Event-Based Microscopic Analysis of the FX Market

Shaimaa Masry Hussein

A thesis submitted for the degree of
Doctor of Philosophy

Centre for Computational Finance and Economic Agents
School of Computer Science and Electronic Engineering
University of Essex

April 2013

Abstract

The foreign exchange (FX) market is the largest and most liquid financial market in the world. Like the centre of a spider web, the foreign exchange market connects to all other financial markets around the world. It is a global network that allows its participants to trade 24 hours 5 days a week from different geographical locations. Given this unique nature of the FX market, millions of daily tick data, referred to as high frequency data (HFD), are generated as a result of market participants' decisions and interactions.

To understand market dynamics, our approach is to explore the microscopic world of the FX market by analysing in depth the millions of daily tick-by-tick prices and the micro-behaviour of FX participants, which in turn formulate a collective market macro-behaviour. This thesis conducts its analysis using an event-based approach. Events are actions taken by traders in the market. We carry out three studies with the aim to get an insight into how these events drive the FX market. With these studies, we aim to make general inferences about market behaviour. The first two studies are empirical research based on analysing a unique high frequency real transaction data set of FX traders, whereas the third study formalises the market micro-dynamics.

To prepare for our empirical studies, we have produced, to the best of our knowledge, the biggest set of HFD ever, which comprises tick transactions carried out by over 45,000 FX traders on an account level for over 2 years. In addition to cleaning the data set from any erroneous observations and validating the quality of the data, we provide strong indicators that the data set is representative of the of the global FX market. This confirms the reliability and validity of this research results. This data set is invaluable to future researchers.

The first empirical study tracks and analyses the FX market seasonal activity from a microscopic perspective, using the tick transactions of the HFD produced. We provide empirical evidence that the unique signature of the FX market seasonality is indeed due to the different time zones market participants operate from. However, once normalised using our custom-designed procedure, we observe a pattern akin to equity markets. Thus, we have revealed an important FX market property that has not been reported before.

The second empirical study conducts a microscopic analysis of FX market activity of the produced data set along price movements. Given the high frequency

and irregular nature of FX tick data, we adopt an intrinsic time scale approach proposed by Olsen Ltd. Intrinsic time is defined by exchange rate turning points of a pre-specified threshold, which are called directional change events. We provide empirical evidence for decaying market liquidity and price ticks changes at the end of the price movement, the overshoot period. We find that a price overshoot stops due to more participants placing counter trend trades. The overshoot period is of special importance as it measures the excess price move of a given threshold and indicates the extent of imbalance in the market for the specified threshold. To our knowledge, this is the first study that deciphers FX market activity during price overshoots. It lays the foundations for understanding how FX market activity changes as the price movement progresses and how small imbalances of market activity in large overshoots can alter the price trajectory.

The third study formalises market dynamics using calculus. In this approach, we define the different market states mathematically and demonstrate the consequences of placing an order into the market. This calculus enables us to analyse market dynamics and properties scientifically. For example, it allows us to study feedback loops, which account for the full effects of cascading margin calls. It also allows us to compute how big a sell order has to be to cause the market to fall by a certain percentage in a simple double auction market model. This work demonstrates how market dynamics and properties can be studied rigorously. It lays a solid foundation for extensive scientific analysis of complex market models.

To my loving Mum and to my loving Dad who left us early...

Acknowledgements

First and foremost, I praise God Almighty for giving me the strength and the capability to pursue this program and successfully complete it.

It is a pleasure to express my sincerest gratitude to everyone who offered me guidance and support throughout my PhD journey.
I am heartily grateful to Prof. Edward for his professional guidance, valuable advice, continuous support and motivation. From him, I have learnt most valuable research and writing skills. I will never forget the amount of support he gave me to overcome all the obstacles I have had throughout my study. Prof. Edward has been simply the best mentor and advisor I could have.
My sincere thanks also go to Prof. Richard Olsen for sharing his immense knowledge and inspirational insights about the foreign exchange world. I am truly thankful for his valuable guidance and his steadfast encouragement to complete this research. He really made me learn how to enjoy research.
I also owe sincere and earnest thankfulness to Dr. Alexandre Dupuis for his very much appreciated valuable assistance and all the bright research ideas he gave me throughout my PhD study. Working with him was really a pleasant experience.
Special thanks go to all my colleagues and friends in Colchester, whose company I really enjoyed, in particular my dear friend and my office mate, Monira Aloud.
Finally, I am truly indebted and grateful to my loving mum, my great dad and my caring brothers for every moment of the love and support, they have provided over the past years. Thank you for allowing me to realize my own potential. You are the greatest gift ever. I would also like to thank my fiancé, whose continuous motivation at the end of my PhD journey was a real incentive for me.
It has been a long journey but also the most enjoyable experience with the support of all these people.

Contents

List of Figures

List of Tables

List of Algorithms

Part I

Introduction

Chapter 1

Introduction

This chapter is an introduction to the thesis. It provides a general overview of the research work done, and discusses its aim and main objectives. The thesis structure is then described in detail. Finally, all the published thesis results are listed.

1.1 Overview

The behaviour and dynamics of financial markets have been analysed by researchers since the 1770s, starting with Adam Smith's notion of the "Invisible Hand", which assumes homogeneous and rational behaviour of investors and has developed into classical economics [166]. However, many of the phenomena observed in market microstructure, as also in the behaviour of high frequency traders and the intricacies of electronic trading platforms are highly complex and are outside the scope of classical economics. Studying the impact of traders' psychology, heterogeneous expectations and technical trading on the market [26, 143, 189, 193]; exploring the market microstructure [42, 71, 95, 133, 170, 199]; applying agent based modelling and artificial intelligence techniques [102, 118, 119, 182], are all attempts to gain insights into traders' behaviour and market dynamics.

The foreign exchange (FX) market, focus of this thesis, is the largest financial market worldwide, with a daily average turnover of 4 trillion USD and an average daily transaction volume for spot currencies of 1.4 trillion USD or equivalent of 10 percent of the GDP of the US [20]. The FX market determines the exchange rates of global trade and is an important factor in determining the relative wealth of countries. The importance of FX markets has developed with increased global trade, capital flows and investment, and with real-time high frequency trading brought about by advances in technology and the spread of electronic trading platforms [136]. The participants in the FX market are diverse and include financial institutions, hedge funds, corporations and retail traders. Like the centre of a spider web, the foreign exchange market connects to all other financial markets around the world [159]. Therefore, it has become essential to study the dynamics and understand the behaviour of the FX market, which in turn affect the global economy.

The financial crisis that started in 2007 has stimulated new research to understand the market dynamics, reduce unforeseen losses, increase market efficiency and develop an early warning system to prevent market crashes. The market dynamics cannot be explained by macroeconomic data only. We need to dig deeper and analyse the buy and sell flows in financial markets. We have to analyse the millions of daily tick-by-tick market prices (called high frequency data, HFD) and the transactions of market participants and their interactions [41, 165]. The market consists of a network of people who are the real decision makers and who interact and decide in human ways, pursuing human goals and satisfying their human needs. *"If you understand what everyone else is doing, and why everyone else is doing it, it makes it very easy to understand what is going to happen" [159].*

Thus, drawing analogies from physics, where findings are being generalized from microscopic observations [195], we believe that exploring the microscopic world of financial markets in general and the FX market in particular would provide an insight into the collective macro-behaviour of financial traders and explain many of the unjustified market phenomena.

However, there are three main challenges associated with adopting this research approach on the FX market:

First, we must have access to a rich data set of real tick transactions of FX traders, whose micro-behaviour can be tracked and analysed to make general inferences about FX market dynamics. Access to FX transaction data has only been recently enabled. In addition, it has been always restricted to a few currencies and to an aggregate level of transactions that does not differentiate between the different participants placing the trades. Accordingly, activity of individual traders cannot be tracked and the impact of the different participants on the market cannot be analysed.

Second, the foreign exchange market generates millions of daily tick data, referred to as high frequency data (HFD), as a result of market participants' decisions and interactions. To understand the FX market behaviour, we need to analyse and dig into this HFD. However, there are several challenges associated with its analysis. It has no standard structure, as it depends on the institution's policy with regard to the production and storage of data[39]. Moreover, HFD may contain a great deal of erroneous or misleading observations that may result from computer system errors or internal system procedures[49, 204]. Therefore, prior to data exploration, an important laborious step is to clean the data from any outliers and produce a reliable data set that can be used to analyse and make general inferences about traders' behaviour.

The third challenge is the analysis of the high frequency data. There is no established framework or methodology to follow from previous research, especially given that both the type of data and the research sought in this work are unique.

1.2 Research Aim and Objectives

The aim of this research is to conduct a microscopic analysis of FX market traders' behaviour, using event-based approaches, in order to lay a solid foundation for accurate modelling of FX traders and their behaviour. This comprises the following objectives:

1. Study the complex structure of the FX market and its unique nature of trading that would impact FX traders' behaviour.

2. Survey the various approaches adopted in the literature for studying the FX market behaviour and explore their limitations.

3. Explore and understand the high frequency data set supplied by the OANDA Corp [156], an online FX trading platform, for our research work. The data set covers over 140 million tick transactions on an account level, carried out by more than 45,000 accounts on the OANDA platform for over two years. A thorough understanding of this large unique high frequency data set is a vital step to confirm the data set representativeness of the FX market and its potential usefulness for understanding FX traders' behaviour.

4. Prepare the data set for analysis by cleaning the data set from any erroneous observations and validating the quality of data. This involves tracking the transactions

of each individual trader from positions opening to closure in each currency pair to ensure the consistency and reliability of trades. A proper data set preparation process is crucial to ensure the validity of our research results.

5. Given the complex nature of FX trading and its irregular spaced high frequency data, we need to develop new approaches for studying traders' micro-behaviour and for making general inferences about market dynamics, with an ultimate goal of laying a solid foundation for modelling market dynamics. Following are three critical issues in the FX market, which are studied in this thesis from a microscopic perspective:

    a) One of the most common conventions in the literature, is that the FX market's unique and complex nature, as compared to equity markets, is due to the existence of different overlapping world trading sessions, from which traders can operate 24 hours 5 days a week. Therefore, the different geographical locations of FX traders, is one of the parameters of primary importance in analysing and modelling their behaviour. Consequently, the objective is to develop an approach for normalising traders' time zones and to examine if new market dynamics are revealed after normalisation.

    b) The presence of a large variety of traders with different sleeping hours and trading strategies causes a discontinuous flow of physical time. This means that each clock tick of the uniformly progressing physical time does not necessarily correspond to a transaction in the FX market. Traders' actions are usually associated with exchange rate price movements, which in turn are determined through a complex market making process. Thus, one of the controversial issues in the FX market is traders' behaviour with exchange rate price movement. Both traders' transactions and exchange rate series are discontinuous and irregularly spaced. The objective is to develop an approach for synchronizing market activity along price movements and derive rhythmic behavioural patterns, from which we can possibly understand new market dynamics.

    c) The FX market attracts many traders, because the liquidity of the market enables market makers to offer their customers a high degree of leverage.

Traders buy a currency by selling another currency using a loan facility provided by the market maker on the basis of a collateral, referred to as margin, that the trader has deposited. The degree of leverage might be as high as 50 times the underlying capital of the trader. Foreign exchange trading seems to offer fast riches, but is in fact highly risky. Traders can easily get trapped in losing positions. If their equity drops below a critical threshold, margin calls are triggered forcing traders to close out their positions, possibly at the worst possible moment, thus incurring large losses. Since many traders may be trapped in the same situation, market price can be affected, which in turn can lead to further positions close-outs and losses. Hence, the objective is to develop an approach to formalise the consequences of margin trading and study the impact of such feedback loops, through which we can provide early market warnings.

1.3 Thesis Structure

The thesis structure is based on the aims and objectives discussed in the previous section. It begins with a comprehensive background and literature review in chapter 2, describing the FX market structure, trading nature and its main stylized facts. The greater part of the chapter is then devoted to the main research approaches adopted in previous studies to understand the FX market behaviour and some of its dynamics. These are then contrasted with the proposed microscopic analysis approach of the thesis.

Chapter 3 prepares the high frequency data set provided by [156] for the microscopic analysis of traders' behaviour. A cleaning procedure is carried out to remove all erroneous or misleading observations from the set that would affect the validity of any forthcoming results. The quality of data is then validated, ensuring a clean and consistent transaction data set. Moreover, evidence is provided for the representativeness of the data set of the global FX market. This makes the produced data set an invaluable source for understanding FX traders' behaviour and any future research. The produced data set is used in chapters 4 and 5.

Chapter 4 investigates the nature of the FX seasonal activity against the one in equity markets and questions whether the different signature in the two markets is due to the overlapping FX world trading sessions. In an attempt to answer this question, the chapter introduces a new event-based algorithm to normalise the time zones of the different traders according to their daily sleeping hours. The results of implementing the normalisation procedure is then reported, and a detailed description is given of how the normalized seasonal activities of all traders can be synchronized and aggregated along one time window. The resulting new signature of the normalized FX seasonality is then depicted and discussed. Finally, the chapter describes how the findings of this empirical study can be used to make general inferences about the position management of FX traders during the trading day.

Similar to chapter 4, chapter 5 tracks and analyses traders' activity but from a different event-based microscopic perspective. It studies the FX market activity patterns with price movements. Given the high frequency and discontinuous nature of the FX data, the chapter highlights the problems associated with using the traditional physical time scale for market activity synchronization. The chapter then proposes an alternative directional change event-based time scale, the so called intrinsic time, which is defined by price turning points of a pre-specified threshold. An algorithm for defining the intrinsic time scale is then outlined and applied on several price series of different currencies supplied by OANDA Corp. The chapter then expands into a further in-depth analysis by discretising each price movement into smaller sub-events to synchronize and decipher market activity on a microscopic level. Empirical results are then reported and evidence is provided for new market dynamics that are valid across all magnitudes of price events.

A third microscopic angle, from which traders' behaviour can be also analysed, is presented in chapter 6. It introduces the market-calculus as a new event-based microscopic approach to formalise and analyse market dynamics and feedback loops, like cascading margin calls. Due to the lack of knowledge of the FX market maker model in the literature, the chapter illustrates the high potentiality of the market-calculus approach in the context of a simplified version of the popular double-auction market. The chapter starts by explaining how the proposed calculus approach can formalise market mechanisms and enable rigorous reasoning. Then it formalises two models for market clearance in a double auction market. Given this definite formalism, the chapter proposes methods for market risk assessment and illustrates how the impact of an order of any size can be computed; or how big a sell order has to be to cause the market to fall by a certain percentage. Finally, using the same logic, a demo simulator is implemented, to demonstrate market clearance dynamics, margin trading feedback loops and market risk assessment.

The thesis is concluded with chapter 7. It summarizes the work done in the thesis, lists its main contributions, and discusses possible further work.

1.4 Publications

Some of the original work done in this thesis has been published in the following peer-reviewed papers:

1. Masry, S., Dupuis, A., Olsen, R., Tsang, E., 2013. Time zone normalisation of FX seasonality. Journal of Quantitative Finance, in press .

2. Tsang, E.P.K., Olsen, R., Masry, S., 2013. A formalization of double auction market dynamics. Journal Quantitative Finance, in press .

3. Masry, S., Dupuis, A., Olsen, R. & Tsang, E., April 2013. Deciphering FX market activity in price overshoots. Journal of Empirical Finance, under review.

4. Masry, S., Tsang, E.P.K., 2011. Simulating market clearance dynamics under a simple event calculus market model. In: Proceedings of the 3rd Computer Science and Electronic Engineering Conference, Colchester, ISBN: 978-1-4577-1300-2 ,United Kingdom, IEEE.

5. Masry, S., ALOud, M., Dupuis A., Olsen, R., Tsang, E. K., 2010. High frequency FX market transaction data handling. In: 4th CSDA International Conference on Computational and Financial Econometrics, London, United Kingdom.

6. Masry, S., Tsang, E.P.K., December 2010. Monitoring the financial market using event calculus. Workshop in Accounting, Finance and Management, BCS SGAI International Conference on Artificial Intelligence, Cambridge, United Kingdom.

7. Masry, S., Aloud, M., Tsang, E., Dupuis, A., Olsen, R., 2010. A Novel approach for studying the high frequency FX market. In: Proceedings of the 2nd Computer Science and Electronic Engineering Conference, Colchester, ISBN: 978-1-4244-9029-5, United Kingdom, IEEE.

Part II

Background and Literature Review

Chapter 2

Literature Review

In this chapter we describe the nature of the foreign exchange (FX) market from different aspects and provide a synthesis of the different approaches used in studying the FX market behaviour.

2.1 Introduction

The fast growth of FX and the unique characteristics of its trading environment have drawn the attention of many researchers and experts. FX is the largest financial market in the world, opening 24 hours a day, 5 days a week. The Bank for International Settlements (BIS) Triennial Survey in 2010 [20] reports an average daily turnover of 4 trillion USD in the FX market, which is 30 times larger than the turnover of both NASDAQ and NYSE [156]. In addition to financial institutions, companies and hedge funds engaged in FX trading, the evolution of internet trading has also attracted individual currency traders. In general, the ability of trading at any time and responding to currency fluctuations on the spot, having high level of liquidity, narrow spreads and low margin requirements, has attracted various types of traders to the FX market. Hence, studying FX market dynamics and behaviour is a fertile research area and provides a broader perspective for behavioural analysis.

The fact that classical economics has not captured many of the exchange rate features has opened the gate for various research approaches to study the FX behaviour. Surveying the FX literature reveals two main streams of research studies of the FX market, the first one exploring the market microstructure and its stylized facts while the second one studies the market behaviour leading to these stylized facts.

Digging further into the second research area, we find three main perspectives from which market behaviour has been analysed. We discuss each one in the following sections.

The first perspective is the behavioural finance approach, which is mainly concerned with how traders' psychology contributes to their over- or under-reaction to new information and price movement. The literature shows a long strand of researchers providing empirical evidence for several psychological biases in the FX market. Traders' heterogeneous expectations and beliefs [81, 82, 106, 132, 142, 157], overconfidence [22, 86, 160], loss aversion [155], disposition effect [188] and feedback trading [3, 32, 117] are all examples of biases that have been reported and empirically proven in the FX market.

A different perspective for explaining the market dynamics has been adopted by the FX market microstructure researchers, pioneered by Evans and Lyons. Using empirical models, they analyse the effect of cumulated order flow [29, 63, 66], news arrival [7, 10, 44, 62, 67], fundamentals [63, 67], feedback trading [59, 63, 83] and institutional interventions [149, 173] on exchange rate movement and the general market state .

The third and most recent approach for explaining the FX market is agent-based modelling (ABM), which models the features of a real financial market. ABM represents the financial market as a group of interacting heterogeneous agents, the basic units of the model, which can adapt to their environment, learn from the information they obtain, and interact dynamically with each other and with their environment [118, 182]. The aim is to make inferences about market behaviour and the causalities of the emergence of the different market anomalies.

The rest of this chapter is organized as follows: Section 2.2 describes the FX market structure. Section 2.3 discusses the most important FX market stylized facts discovered so far. Section 2.4 introduces the different research approaches adopted to understand the emergence of these stylized facts and the market behaviour. Conclusions and discussion of the main limitations of the various approaches used in the literature are given in section 2.5.

2.2 Foreign Exchange Spot Market Structure

The FX market consists of two dominating parts, the FX spot and the continuously growing FX derivatives (forwards, swaps, options and futures). The operations of the

FX spot market account for 90% of the FX market volume. According to the BIS Triennial Survey [20], the average daily transaction volume for spot currencies reached 1.4 trillion USD in 2010. The high frequency data produced by this market serve as a role model for the exploration and analysis of high-frequency data, leading to inferences about market microstructure and traders' intraday behaviour [49]. This research tackles only the FX spot market. The FX spot market structure can be generally characterized by the following main aspects:

First, FX markets by their nature are decentralized and distributed in several financial centres and thus not controlled by any central authority.

Second, the participants in FX can be categorized into FX dealers, brokers and customers. FX dealers, representing mainly major commercial banks, trade among each other and with external customers, who in turn, may be large corporations, financial institutions or retail investors.

Third, the transactions between the different FX participants can take place through different channels: i) Inter-dealer trades, which accounts for 43% of the FX market turnover [19] , can be executed either directly (over-the counter) or indirectly via voice brokers (for low volume currencies), or the now prevailing electronic brokerages like (EBS) and Reuters Dealing 3000; ii) Dealer- institutional customer trades can be either directly via dealer banks, or indirectly via electronic brokerages or now the smaller internet trading platforms. These can be either platforms developed by banks or consortiums of banks, also called electronic communication networks, like FxConnect and Hotspot FXi. Financial institutions and large corporations can also now trade using non-banking trading platforms such as OANDA FXTrade. While electronic brokerages can only be used by inter-bank foreign exchange traders, internet trading platforms can now be used by all types of customers; iii) the third channel is individual or retail trades which are settled through internet trading platforms [120, 170].

Finally, it is important to differentiate between the two main trading mechanisms in the FX market; the direct (over-the-counter) market and the indirect (brokered) market via voice brokers or electronic brokers. The direct market is quote driven, as the dealer gives quotes on request, and the initiator decides on the time, quantity and direction of the trade. Transactions are completed through private bilateral deals among traders and are not disclosed to other market participants. However, direct trading is non-anonymous, as the client's identity is revealed to the dealer. In contrast to the direct market, the indirect market, which is nowadays dominated by electronic brokerage systems like EBS and Reuters Dealing 3000, is order-driven. Transactions are mediated by brokers and prices and quantities are set altogether. Although traders' identity is kept anonymous during trading, prices, quantities and order types of the different transactions are communicated to the rest of the market [32, 198].

2.3 FX Stylized Facts

Since the end of the Second World War, a long strand of researchers have enriched the FX market exchange rate theory, starting from the simple model of exchange rate determination (in the 1950s), through the intersection of the demand and supply curves for FX, to the portfolio balance model and the Dornbush's stick price monetary model in the late 1970s and early 1980s. All these models were mainly based on the efficient market and rational expectations concept. The most important notion of the rational expectation model is that exchange rates follow fundamentals (inflation, output growth, interest rates, GDP, etc.) and that changes occur only due to unexpected news or shifts in fundamental variables. Despite the popularity of all these theories, their empirical failure has been proven by the work of Meese and Rogoff in 1983 [139]. They showed that there is no stable relationship between exchange rate and fundamentals in the short run, referred to as the exchange rate disconnect puzzle by [161]. This has led to a large empirical literature exploring several empirical puzzles and stylized facts of the FX spot market, which could not be explained by former models [144]. The list of exchange rate stylized facts and puzzles is long, but we cite here the most popular ones discovered so far in the literature for different data frequencies, ranging from ultra-high frequency(tick-by-tick) to low frequency (years).

One of the most examined stylized facts is exchange rate volatility. A long standing research shows that exchange rates generally exercise both excessive volatility, exceeding the volatility of most of the fundamentals and economic variables thought to affect them, and volatility clustering, where exchange rates exhibit periods of high and low volatility [73, 76, 79, 128].

There is also evidence of seasonal heteroskedasticity in the form of intra-daily and intra-weekly, but not for lower frequency, FX price volatility clusters [16, 17, 105]. According to [49], this puzzle can be partially explained by fat tailed return distributions, which is another stylized fact of exchange rates. Despite the various views regarding the distributions of FX returns, all researchers agree that daily returns are fat tailed for short-term horizons and deviate from a Gaussian random walk model [49].

A further statistical property is the positive relationship of exchange rate volatility with trading volume across different currency markets, as proved by [35, 90] and more recently by [33]. They all confirm the strong positive relationship between volume and trading for long horizons. For high frequency (minute data), [29] show that volume and volatility rise sharply at certain times of the day and exhibit a double U-shape pattern, where peaks occur during the overlapping sessions of the FX market.

The behaviour of spreads with volume and volatility is another interesting stylized fact. According to [170], bid-ask spreads almost perfectly mirror the pattern of volume and volatility. FX spreads also exhibit a clear weekend pattern, as foreign exchange trading largely ceases from about 21 GMT on Fridays until 21 GMT on Sundays. FX spreads on Saturdays and Sundays can have double and more the size of those on weekdays. Other studies have also shown that this week end effect can last till Monday morning, where spreads can be exceptionally wide in some markets like Tokyo (see [49, 170]).

Returns autocorrelation and correlations with other currencies have been also spotted in the FX market. Several studies on Reuters FX indicative quotes find strong signs of a first order moving average negative auto-correlation at the highest frequency, disappearing as the information process is over [17, 88–90]. Examining USD/DEM quotes for the period (1987-1993), [49] show that there is significant autocorrelation up to a 4 minutes lag, whereas insignificant for longer lags. The autocorrelations mainly lie within the 95% confidence interval of an i.i.d. Gaussian distribution. Daily returns of different currencies are also correlated as a result of the news effect. The correlation between daily euro-dollar and sterling-dollar returns, for example, is 70 percent, while correlations between these European exchange rates and dollar-yen are smaller: both are 46% [170].

Profitable speculations based on exchange rate movements are another property of exchange rates. Research has long shown that profits can be earned in the FX market by speculating relatively on a short term basis via the use of various types of technical trading rules (see [143, 144]). Moreover, several studies show that foreign exchange dealers rely heavily on speculative profits, and that their main profit is from exchange rate movement rather than from spread. Their speculative positions are based on information gathered from customers, from professional colleagues at other banks, and from real-time news services [170].

[170] also shows that some dealers also engage in arbitrage across markets, such as triangular arbitrage or covered interest arbitrage. Arbitrage opportunities, though typically short-lived, arise frequently and occasionally provide sizeable profits [6]. A more recent study by [75] reveals that triangular arbitrage opportunities exist in the spot foreign exchange market with short durations and small magnitudes, but generally have decreased in recent years, implying an increasing price efficiency.

Finally, scaling laws have been found for mean absolute returns and mean squared returns over varying time intervals, from few minutes to few months. It states that absolute or squared returns are proportional to a power of the interval size [49]. A comprehensive list of more recent scaling laws for exchange rates can be found in [87], where they have been able to unveil 12 new scaling laws holding across 13 exchange rates for up to three orders of scaling magnitude.

2.4 Research Approaches for Explaining the FX Market Behaviour

As mentioned earlier, the second stream of research in the FX market aims to understand the FX market behaviour and the emergence of its stylized facts. We introduce here the three main approaches adopted to understand market and traders' behaviour. These approaches are not restricted to FX markets only; on the contrary, they were first applied on stock markets and other financial markets, and, as a natural evolution for research, they have been also adopted in studying the FX market. Although referring sometimes to studies made on other financial markets, we mainly focus here on the FX research. In this section, we will briefly cite some of the scholarly works done in each approach.

2.4.1 Behavioural Finance Approach

The study of traders' psychology has been advocated for by the behavioural finance group many years ago, seeking for psychology-based theories to explain stock market anomalies. According to [186], behavioural finance attempts to show how psychological factors can contribute to traders' under- or over-reacting to new information and the consequent price movement. Within the behavioural finance paradigm it is assumed that the information structure and the characteristics of market participants systematically influence the investment decision as well as market outcomes. The roots of behavioural finance go back to the 18th century, where Adam Smith was the first proponent of a link between economics and psychology describing some of the psychological principles of individual behaviour in his book 'The Theory of Moral Sentiments'. While, the 1850's, witnessed strong opposition from the proponents of natural sciences and the developing theories of rationality and homo economicus, John Maynard Keynes and Irving Fisher, in the 20th century, based much of their theories on psychological factors influencing economics and market behaviour. The rationalists then took over in the 1960's and 1970s with the hypothesis of efficient markets and rational agents. However, the failure of the theory of efficient markets and the supposedly rational behaviour of investors to explain many of the stock market anomalies has directed many researchers to behavioural finance again in the 1980s. In classical finance it is argued that behavioural finance has no unified theory similar to utility maximization and rational beliefs. Even though this critique is valid, behavioural finance is based on actual observations of how people behave, using extensive experimental evidence. This is different from rational theories that make assumptions of how people behave, introducing a priori constraints in our understanding of the phenomena of financial markets [189, 193].

Studies in the field of traders' psychology are enormous, covering academic work, experiments, professional and institutional articles. For detailed reviews, see [26, 189, 193].

In the following sub-sections, we cite only the main scholarly works that provide empirical evidence and experiments of the most popular traders' biases and how their psychology can affect their trading behaviour.

2.4.1.1 Traders' Decision Making Process

One of the major researches carried out to understand traders' psychology is studying their decision making process and linking it with their psychological traits. For instance, [122] describe an experiment undertaken with a sample of ten professional stock traders where they attempted to verify the link between decision making and emotions by measuring the real time psycho physiological characteristics. They suggest that experts judgements are often based on intuition not explicit analytical processing. [99] also introduced a new understanding of the individual's decision making process and refers to 'seven well-known biases in human decision making', titling them as 'The Seven Deadly Sins in Financial Decision Making', which apply to both open outcry and electronic traders. These sins are: confirmation bias, optimism bias and illusion of control, overconfidence in predictions, mistaken beliefs, risk aversion, mental accounting biases; the construction of values and preferences; and finally overreaction.

Steenberger, in his book 'The Psychology of Trading' [192], explains the challenges of being a professional trader and the ways in which traders can benefit from understanding the information contained within their emotional, cognitive and behavioural patterns. He had the opportunity to analyse the psychology of a group of traders and was successful in getting worthwhile insights into their minds. He finds certain behavioural patterns that can be observed with individuals who are trading for living.

As for FX traders, one of the most interesting and comprehensive studies was provided by [159] in his book, 'The Psychology of the Foreign Exchange Market'. He describes and explains many of the decisions made affecting FX market dynamics. He engages in his analysis a variety of psychological perspectives ranging from social market dynamics to traders' personality characteristics and cognitive biases and intuitions.

2.4.1.2 Expectations, Heterogeneity of Beliefs and Technical Trading

Two important strands of literature, based on traders' psychology, that have emerged to explain the high levels of trading volume in financial markets are the " heterogeneity of beliefs" literature and the "overconfidence" literature [86]. Heterogeneity of expectations can arise from the way participants respond to news signals, the way they react to past exchange rate movements, the way they forecast future exchange rate values and the way they perceive risk.

The heterogeneity of traders' beliefs and their various expectations and behaviour to news events, creating a complex dynamic trading environment, has been extensively surveyed and evidenced by several authors including [9, 55, 61, 80, 106, 158, 159, 192, 194].

The heterogeneity of traders' beliefs can be also shown in the way they judge exchange rate movements and upon which they develop their trading strategies. It has been shown that the majority of traders do not only base their strategies on fundamental information but also on technical trading, especially over short and medium horizons, where exchange rates are more manipulated by speculative behaviour, technical trading and bandwagon effects rather than macroeconomic fundamentals. For a good review of technical analysis in FX see [143].

One of the earliest studies conducted in the eighties by [79] tried to estimate the proportion of chartists and fundamentalists in the aggregate exchange rate expectation of the US Dollar. The reliance of FX traders on technical analysis and their heterogeneity in exchange rate forecasts has been further studied by a long strand of researchers such as [4, 46, 81, 82, 98, 121, 126, 132, 140, 141, 148, 150, 157, 158, 176, 194, 197]. A more recent empirical study of technical trading in the FX market has been carried out by [142], showing that foreign exchange traders rely heavily on technical analysis. It is of high importance among foreign exchange professionals such as dealers and fund managers, a finding that has held from the 1970s to the present day. They explore a huge data set, covering 15 years of exchange rate expectations. They find that influences arising from chartists' and fundamentalists' behaviour are most useful in explaining heterogeneity.

2.4.1.3 Overconfidence

As mentioned earlier, investors' overconfidence is the second strand of literature beside "heterogeneity of beliefs" that is assumed to cause the high trading volumes in financial markets [86].

According to [55], overconfidence is *"the key behavioural factor needed to understand the trading puzzle"*. Overconfidence is sometimes referred to as "miscalibration" because perceived probabilities are poorly calibrated to true probabilities [160]. As described in [26], overconfidence may in part stem from two other biases, self-attribution bias and hindsight bias. Self-attribution refers to a person's belief that any success is due to his talents and any failure is just bad luck, leading the person to be overconfident after having some successive successes. Hindsight bias refers to "*the tendency of people to believe, after an event has occurred, that they predicted it before it happened* ", giving them the illusion of being able to accurately predict the future.

Empirical evidence of overconfidence, using different proxies, have been shown by several authors such as [21, 86, 110, 191], and more recently for currency markets [155, 160]. The latter show that currency dealers exhibit two types of overconfidence. They tend to overestimate both the precision of their information and their personal success. They also show that there is no difference between junior and senior currency dealers, and that overconfident dealers are not driven out of the market as would have been expected due to their irrationality.

In general the overconfidence bias has been used to explain several market anomalies like short-run momentum and long run reversals and the resulting profitability of trend-following technical strategies in currency markets[160]; excessive trading volumes, high volatility in exchange rates and the engagement of market participants in excessive levels of risk [21, 54, 77, 86] ; and also intensive diversity among price forecasts and heterogeneous beliefs [26, 30, 104].

2.4.1.4 Feedback Trading

Another important behavioural aspect of traders that is related to the previous two biases is feedback trading. Feedback traders are traders whose demand is based on the history of past returns rather than expectations of future fundamentals. Based on traders' beliefs, they can be negative or positive feedback traders. In the case of foreign currency, positive feedback trading exists when traders buy (sell) after the exchange rate has depreciated (appreciated), and negative feedback trading exists when traders buy (sell) after the exchange rate has appreciated (depreciated).

A simple factor motivating positive feedback trading can be extrapolative expectations, where investors' expectations of future returns are based on past returns [26]. Other possible reasons for positive feedback trading are portfolio insurance strategies, margin call-induced selling after periods of low returns, extensive use of stop-loss orders, and also central bank interventions in the foreign exchange markets. Negative feedback trading, on the other hand, is mainly a result of trading strategies seeking profits with the currency appreciation (depreciation). Thus, while positive feedback stabilizes the currency exchange rate, negative feedback exercise further pressure on the exchange rate to either appreciate or depreciate and hence deviate the currency from its long-term value [3].

The literature has shown for a long time the role of feedback trading in stock markets and other types of assets. For instance, empirical evidence of feedback trading in the US stock market has been found by [184] and in several developed European markets by [34, 113] and Asian stock markets [113].

The first study on the role of feedback trading in exchange-rate markets was carried out by [3]. Contrasting [3] results against the stock market literature, they find that, similar to stock markets, feedback trading in foreign currency markets is more present during high volatility causing market instability. However, there is only feedback trading in less developed currency markets, whereas feedback trading exists in both developed and emerging stock markets. Besides, negative feedback trading is more prevalent in foreign currency markets. More recent empirical evidence of both positive and negative feedback trading in FX has been found by several other authors. In contrast to [3], [117] show a presence of negative/positive feedback trading in several industrial and emerging economies' exchange rates with respect to the US dollar but an absence of such behaviour for the Euro, indicating the credibility of the Euro currency in the eyes of foreign currency traders. Other empirical studies of feedback trading in FX have been carried out by [32, 137] who show evidence for negative feedback trading in semi-daily commercial-customer order flow but not in corresponding financial-customer order flow. [50] also find evidence of feedback trading in transaction-level inter-dealer trading data. More recently, [180] find evidence for positive feedback trading in inter-dealer order flow using Granger-causality tests applied to the Evans and Lyons [70] daily data.

2.4.1.5 Disposition Effect

One of the most widely reported biases in traders' behaviour is the disposition effect, namely "*sell winners too early and ride losers too long*", a phenomenon documented by [188]. This behavioural bias is mainly based on the prospect theory of Kahneman and Tversky (1979), where gains and losses are often judged according to a reference point, a function of the asset purchase price. Investors then become risk averse in the case of profits and risk seeking in the case of losses. As a result, they hold onto the losing stock for a chance to break even [188].

Empirical evidence for the disposition effect in stock (and other assets markets) has been provided by many researchers for both retail [48, 56, 123, 162, 175, 201] and institutional investors [38, 92]. In contrast to [162, 175, 201], who demonstrated the disposition effect by aggregating across individual investors, [56] analyse the disposition effect at an individual level, trying to link the underlying investors in their used data set with their characteristics. They show that wealthier investors, professional investors, and investors trading more frequently exhibit a smaller disposition effect. Many of them actually show "reverse disposition". [175] shows that the disposition behaviour also varies with the firm size, i.e., it is concentrated in stocks among the top 60 percent of the market capitalization distribution.

For the FX market, empirical evidence for the disposition effect seems to be not in complete conformity with the existing literature for stock markets. [155] use a proprietary currency traders' database (State Street Corporation) of large institutional investors. Following the same methodology of [162] and the methodology of [92] in measuring the disposition effect, they find that institutional currency traders are not susceptible to the disposition effect. On the contrary, they aggressively reduce risk in the increase of losses, and mildly increase risk in the wake of gains. Moreover, any increases in risk-taking are short-lived, reversing themselves within a calendar quarter. These basic results are pervasive across the major currencies. [155] argue, in contrast to [92], that institutional currency traders are sophisticated enough to fall into a common pitfall like the disposition effect. As for FX retail traders, a study has been carried out by [152, 153] examining the disposition effect among retail foreign exchange traders, by applying a panel duration approach on a data set acquired from OANDA FXTrade. The authors find no general disposition effect for the different roundtrips as it varies with the size of profits or losses obtained over a round trip. While the usual disposition is reported for large profits and losses, an inverse disposition effect can be found for small profits and losses. The study also shows that overconfident traders (captured by high trading volumes and past trading success) exhibit a more severe form of disposition effect.

2.4.1.6 Loss aversion

Another important behavioural aspect introduced under the prospect theory framework by Tversky and Khaneman (1979) is loss aversion. It indicates that people are more sensitive to losses than to gains. An investor can be loss averse over changes in his total wealth, or over changes in the value of his portfolio or individual stocks[189].

Many experiments show that individuals usually exhibit narrow framing, i.e. consider narrowly defined gains and losses [23]. Theoretical studies and models of investors' loss aversion and narrow framing to understand stock market behaviour have been carried out by researchers like [23, 24, 28].

The first study on the FX market, showing empirical evidence of the model introduced by [24], was carried out by [155] on a large data set of currency institutional traders. They show that investors are loss averse; they are more sensitive to losses in financial wealth than to gains, and they become more loss averse if their previous investment results have been poor. According to [155], institutional fund managers are more likely to be risk averse and to derive utility from past performance than retail or professional investors, as they manage other people's accounts and any bad performance would affect their reputation.

Having all this empirical evidence literature of traders' behavioural biases in the FX market, many models, both theoretical and computational, have been developed to replicate the features of exchange rate series based on agents' heterogeneous expectations and many of the above-discussed behavioural aspects. This is discussed further in the following sections.

2.4.2 FX Market Microstructure Approach

The last two decades have witnessed tremendous research in an important branch of economics and finance, the market microstructure. O'Hara states that it is " *the study of the process and outcomes of exchanging assets under a specific set of rules*" [163]. The access to rich data from a variety of sources in recent years has rapidly increased knowledge of microstructure and has allowed the development of many complex empirical models. The FX market microstructure research is pioneered by Evans and Lyons. For excellent reviews of the most important empirical models of market microstructure see [63, 95, 133, 199]. In the following subsections, we briefly cite the main scholarly work in market microstructure that aims to explain the market behaviour.

2.4.2.1 Order Flow Empirical Evidence and Modelling

The complexity of exchange price behaviour and the strong need to understand the market dynamics has led to a long strand of empirical research proposing different approaches to exchange rate determination, trying to link it with order flow, news arrival, fundamentals and feedback trading. The empirical microstructure literature has been substantially examining the role of foreign exchange order flow. Order flow is defined as the cumulative flow of signed transactions, where each transaction is signed positively or negatively depending on whether the initiator of the transaction (the non-quoting counterparty) is buying or selling, respectively [131]. It is important to differentiate order flow from trading volume, as the order flow sign conveys information. Positive (negative) order flow indicates to a dealer that, on balance, their customers value foreign currency more (less) than his asking (bid) price. By tracking who initiates each trade, order flow provides a measure of the information exchanged between counterparties in a series of financial transactions [65, 69].

Since the FX market has inter-dealer trading and dealer-customer trading, we have two types of order flow, namely customer order flow and inter-dealer order flow. Empirical evidence for the explanatory power of the inter-dealer order flow on exchange rate returns has been provided by several researchers on different currencies. For instance, [62] use a 4 month period of inter-dealer transactions for DM/Dollar and Yen/Dollar currency pairs; [63] use a four month period of interbank order flow for seven currencies against the dollar, showing an increase in the predictive power of order flow for exchange rate returns with an increase in the currencies; [172] use a large number of transactions from the indirect spot market; [29] use a long data set of 6 years of inter-dealer and exchange rate data for Euro/Dollar and Dollar/Yen currency pairs from EBS; and [178] use one year of high frequency data of three major exchange rates (USD/EUR, USD/GBP, USD/YEN) collected from Reuters (D2000-2) inter-dealer trading platform.

The relation between customer order flow and returns has been also examined in some studies, showing a difference in the results among the different customer groups, indicating that financial customer order flow is positively related to returns, while commercial-customer order flow is negatively related to returns. According to [65], customer orders should also have explanatory power for depreciation rates, as customer orders are originally motivated by information, which are in turn aggregated and carried in the inter-dealer order flow. Similarly, using a unique data set covering almost all transactions in the SEK/EUR market on a day-to-day basis, [32] find that a positive correlation between net purchases of currency of financial customers and the exchange rate is matched by a negative correlation between the net purchases of non-financial customers and changes in the exchange rate. Thus, when financial investors buy SEK, the SEK tends to appreciate. The correlation becomes stronger for lower frequencies. [72, 83] show similar results in the sense that they find a positive correlation between financial customers net purchases of foreign currency and changes in the foreign exchange rate. [72] find that extreme exchange-rate movements at high-frequency are generally associated with large net flows from financial institutions; in contrast, low frequency trends are associated with secular net flows from non-financial corporations. According to [170], this pattern suggests that financial customers are typically net consumers of overnight liquidity while commercial customers are typically net suppliers.

When it comes to modelling, one of the most influential order flow models, always discussed in literature, is the Evans and Lyons model [62, 63], which is based on Kyle's (1985) sequential equilibrium model. In contrast to Kyle's model, Evans-Lyons model does not assign private information to a single participant, but it is scattered among traders and then aggregated through the trading process. The main idea of this microeconomic approach for exchange rate modelling is that the equilibrium is a function of the dealers quoted prices at certain point in time, and information will only impact exchange rates when it affects dealers' quotes. Dealers revise their quotes either in response to direct information captured from public information and macroeconomic announcements or in response to indirect dispersed information reflected in orders received from customers and other dealers. According to Evans and Lyons [62, 63, 70], the trading day includes three rounds, where dealers trade with customers in the first round, then trade among each other in the second round and finally trade again with customers in the third round to absorb dealers' imbalances, leaving them with no open positions overnight.

Discussions of the Evans-Lyons model and its contrast with the Kyle's model can be found in [63, 65, 70, 181]. But to sum up the key features of the Evans and Lyons model [63] is that order flow has contemporaneous impact on prices at very high frequencies, but the inverse causality (through feedback trading) is not true. This is in contrast to [50], who show that when aggregating order flow over time, a simultaneous impact can be observed, a phenomenon called *"contemporaneous feedback trading"* by [50]. This is however ruled out in order flow models. Hence, they propose a VAR model that allows contemporaneous feedback trading, showing that the price impact of trades is much stronger when such a feedback trading strategy is incorporated.

According to [181], despite the usefulness of the Evans-Lyons model, it should not be referred to as the model accurately describing the FX market, as the underlying assumption of the Evans-Lyons model that customers absorb daily inventory imbalance does not totally reflect real market practice. For instance, hedge fund traders "*will often trade against their own positions in small size in order to encourage the inter-dealer community to adopt similar positions, subsequently reversing this tactic with substantially greater volumes to effect a short squeeze on the market*" [181]. Another example is the currency overlay managers, who split their large orders into smaller ones distributed amongst different dealers, sharing risk among the different dealers and decreasing their trading profit. Accordingly, [181] propose a primary push-pull trading model. They classify customers into push or pull customers. Large active traders or traders who are assumed to be informed are expected to initiate price movements and hence are called *"push customers"*. Less informed, passive traders or traders who mainly rely on technical analysis, are pulled into the market by price movements and hence are called *"pull customers"*. Therefore many assumptions of the Evans-Lyons model like customers' absorption of daily inventory and risk neutrality will be released [181].

The significance of signed transactions flows has been also addressed in several other models. [97] provide an integrated analysis of exchange rates, equity prices and equity portfolio flow trying to capture the interaction between optimal portfolio choice under market incompleteness and exchange rate dynamics in a simple model. They show that equity flows affect exchange rates via induced currency transactions and focus on how shocks in equity markets and dynamics rebalancing of equity portfolios initiates FX order flow, which in turn affect exchange rate movements.

Another empirical model on the impact of order flow on exchange rate has been also introduced by [15]. The fact that most short-run exchange rate volatility is related to order flow [63], which in turn is associated with investor heterogeneity, was the primary motive for Bacchetta and Wincoop's model [15]. They introduce heterogeneous/disperse information into a standard rational dynamic monetary model that can explain the exchange rate determination puzzle and the evidence on order flow. They show that the model can explain important stylized facts like disconnect puzzle for exchange rates from fundamentals over short and medium horizons, but the close relation of exchange rate changes to fundamentals over long horizons, the weak predictive power of exchange rate changes to observed fundamentals and finally the close relation of exchange rate to order flow.

A new framework for modelling exchange rate dynamics has been also introduced by [67], where they connect the dynamic general equilibrium (DGE) and microstructure approaches to exchange rates, combining risk-averse decision making, heterogeneous information and social learning in the sense that agents learn from the equilibrium actions of others. Their model accounts for many exchange rate stylized facts like persistent gaps between exchange rates and macro-fundamentals, excess volatility, exchange rate movements due to trades/order flow without the existence of macro news, and also little exchange rate movements with the occurrence of macro news.

2.4.2.2 Order Flow, Fundamentals, Information and Feedback Trading

While in the previous section, we have shown that order flow has an explanatory power for exchange rate movements, in this section and the following, we address the studies explaining the source of the order flow and how its explanatory power comes into effect.

[178] use data for three main exchange rates from Reuters, and documents that information in order flow is intimately related to a broad set of economic fundamentals. They show the importance of macroeconomic information in driving trading and asset allocation decisions of foreign exchange market participants and, as a consequence, in moving exchange rates. They even provide evidence that order flow has a predictive power of future fundamentals, a finding that has been previously supported by [67], who argued that order flows convey information about fundamentals that is not yet common knowledge to all dealers.

Another interesting finding by [67] is that the pace of information aggregation is slow, which may be a possible explanation for the longstanding puzzles in exchange rate economics; the disconnection between spot exchange rates and fundamentals over short and medium horizons. One factor that might contribute to the slow pace of information aggregation is the presence of price-contingent order flow generated by feedback trading. Stop loss orders, for example, represent a form of positive feedback trading, in which a fall in the FX price triggers negative order flow from customers wishing to insure their portfolios against further losses. In line with Evans and Lyons, [169] also argues that feedback trading will be an important component of order flow when quotes approach the points at which stop-loss orders cluster. A fall in FX quotes at these points can trigger a self-reinforcing price-cascade where causation runs from quotes to order flow.

The effect of feedback trading has been also shown by [83] in a study assessing the relation that exists between order flow, exchange rate returns and fundamentals. Their results indicate that some traders employ positive feedback rules over short-term horizons and then unwind their speculative position. As a consequence, they appear to follow negative feedback trading rules over long horizons. Their results show insignificant values for the covariance between fundamental news and flow surprises, expected short-term innovations in order flow, expected long-term innovations in order flow and total innovations in order flow. They conclude that the positive impact of order flow on exchange rate is a transitory phenomenon not necessarily related to fundamental information. The latter argument, that the information carried by the order flow is not necessarily fundamental, is in line with [29] findings. They find that inter-dealer foreign exchange order flow has a strong impact on short- and medium-term exchange rate returns, but much less explanatory power for long-term exchange rate fluctuations. The fact that order flow may carry non-fundamental information has been nicely discussed in [170], citing evidence for the relevance of non-fundamental information and arguing that the large scale speculative high frequency trading in the FX market may not be based on future views of fundamentals. Dealers consider non-fundamental information relevant, even though it has a transitory impact on the market.

2.4.2.3 News Arrival and Order Flow

The analysis of the effect of news arrival on exchange rates goes back to the series of studies started by Meese and Rogoff in the 1980s [139], who revealed the disconnection between exchange rate forecasts and fundamentals or macroeconomic variables. Since then, several researchers have tried to examine the impact of macroeconomic variables on exchange rates by adopting a new approach, namely measuring the effect of fundamental shock and macro announcements on daily and intra-daily exchange rates as well as on order flow.

Generally, we can find two strands of research concerning the effect of news arrival on exchange rates. The first strand examines the direct effect of news on exchange rate movements, while the second strand examines the channels (e.g. order flow) through which this effect is transmitted. Researchers like [7, 58, 88] have shown that macroeconomic news announcements lead to short lived jumps in exchange rates. One of the most cited studies on the effect of news arrivals, is a study conducted by [10], who use a large data set consisting of six years of real-time exchange rate quotations, macroeconomic expectations, and macroeconomic realizations (announcements.). They show that unexpected fundamental shocks significantly affect exchange rates, whereas expected macroeconomic information has no effect on exchange rates. A more recent study by [43] uses a surprise announcement of an increase in German interest rates coupled with concurrent transactions of DM/USD data, trying to examine dealers' reactions to this news. From 30 seconds after the announcement a huge volume of seller-initiated trades takes place. Dealers holding long positions in Dollar, place market orders at existing bid prices, whereas dealers having short position in Dollar, and have open buy limit orders for dollars, book their profit and let the market hit their bids. [43] attribute the continued high volume and price volatility of the DM/USD to the speculative activity when the news was announced.

Studying the channel through which news is transmitted to exchange rates, [178] argue that order flow can be seen as a vehicle that carries aggregate interpretations of news and expectations of fundamentals to exchange rate. [124] analyse the effect of news arrival on exchange rates by studying the interrelation between order flow, spot rates and macro news. They study the effect of news arrival on exchange rates and order flow separately; the impact of order flow on exchange rates around announcement dates; and finally the effect of news arrival and order flow on exchange rates simultaneously. Love and Payne [124] find that exchange rates react immediately to news arrival even at high frequencies (1-minute). They also show that news arrival affects order flow, with both immediate and delayed effects, where 30-60% of the simultaneous impact of news on exchange rates is transmitted by order flow. They conclude that nearly 50% of public information simultaneously released to all market participants is impounded into exchange rates via order flow.

In contrast to [124], [64, 67] argue that public information/news does not affect order flow and is immediately incorporated in currency values. According to them, the order flow is only affected by the private component of information and changes exchange rates consequently, noting that private information here refers to possible different traders' interpretations to macro news. Evans and Lyons calculate that roughly 70% of daily exchange rate variance due to news arrival is via order flow and 30% is via the direct effect.

Trying to provide interpretations for the fact that common news can be a source for customer order flow, [67] argue that this news may imply heterogeneous views about future returns, which in turn induces portfolio adjustment. They find evidence that the exchange rate effects of macro announcements operate via both a direct channel (i.e., when the announcement contains common knowledge (CK) news on GDP, unemployment etc.) and an indirect channel, which operates via order flow and conveys dispersed information about fundamentals to dealers. [69] address again the presence of an indirect channel through which public news affects prices. Their intraday data shows that order flow contributes more to changing FX prices in the period immediately following the arrival of news than at other times. With both the direct and indirect channels operating, they estimate that macro news accounts for 36 percent of total FX price variance in daily data.

2.4.2.4 Foreign Exchange Intervention

Foreign exchange intervention is the practice of monetary authorities buying and selling currency in the foreign exchange market to influence exchange rates. Two approaches have been recently adopted to understand the effect of interventions: event studies and high frequency data.

Event studies examine the market's reaction to the event and its effect on the price. High-frequency data use both exchange rates and intervention data to better understand the behaviour of exchange rates immediately around intervention. Most central banks routinely "*sterilize*" their foreign exchange operations; that is, they buy and sell domestic bonds to reverse the effect of the foreign exchange operation on the domestic money supply [149].

Theoretically, sterilized intervention affects the value of currencies and the level of activity in FX markets, either through a portfolio-balance channel, supplies of bonds denominated in different currencies, or via a signalling channel, communicating information about future monetary policy or the long-run equilibrium value of the exchange rate. News of an intervention operation would quickly spread in the market through the inter-dealer section of the FX markets. Usually, the diffusion of news of foreign exchange intervention is so rapid that newswire services would report intervention activity in the space of few minutes [198].

Some researchers suggest that foreign exchange intervention can be effective because it carries information. This is because central banks are the most informed agents in the market, which trade on superior information and consequently alter securities prices. This way, order flow in foreign exchange markets affects exchange rates as it transmits private information possessed by central banks. For a comprehensive list of central bank intervention in FX, see surveys by [149, 198].

2.4.2.5 Price Discovery in the FX market

Another important line of models in the FX microstructure research focuses mainly on the dealer behaviour with regards to price setting and risk management. This is in contrast to general equilibrium Evans-Lyons models [63, 68], that mainly focus on how the entire market builds such a consensus and settles on an exchange rate [179] . The first study addressing the price discovery process by FX dealers was Lyons [130] based on the Madhavan-Smidt model [134]. They find that the dealer increases his spread with trade size to protect against private information, and adjusts the midpoint in the spread to induce trade in a preferred direction to adjust inventory. According to [171] this adverse selection framework, which focuses on the information advantage of customers over dealers, for the price discovery process has been adopted in most of the foreign exchange rate microstructure studies.

However, empirical evidence shows violation of this theory. Using the same model as Lyons, [32] provide empirical evidence of no support for such information or inventory effects. They suggest that this might be due to the change in the trading environment caused by the introduction of electronic brokers. [171, 174] also show that customer spreads are inversely related to trade size and are narrower for the customers that dealers assume to be most informed.

Another alternative price discovery framework is proposed by [171] for which they provide some empirical evidence. The pricing mechanism involves the way dealers trade and the type of orders (market or limit orders) they place in the inter-dealer market after individual customer trades. The framework is based on the condition that price discovery must take place in the inter-dealer market, since customers' information is not immediately reflected in the prices they pay. Since dealers trade with informed customers, they will tend to place market orders, motivated by their inventory and their newly acquired information. In this way the information from customer trades will be reflected in inter-dealer prices. In contrast, after trading with uninformed customers, dealers "*will be more likely to place parallel limit orders or to wait for incoming calls, leaving price relatively unaffected* "[171] . For a good review of the price discovery process in currency markets see [111, 171].

2.4.3 Heterogeneous Agent Based Modelling Approach

As described in section 2.4.2, the order flow model can account for 50% of the explanation of the exchange rate behaviour[131]. However, an alternative explanation for the discrepancy between short- and long-term exchange rate expectations is the heterogeneous expectations and forecasting rules of the FX participants. Hence, an important route of microstructure modelling focuses on capturing traders' heterogeneous beliefs and strategies. In this recent behavioural agent-based approach, computer simulation models are the main modelling framework. For an excellent review on heterogeneous agent-based modelling in economics and finance, see [102] . In this section, we briefly introduce heterogeneous agent-based models of the FX market.

The robustness of agent-based models is attributed to their ability to replicate many of the exchange rate stylized facts, such as high volatility, fat tails, volatility clustering, exchange rate disconnect and large trading volumes. One main source for this heterogeneity, and upon which most models in the literature are based, is the chartist-fundamentalist (CF) trading approach. The first foreign exchange rate heterogeneous agent model with a CF framework was originally suggested by [79], which was then extended by several other researchers. According to the CF approach, agents are classified into chartists, fundamentalists and portfolio managers. According to [79, 81] the log of the exchange rate is determined by portfolio managers, who buy and sell foreign currency in response to changes in the expected exchange rate and a set of other contemporaneous variables. The expected exchange rate is in turn generated using a weighted average of chartist forecasts and a weighted average of fundamental forecasts. In this model, forecasts are generated by fundamentalists using a long run equilibrium, while chartists predict the future rate using spot rate. Using simulation, the model reveals that portfolio managers keep learning the model by subsequently revising the weights given to chartists and fundamentalists, leading first to a bubble in the exchange rate that smooths out after a while, as the portfolio manager model leans to the correct one and fewer revisions are made [4].

Extending the Frankel-Froot CF model, different types of chartist-fundamentalist set-ups have been implemented into a wide range of different exchange rate models and have been introduced by several researchers. For instance in [203], exchange rate is modelled with fundamentalists having regressive expectations, where the perceived fundamental value of the exchange rate follows changes with the arrival of news, and chartists form their expectations using a variety of technical rules. The simulation of [203] model displays many of the exchange rate dynamics. Another more general heterogeneous agent asset-pricing model with a CF setup, has been developed by [101]. The model was originally developed for the stock market but can also be generalized to the FX market. [101] calls his model an "adaptive belief system for financial markets", involving evolutionary competition between trading strategies of chartists and fundamentalists. [53] also propose another non-linear exchange rate model including transactions costs and heterogeneous agents who switch between their trading rules based on the past profitability of these rules. They show that their model displays many of the exchange rate dynamics.

Another group- but few- of exchange rate models combine target zone models, first introduced by Krugman [116], with chartist-fundamentalists beliefs. A target zone implies setting a central parity around which the exchange rate is allowed to fluctuate within an interval. When the exchange rate approaches the limits of the interval, the monetary authority intervenes in the domestic monetary market and prevents the exchange rate from deviating outside the interval. Thus, an undervalued (overvalued) currency is bought (sold) when the distance between the exchange rate and its fundamental value exceeds a critical threshold value. There are very few target zone models including the chartist-fundamental behaviour as a source for heterogeneity of traders' beliefs. Two recent studies by [27, 177] combine target zone type intervention models with chartist-fundamentalist setup to model exchange rates, showing that these interventions have a stabilizing nature on the exchange rate prolonged by both chartists and fundamentalists and that the impact of speculative activity declines within the target zone regime.

The heterogeneity of beliefs is not the only behavioural aspect embedded in market modelling; there are other behavioural biases like feedback trading, overconfidence, loss aversion, and disposition effect, for which empirical evidence has been described in section 2.4.1. In contrast to the stock market, where we find several market models and frameworks embedding these behavioural biases [18, 23, 25, 54, 103], we find few evidence of integrating these behavioural aspects, except for feedback trading, in foreign exchange market agent-based modelling. The case for feedback trading is different, as it is already implied in technical trading rules, in which traders depend on past performance [102]. As for other biases like overconfidence and disposition effect, we find a very recent FX framework addressing them, introduced by [154]. The authors find that the standard disposition effect is complemented by the impact of the total portfolio performance on the length of investment periods. The model also delivers detailed insights into investors' overconfidence, leading to clustering of trades and to longer holding periods.

With the availability of large machine-readable market data sets and the advances of computational technologies, many researchers have adopted a recent approach to understand and replicate market properties, namely the "Artificial Financial Market" model. The dynamics of an artificial market model (computational agent based model) are solely determined by either naive or more intelligent computer programs that model the various behaviours of the market and its interacting agents [196]. The Santa Fe Institute Artificial Stock Market was the first and most influential multi-agent artificial financial market, introduced by [12], that motivated further research on artificial financial markets. It consists of various types of agents who choose from among many investment strategies. Agents' investment decisions are affected by their expectations about future prices, which in turn affect realized prices, leading to new beliefs and so on and so forth. Thus, agents keep adapting their behaviour with every new observation filtering out less successful strategies. Hence, the market co-evolves over time [102]. When constructing an artificial financial market the researcher is faced with a large number of basic design issues, including agent preferences that can vary from very simple to very sophisticated ones, price determination, evolution and learning that can be based on several artificial intelligence techniques, information representation, social learning rules determining how agents interact and learn from each other, and finally benchmarks (calibration) which concerns validating the agent based market against well defined market dynamics [118, 119].

Generally, the literature cites several agent based models for the FX market behaviour ranging from simple and partially computational models like [4, 52, 197, 202] to more rich extensive computational multi-agent exchange rate models like [11, 57, 115, 129], where some of them use genetic algorithms to evolve decision rules that were used to determine the composition of the agents' portfolios in a foreign exchange market. For a detailed overview of various agent based financial market models see [118, 119, 182].

2.5 Conclusions

Having surveyed the various approaches for studying the FX market behaviour, we find that the new paradigm is towards studying the behaviour of markets, either through behavioural biases of traders, or through micro-structural studies of market dynamics, or through agent-based modelling and simulation. The latter can also combine the second approach by embedding traders' behavioural biases in modelling trading agents. In the end, the same objective is always sought, namely to explain and model one or more of the foreign exchange rate puzzles. For instance, we have shown in the previous sections, that many of the above cited models, aim to explain and replicate the exchange rate disconnect puzzle from fundamentals over short and medium horizons. In the micro-structural approach, this puzzle has been explained and replicated by the aggregate order flow model and the associated information [14, 15, 63, 67, 83]. As for agent-based modelling, the disconnect puzzle in addition to other exchange rate stylized facts have been replicated after embedding some behavioural biases of traders like the chartist-fundamentalist trading approach, overconfidence and feedback trading [102, 118, 119], for which empirical evidence has been already provided by behavioural finance researchers.

Despite the success of these models in replicating an important stylized fact like the disconnect puzzle, however, there is lack of information about the real high frequency individual traders' activity, which in turn is the main driver for the market behaviour [159, 165].

This is mainly due to two reasons. First, the analysis of previous studies has been restricted to either large samples of low frequency data or very small samples of high frequency data. However, it has been empirically proven that most of the FX transaction volume is intraday, as dealers can hold only limited overnight positions. Hence, they provide only intraday liquidity in FX markets and fulfil buy and sell orders at their posted bid/ask quotes throughout the trading day, but are less likely to do so for longer horizons [32, 33, 130]. Also the fact that trading in the FX market is extremely frequent with orders being fulfilled on the spot, the market volume has a very short duration. Hence, to understand the FX market dynamics, we need to analyse and dig further into its high frequency tick transactions. Therefore, this research is empowered by a unique high frequency data set of more than 140 million real tick transactions carried out by over 45,000 FX traders for over 2 years on the OANDA FX trading platform [156]. The production process of this high frequency data set is discussed in details in chapter 3.

Second, the transactions analysed by previous research are not on an account level, giving no possibility for tracking individual traders' transactions and portfolio management along with exchange rate movements. As a result, the impact of traders' activities and trading positions on the market has not been actually studied in depth before. For instance, all previous studies dealing with transaction data focus on opening trades and do not take into sufficient account that part of the of the volume is closing trades. However, traders have both opening and closing positions. Studying when traders close their positions and defining at what point they accept to clear their positions is missing in the literature. Understanding such a behavioural aspect can reveal important information about market dynamics and lay a solid foundation for market modelling. Given these facts, we use our unique data set to track and analyse FX traders' activity in chapters 4 and 5.

In chapter 4, we track individual traders' seasonal activity, which enables us to normalise the different time zones from which FX traders operate. The resulting normalised seasonality conveys new information about FX traders trading intensity and positions closure during the day.

In chapter 5, we analyse and synchronize traders' activity with price movements. Given the second by second transaction volume in the FX market, the BIS triennial study states that 1.4 trillion USD spot are traded per day. This represents a microscopic transaction volume of 16.2 mill transaction per second; 8.1 mill buying volume and 8.1 mill selling volume[20].With such a small volume per second in the global FX market, any small imbalance in the buying or selling volume can trigger disproportionate price moves. In an attempt to understand the impact of microscopic activity on the global FX price move, we decipher the FX market activity in chapter 5 along an intrinsic time scale, based on exchange price turning points of a pre-specified threshold, and provide empirical evidence for further new market dynamics.

The microscopic second by second liquidity of the FX market, along with its 'random price spikes', has also a great impact on market movement when associated with leverage trading. This is an important aspect that has not been discussed in the literature. Leverage is the factor of the actual position size relative to the assets deposited in the account that is the underlying collateral. Investors can open positions larger than their assets, because brokers provide them with a loan facility to buy currencies. Some brokers are ready to provide leverage up to 50 times the underlying capital or even more. Leverage trading is also referred to as margin trading, where margin is the minimum amount required to be deposited and held by the investor in reference to the total value of his/her trades . In contrast to other financial markets, FX offers a much higher leverage, increasing its traders' purchasing power[156]. High leverage is always attractive as long as traders make profits. However, once they start losing, their losses are multiplied, leading to a closure of their positions either willingly through stop loss orders or forcefully through margin calls, which are triggered once the account balance is below the specified margin requirement. Since many traders may be trapped in the same situation, there may be a cascade of margin calls to close out their losing positions. This amplifies the flow imbalance of buyers over sellers or vice versa. If a trader was short and the price moves up to an extent that a margin call is triggered, then the trader has to close the short position by buying the underlying traded currency, which in turn increases the price imbalance. This is a destabilizing force in the market. We can only understand how such situations come about, if we understand the trader behaviour. Such micro-behaviour has not been published before. Chapter 6 introduces a new event-based microscopic approach to formalise and analyse market dynamics and feedback loops like cascading margin calls, existing in both FX and stock markets, using a simplistic double auction market model.

Hence, despite the various contributions of the different approaches studying the FX market, they still leave several open questions about traders' and market behaviour. As briefly described, in the next chapters, we propose new microscopic approaches for studying and analysing the FX market behaviour, using a large sample of high frequency transaction data of FX traders on an account level. The size and the level of details of the data set allows us to carry out an in-depth analysis of traders' behaviour.

Part III

Thesis Contributions

Chapter 3

High Frequency FX Data Set Production[1]

The foreign exchange market generates millions of daily tick data, often referred to as high frequency data (HFD), as a result of market participants' decisions. By analysing these data, one could reveal many of the market properties. However, HFD may possibly contain observations that are not reliable in terms of actual market activity. In this chapter, we manipulate a real data set storing the full transaction history of more than 45,000 traders on an account level for 2.25 years. Prior to exploring the data to discover basic mechanisms of the market dynamics, we perform a cleaning and preparation process to remove any erroneous or misleading observations from the set, and to validate the consistency of transactions that would affect the reliability of any forthcoming results. We can confirm a clean transaction data set allowing for a better understanding of the financial market and invaluable future research

.

3.1 Introduction

The high liquidity of the FX market makes it an important source for high frequency data (HFD). HFD refers to a large amount of data, irregularly spaced in time, representing the

[1]Some of the work presented in this chapter has been taken in parallel by myself and Monira AlOud, a PhD student at the College of Computing at the University of Essex. However, we adopted different methodologies and accordingly developed different code. Fortunately, our results for the cleaned data set are very similar, which indicates that both of us successfully eliminated all erroneous observations from the data set. Furthermore, such similarity of results confirms the validity of the produced data set.

full record of transactions and their associated characteristics at tick-by-tick frequencies, where a tick can be a second or less [49, 60]. To understand the FX market behaviour, we need to explore the millions of tick data generated each day as a result of the market participants' decisions.

Research in financial markets microstructure has been recently using HFD to detect trends and patterns of different magnitude in price time series. The objective is to study financial markets stylized facts and to model real time market dynamics [204]. Even though advances in databases and computing power have made these data sets increasingly accessible, several challenges are associated with its analysis. There is no standard structure for HFD. It depends on the institution's policy with regard to the production and storage of data [39]. In addition, high frequency data sets can have a variety of erroneous or misleading observations and data gaps that may result from computer system errors or internal system procedures using dummy ticks [49, 204]. A failure to recognize erroneous data may affect the research validity. Hence, cleaning and filtering an HF data set is an inevitable laborious step prior to data exploration. There are some contributions [37, 39, 49, 168] in the literature dealing with HFD filtering issues. However, each study presents a different approach adjusted to the data set being processed. The criteria for filtering HFD differ with reference to the structure and types of errors of the underlying data set. The key is that handling a high frequency data set cannot be standardized and must be customized according to the nature of the available data. Generally, to handle a HF data set one should first understand the data set and its potential usefulness, then state the objective from data cleaning and finally decide on the cleaning process.

In this chapter, we deal with a unique high frequency data set representing the transactions history of thousands of FX traders on an account level. The data set is provided by OANDA Corp., an online trading platform. The data set is invaluable and its nature promotes for close monitoring and a microscopic analysis of individual traders' activities, which in turn can possibly explain several market dynamics. Therefore, the objective is to filter the data set from any possible non-reliable transactions in terms of traders' actual activity. This chapter describes all the steps carried out to clean the data from all erroneous or misleading observations as well as to validate its cleanness and prepare the data set for the microscopic analysis. Section 3.2 describes the OANDA data set and its potential usefulness. Section 3.3 states the objective from the cleaning process. The data set cleaning and preparation process is then discussed in section 3.4. Main conclusions are then given in section 3.5.

3.2 The Data Set

3.2.1 The OANDA FxTrade Platform

In this chapter, we handle a high-frequency data set of FX traders' transactions supplied by OANDA Corp [156]. The OANDA corporation, short for "Olsen And Associates" is a financial services provider of currency conversion, online retail foreign exchange trading (OANDA FxTrade), online foreign currency transfers, corporate hedging consultancy, and FX information.

OANDA FxTrade (Figure 3.1) is a market maker and one of the largest non-bank Futures Commission Merchants (FCMs) that specializes in spot FX trading. Over 20% of the world's online spot FX transactions take place on its servers. The platform serves a wide range of traders, from novice individual traders to corporations and financial institutions. It also offers several innovative and unique features relative to other online trading platforms. Traders are not subject to any lot size restrictions. That is, they all trade under the same terms and conditions, in particular the same prices with competitive tight spreads for small and large orders. Tight spreads are enabled through fast seamless execution. Using Straight Through Processing (STP), all transactions are fully automated and instantly settled. Another unique feature of the platform is the application of continuous second-by-second interest. This is in contrast to the standard method of interest calculation by brokers, who apply interest only once a day on open positions. This may create an artificial bias toward shorting weaker currencies (with higher interest rates) and potentially reward buyers of stronger currencies with lower rates of interest. Finally, OANDA is of one the few market makers that tries to achieve an open transparent market by publishing the underlying market open short/long positions to help traders judge the depth of the market for a given currency pair. In contrast, most banks and FX market makers either do not share their data on open positions and pending orders or give it only to select clients.

Figure 3.1: Screenshots of the OANDA online trading platform (FxTrade). Source: http://fxtrade.OANDA.com

3.2.2 Importance of the Data Set

The available data set covers all tick transactions on an account level, made available on an anonymous basis, carried out on the OANDA platform from January 1, 2007 to March 5, 2009. This makes a total of more than 147 million transactions carried out by 45,845 different accounts trading in 48 different currency pairs. About 64 % of the data set transactions are in 7 out of the 48 currency pairs traded in the data set (Figure 3.2). These are EUR/USD, USD/JPY, GBP/JPY, EUR/JPY, AUD/JPY, AUD/USD and GBP/USD in order, where EUR/USD has the highest trading proportion ($\sim 37.59\%$) and GBP/USD has the lowest proportion ($\sim 7.16\%$). The remaining 36% of the total transactions are distributed among the other 41 currency pairs.

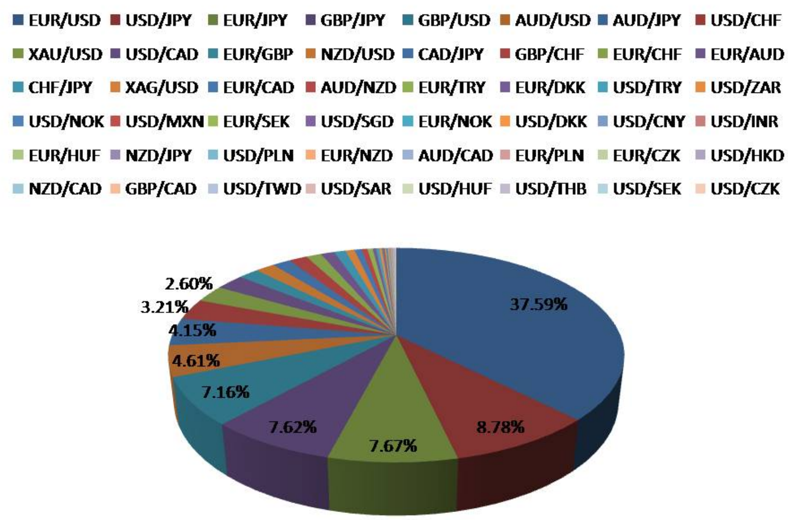

Figure 3.2: Currency pairs traded on the OANDA platform for the 2.25 years of the data set period.

It is worth noting that there is no platform apart from OANDA that stores the details of traders' transactions and their positions' history over several years, representing the worldwide exposure of traders in FX . Understanding the FX market requires exploring its high frequency (tick-by-tick) data, especially given that traders' decisions, formulating the collective market behaviour, are based on observing this high frequency data [49]. Therefore, having access to such a unique and large high frequency data set, from a trusted source like OANDA, can lead to a new generation of research on the high-frequency FX market.

Access to transaction data in FX markets has been only recently enabled. Having surveyed the various approaches for studying the FX behaviour in chapter 2, we find that despite the various time horizons and frequencies (large samples of low frequency data or small samples of HF data) of the transaction data explored by previous studies, it has been restricted to few currencies. Most important, the studied market transactions are not on an account level. They do not differentiate between the impacts of the different participants trading in the market, leaving many open questions about the market behaviour.

In contrast, our data set records more than 147 million physical transactions of more than 40,000 traders at an account level in all traded currency pairs, making it the largest set of HFD ever. The size and level of details of the data set, allows for tracking individual traders' transactions and analysing their trading behaviour on a microscopic basis, which in turn formulates the collective market behaviour. In the following chapters, we describe in details how we can conduct a microscopic analysis of FX traders' actual behaviour using the underlying data set. Thus, this chapter represents a major contribution by producing such an invaluable data set, not only for our work but also for future research.

3.2.3 Description of the Data Set Fields

Each transaction record contains detailed information about the transaction number, the account (anonymised) placing the transaction, the timestamp at which the transaction is executed, the traded currency pair, the executional price, the number of units and amount value traded. However, no information is provided about the accounts' balances. Below is a description of each field. Table 3.1 shows a snapshot of some transactions records of one account.

Table 3.1: Snapshot of a random account transactions

TransId [a]	AccountId[b]	Timestamp[c]	CurrencyPair[d]	Type[e]	Price[f]	Units[g]	Amount[h]
59697811	4	1196438200	EUR/USD	BuyOrder	1.46974	5	7.3487
59697812	4	1196449156	EUR/USD	CloseTradeB	1.46738	5	7.3369
65085899	4	1197982896	EUR/USD	SellOrder	1.44068	10	14.4068
65085900	4	1197983022	USD/JPY	BuyOrder	113.438	50	5671.9
65085901	4	1197987388	EUR/USD	CloseTradeS	1.44294	10	14.4294
65085902	4	1197988337	USD/JPY	CloseTradeB	113.252	50	5662.6
74388724	4	1205220700	EUR/USD	BuyOrder	1.53604	7	10.75228
74388725	4	1205221715	GBP/USD	BuyOrder	2.00877	10	20.0877
74388726	4	1205221716	GBP/USD	BuyOrder	2.00877	20	40.1754
74388727	4	1205222761	GBP/USD	ClosePositionB	2.00688	30	60.2064
74388737	4	1205222784	EUR/USD	SellOrder	1.53636	7	10.75452

[a] TransId: unique transaction identity

[b] AccountId: anonymised account identity

[c] Timestamp: timestamp of executed transaction

[d] CurrencyPair: pair of currencies exchanged

[e] Type: type of transaction (order)

[f] Price: execution price.

[g] Units: number of traded units expressed in base currency.

[h] Amount: transaction value in account's home currency

Trans ID: The transaction id (Trans ID) is a unique identity, stored once a transaction is carried out on the OANDA FxTrade platform. A transaction is an action performed on the customer account. The different types of transactions are explained in the data set field 'Type'. A transaction is equal to one record in the data set. More than one transaction can be executed at the same timestamp.

Account ID: The account id is a unique and anonymised number characterising each account. The account numbers do not represent the real identities of OANDA traders. They have been changed by OANDA for privacy issues. One trader may have different accounts.

Timestamp: The timestamp indicates the time at which the transaction has been executed. It is stored in Unix format, which is measured as the number of seconds between the date of interest and Jan 1 1970.

Currency Pair: The currency pair (Base/Quote, e.g. EUR/USD) indicates what one buys (sells) in base and sells(buys) in quote.

Type: This is the type of the executed transaction. A transaction can be of type:

BuyOrder (SellOrder): buying (selling) a number of units in the base currency and selling (buying) the same number of units in the quote currency of the currency pair. A BuyOrder (SellOrder) can occur in the data set as more than one transaction, closing previous open sell (buy) orders previously placed by the trader.

CloseTradeB (CloseTradeS): Closes a single previously made buy (sell) order.

ClosePositionB (ClosePositionS): Closes a sequence of open buy (sell) orders, equivalent to the customer's total long (short) position in a certain currency pair. ClosePositionB (ClosePositionS) is recorded in the system as different transactions. Each one of these transactions has a number of traded units equivalent to the total size of the open buy (sell) orders being closed.

Units: The number of traded units expressed in base

Price: This is execution price of the transaction.

Amount: The amount is calculated as (Price * Units), converted in home currency, the currency of the account.

3.3 Objective of the Cleaning Process

As described in the previous section, the available high frequency data set is considered an invaluable source for tracking and analysing FX traders' behaviour on a microscopic basis, from which general implications about the market behaviour can be made. Hence, our objective is to clean the data from any erroneous or misleading observations not corresponding to traders' actual activity.

To achieve this objective, we need first to carry out an exploratory data analysis and get a 'look and feel' of the underlying data and its complexity. We are handling millions of transactions in different currency pairs carried out by thousands of traders, who have different trading behaviours and different trading strategies. The heterogeneity of traders and the complexity of their transactions can be visualised by Figure 3.3, which depicts the tick transactions of only 10 sample accounts, carrying out over 9000 transactions over the 2.25 years of the data set period.

As illustrated by Figure 3.3, traders vary in terms of trading frequency, trading volume (number of units) and even in their trading time window, where some traders are only active for a few months and others trade over the whole data set period.

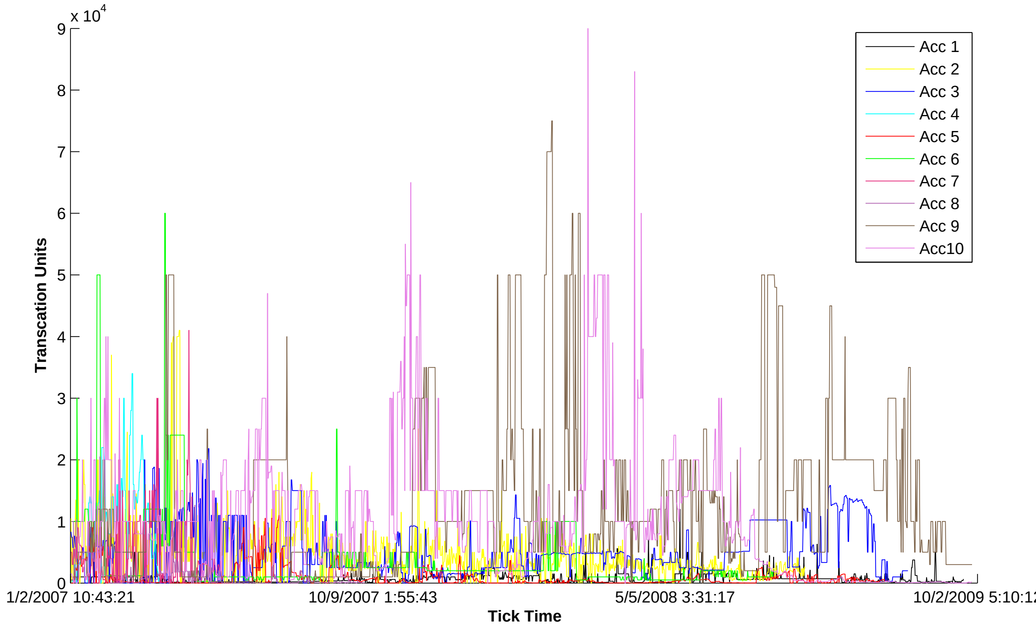

Figure 3.3: Tick Transactions of 10 sample accounts

What makes it more complex is that this heterogeneity is not only among the different traders, but also within the trading flow of each account. This is illustrated in Figure 3.4, where we take one sample account from Figure 3.3 and aggregate its number of transactions on a minute, hourly and daily basis. In each time span we can observe some outliers at certain points relative to the whole trading period. The question is whether these outliers are due to a sudden change in the account's trading frequency, or are misleading and erroneous observations. These can be due to several factors, discussed in the following sections, like dummy transactions of internal system procedures, spurious transactions entered by OANDA's accounts for testing purposes, or even platform dependent storage policies of executed transactions, where one trade can be stored in many transaction records in the database. A failure to understand the origin of these outliers and to recognize erroneous transactions can distort the research analysis. Therefore, we need not only an aggregate level but also an in-depth data exploration to ensure the cleanness and reliability of the data on an account level.

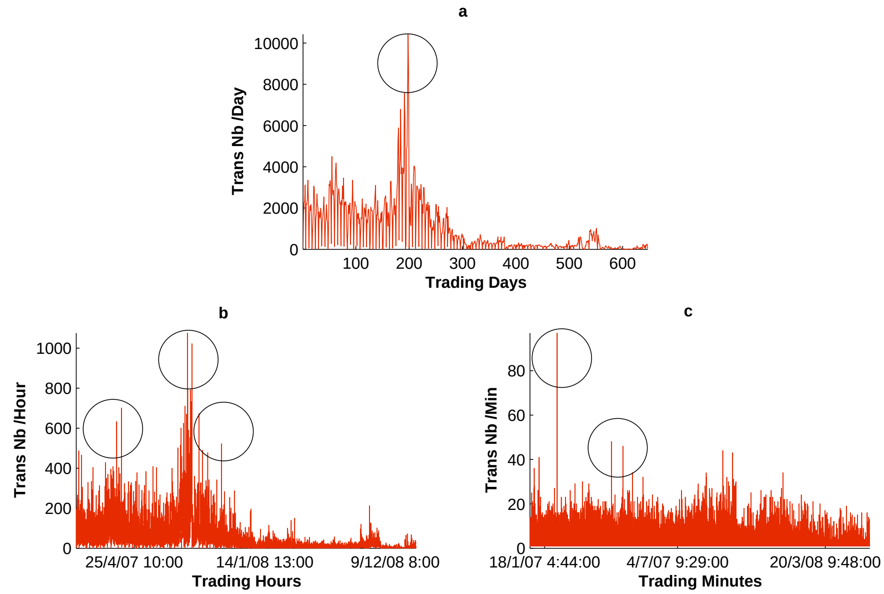

Figure 3.4: Aggregated number of transactions for one sample account on a a) daily b) hourly and c) minute basis. Circled data points represent possible outliers resulting from erroneous observations

3.4 The Data Cleaning and Preparation Process

After identifying the objective from the cleaning process, we carry out in this section a step-by-step procedure, depicted in Figure 3.5, to remove non-reliable transactions in terms of traders' actual activity and to prepare the data set for analysis. First, we study OANDA's internal system storage procedures causing any misleading observations and affecting the data reliability. Going a step further to understand and produce a validated clean data set for research analysis, we explore the aggregate trades' flow of all accounts over the whole data set period (2.25 years). Then, we dig further into the traders' activities and track each account's transactions and positions, to confirm their consistency in terms of the types and size of orders placed and the consequent open and closed positions. Finally, we analyse traders' intraday seasonality (activity pattern) and compare it with the benchmarked intraday FX market seasonality reported in the literature.

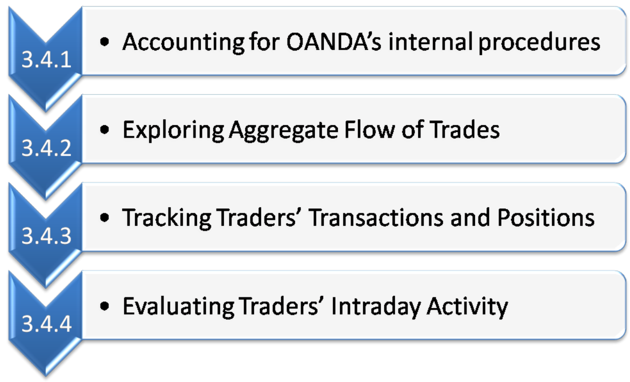

Figure 3.5: Data cleaning and preparation process

3.4.1 OANDA's Internal Procedures

3.4.1.1 Execution and Storage of Transactions

To understand the execution and storage of transactions on the OANDA platform, we use a demo account to carry out several trading experiments on OANDA's 'FxTrade Practice' platform (see Figure 3.6). This is a real-time version of OANDA's foreign exchange trading platform, FxTrade, where the trader can trade under real market conditions with live prices and spreads for any period of time [156].

The fact that the system stores all executed transactions has allowed us to view all activities performed on the experiment account over any time span. Studying the experiment's transactions, retrieved from the 'FxTrade Practice' database system along with the underlying data set, we find that one trade can be stored in many transaction records in the database. Possible reasons include placing a buy (sell) order that closes several previous open sell (buy) orders; closing a position using order types 'ClosePositionB' or 'ClosePositionS'; and placing a stop loss order to automatically close an open position when the price moves against the trader and reaches a specified threshold. Another reason can be forced close-outs (margin call) in case of margin trading. Margin trading allows the trader to enter into positions larger than his/her actual account balance, given that the trader has sufficient collateral referred to as margin. Once the trader's account value depresses to a value specified by OANDA, a margin call takes place, which automatically closes out all of the trader's open positions.

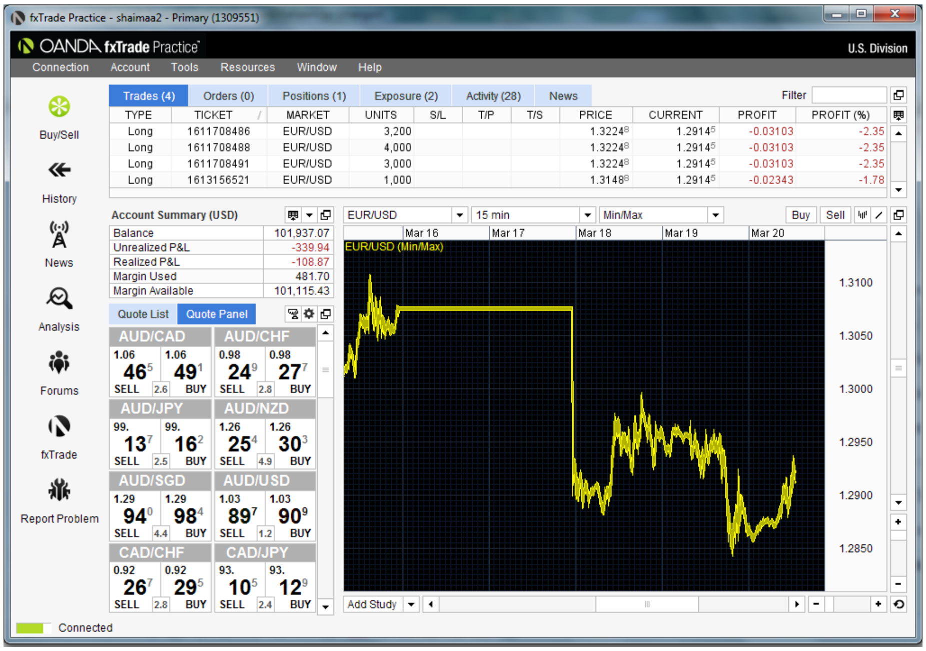

Figure 3.6: Snapshot of the demo account trades on the OANDA 'FxTrade Practice' platform. Source: https://fxtrade.OANDA.com

Analysing the transactions in the current status can provide misleading information about the traders' behaviour in terms of frequency of trading and adopted strategies. Therefore, we carry out a merging procedure on all related transactions. These can be generally categorized into three groups illustrated in Table 3.2.

Table 3.2: Data set transactions groups

<table>
<tr><th>AccountId</th><th>Date Timestamp</th><th>CurrencyPair</th><th>Type</th><th>Price</th><th>Units</th><th colspan="2">Group</th></tr>
<tr><td>8</td><td>1/1/2007 9:57:07</td><td>EUR/USD</td><td>ClosePositionB</td><td>1.3197</td><td>500</td><td rowspan="3">A</td><td rowspan="4">B</td></tr>
<tr><td>8</td><td>1/1/2007 9:57:07</td><td>EUR/USD</td><td>ClosePositionB</td><td>1.3197</td><td>1000</td></tr>
<tr><td>8</td><td>1/1/2007 9:57:07</td><td>EUR/USD</td><td>ClosePositionB</td><td>1.3197</td><td>2000</td></tr>
<tr><td>8</td><td>1/1/2007 9:57:10</td><td>EUR/USD</td><td>ClosePositionB</td><td>1.3197</td><td>200</td><td></td></tr>
<tr><td>8</td><td>1/1/2007 10:50:00</td><td>EUR/USD</td><td>BuyOrder</td><td>1.3195</td><td>100</td><td></td><td rowspan="2">C</td></tr>
<tr><td>8</td><td>1/1/2007 10:50:00</td><td>EUR/USD</td><td>BuyOrder</td><td>1.3196</td><td>100</td><td></td></tr>
</table>

1. **Group A** includes transactions placed by the same account in one currency pair at the same timestamp and price using one type of order. Each set of related

transactions are merged into one trade. For each resulting trade, we aggregate the total units and amount traded as well as the number of merged transactions.

2. **Group B** is restricted only to transactions of order types 'ClosePositionB' and 'ClosePositionS'. It is similar to Group A except for one field. The transactions have different timestamps, usually a gap of few seconds. To know whether these transaction records belong to one trade or not, we check whether the trader has opened a position in this currency pair in the time gap between them. If no orders took place, then the transactions are merged. Total units and amounts traded and the number of merged transactions are then aggregated. Algorithm 3.1 explains the merging procedure for Group B transactions.

3. **Group C** includes transactions placed by the same account in one currency pair at the same time timestamp but at different prices. Group C transactions are either buy or sell orders. Despite the fact that these transactions belong to one trade, we decided not to merge them, as this would require finding an average price for the trade. A change in the original price can lead to an information loss of actual executional prices and the profitability of each transaction.

Algorithm 3.1 Merging Group B Transactions

Input:
D */* data set transactions (tuples). D is ordered ascending wise by the fields : Account No, Currency Pair, Timestamp. An identity field ID is then added to the ordered database. D={ T/T=(id,account,currency,timestamp,type,price,totalUnits,totalAmount)} */*
Output:
Merged Transactions of type 'ClosePositionS'/ 'ClosePositionB'

```
Begin
For i ←1 to n // number of D tuples
    fetch T_i from D ; //retrieve tuple (i) from D
    If (T_i.type=('ClosePositionB' or 'ClosePositionS')) Then {
        id ←T_i.id ; ac ←T_i.account ; curr ←T_i.currency ;
        type ←T_i.type; price ←T_i.price;
        While (ac=T_i.account and curr=T_i.currency and type=T_i.type
                and price=T_i.price)
          gap←T_i.id - id;
          If (gap≤ 1) Then {
            /*merge current tuple(T_i) with prev. tuple and
            update totalAmount and totalUnits*/
            If (gap=1) Then mergeTuples(T_i, id) ;
            id ←T_i.id ; i← i+1 ;
            fetch T_i from D
          }else break ;
         End While
    }//end of If
End
End
```

3.4.1.2 OANDA's Interest Payment Internal Procedure

As mentioned earlier, OANDA applies a continuous second-by-second interest on its trading platform. The data set does not contain any information about interest payments. However, OANDA's internal procedure for the interest calculation must be considered in the data cleaning process. The system enters spurious (dummy) transactions having an amount value of zero referring to incurred interest payments. These transactions are then merged once a day, though frequency may vary, to calculate the interest for the open positions. The reason for doing this is to speed up daily interest payment calculation. To account for possible erroneous analysis as a result of these spurious transactions, they are completely removed from the data set. Figure 3.7 shows the number of dummy transactions removed every month.

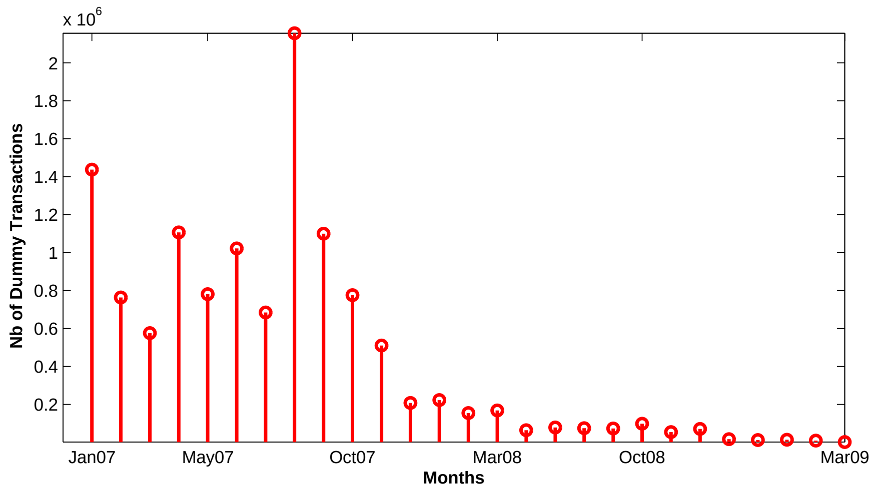

Figure 3.7: Monthly dummy transactions removed from data set

3.4.2 Aggregate Flow of Trades

After handling the system's internal procedures, we start exploring the aggregate number of trades per month. We use now the terminology 'trade' after carrying out the merging process in the previous section . Figure 3.8 shows the monthly evolution of the transactions and trades before and after accounting for the internal system procedures respectively against the number of accounts.

The variation of the number of accounts from month to month is natural, as new accounts keep being created, while others stop trading either for some time or up till the end of the data set period.

The evolution of the number of trades at OANDA mirrors market volatility, especially in August 2007, where foreign exchange markets experienced an extensive increase in volatility as a consequence of major dislocations in other financial markets. This rise in volatility was accompanied by higher turnover in the foreign exchange spot market.

According to the Bank for International Settlements 78th annual report [19], the spread of electronic trading platforms has generally contributed to the unprecedented rise of FX turnover in 2007, in part because it has enabled large financial institutions to set up algorithmic trading systems, and has provided trading facilities to retail investors.

A questionable observation in the data is that, even though the number of accounts increased, their trading activity dropped from September 2007 onwards. It is worth noting that the drop in the number of accounts in March 2009 is a data artefact and must be ignored, as we have data for only 5 days in March 2009.

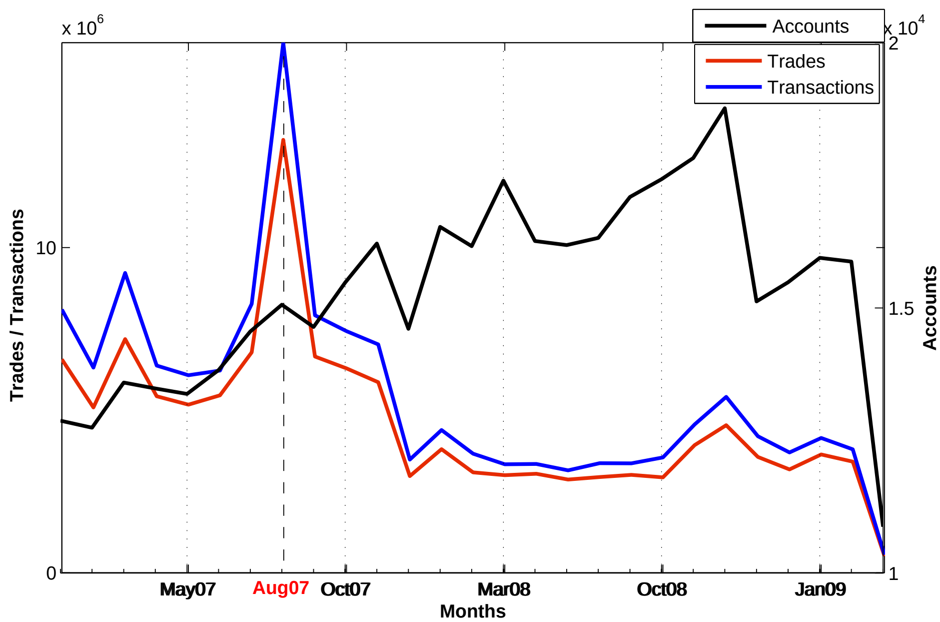

Figure 3.8: Monthly evolution of number of accounts vs. number of transactions/trades (Jan07-Mar09). Traders' activity is referred to as 'transactions' before handling internal system procedures, and as 'trades' after the handling.

In an attempt to understand the current trades flow relative to the number of accounts evolution, we examine the trading frequency of all accounts over the whole data set period, to decide on whether the different behaviour from 2007 to 2008/2009 is due to traders' varying trading frequency or due to spurious transactions. We define trading frequency as the maximum number of daily trades executed by each account per month

(30 days period). To calculate the daily number of trades, we use a moving average of a 30 days window. The n day simple moving average for day d is computed by:

$$MA_d = \frac{\sum_{i=1}^{n} T_{(d-i)+1}}{n}, \; where\, n = 30 \tag{3.1}$$

We find that accounts daily trading activity can largely vary from one trade to thousands of trades per day. Figure 3.9 shows the moving average daily trades of two sample accounts having different trading frequencies.

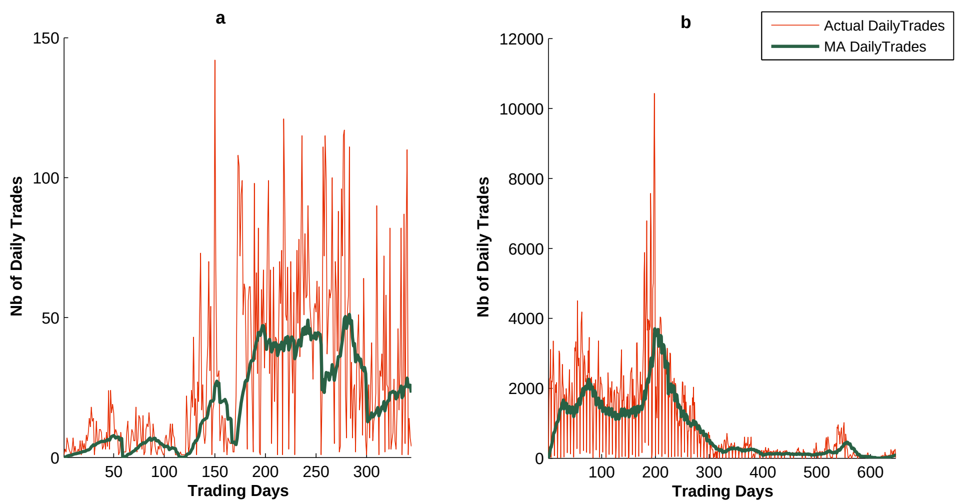

Figure 3.9: Actual and moving average daily trades for two sample accounts having different trading frequencies

Categorizing these trading frequencies into classes, we come up with eight different trading frequency classes (AA,A,B,C,D,E,F,G,L), as shown in Table 3.3.

Table 3.3: Range of trading frequency classes

Trading Frequency Class	Max Nb of Trades Per Day
AA	>1000
A	>500 && <=1000
B	>300 && <=500
C	>200 && <=300
D	>100 && <=200
E	>50 && <=100
F	>20 && <=50
G	>10 && <=20
L	>0 && <=10

It is a complex task to assign an account to a certain trading frequency class, as accounts have various trading frequencies over the different trading periods. Furthermore, accounts have different trading windows. Some accounts trade for a few months whereas others trade for the whole data set period. In addition, many accounts have gaps in their trading windows. For instance, an account might be trading for a total of six non-consecutive months, which are spread over the whole data set time period. So, taking an average of an account trading frequency over the whole data set period (27 months) can be misleading. Also, comparing accounts belonging to the same class can be misleading due to the different length and order of the trading windows. Considering all these complications, we have classified the accounts into the different classes over various time spans; yearly, quarterly and monthly. The accounts considered in each time window, are only those trading in each month of the window. This way, accounts belonging to the same class are comparable as they have the same trading window. The trading frequency of each account has been calculated by taking the average trading frequency over the considered time span. The shorter the time window, the less variance we have in the trading frequency of each account over the different months and the more accounts we have trading for the whole window. For all trading windows, class 'L' is the dominant class, showing that the majority of accounts exercise a low trading frequency. Figure 3.10 shows the percentage number of accounts in each trading frequency class over yearly, quarterly and monthly time spans.

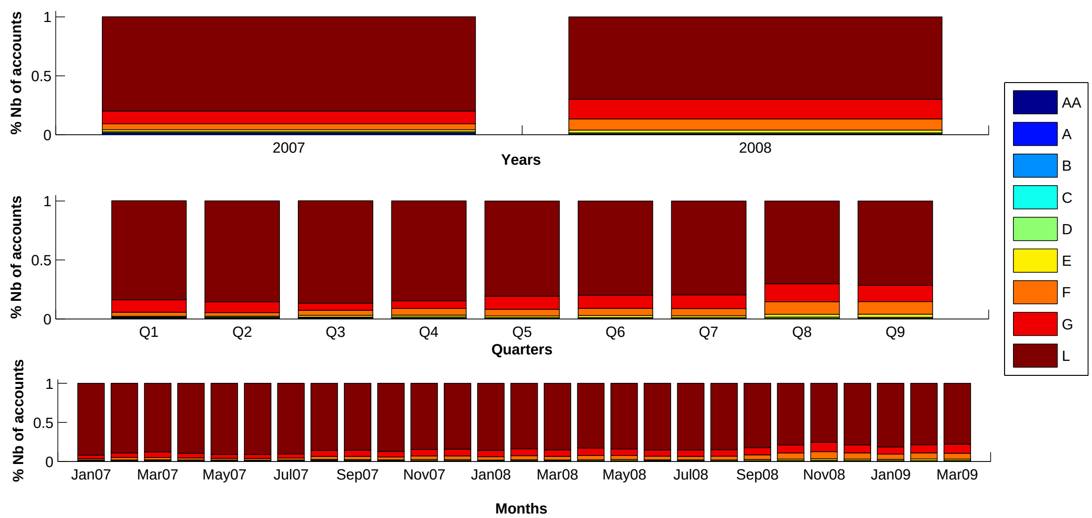

Figure 3.10: Yearly, quarterly and monthly trading frequency of all accounts

Accounts trading more than 1000 trades at least for one day in a month are defined as ultra high frequency (HF) accounts. An account that is defined as HF in one month may not trade high frequency in the next one. Searching the data set for HF accounts, we find 114 HF accounts of which some accounts trade in all 27 months and others trade high frequency for only a few months in 2007. Figure 3.11(a) displays the number of HF accounts per month, showing more high frequency traders in 2007 relative to the remaining data set time period. To measure the impact of HF accounts on the aggregate trades flow, we use linear correlation (eq. 3.2) between HF accounts daily trades and both the absolute (r=0.8915) and moving average (r=0.9324) total number of daily trades in the data set (Figure 3.11(b,c)). We find a strong positive correlation, indicating that the increase in the number of trades in year 2007 is not only due to market volatility but also due to HF accounts.

$$r_{X,Y} = \frac{1}{n}\sum\left(\frac{X_i - \mu_x}{\sigma_x} * \frac{Y_i - \mu_y}{\sigma_y}\right) \tag{3.2}$$

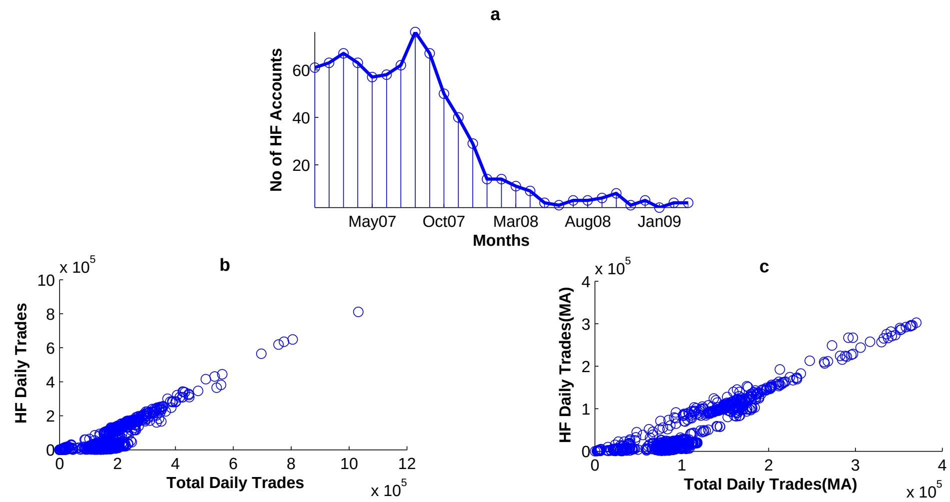

Figure 3.11: a) Number of HF accounts per month b) Correlation of the HF accounts' actual daily trades and the data set total daily trades. (r=0.8915) c) Correlation of the HF accounts' moving average daily trades and the data set total daily trades (r=0.9324).

The extracted HF accounts contribute largely (~44%) to the number of trades of the whole data set. However, the volume of their trades represents only ~0.68% of all accounts' trades volume. The volume is the sum of the total traded units over a specific period of time. Having such an extremely small trading volume relative to the number of trades, the identified HF accounts may be spurious accounts. Questioning OANDA quants about this trading behaviour, they confirm that such accounts are indeed spurious and are only entered into OANDA's trading platform for testing purposes. Therefore, we remove 114 HF accounts and their related trades. Figure 3.12 shows the number of trades after removing the HF accounts. We can still observe the sudden increase in August 2007 resulting from the high market volatility during this period, but the evolution of the number of accounts goes now in line with the monthly flow of trades.

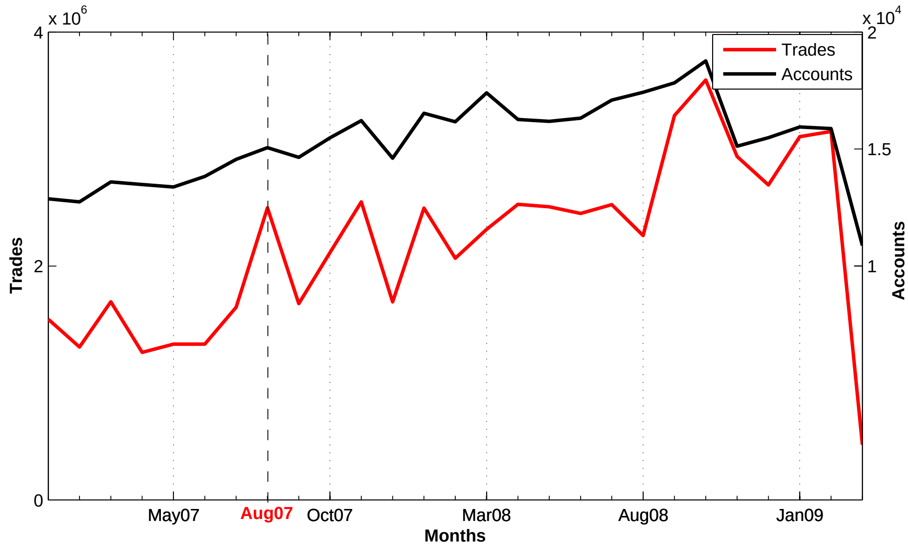

Figure 3.12: Monthly evolution of number of accounts vs. number of trades (Jan 07-Mar09) after removal of HF Accounts.

3.4.3 Tracking Individual Traders' Positions

In this step, we dig further into the data by tracking individual traders' activities from opening to closing their positions. The objective is to validate the consistency of the accounts' trades in terms of the number of units bought and sold and the resulting opened and closed positions in the market. Given a closed currency position, the account's exposure in this currency must be equal to zero. The exposure Ex (eq. 3.3) represents the aggregated holdings of the traded currency pair, by computing the net sum of all traded units since the last flat position up till the current timestamp (t). The units are signed according to the trade type. If the trade (eg. buy order) increases the account's holdings, the traded units $u(t_i)$ are positive. If the trade decreases the holdings (e.g. sell order), $u(t_i)$ is negative. Eventually, the resulting exposure is also signed. A positive exposure means an open long position, while a negative one means an open short position in the underlying currency. Consequently, a zero exposure indicates a closed (long/short) position.

$$Ex_{(t)} = \sum_{i=1}^{n} u(t_i), where\, t_1 < t_2 < ...t_n < t \tag{3.3}$$

To confirm that the exposure of closed positions is zero, we need first to identify when each account opens and closes its positions in the different currency pairs. Generally, a trader makes a trade to open a new long (buy) or short (sell) position in a currency pair; switch from one position to another; manage the current open position by decreasing or increasing its size; or just close an existing open position. Following is a definition of the different functionalities of an account's trades:

Opening Trade: This is any trade of type 'BuyOrder' or 'SellOrder' occurring after an account's position has been closed in a certain currency pair. A 'BuyOrder' opens a long position where a 'SellOrder' opens a short position in the underlying currency pair.

Closing Trade: This is any trade that closes the long (short) position of an account in a certain currency pair. There are three possibilities for a closing trade:

- ***ClosePositionB (ClosePositionS) trade***: This trade closes the already open total long (short) position of the account in a certain currency pair.
- ***CloseTradeB (CloseTradeS) trade***: Once a CloseTradeB (S) occurs, then we know for sure that the account has a zero short (long) position, as an account can hold only either a short or a long position in a currency pair. A CloseTradeB (CloseTradeS) decreases the long (short) position. However, a CloseTradeB(S) can close the long (short) position, in case it closes the last open BuyOrder (SellOrder)
- ***SellOrder (BuyOrder) trade***: A SellOrder (BuyOrder) decreases the long (short) position. However, a SellOrder (BuyOrder) can close the long (short) position, in case it matches against all the units of the remaining open buy orders (sell orders) of the long (short) position.

Switching Trade: This is any trade that closes one position and opens another position simultaneously. This means switching from long (short) to short (long) position in a certain currency pair using a Sell Order (Buy Order). For instance, say an account holds a long position of 100 units. Placing a sell order of 150 units will close the long position and open a short position of 50 units.

Managing Trade: This is any trade that increases or decreases the long or short position of an account in a certain currency pair. Trades managing the long position are ' BuyOrder', 'SellOrder' and 'CloseTradeB', while trades managing the short position are ' BuyOrder', ' SellOrder' or' CloseTradeS'.

To extract these trades, we implement Algorithm 3.2. We start by identifying the first timestamp, where the account's position in the underlying currency pair is flattened. This is identified by either locating the first occurring 'ClosePositionB(S)' or 'CloseTradeB(S)', whichever comes first. In the case of 'ClosePositionB(S)', we know for sure that the account's total position is closed and flattened and any consequent trade must be either a 'BuyOrder' or 'SellOrder' to open a new long or short position respectively. In the case of 'CloseTradeB(S)', we do further calculation. For instance, if we locate a 'CloseTradeB', then we are certain that the account is holding a long position that is being either managed (decreased) or closed by the underlying 'CloseTradeB' order. To decide on whether the long position is being closed, we calculate the size of the position by tracking back all previous trades. We add up the units of all previous buy orders that increase the long position while we deduct all units of any previously made sell orders. If the resulting net sum is equal to the size of the 'CloseTradeB', then the position is flattened, otherwise, we continue tracking the account's trades after the placed 'CloseTradeB' order, until the position is flattened.

Once the first flattened position for the specified currency and account is located, we start tracking all the account's subsequent currency trades. As illustrated in Algorithm 3.2, we consider the different order types ('BuyOrder', 'SellOrder', 'CloseTradeB', 'CloseTradeS', 'ClosePositionB', 'ClosePositionS') offered by the OANDA platform to identify when the trader opens, increases, decreases or closes a (long/short) position.

Having this information , we can easily calculate an account's exposure Ex, defined in eq. 3.3, from opening up till closing a position in a certain currency pair. Figure 3.13 shows the exposure of one sample account after tracking its trades in EUR/USD during one hour. The depicted account has opened and closed four positions during one hour and consequently has a zero exposure, each time the account's underlying long/short position is closed.

Tracking all accounts trades in the different currency pairs, we are able to confirm the reliability of the data set and its consistency in terms of the number of units bought and sold and the resulting opened and closed positions in each currency pair.

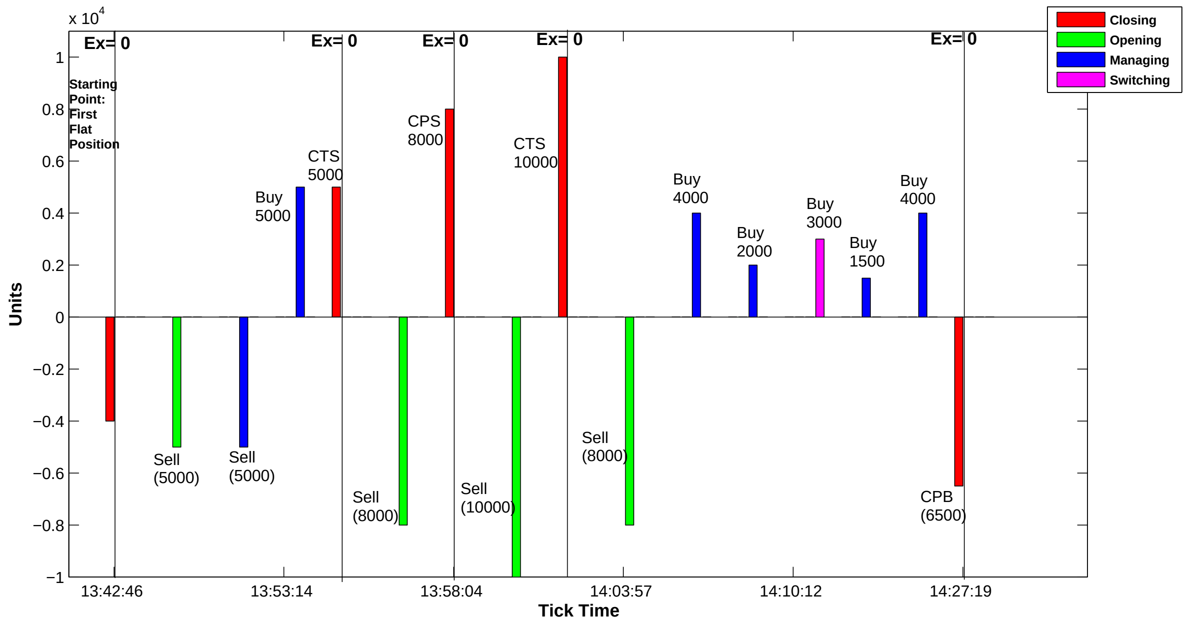

Figure 3.13: Closing, opening, managing and switching trades of one sample account during one hour. Each bar represents an order placed by the trader: Sell, Buy, CTS (CloseTradeS), CPS (ClosePositionS), CPB (ClosePositionB). The exposure (Ex) is zero when the account's underlying long/short position is closed. This account has opened and closed four positions during one hour after the starting point.

Algorithm 3.2 Tracking Traders' Positions

Input:

Ac *//Account ID*

C *//Currency Pair*

Output:

Identify the closing, opening, switching and managing trades

Begin

//Return the timestamp of the first closing trade (ClosePositionB(S),CloseTradeS(B)) flattening the position of Ac in C

startTime←getFirstClosingTradeTimestamp (Ac,C);

//Return trades of Ac in currency pair C

Trades←getTrades (Ac, C, startTime);

*/*Trades={T/T=(time, type ,units ,long, short ,E)}, where T.time ≥startTime, long(short)=size of long(short)position,*

E=trade effect='Closing' /'Opening' /'Managing' /'Switching' .Up till this point the fields (long, short ,E) have no values./*

For i ←1 to n *// number of Trades*

If (T_i.time= startTime) Then**{**
T_i.short← 0; T_i.long← 0; T_i.E←'Closing';
}Else if (T_i.time>startTime) Then **{**
If (T_{i-1}.E='Closing') Then**{**
If T_i.type='BuyOrder' Then {T_i.long←T_i.units; T_i.short← 0;}
Else if T_i.type='SellOrder' Then {T_i.short←T_i.units; T_i.long← 0;}
$T_i.E$ ←'Opening';
}Else If (T_{i-1}.E='Opening' or T_{i-1}.E='Managing') Then{
If (T_{i-1}.long= 0 and T_{i-1}.short> 0) Then
*/*increase short position if type='SellOrder' and*
*decrease if type='BuyOrder'/ 'CloseTradeS'/ 'ClosePositionS' */*
adjustShortPosition(T_i.type,T_i.units);
Else if(T_{i-1}.long> 0 and T_{i-1}.short= 0) Then
*/*increase long position if type='BuyOrder' and*
*decrease if type='SellOrder'/ 'CloseTradeB'/ 'ClosePositionB' */*
adjustLongPosition(T_i.type,T_i.units);
//Update trade effect E
If (T_i.long= 0 and T_i.short= 0) Then T_i.E←'Closing' ;
Else if (T_i.long > 0and T_i.short= 0) or (T_i.long= 0 and T_i.short>0)
Then T_i.E←'Managing' ;
Else if (T_i.long < 0 and T_i.type = 'SellOrder') Then
{T_i.short←-T_i.long; T_i. long← 0; T_i.E←'Switching'};
Else if (T_i.short< 0 and T_i.type='BuyOrder') Then
{T_i.long←-T_i.short; T_i.short← 0; $T_i.E$ ←'Switching'};
}}

End *//end of for*

End

3.4.4 Traders' Intraday Seasonality

Analysing traders' intraday activities and comparing this with the benchmarked FX market seasonality reported in the literature is a further validation procedure of the data set cleanness. Seasonality is periodical (intra-daily, intra-weekly, intra-monthly) patterns of market activity [49]. We define intraday seasonality as the sum of trades in each hour (h=0,1,2..23) of the day, divided by the total number of trades over the whole data set period. The intraday seasonality is revisited in more detail in chapter 4.

$$Seasonality(h) = \frac{\sum trades(h)}{totalTrades} \tag{3.4}$$

Several empirical studies provide evidence for the FX market intraday double U-shape or 'camel-shape' pattern [49, 108]. Analysing the intraday activity of the dirty and cleaned data set, we can observe a clear difference in the trading pattern. As depicted by Figure 3.14, the unprocessed data set shows a noisy pattern due to the spurious dummy transactions resulting from OANDA internal procedures, whereas the cleaned data set reveals a double U-shape pattern, going in line with the reported for FX market activity in the literature. It is generally assumed that activity peaks occur as a result of the overlapping trading sessions of the different markets around the world [49, 108]. Hence, the first peak occurs when the Tokyo and London sessions overlap, while the second one takes place when the London and New York sessions overlap (Figure 3.14(b)). With this analysis, we can confirm the representativeness of our cleaned data set of the global FX market.

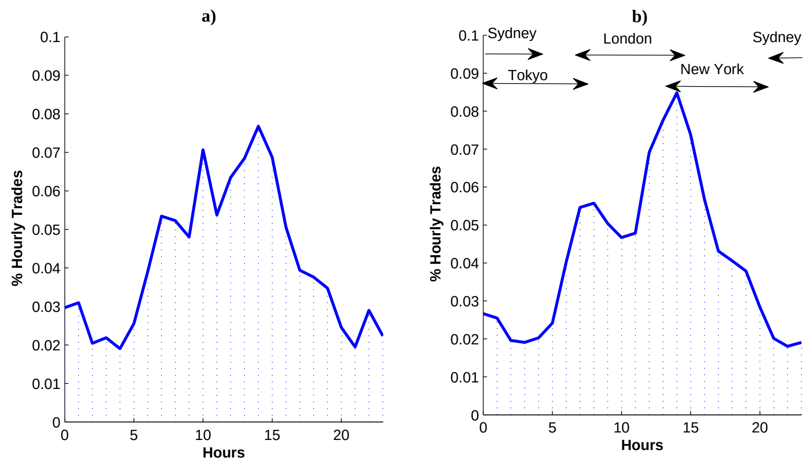

Figure 3.14: FX traders intraday seasonality a) before and b) after the cleaning and preperation process of the data set.

3.5 Conclusion

In this chapter, we have produced, as far as we know, the biggest set of HFD ever, of the tick transactions carried out by over 45,000 FX traders on the OANDA FX trading platform for 2.25 years. There is no platform apart from OANDA that stores the details of traders' transactions and their positions' history over several years. This makes our data set unique, as it includes detailed information about each transaction on an account level. To prepare the data set for analysis, we have explored the structure and the unique properties of the data set dictated by the specific context of the OANDA platform. Taking this into consideration, we have cleaned the data set from any erroneous or misleading observations that would affect the research validity. We have also validated the quality of data by tracking traders' activity from positions opening to closure to confirm the consistency and reliability of traders' transactions.

In addition to cleaning and validating the quality of the data set, we have provided strong indicators that the data set is representative of the whole market. We show how the trades flow of the cleaned data set mirrors market volatility in August 2007 and how its intraday pattern mirrors the double U-shape of the global FX market activity reported in the literature. This goes back mainly to the nature of the OANDA platform, which serves a wide range of traders under the same terms and conditions, in particular the same prices with competitive spreads, representing a worldwide exposure of FX traders.

With such a unique and detailed data set, researchers can study trading patterns, as demonstrated in chapter 4, and trading behaviour in microscopic perspective, as demonstrated in chapter 5. Our use of this data set is limited by our knowledge of the market micro-dynamics, which chapter 6 attempts to formalise. However, the production of such a unique and invaluable HF data set is a major contribution for future research.

Chapter 4

Time Zone Normalisation of FX Seasonality

This chapter tracks and analyses the FX market seasonal activity from a microscopic perspective, using the tick transactions carried out by FX traders on the OANDA FX trading platform. To account for the FX market continuous trading process around the world and its different trading sessions, we design an event-based procedure to normalise the traders' time zones. Our study provides empirical evidence that the particular intraday seasonality observed in the foreign exchange market is indeed due to the different geographical locations of its traders. The resulting normalised intraday seasonality shows a pattern akin to the ones observed in regulated exchanges, where traders are more active at the beginning and at the end of their trading session. We also provide empirical evidence that the trading intensity at the end of the day is due to traders closing their positions.

4.1 Introduction

The Foreign Exchange (FX) market is a sleepless market, operating 24 hours 5 days a week as earlier described in chapter 2. Unlike stock markets, foreign exchange trading is not dealt across a trading floor during a fixed period of eight hours a day. Instead trading is done online via telephone and computer networks between dealers in different trading centres around the world. These trading centres cover most of the world's time zones, allowing traders to respond to market dynamics day and night. FX trading begins on Sunday evening when the market opens in Sydney for the start of the trading week. As the clock moves, other trading centres around the globe come online, starting with

Tokyo, followed by London and then New York. At the time the New York trading session closes, 4:30 PM EST, Sydney reopens for the next trading day. Traders can continue trading around the clock until the New York session closes on Friday.

Having different trading natures, foreign exchange and equity markets reveal different intraday activity, the so-called seasonality. Several empirical studies show that equity markets exhibit a single U-shaped intraday seasonality, where high activity takes place at the beginning and end of the day [2, 78, 85, 96, 109]. This is mostly due to market participants opening fresh positions in the morning and closing positions at the end of the working day to prevent them from carrying the risk overnight. In contrast, the FX market shows a very different signature, a double-U shape pattern [8, 9, 49, 107, 108], where peaks occur at 8am and 2pm GMT (see Figure **4.1**). It is generally assumed that the reason for the different seasonality between the equity and the FX markets is the overlapping trading sessions around the globe. Supposing the assumption to be correct, the literature does not report any empirical evidence on the shape of the activity of any session.

Exploring previous empirical studies using high frequency FX data, such as [49, 51, 107, 108, 185], we find that each study focuses only on the pattern of one or few currencies, where some currencies may differ in their properties in the different trading sessions. For instance, evidence is provided by [49] for seasonal heteroskedasticity in volatility and quote frequency, showing a double U-shape for USD/DEM , USD/JPY and USD/CHF. They argue that these patterns may be explained by the behaviour of three partially overlapping main markets; America, Europe, and East Asia. Similar results have been reported by [107] for the activity (measured by number of quote entries and number of deals) in USD/JPY. Using the same explanation as [49] , they argue that the double U-shape is due to high activities taking place at the opening of the market. However, after adjusting for winter and daylight saving time of the three major markets (Tokyo, London, NY), they find no U-shape intraday activity pattern in the USD/JPY market. The activities are high during the opening hours but not the ending hours of the three major markets, even during the closing hours of New York on Friday. Another study carried out by [108] on USD/JPY and EUR/USD shows again double U-shape for deals, price change and return volatility. Examining the Tokyo, London and New York participants separately, they show that the U-shape of intra-day activities (deals and price changes) and return volatility is confirmed for Tokyo and London but not for

New York participants.

Overall these previous studies decompose the well-known double U-shape pattern into trading sessions. To the best of our knowledge, none has been able to empirically demonstrate that the double U-shape pattern is indeed an overlap of single U-shape patterns (characterizing single participants trading in different currencies) amid traders' behaviour in regulated exchanges.

In this chapter, by normalising for the different time zones in which FX market participants evolve, we show that the overlapping sessions assumption is indeed valid. We also demonstrate that the FX intraday seasonality, once normalised, exhibits a signature similar to the one measured in equities [2, 78, 85, 96, 109]. Using this finding, we have been also able to identify how traders manage their positions during the trading day. The results reported in this chapter reveal important features of the FX market dynamics.

The rest of the chapter is organized as follows: Section 4.2 defines the market seasonality and questions the reported FX market activity pattern in the literature. Section 4.3 addresses the time zone normalisation and the microscopic analysis of the individual FX traders seasonal activity. Traders being excluded from the analysis are discussed in section 4.4. Section 4.5 contrasts the aggregate original FX seasonality with the aggregate time zone normalised one. In section 4.6, we decompose the FX traders' seasonal activity in attempt to explain the trading pattern during the day. Main conclusions are then given in section 4.7.

4.2 FX Market Activity Pattern

The FX Market activity exhibits seasonal patterns over various time scales: intraday, day, week, even month. The activity can be measured in different ways during a fixed period of time, e.g. number of quotes, trading volume or trading value. Here we use an alternative activity measure A_{h_i} as the hourly share of daily number of trades

$$A_{h_i} = \frac{\sum_d^{n_d} N_{h_i,h_{i+1}}^d}{\sum_d^{n_d} N_d} \tag{4.1}$$

where h_i is the i^{th} $(i = 0, 1, ..23)$ hour of the day, n_d is the number of days over which the activity is measured, $N_{h_i,h_{i+1}}^d$ is the number of executed trades on day d between hour h_i and h_{i+1} and $\mathcal{N}_d$ is the total number of executed trades on day d. In what follows, for the sake of simplicity, we refer to A_{h_i}as A when no particular time tranche h_i is considered.

To analyse the activty of the FX market, we use our unique high frequency transaction data set supplied by OANDA Corp. As illustrated in the previous chapter, OANDA is a one of the largest non-bank Futures Commission Merchants running over 20% of the world's online spot FX transactions. It serves a wide range of traders, from novice individual traders to corporations and financial institutions with exactly the same terms and conditions, making our data set a good representative of the whole market. The activity A of OANDA traders is shown in Figure **4.1**(a) where we observe a double U-shape pattern in agreement with the literature [49, 108]. A similar pattern can be observed when computing the intraday tick seasonality of the FX market currency pairs, calculated the same way as in eq. 4.1 but measured using the hourly share of daily number of price ticks. Figure **4.1**(b) shows such a double U-shape pattern for EUR/USD and EUR/CHF. In addition to conclusions already drawn in chapter 3, these graphs stress that the present data set exhibit the same behaviour as the global market.

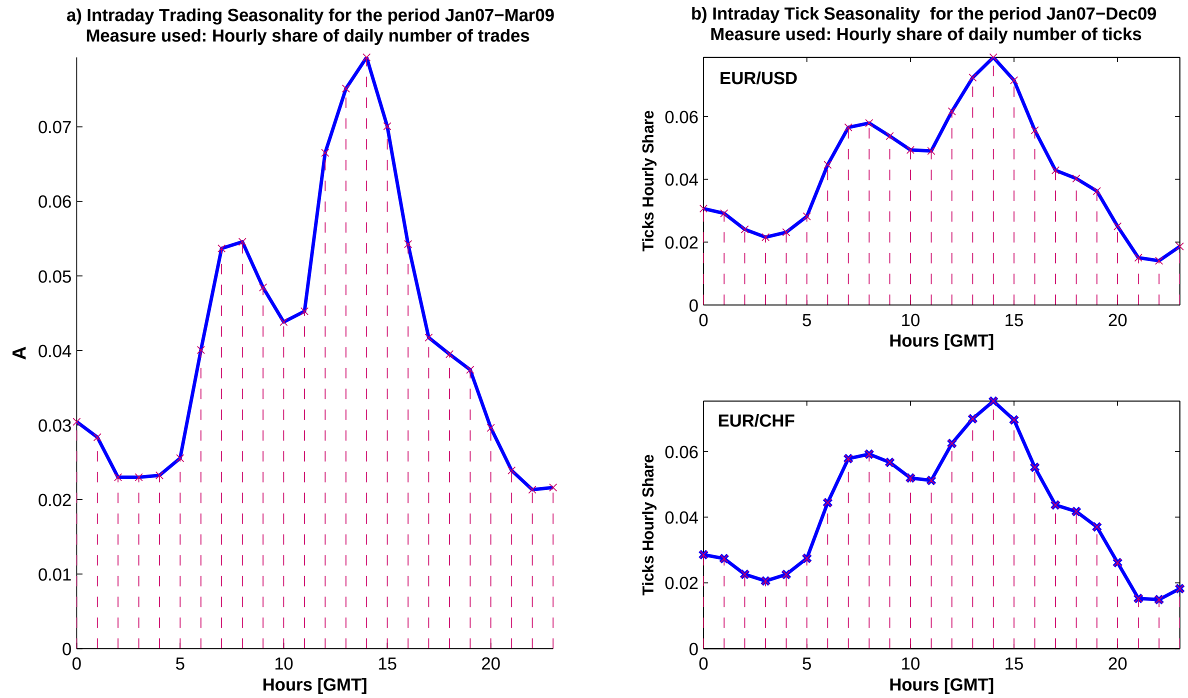

Figure 4.1: Intraday seasonality patterns using different activity measures. a) Hourly share of daily number of trades A computed from trades placed at OANDA between Jan 2007 and Mar 2009. b) Hourly share of daily number of ticks in (top) EUR/USD and (down) EUR/CHF between January 2007 and December 2009.

4.3 FX Seasonality Normalisation

In this section, we develop a new procedure to normalise the time zones of FX seasonality. The time window in Figure 4.1 is based on the actual timestamps (converted to GMT) at which market participants from the different world regions are trading. So, it may be morning (i.e. beginning of a trading day) for one trader and night for another trader (i.e. end of his/her trading day). This way, it becomes a complex task to contrast the behaviour of different traders as their activity is not aligned over the same time horizon, as in the case of stock markets, which have one constant trading time window of 8 hours a day.

Therefore, the objective is to define a new trading time window that is independent of the world time zones and that allows for synchronizing traders' seasonal activity. This means we cannot depend on the various physical timestamps of accounts' transactions. Therefore, we develop a new event-based approach to align traders' activity along the same trading time window. It is based on synchronizing the different traders according to their sleeping (inactive) hours. We explain below our approach in detail and the consequent normalised FX seasonality.

4.3.1 Normalised Intraday Seasonality

We identify a trading time window by first computing the hourly share of daily number of trades A, defined in eq. 4.1, for each of the 45,845 accounts over a specified period of time. Using this information, we define the hours during which an account is asleep (inactive) over the underlying period. Consequently, the activity period of this account is also identified. Accounts having the same number of sleeping hours are then grouped together and go through a time zone shifting (normalisation) process, so that the trading window (the first and the last hour of the active trading period) is identical for every participant in the group. Following this, we explain how we define accounts' sleeping hours and how we normalise their trading time zones according to their sleeping groups.

We note that one hour with no trading activity is not enough to be elected as a sleeping hour. Indeed it has to be consistent throughout the defined time period and element of a series of n consecutive sleeping hours$\{s_1, s_2, ..s_n\}$, where $3 \leq n \leq 17$. Accounts sleeping more than 17 hours a day are being ignored as non statistically significant, indeed in these accounts not enough trades were made. Accounts sleeping for less than 3 hours are also not considered, due to the high possibility of them being the hosts of algorithmic traders. The excluded accounts are discussed in section 4.4.

Similarly, we note that an hour with some trading activity is not enough to be defined as an active hour. To reduce any possible noise, we apply a filter given by:

$$A_{h_i} < r \text{ and } D_{h_i} < s = \begin{cases} true & A_{h_i} \text{ is set to zero} \\ false & A_{h_i} \text{ keeps its numeric value} \end{cases} \tag{4.2}$$

$$D_{h_i} = \frac{n_{h_i}}{n_d} \tag{4.2a}$$

The account is considered inactive (zero activity) in the i^{th} hour (h_i), if the two conditions of the filter are satisfied. The first condition checks if the account's hourly share activity A_{h_i} is smaller than threshold r. The second conditional statement checks if D_{h_i}, the percentage number of days an account has traded in h_i is below threshold s. $D_{h_i}$is defined in eq. 4.2a by the number of days n_{h_i}, in which the account has traded in h_i relative to the total number of the account's active days n_d. The thresholds values ($r = 1\%$ and $s = 2\%$) have been heuristically chosen with the objective of having the minimal information loss. This way, we alleviate insignificant active hours, while making sure that the small hourly share activity is not due to the account's short trading period.

Defining the sleeping period for each account, we find 15 different groups of accounts where each group sleeps n hours during the day over different time spans: monthly, yearly, 2.25 years. As described in chapter 3, the number of accounts in OANDA is not constant along the 27 months. As a result the number of accounts per sleeping period differs according to the time span used. Moreover, the behaviour and the trading frequency of traders can differ from one period to another. This means, that a trader might be in a sleeping mode for 17 hours a day during one month, and in another month he/she can be inactive for only 7 hours. This confirms the idea that in addition to the time zones variant, traders' trading time preferences also play an important role. Figure 4.2 shows the proportion of accounts in the different sleeping groups for different time spans. An account can belong to one sleeping group in one time span and then switch to another in a different time span. The active trading period of an account can be equal to or less than the defined time span.

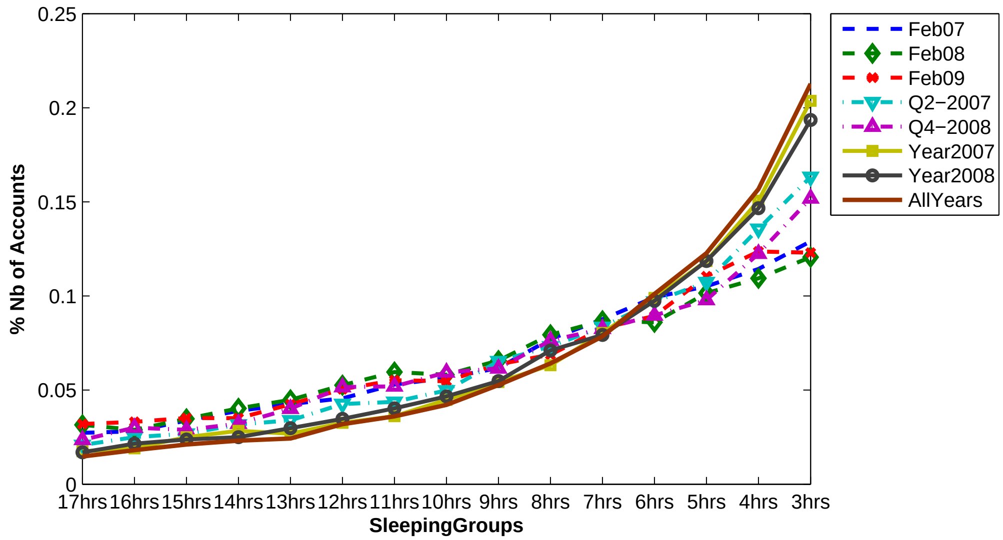

Figure 4.2: Percentage number of accounts per sleeping group over different time spans: monthly (Feb07, Feb08, Feb09), quarterly (Q2-2007, Q4-2008), yearly (2007, 2008), and all years (Jan 07- Mar 09).

Having categorized the accounts into sleeping groups, we set up a procedure to shift real trading time and align the traders in each group by their sleeping hours. This is done by a time zone normalisation procedure explained in Algorithm 4.1. For each sleeping group, we define a reference sleeping time period during the day, according to which the group accounts are shifted. An account's intraday activity can be shifted by 1 to 23 hours along the day, so that all participants in the same group start and finish trading at the same hour.

Algorithm 4.1 Normalised Intraday Seasonality

Input:
D/ /database accounts' trades
$Ac = (ac_1, ..ac_k)$ *//Distinct accounts in D*
P //the time period (monthly, quarterly,yearly,..) over which to implement the algorithm
Output:
Normalised Intraday Seasonality

Begin
For each $ac \epsilon Ac$ **do**
  //return the intraday seasonality array for each account's trades using eq. 4.1
  $A[]$←getAccountIntradaySeasonality(ac,D,P);
  //Apply filter on each account intraday seasonality using eq. 4.2
  For $j \leftarrow 0 : 23$
    If $(A[j] < 0.01 \,\text{and}\, D_{h_j} < 0.02)$ **Then** $A[j] = 0$;
  End
  //search for ac sleeping window of n successive sleeping hours,$3 \leq n \leq 17$
  n←searchSleepingHours($A[]$);
  //group ac with other accounts sleeping n hours
  joinSleepGrp(n,ac,$A[]$))
End
/ Normalize the the intraday seasonalities of the accounts in each sleeping group (*$SleepGrp_n$*), with* $n = \{3, 4, .., 17\}$*, to align them on the same trading time window/**

For $n \leftarrow 3 : 17$

  */*define a reference sleeping time window* $T_n[]$ *for* $SleepGrp_n$*, to which the sleeping window of each account in the group must be aligned*/*
  $T_n[] \leftarrow \{h_s, h_s + 1, ..h_s + n - 1\}$, where $h_s = starting\, sleeping\, hour$,$0 \leq h_s \leq 23$
  //hs is fixed so that all traders start sleeping at the same hour. We choose $hs = 22$

  For each $ac \epsilon SleepGrp_n$ **do**
    $Sleep_{ac}[]$←getSleepingWindow(ac); *//return the account ac sleeping window*
    If $Sleep_{ac}[0] \neq h_s$
      //calculate the hours lag (shift) btw $T_n[]$ *and* $Sleep_{ac}[]$
      Then $shift$ ←getLag($T_n[]$, $Sleep_{ac}[]$);

    //normalise time zone of ac by shifting its intraday seasonality in the relevant group
    $\bar{A}[]$←normaliseAccountSeasonality($SleepGrp_n$,ac,$shift$);
  End
End *//end of for*
End

The result of the normalisation process is illustrated in Figure 4.3 where we show the intraday trading activity of a sample of 10 accounts from three different sleeping groups: 17, 12 and 8 hours sleeping groups. Figure 4.3(a) shows accounts having the same number of sleeping hours. Their different sleeping time windows makes difficult a visual inspection of their similarity. Figure 4.3(b) shows the same accounts where the time zone normalisation has been applied.

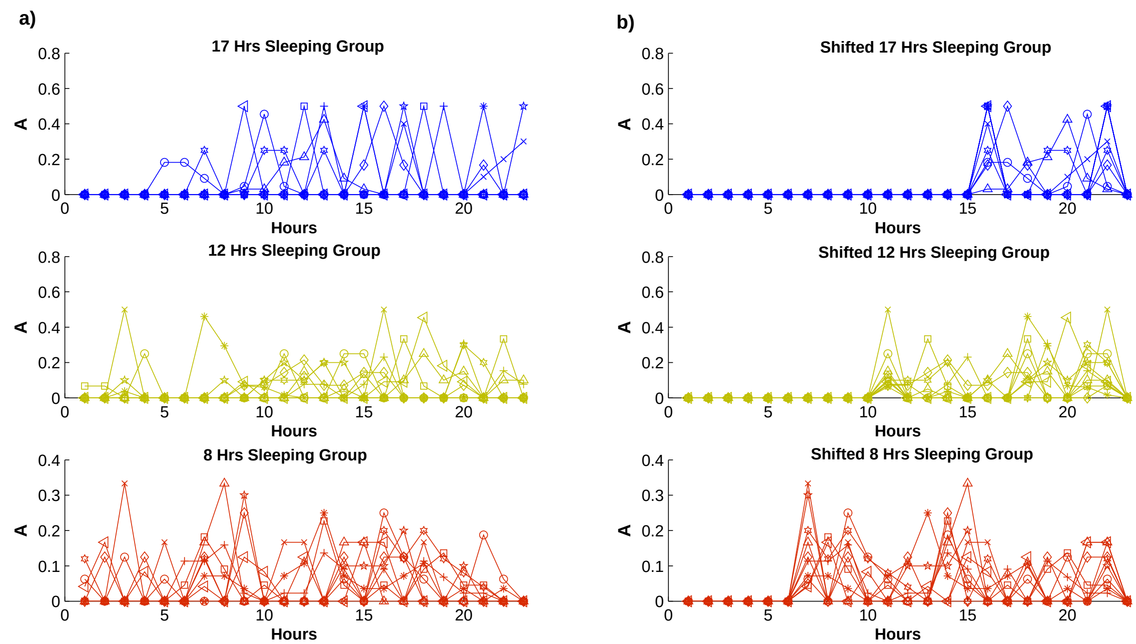

Figure 4.3: Intraday activity of 10 sample accounts from three sleeping groups a) before and b) after the shifting process. Before the shifting process, accounts belonging to each sleeping group have the same number of sleeping hours but different sleeping time windows. After shifting (b) accounts of each sleeping group have the same sleeping time horizon.

Aggregating shifted activities from each account in each sleeping group, we highlight the main result of this chapter: accounting for the different time zones, participants of the over the counter FX market behave in a similar way as the equity market participants. Figure 4.4 depicts that result, showing the normalised intraday seasonality of the different sleeping groups for the whole data set period (2.25 years). All curves are centralised to 12.00 pm for illustration purposes. We do not observe the double U-shape pattern of the FX intraday activity any more. Instead, traders exhibit a single

U-shape pattern with intensive trading activity peaks at the beginning and end of the day. These peaks however are less visible for those accounts sleeping 3 to 5 hours. They usually show a higher trading activity in the middle of their trading day. This intraday behaviour and the observed peaks are consistent over shorter time spans: yearly, quarterly, monthly and even daily (Figure 4.5).

It is worth noting that the minimal period for observing these patterns is strongly associated with the minimal period of observing the double U-shape seasonality of the FX market. We find that at least one trading day covering all world sessions in FX, or 2 days in physical time, will show a double U-shape pattern or a single U-shape in case of time zone normalisation (see Figure 4.6).

Our results of the FX traders' behaviour go in line with previous empirical studies on security traders intraday activity. They show a relative intensive intraday trading activity in the opening and closing hours and more quiet activity in the noon [2, 78, 85, 96, 109, 200].

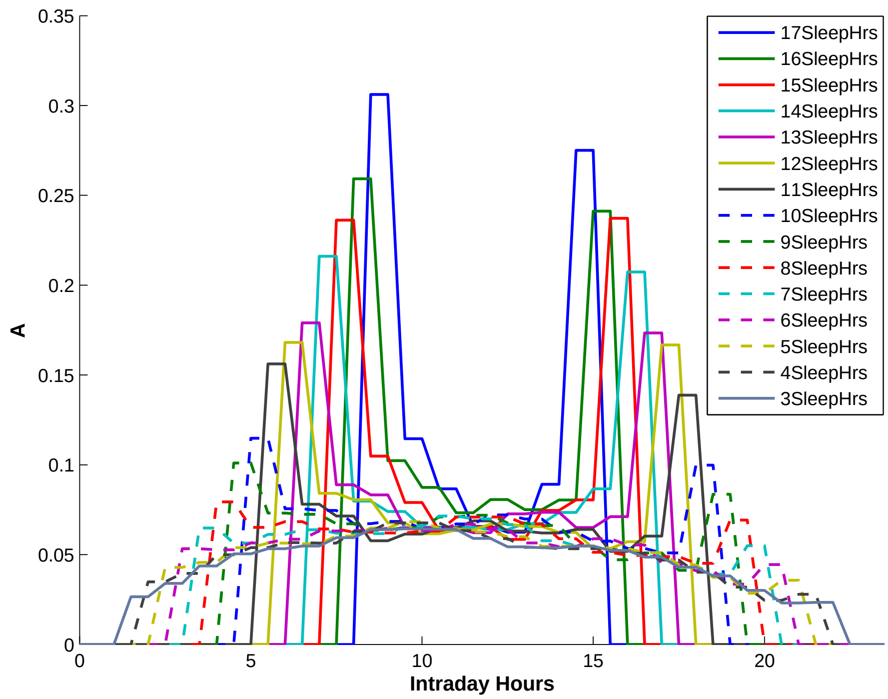

Figure 4.4: Time zone normalised FX intraday seasonality of the 15 sleeping groups over the 2.25 years of the data set period. Each curve represents the aggregated shifted intraday activities of all accounts in one sleeping group. Most of the sleeping groups show a similar behaviour, where traders trade intensively at the beginning and end of their trading days. A single U-shape can be observed in all groups except for traders sleeping from 3 to 4 hours a day.

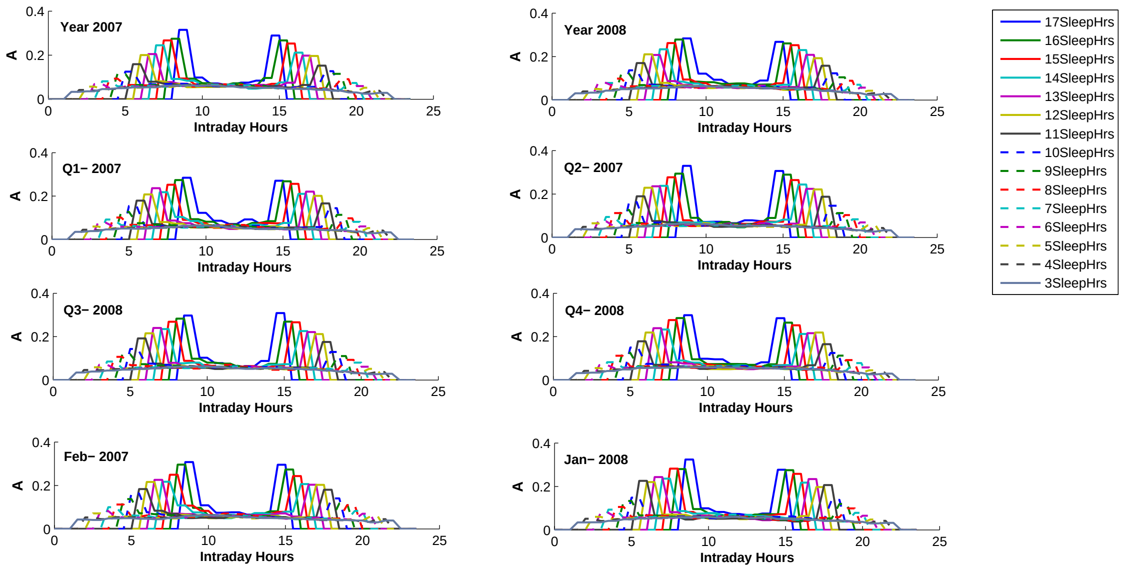

Figure 4.5: Time zone normalised FX intraday seasonality of the 15 sleeping groups over different time spans: yearly, quarterly, monthly. Each curve represents the aggregated shifted intraday activities of all accounts in one sleeping group. Most of the sleeping groups show a similar behaviour, where traders trade intensively at the beginning and end of their trading days. A single U-shape can be observed in all groups except for traders sleeping from 3 to 4 hours a day.

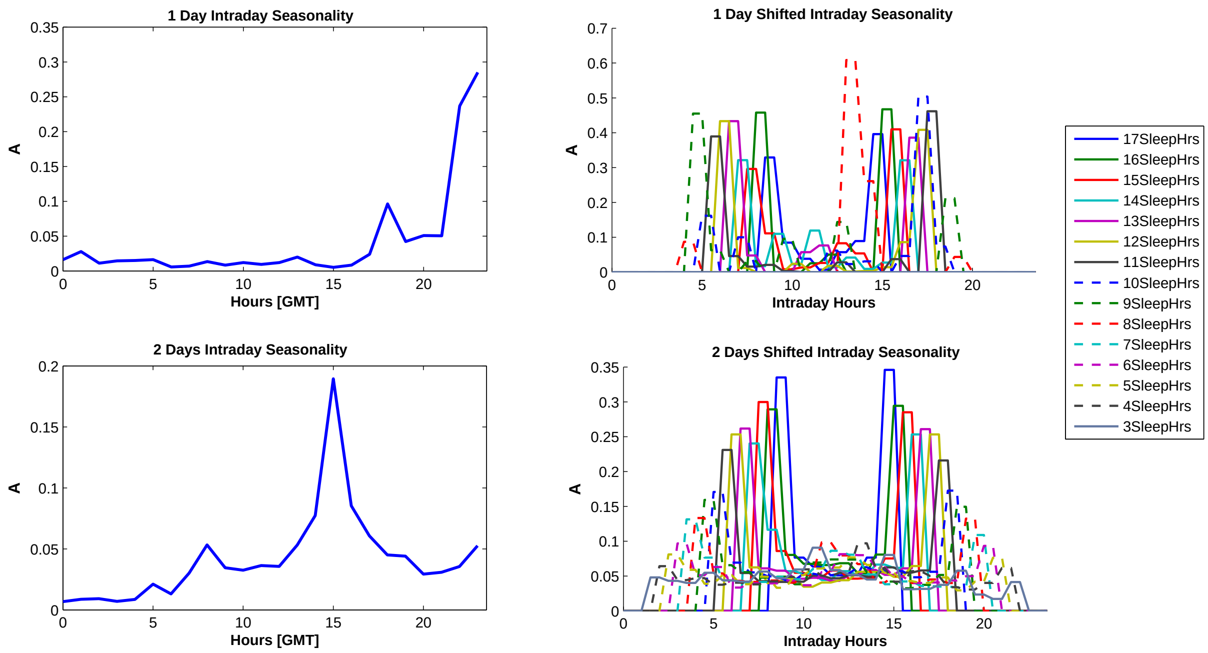

Figure 4.6: Minimal period to observing the intraday U-shape pattern. A period of one day in physical time (top) does not show the camel-shape or the normalised U-shape pattern. A minimum of two days period in physical time (down) is needed to observe these patterns.

4.3.2 Normalised Intraweek Seasonality

Similar to the intraday seasonality defined in eq 4.1, the intraweek seasonality represents the average number of all accounts' trades that occur in each hour of the week ($wh = 0, 1, 2, 3, 4, ..., 167$) over a certain period of time. The week starts on Sunday ($wh = 0 : 23$) and ends on Saturday ($wh = 144 : 167$).

To normalise the time zones of FX traders' intraweek seasonality we use the findings in section 4.3.1. We now have a priori knowledge about the trading activity behaviour of the different traders over different periods. First, we know the sleeping group of each account, and second we know the number of hours it has to be shifted to normalise its intraday activity. Using this information, we shift the original timestamps of each account's tick transactions (trades) by the required number of hours. After this time zone normalisation (shifting) process, we calculate the intraweek seasonality of each account. The normalisation of the intraweek traders' activities is described in details in Algorithm 4.2. Figure 4.7 displays the aggregated traders' intraweek seasonality in each sleeping group after time zone normalisation. Again, we observe two trading peaks at the beginning and end of each day during the week, confirming and validating the results in the previous section. The low activity on Saturday and on Sunday is due to the FX market operating hours described earlier, where FX trading starts Sunday evening in Sydney and ends Friday evening in New York.

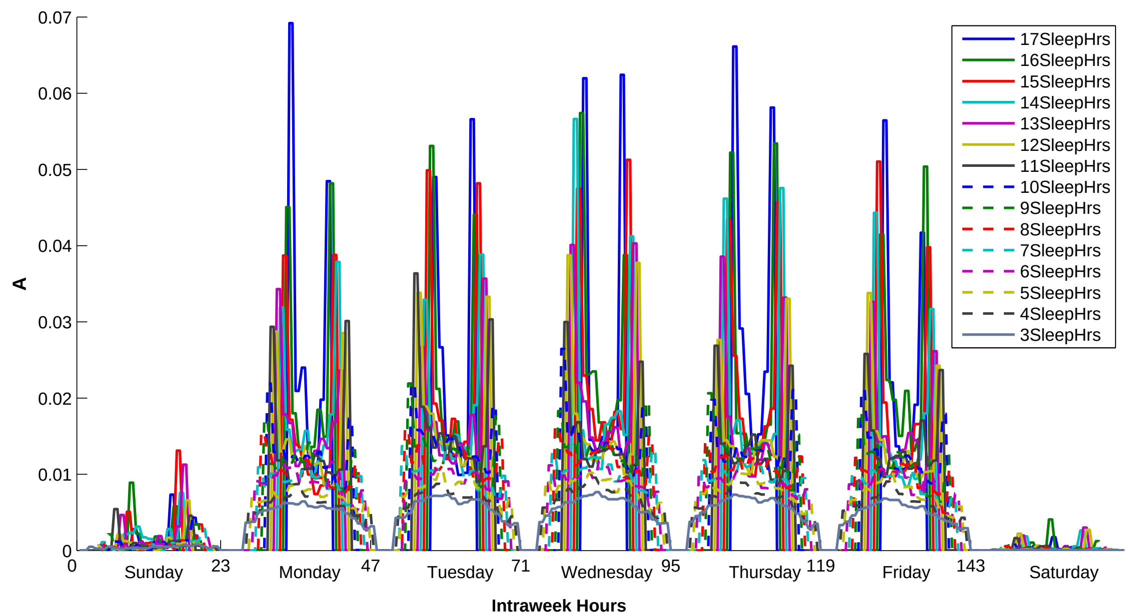

Figure 4.7: Normalised FX intraweek seasonality over the 2.25 years of the data set period. Each curve represents the aggregated shifted intraweek activities of all accounts in one sleeping group. A single U-shape can be observed in most of the sleeping groups. They show a similar behaviour, where traders trade intensively at the beginning and end of their trading days, except for traders sleeping from 3 to 4 hours a day. Low activity is observed on Saturday and Sunday due to the FX market operating hours.

Algorithm 4.2 Normalised Intraweek Seasonality

Input:
D //database accounts' trades records
$ShiftedAccounts = \{SA|SA = (ac, shift, sleepingGrp)\}$ */*extracted from Normalised Intraday Seasonality. Each SA record contains information about Account Id (ac), the nb of hours the account must be shifted to normalize its time zone (shift), and the sleeping group it belongs to (sleepingGrp).*/*
P //the time period (monthly, quarterly,yearly,..) over which to implement the algorithm
Output:
Normalised Intraweek Seasonality

Begin
For each $SA \epsilon ShiftedAccounts$ **do**
  $Trades_{ac} \leftarrow$ getAccountTrades($SA.ac$,D,P),
  where$Trades_{ac} = \{T|T = (id, timestamp, weekHour)\}$
  */*weekHour has a no value and is computed after time zone normalization.*
  weekHour={0,1,..167} represents the hour of the week at which a trade is placed./*

  //shift all trades timestamps to normalize time zones
  For each $T \epsilon Trades_{ac}$ **do**
    $\overline{timestamp}\leftarrow$shiftTradeTimestamp($T.timestamp$,$SA.shift$);
    */*compute the week hour after the shift. The shift takes place during the same original day in which the transaction takes place.*/*
    $SA.weekHour \leftarrow \text{getHour}(\overline{timestamp}) + ((\text{getWeekDay}(T.timestamp) - 1) * 24)$;
  End

  //calculate account's normalised intraweek seasonality using eq. 4.1
  For $wh \leftarrow 0$:167 *//week hours*

$$A[wh] = \frac{\sum_{d}^{n_d} N^d_{wh_i,wh_{i+1}}}{\sum_d^{n_d} \mathcal{N}_d}$$

  End
End *//end of for each*

End

4.4 Traders Excluded From Time Zone Normalisation

As mentioned in the previous section, we have excluded accounts sleeping more than 17 hours or less than 3 hours. To identify the impact of these accounts, we analyse their trading frequency, discussed earlier in chapter 3.

We generally find a strong relationship between the sleeping hours and the trading frequency of the different accounts. The lower the trading frequency of the accounts, the more sleeping hours they show. Figure 4.8 shows the percentage number (density) of the different sleeping groups accounts in each trading frequency class over a one-month trading period (January 2008). Class 'A' means high trading frequency, while class 'L' represents low trading frequency. The ranges of the different trading frequency classes have been previously described in Table 3.3. Those accounts exhibiting higher trading frequency sleep less than those accounts exhibiting low trading frequency. For instance, Figure 4.8 shows that 92% of the high frequency trading accounts (class 'A') belong to the '0SleepHrs' group in January 2008. The trading frequency class 'A', defined earlier in 3.3, contains high frequency accounts trading more than 500 but less than 1000 trades per day. However, the inverse relationship is not always true as shown in Figure 4.9. Each sleeping group can have a variety of accounts belonging to different trading frequency classes. For the '0SleepHrs' group, for instance, we can see other accounts of many other trading frequencies than class 'A' (Figure 4.9). This relationship between accounts' sleeping hours and their trading frequencies is the same across all months.

Relating this trading frequency to the accounts excluded from our analysis, we find that those accounts sleeping more than 17 hours represent only5% of all accounts, having very low trading frequency per day and trade for very short term periods; a couple of days during the month. Hence, the analysis of their behaviour is statistically irrelevant and ignored.

The second set of traders, sleeping less than 3 hours, has been also excluded from our analysis. Despite the diverse trading frequency of the accounts sleeping zero to two hours per day, they have been excluded due to their likelihood of being algorithmic traders. Identifying algorithmic traders from traditional traders requires heavy computation, as each algorithmic trading program has its own trading rules that need to be discovered. Although the identification of algorithmic traders is out of our research scope, however, we can easily extract a considerable amount of algorithmic trading accounts after implementing our 'Sleeping Hours' approach. We argue that traders having less than 3 sleeping hours during the day throughout the 2.25 years are algorithmic traders. The percentage number of accounts trading for 6 to 27 months with either zero, one or two sleeping hours during the day represent $\sim$ 46% of the total number of accounts in our data set. Thus, the effect of algorithmic traders on the aggregate behaviour of the FX market can be immense.

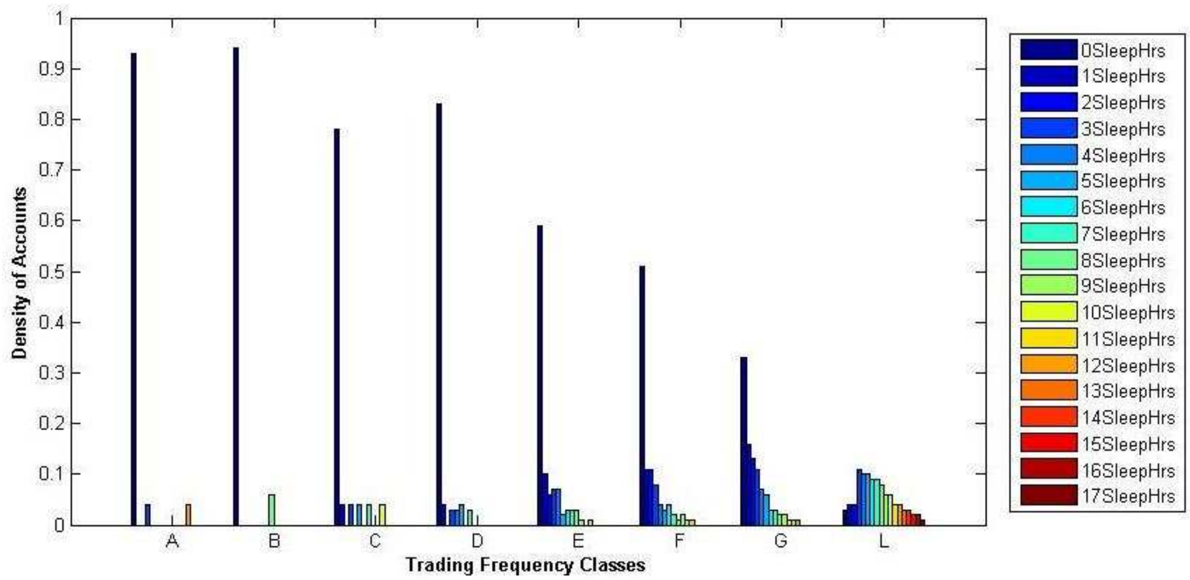

Figure 4.8: Density of sleeping groups accounts in the different trading frequency classes over a one-month period (January 2008). The same pattern can be observed across all months. The range of each trading frequency class (A,B..,L) is defined in Table 3.3. 'A' means high trading frequency; 'L' means low trading frequency. Class 'AA' is not used here, as it have been removed in the data set cleaning process in chapter 3.

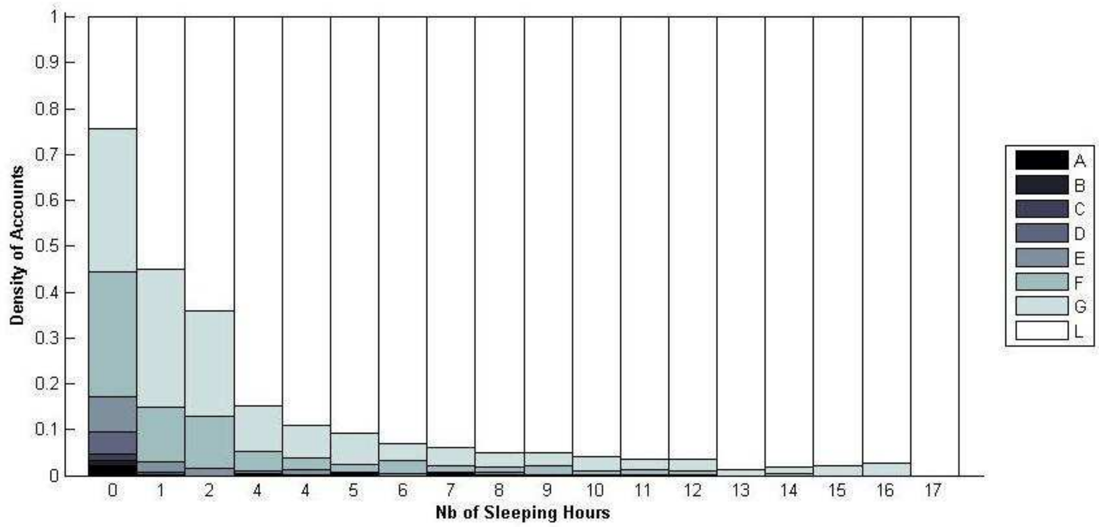

Figure 4.9: Density of trading frequency classes accounts in the different sleeping groups over a one-month period (January 2008). Class 'L' dominates most of the sleeping groups. The same pattern can be observed across all months

4.5 Aggregate Traders' Behaviour

So far we have examined the seasonal activity patterns of the individual accounts from a microscopic perspective. We have categorized them into sleeping groups and we have studied the time zone normalized FX intraday and intraweek seasonality of each group. We show that they do not differ from equity traders in terms of their U-shape trading behaviour, with high trading intensity at the beginning and at the end of the day.

In this section, we study the aggregate intraday (intraweek) seasonality of all sleeping groups accounts before and after time zone normalisation over the 2.25 years of the data set period. The aggregate intraday trading pattern is computed as the average seasonality (hourly share of activity A) of all sleeping groups intraday activities depicted in Figure 4.4. The aggregate intraweek pattern is computed in the same way from sleeping groups intraweek seasonalities illustrated in Figure 4.7. For the non-normalized aggregate seasonality, we aggregate the sleeping groups accounts intraday (intraweek) activities before time zone normalisation.

Figure 4.10 depicts the aggregate intraday and intraweek activity patterns of all sleeping groups before and after time zone normalisation. Taking first the non-normalised seasonalities, we can clearly see the double U-shape pattern for the intraday as well as for each day of the intraweek aggregate seasonality. Saturday and Sunday in the intraweek seasonality are exceptions and show low activity due to the FX market opening hours, where FX trading starts on Sunday evening in Sydney and ends on Friday evening in New York.

However, if we aggregate sleeping groups seasonalities after time zone normalisation, we find a single U-shape pattern for both intraday and intraweek activities. Thus, we empirically demonstrate that the double U-shape pattern is indeed an overlap of single U-shape patterns characterizing single market participants trading in different currencies.

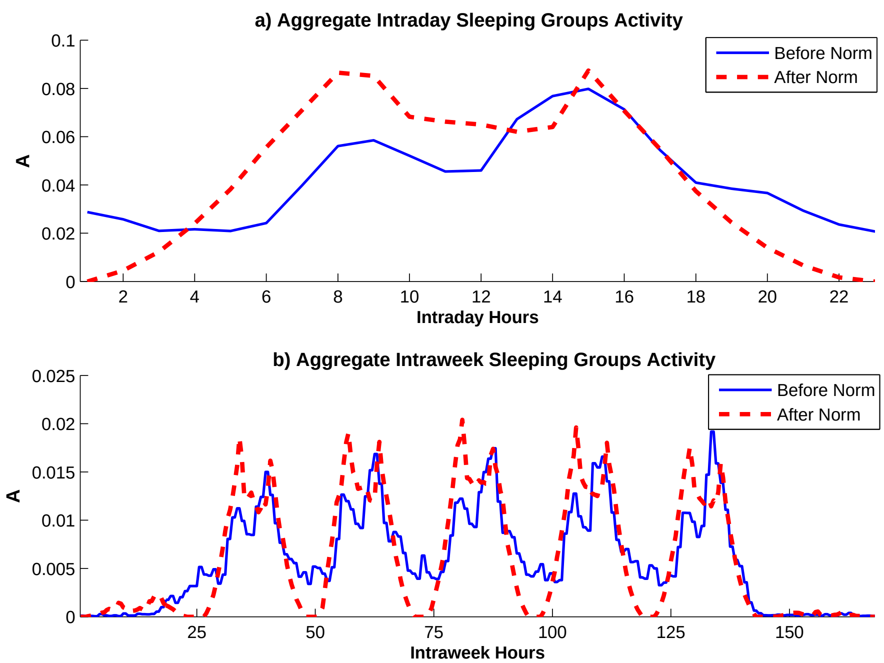

Figure 4.10: Aggregate seasonality of all sleeping groups before and after time zone normalisation over the 2.25 years of the data set period. Both a) intraday and b) intraweek aggregate sleeping groups activity exhibit a double U-shape before and a single U-shape after time zone normalisation. Saturday and Sunday in the intraweek seasonality are exceptions due to the FX market opening hours.

4.6 Decomposing Traders' Activity Patterns

In the previous chapter, we have tracked individual traders' activities from opening to closing their positions to validate the data set quality. This required identifying when each individual trader opens, manages or closes a position or just switches from one position (long/short) to another one.

In this section, we decompose FX traders' intraday seasonality over the 2.25 years into these four types of trades: opening, closing, managing and switching trades (defined in section 3.4.3). We follow the same intraday FX seasonality normalisation procedure explained in section 4.3.1 and described in Algorithm 4.1. The only difference here is that we first differentiate between the types of trades of each account before calculating the intraday seasonality and applying the time zone normalisation process. So, instead of computing the seasonality of all trades placed by the individual account, we calculate and normalise the seasonality of the account's opening, closing, managing and switching trades separately.

The decomposition of the intraday activity into these four types is illustrated by (Figure 4.11) revealing more information about traders' behaviour.

For all sleeping groups, we can observe the different types of trades. Switching trades however, represent the smallest percentage for all groups. As for the managing trades, traders belonging to short sleeping groups (3-5 sleeping hours) seem to use them more often. In these groups, opening, closing and managing trades move closely together, indicating how traders in these groups must be closely monitoring and manipulating their trading positions. This can also explain the relatively high trading frequency of the accounts belonging to short sleeping groups. In contrast, traders with longer sleeping hours rely more on opening and closing trade types.

Similar to the composed activity in Figure 4.4, we find intensive trading activity (peaks) at the beginning and end of the trading day and a relatively lower activity in the middle of the day for closing, opening and managing trades. However, this intensity declines with accounts having shorter sleeping periods, and an opposite pattern appears, where the highest trading intensity takes place in the middle of the day. Despite the different intensity of the trading activity between the short and long sleeping groups, the pattern of the last trading hour is consistent for all traders. Traders go more for closing their open positions rather than opening new ones at the end of the day.

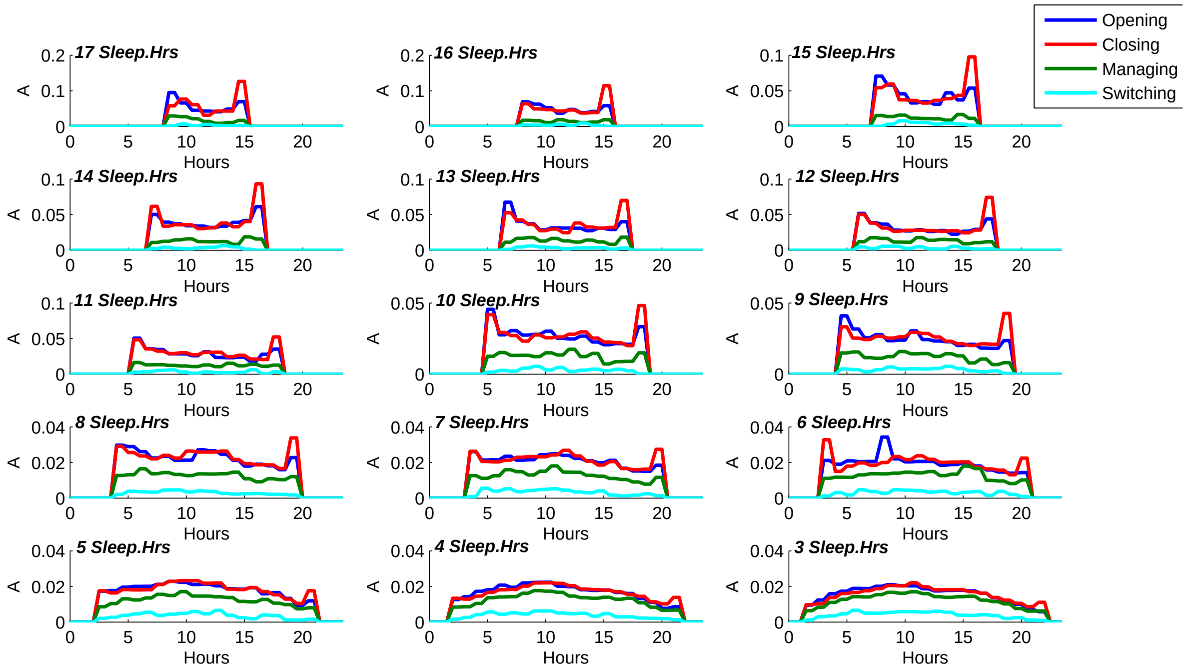

Figure 4.11: Intraday activity of closing, opening, managing and switching trades for the 15 sleeping groups

4.7 Conclusion

In this chapter, we have conducted an empirical microscopic analysis of the activity of FX traders. Using the data set produced in chapter 3, we have studied the intraday and intraweek seasonalities of more than 45,000 individual accounts.

It is well known that FX and equity markets have different seasonalities. The assumption for these different behaviours has been referred to the unique nature of the FX market, where traders can access the market 24 hours a day from different geographical locations. This is very different to the typical equity exchange that has specific opening hours in a particular time zone.

We provide empirical evidence on the validity of this assumption. We normalize traders' time zones by applying an event-based algorithm that disregards the transactions physical time, and synchronize traders' seasonal activities according to their inactive (sleeping) hours during the day. We show that the well-accepted double U-shape pattern of FX market intraday seasonality is mainly due to the different time zones market participants operate from. However, once normalised, a remarkable single U-shape pattern reveals, similar to the one observed in equity markets, where traders trade intensively at the beginning and end of the trading day. To the best of our knowledge, no one has been able to empirically demonstrate that the double U-shape pattern is indeed a result of an overlap of single U-shape patterns characterising the behaviour of single participants operating from different time zones. Similar results are shown for the intraweek analysis.

Using this finding, we further analyse the reason for the trading peaks at the beginning and end of the trading day by decomposing traders' intraday activities into opening, closing, managing and switching trades. We find that regardless of the number of hours traders are inactive during the day, they all tend close their positions at the end of the trading day. As for the peak at the beginning of the day, it is a result of both opening and closing trades.

Possible future work can be exploring mathematical functions relating single with double U-shaped seasonalities. Generally, adopting a bottom-up perspective where an accurate modelling of traders, and thus their geographical location, are of primary importance, we believe that the results reported here represent an additional step in the quest of unravelling the dynamics of the FX market.

Chapter 5

Deciphering FX Activity Along Intrinsic Time

In this chapter, we conduct a microscopic analysis of FX market activity patterns along price movements. Given the high frequency and discontinuous nature of FX data, we use an intrinsic time scale for market activity synchronization. This time scale is defined by exchange price turning points of a pre-specified threshold. Using this event-based time approach, we are able to decipher market activity during price movement events of any given scale. We provide empirical evidence of diminishing market liquidity at the end of the price movement, the overshoot period, for all currency pairs studied. We find that a price overshoot stops due to more participants placing counter trend trades, a finding that is valid across all magnitudes of price movement events. We also find that this market liquidity pattern has an impact on price ticks moves, which follow the same decaying pattern during overshoot periods. The results in this study have been only enabled through utilizing the intrinsic time scale. We show that intrinsic time is a powerful tool that can discover recurring market behavioural patterns, which are otherwise hidden, if we use physical time only

.

5.1 Introduction

As described earlier in chapter 2, the high frequency FX market is the world's most liquid financial market with a daily volume of 1.4 trillion USD of spot transactions, representing a microscopic transaction volume of 16.2 million transaction per second [20]. With this second by second small transaction volume, any small imbalance in the buying or selling volume can trigger disproportionate price moves affecting the global FX market. In an attempt to understand the impact of microscopic activity on the global FX price move, in this chapter we conduct a microscopic analysis of FX traders' activity with price movements. To carry out this study, we use the unique high frequency transaction data set produced in chapter 3. Given the unevenly spaced nature of high frequency data, we base our analysis on an event-based time scale, the so called intrinsic time, which is determined by the price time series itself. The notion of time here is simply an index that determines the position of an event occurrence within a set of asynchronous events [112, 146].

Following the same definition of intrinsic time adopted by [31, 49, 87, 146, 151], we use the directional change events approach. It focuses on exchange price turning points at a pre-specified fixed threshold of any size. A price overshoot (*os*) is usually associated with each directional change event. An overshoot is a continuation of the price movement in the same direction as the directional change event, but in excess of the targeted threshold size. Hence, the total price move is composed of a *dc* section and an *os* section.

One of the interesting stylized facts reported by [87], is the scaling law of the relationship between the directional change and the overshoot sections of the total price move. A directional change is followed by an overshoot of the same magnitude, making the total move double the size of the directional-change threshold. Given this information, [31] used the discovered scaling law as a metric for the dynamic market equilibrium in terms of rates of price change. The market is off equilibrium, once the overshoot exceeds the expected size. Thus, the overshoot becomes an indicator of the balance of supply and demand. By analysing the extent of market price deviation from its equilibrium, [31] design the so called Scale of Market Quakes (SQM) measure. It quantifies the impact of price events in the FX market, occurring in response to news announcements.

Despite the work done by [31, 87], the overshoot period has not been explored in depth yet; leaving many open questions about market behaviour during this period. The literature does not report any empirical evidence on when or why a price overshoot stops, and whether this can be related to a certain activity of market participants that is consistent across all magnitudes of price events .

To approach this question, we decipher the overshoot period by synchronizing FX traders' activity with price movement events of different magnitudes across several currency pairs. Our study reveals new results about FX market dynamics. We provide empirical evidence of decaying market liquidity along the price movement event. Market liquidity reaches its minimum at the end of the overshoot period, a finding that is valid across all scales and for all studied currency pairs. Both counter trend and with trend traders decrease their activity at the end of the overshoot period of the price movement event. However, counter trend trades (CTT) have a higher intensity relative to with trend trades (WTT) at the end of the price move. Thus, we can claim that a price overshoot period stops due to more participants placing counter trend trades, representing a stabilizing force. In other words, given this diminishing market liquidity, placing a small WTT order can drive the market far away in terms of fueling the underlying price movement, and hence making its overshoot period longer.

We also provide empirical evidence for the strong impact of liquidity on the patterns of price ticks changes along the intrinsic time scale. We find that price ticks follow the same decaying pattern during the overshoot period. This is similar to recent findings in equity markets [74, 183], which show that price changes are affected by liquidity fluctuations, and not by the volume of orders, which has been a common conjecture adopted in previous research [36, 47, 84].

The rest of this chapter is organized as follows: Section 5.2 describes the data set used for this study. Section 5.3 defines the intrinsic time scale. Section 5.4 discusses the discretisation of the internisic time scale into further events of smaller thresholds. Section 5.5 describes the methodology of synchronizing market activity with intrinsic time, which is then implemented in section 5.6 to decipher FX market behaviour during price overshoots. Section 5.7 relates price ticks patterns with market liquidity. Section 5.8 highlights the main conclusions.

5.2 The Data Set

To analyse market activity along price movement events, we use two data sets. The first is our unique high frequency transaction data set produced in chapter 3, which covers the transaction history of more than 45,000 accounts, on an account level, trading on the OANDA FX trading platform over a 2.25 year period.

The second are the price series which are used to compute the different price movement events to build an intrinsic time scale, along which we synchronize market activity.

Thanks to OANDA, we have been provided access to the tick price series of four major currency pairs in the FX market: EUR/USD, AUD/USD, GBP/USD and USD/JPY spanning three years from June 2006 up to June 2009. However, since our transaction data set covers only the time period from January 2007 to March 2009, we only focus on the same time span for the price series. For each price series, we know the tick timestamp, the bid price (trader's sell rate) and the ask price (trader's buy rate). Figure 5.1 depicts the price series of the four currencies for a sample period of one month.

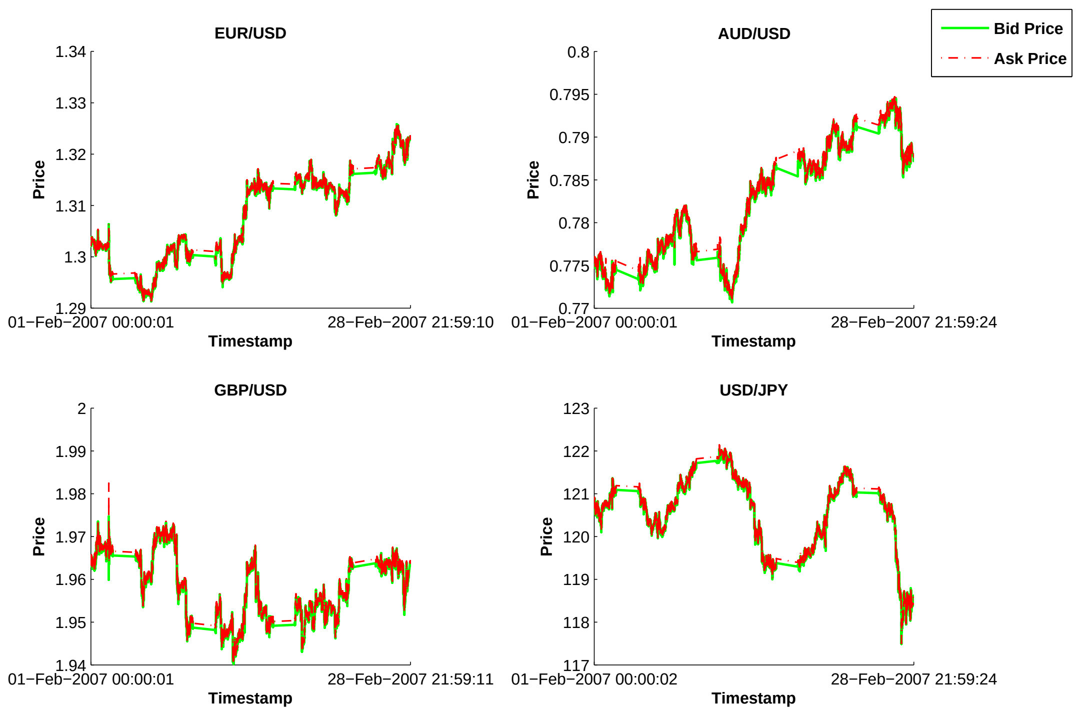

Figure 5.1: Tick price series for the currency pairs: EUR/USD, AUD/USD, GBP/USD and USD/JPY for a sample period of one month from 1-Feb-2007 to 28-Feb-2007.

In chapter 4, we show that both the hourly share of OANDA traders daily activity and the hourly share of the daily number of price ticks go in line with the previous literature [49, 108] and exhibit a double U-shape pattern. Peaks occur during the day as a result of the overlapping trading sessions around the globe (Figure 4.1). Generally, with the conclusions already drawn in chapter 3 and chapter 4 that the data set is representative of the FX market, we can confirm that the results of this study describe the real behaviour of the FX market.

5.3 Intrinsic Time

Time has been always a major topic of research in philosophy, physics and all other disciplines of the natural sciences. The traditional notion of time is defined as "*..a universal clock ticking independently of matter and space, just tracking all processes in the universe*" [5]. However, the universe has different systems, where each has its own processes and events. The occurrence of these events is the cause for time motion, making the system change from one state to another [5]. Using a global variable like physical time, that has exactly the same value at any level of the system, can totally distort the analysis of a dynamic stochastic discontinuous system like the high frequency FX market. The presence of a large variety of traders sleeping at different times, having different objectives and trading strategies causes a discontinuous flow of physical time. In other words, each clock tick of the uniformly progressing physical time does not necessarily correspond to a transaction event in the financial market. Therefore, another approach has been introduced by [146], in which the time scale in financial markets is defined via events, the so called intrinsic time.

Using this convention, we use intrinsic time to discretise and understand the high frequency FX market. Adopting the same definition introduced by [31, 49, 87, 93], we define intrinsic time in terms of price movement events.

An event represents the total price movement (TPM) between two extreme price levels (extrema), expressed as a relative price jump of a pre-specified threshold size. This movement can be decomposed into a directional change section (dc), price reversion, and an overshoot section (os). Algorithm 5.1 explains how directional change events are identified. A directional change occurs once the current price changes by a specified threshold $\triangle x_{dc}$ from the last extreme price level x^{ext}(extrema). The extrema is the last high price in the up mode or the last low price in the down mode. Once a directional change occurs, an overshoot is usually associated. It is a continuation of the price move in the same upward or downward mode of the underlying dc, but one that exceeds the specified threshold $\triangle x_{dc}$. The overshoot is determined as the difference between the price level at which the last directional change x_{LastDc} occurred and the current extrema at which the new directional change occurs, i.e., the last high when in the up mode or the last low when in the down mode. The last high and last low price levels are then reset to the current price and the up/down mode swaps. To ensure the accuracy of our algorithm, we do not use the mid price $((bid+ask)/2)$ as the current price. Instead, we use the inner price , which alternates between bid x_{bid} and ask x_{ask} prices based on the existing direction mode. Figure 5.2 illustrates the implementation of an intrinsic time scale, defined by directional change events of $\triangle x_{dc}$= 0.4% for the one month projection of EUR/USD price sample depicted in Figure 5.1.

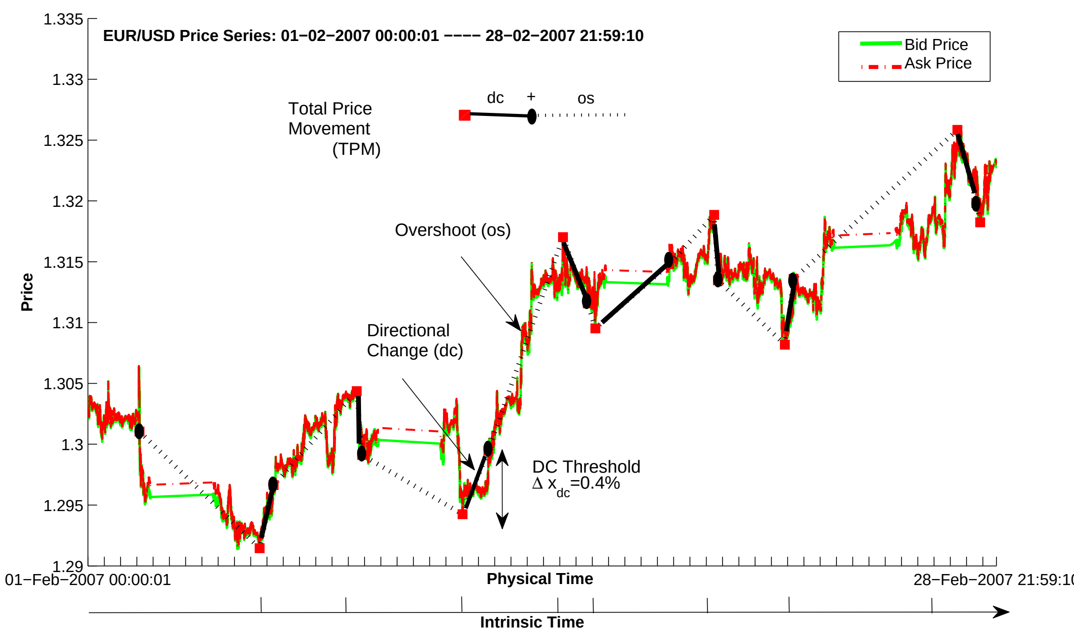

Figure 5.2: Projection of a one month EUR/USD price sample and its associated price movement events defined by a threshold of $\triangle x_{dc} = 0.4\%$. Each price movement (TPM) is composed of a directional change (dc) and an overshoot (os) section. The start of the dc section is denoted by the red rectangle, while the start of the os section is denoted by the black ellipse. The price of extreme price levels alternates between the bid and ask prices depending on the direction mode. See Algorithm 5.1 for more details. The intrinsic time scale triggers with the start of each directional change event (red rectangle), whereas the physical time ticks evenly across the price curve independent of any scale or price pattern.

Algorithm 5.1 Directional Change Events

Input:

//initialize variables

$os \leftarrow 0$; $dc \leftarrow 0$; $\triangle x \geq 0$ (fixed); mode is up ;

$x^{ext} \leftarrow x_0$; $x_{lastDc} \leftarrow x_0$; *//where x_0 is the first price tick (bid or ask)*

Output:

Directional change events dc and associated overshoots os

Begin

//look for new extrema x^{ext}

If (mode is up) Then

If $x_{ask} < x^{ext}$ Then

$x^{ext} \leftarrow x_{ask}$;

Else if (mode is down) Then

If $x_{bid} > x^{ext}$ Then

$x^{ext} \leftarrow x_{bid}$;

//look for a directional change tick dc. //x_p is the inner price, where p is bid if mode is up or ask otherwise

If (((mode is up) and (x_p - x^{ext})/x^{ext} $\geq \triangle x$) or ((mode is down) and (x^{ext}-x_p)/x^{ext} $\leq$-$\triangle x$)) Then {

os$\leftarrow(x^{ext}$-$x_{lastDc})/x_{lastDc}$;

dc$\leftarrow(x_p$-$x^{ext})/x^{ext}$;

If (mode is up) Then {

$x_{lastDc} \leftarrow x_{ask}$; $x^{ext} \leftarrow x_{ask}$; mode $\leftarrow$down;

}

Else if (mode is down) Then {

$x_{lastDc} \leftarrow x_{bid}$; $x^{ext} \leftarrow x_{bid}$; mode $\leftarrow$ up;

}

}

End

The directional change events approach was first introduced by [93], who discovered the scaling law of the fixed relationship between the average price move, or the so called directional change event, and the time interval over which the price move is measured. They related the number of the directional changes $N(\triangle x_{dc})$ to the directional change sizes $\triangle x_{dc}$

$$N(\triangle x_{dc}) = \left(\frac{\triangle x_{dc}}{C_{N,dc}}\right)^{E_{N,dc}} \tag{5.1}$$

where $E_{N,dc}$ and $C_{N,dc}$ are the scaling law parameters. According to this formulation, time varies and the price move threshold is fixed. Hence the same proportional relationship holds for different time intervals ranging from intraday to yearly intervals, indicating the close relationship between short and long term price moves. The work by [93] was the first attempt to overcome the complexity of real time series, focusing only on focal points by defining an activity-based time scale instead of using physical time.

Extending this directional event approach in [87] enabled the authors to unveil 12 new scaling laws holding across 13 exchange rates for up to three orders of scaling magnitude; from intraday to inter-day and longer intervals. Generally, scaling laws have always played a robust role in the natural sciences and have now become an important research area in finance. It is used to facilitate building up and calibrating time series models, characterizing both the short and the long term stylized facts of price time series.

However, the discovery of these scaling laws has only been enabled through the utilization of intrinsic time, or the directional change event approach. It reduces the complexity of time series by focusing on turning points and removing irrelevant details of the stochastic price process. This way, intrinsic time adjusts with the discontinuous nature of the FX market. Moreover, both seasonalities and nonlinearities [49, 107, 108, 138] in the FX market are handled and accounted for by this time scale transformation [40, 94]. Finally, intrinsic time enables a close monitoring of market activity on any given scale and maps how market participants view the market.

5.4 Discretising Intrinsic Time into Sub-clocks

In the previous section, we have defined an intrinsic time scale in terms of price movement (directional change) events. A price movement event TPM, of any threshold size

$\triangle x_{dc}$, is composed of a directional change dc and an overshoot os section. The number of price movements (directional change) events composing the intrinsic time scale depends on the threshold size$\triangle x_{dc}$. The larger the threshold size $\triangle x_{dc}$ of the directional change, the fewer $\triangle x_{dc}$ events we find. This has been already empirically proven and mathematically formulated (eq. 5.1) as a scaling law by [93] and confirmed again by [87]. The latter shows that this scaling law between the count number of directional change events and the threshold size $\triangle x_{dc}$ holds for tick-by-tick price series of 13 currency pairs spanning five years.

Taking this into consideration, we have to decide on the most appropriate $\triangle x_{dc}$ to build up our intrinsic time scale along which we can synchronize the market activity. According to [87], an exchange rate usually exhibits a 10% to 20% movement within a year. However, in order for this to occur, the length of the price curve coastline, which is approximately the sum of all absolute price moves of a given threshold, can be extremely long. For instance, [87] show that the coastline measures on average 6.4% per day given an 0.05% threshold. Meanwhile, a price move of 0.6% is the mean maximal move that can be observed within 24 hours, and it takes on average 220 days for a price move of 6.4% to be measured. This holds across the 13 currency pairs studied in [87], indicating the importance of considering both extreme events as well as the various smaller events preceding and leading to these extremes.

Given this information, using a relatively large or medium threshold size $\triangle x_{dc}$ to build up our intrinsic time scale will hide information about smaller yet important events. Similarly, using a very small threshold size will reveal a great deal of information about price series behaviour even over short time periods. However, it can lead to a blurred picture about the general price trend. This is illustrated in Figure 5.3, where we project a six day EUR/USD price sample onto two different intrinsic time scales. In Figure 5.3(a), the intrinsic time scale is defined by a threshold of $\triangle x_{dc} = 0.4\%$, where the general price trend, with one upward and one downward price move, over the specified time span, is clearly revealed. Figure 5.3(c), in contrast, uses a much smaller threshold size ($\triangle x_{dc} = 0.05\%$) for the intrinsic time. It gives a better zoomed-in picture of the price patterns, showing several upward and downward events. However, it gives no information about the general trend of the price curve.

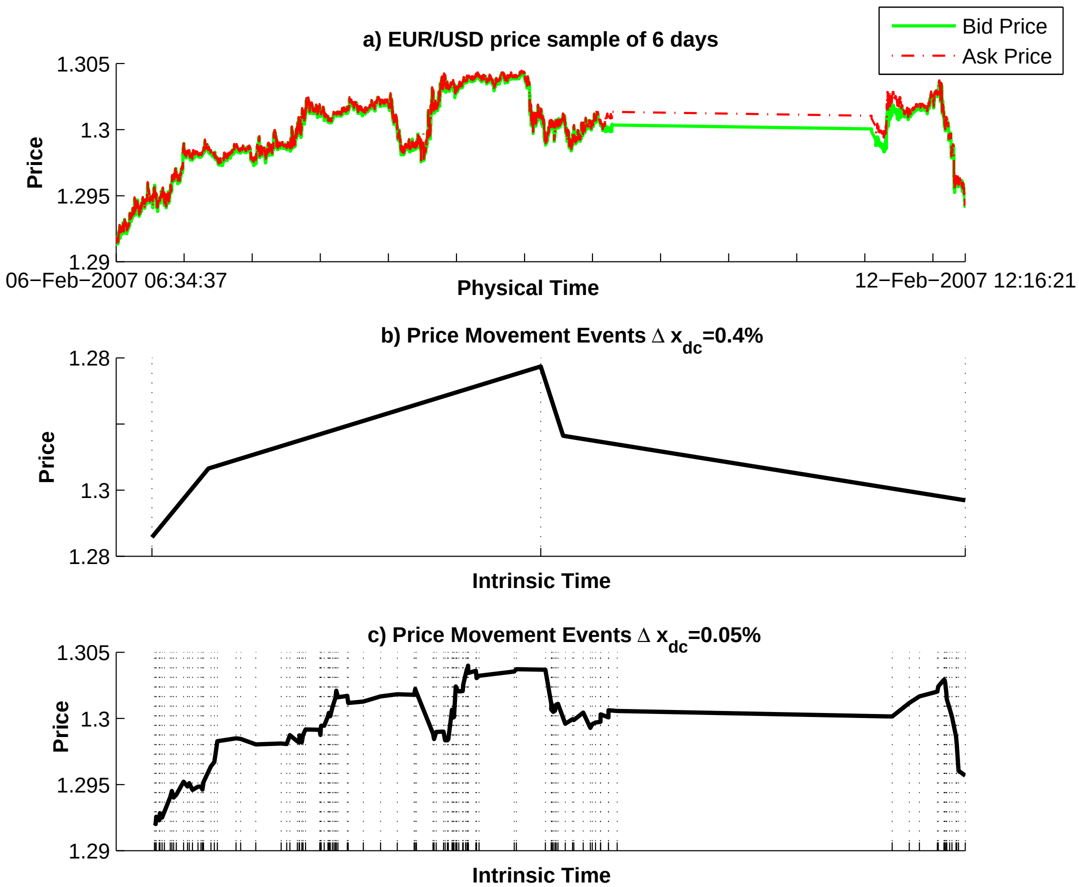

Figure 5.3: Projection of a six day EUR/USD price sample and its associated price movement events of different thresholds. Top graph (a) depicts the bid and ask prices of EUR/USD over a physical time span ranging from 06-02-2007 06:34:37 to 12-02-2007 12:16:21. Graph (b) captures price movements in (a) over an intrinsic time scale defined by a threshold of $\triangle x_{dc} = 0.4\%$. It shows two directional change events (one upward and one downward). Graph (c) captures price movements in (a) over an intrinsic time scale defined by much smaller threshold of $\triangle x_{dc} = 0.05\%$. It provides more information, showing several upward and downward price movement events. However, the general trend of the price series over the six day period is not any more visible as it is in (b).

Therefore, to benefit from different threshold sizes at the same time, we define our intrinsic time scale using an$\triangle x_{dc}$ threshold that is further discretised into smaller events, referred to as sub-clocks. A **sub-clock** can be of any threshold size$\triangle x_{sc}$, where $\triangle x_{dc} > \triangle x_{sc} > 0$. Sub-clock events are computed in exactly the same way as directional change events in Algorithm 5.1. However, we call them sub-clocks as they are aligned with the main total price movement event TPM.

Figure 5.4 shows how we discretise the first upward price movement event (TPM: $\triangle x_{dc} = 0.4\%$) depicted in Figure 5.3 into smaller sub-clocks of different threshold sizes of $\triangle x_{sc}$ (0.2 %, 0.1%, 0.05%). It is important to note that sub-clock events exhibit the same scaling law of eq. 5.1. Hence, the smaller the threshold size of $\triangle x_{sc}$, the more sub-clock events we have during each price movement event. This is illustrated by Table 5.1, which shows the total number of events in EUR/USD of threshold size $\triangle x_{dc} = 0.4\%$ and the total number of associated sub-clocks over the whole data set period (2.25 years).

This way, we can carry out an in-depth analysis and derive rhythmic patterns of market activity on a microscopic level. Moreover, we can synchronize the price movements of different thresholds $\triangle x_{dc}$ using the same sub-clock threshold size $\triangle x_{sc}$.

Table 5.1: Total number of price movement events of $\triangle x_{dc} = 0.4\%$ and the total number of associated sub-clocks that occurred over the data set period (2.25 years). For the sub-clocks events, we use different threshold sizes of $\triangle x_{sc}$(0.2%, 0.1% and 0.05%)

$\triangle x_{dc}$	**Total nb of events**	$\triangle x_{sc}$	**Total nb of sub-clocks**
0.4%	1,762	0.2%	6,430
		0.1%	22,068
		0.05%	69,284

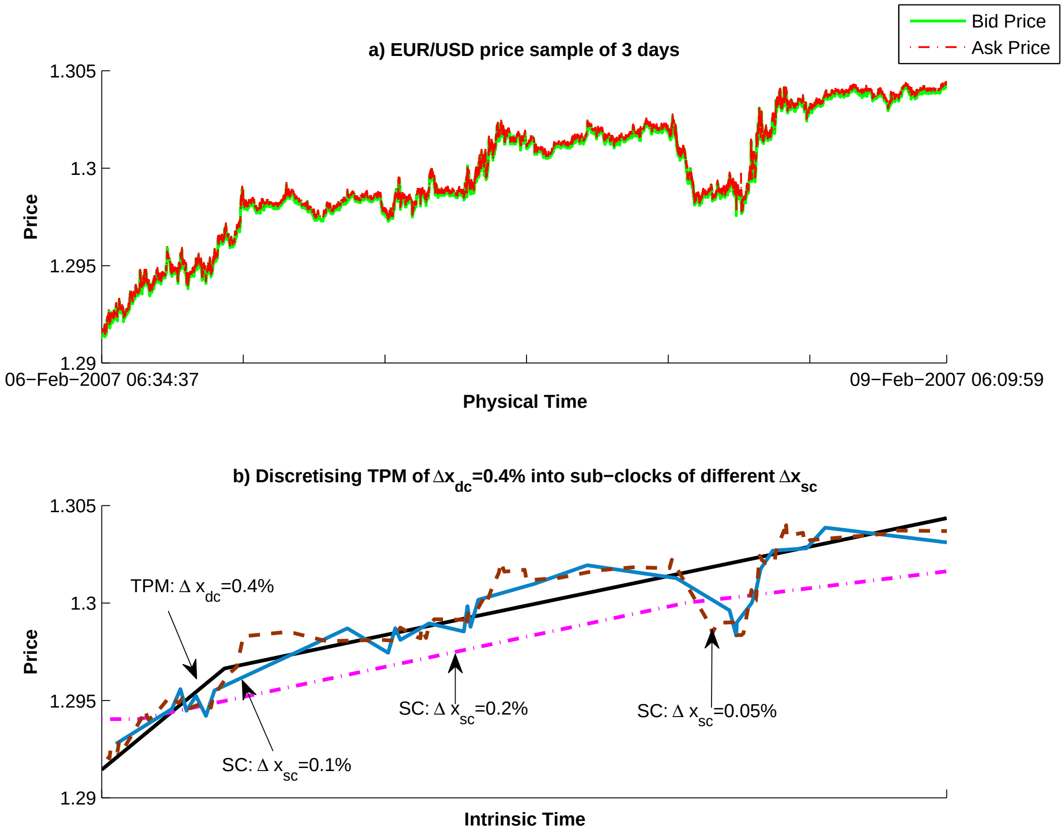

Figure 5.4: Discretising the first TPM (upward move) of Figure 5.3(b) into smaller sub-clocks of different $\triangle x_{sc}$ thresholds. Graph (a) depicts the bid and ask prices of EUR/USD over the physical time span, along which the first TMP $\triangle x_{dc} = 0.04\%$ occurred. Graph (b) captures the total price movement TPM (solid black curve) of threshold $\triangle x_{dc} = 0.4\%$ and discretise it into smaller sub-clocks of thresholds $\triangle x_{sc}$ (0.2%, 0.1%, 0.05%).

5.5 Synchronizing Market Activity with Intrinsic Time

In this section, we explain the methodology for synchronizing market activity with intrinsic time. We first define market activity and classify it into counter trend and with trend trades. We then describe how we synchronize market activity with the discretised intrinsic time scale explained in the previous section.

5.5.1 Defining Market Activity

Since the Eighties, several economists have conducted a number of surveys to study the behaviour of market participants in foreign exchange markets, showing that the majority of traders do not rely only on fundamental information, but also on technical trading. The literature provides a long stream of research studies and empirical evidence on the reliance of FX traders on technical analysis and speculative behaviour, especially over short and medium time horizons [4, 46, 81, 82, 98, 125, 132, 140, 141, 143, 148, 158, 176, 194].

In this chapter we define market activity as the number of trades placed by either contrarians or trend followers, the two major types of technical traders in the FX market. Trend followers are those traders who buy when the price (dc event) goes up and sell when the price (dc event) goes down, whereas contrarians buy when prices are low and sell when they are high. However, since a trader might already be holding a long or a short position in the underlying currency, the type of trade placed to follow or counter the trend depends on the position being held. For instance, a trend following trader who designates the trend as "up" does not necessarily imply that he/she should be holding a long position. It does however mean that he/she should not be holding a short position.

In chapter 3, we have tracked OANDA traders' positions from opening to closure for 2.25 years to validate the quality of the data. As explained earlier in section 3.2.3, OANDA offers its traders the following different types of trades to manage their positions in the market:

- **SellOrder (BuyOrder) trade**: A SellOrder SO (BuyOrder BO) increases the short (long) position while decreases the long (short) position.
- **CloseTradeB (CloseTradeS) trade**: A CloseTradeB CTB (CloseTradeS CTS) decreases a long (short) position by closing one previous open buy (sell) order of the same size.
- **ClosePositionB (ClosePositionS) trade**: A ClosePositionB CPB (ClosePositionS CPS) trade closes out the total open long (short) position of the account in a certain currency pair.

Taking these trades types into consideration, we classify market activity into two groups:

1. **Counter Trend Trades (CTT) :** If the price is moving up (dc>0), trade types that increase the short position (SellOrder) or decrease the long position (SellOrder, CloseTradeB, ClosePositionB) are placed. If the price is moving down (dc<0), trade types that increase the long position (BuyOrder) or decrease the short position (BuyOrder,CloseTradeS,ClosePositionS) are placed.

$$CTT = \begin{cases} SO \cup CTB \cup CPB & if\, dc > 0 \\ BO \cup CTS \cup CPS & if\, dc < 0 \end{cases}$$

2. **With Trend Trades (WTT) :** If the price is moving up (dc>0), trade types that increase the long position (BuyOrder) or decrease the short position (BuyOrder, CloseTradeS, ClosePositionS) are placed. If the price is moving down (dc<0), trade types that increase the short position (SellOrder) or decrease the long position (SellOrder, CloseTradeB, ClosePositionB) are placed.

$$WTT = \begin{cases} BO \cup CTS \cup CPS & if\, dc > 0 \\ SO \cup CTB \cup CPB & if\, dc < 0 \end{cases}$$

5.5.2 Synchronization Methodology of Market Activity

To analyse traders' behaviour along price movements, we need to synchronize their activity with each sub-clock of each price movement event of a given scale. As defined in the previous section, traders' activity is classified into either counter trend or with trend trades. So, for each sub-clock we count the number of CTT/WTT placed by traders, searching for any significant behavioural patterns along the sub-clocks of each total price movement event (TPM). This way, we can also differentiate between traders' activity during the directional change and the overshoot sections of TPM.

However, sub-clocks are not necessarily synchronized with the starting and ending points of the dc and os sections of the total price movement TPM. This is illustrated by Figure 5.5, which depicts only one TPM of $\triangle x_{dc} = 0.4\%$. This is the same price movement event used in Figure 5.4. The price move is discretised into smaller sub-clocks of $\triangle x_{sc} = 0.1\%$, where each sub-clock can have one of the following locations within the TPM :

- Within the dc section of the TPM, such as $SC_{t_1 - t_2}$ in Figure 5.5.

- Between the *dc* section and the *os* section of the TPM, such as $SC_{t_7-t_8}$ in Figure 5.5. As illustrated, the part of the sub-clock that takes place in the dc section is from t_7 to t_8, whereas the part that takes place in the os section is from t_8 to t_9.

- Within the *os* section of the TPM , such as $SC_{t_{13}-t_{14}}$ in Figure 5.5.

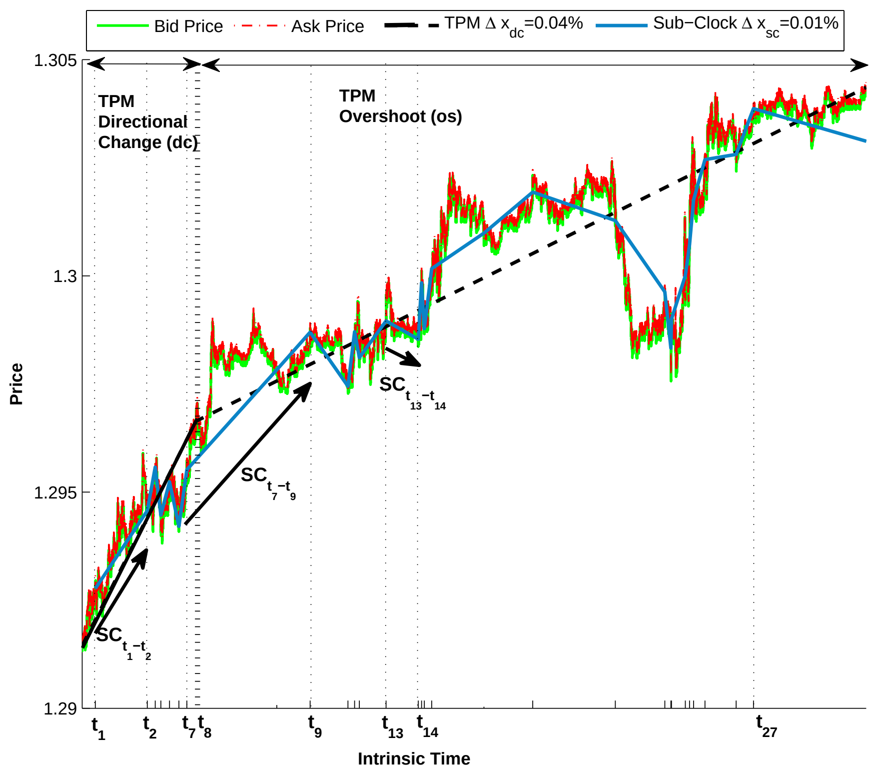

Figure 5.5: Synchronizing sub-clocks with dc and os sections of the total price movement (TPM). This graph captures one TPM ($\triangle x_{dc} = 0.4\%$) of the EUR/USD price series for a period of three days. This is the same TPM depicted in Figure 5.4. The TPM is decomposed into a directional change *dc* (black solid line) and an overshoot *os* section (black dashed line). Discretising this TPM into sub-clocks of $\triangle x_{dc} = 0.1\%$, we find that these sub-clocks are not always exactly synchronized with the starting and ending ticks of the *dc* and *os* sections of the TPM. The ticks of the intrinsic time scale represent the triggers of the TPM and its sub-clocks events.

For sub-clocks that take place totally within the dc or os section, we have no synchronization problem. However, for those sub-clocks that transition within one price movement, from the dc to the os section, we need to account for the right activity proportion in each section of the price movement TPM. The following is an illustrative example of how to compute the synchronized activity over the TPM depicted in Figure 5.5. For simplification we use the following notations:

N_{TPM} for total TPM activity, which is either total CTT or total WTT.

N_{dc} for total activity of the dc section of TPM.

N_{os} for total activity of the os section of TPM.

$n_{sc_{t-t+1}}$ for total activity of one sub-clock starting t and ending $t+1$, where t is the timestamp at which the sub-clock event ticks, and $t+1$ is the timestamp at which the following sub-clock event triggers.

$$N_{TPM=}N_{dc} + N_{os}$$

$$N_{dc} = \sum_{t=1}^{6} n_{sc_{t-t+1}} + n_{sc_{t7-t9}} \cdot \frac{t_8 - t_7}{t_9 - t_7}$$

$$N_{os} = n_{sc_{t7-t9}} \cdot \frac{t_9 - t_8}{t_9 - t_7} + \sum_{t=9}^{26} n_{sc_{t-t+1}}$$

5.6 Deciphering Market Activity Along Intrinsic Time

In the previous sections, we have explained how to define a discretised intrinsic time scale, using price movement and sub-clock events, and how to synchronize market activity along this time scale. In this section, we decipher the FX market activity of our unique high frequency data set along a discretised intrinsic time scale, defined by all price movement events, of a given threshold, and their sub-clocks that occurred over the

2.25 years of the data set period. We analyse the behaviour of counter and with trend trades of our real transaction data set both during the directional change and the overshoot sections of the different price movement events. Our analysis reveals new results about FX market dynamics. We report new empirical evidence about market liquidity behaviour during price overshoots that are valid across all magnitudes of price events. To the best of our knowledge, this is the first study that empirically demonstrates FX activity behaviour during price overshoots.

5.6.1 Visualising FX Activity Patterns Along Intrinsic Time

To visualise FX activity patterns along the discretised intrinsic time, defined in section 5.4, we need to plot the absolute count number of CTT and WTT that took place during each sub-clock of each total price movement event TPM. If we plot such activity over the physical timestamps at which the price movement events and their sub-clocks occurred, no message will be conveyed about the activity pattern. This is illustrated by Figure 5.6, which depicts the number of CTT/WTT over the 0.4% price movement events and their relevant sub-clocks (0.2%, 0.1%, 0.05%) for the 2.25 year data set period.

Therefore, to detect possible market activity patterns, we manipulate the dimensions along which the data are plotted. So, instead of plotting the activity over the timestamps, at which the price movement events and their sub-clocks occurred, we now use intrinsic time scale dimensions. As illustrated by Figure 5.7, we set the x dimension for the price movement events (os and dc sections) , the y dimension for the sub-clocks of each event, and the z dimension for the normalised activity number that takes place at each sub-clock for each event. We normalise the activity by dividing the number of CTT (WTT) of each event sub-clock E_iSC_j by the total number of CTT (WTT) that occurred during the whole event E_i including both the directional change and overshoot periods (see eq. 5.2 and 5.3). This way, we can visualise the activity share of each sub-clock event relative to the total price movement activity.

$$NCTT_{E_iSC_j} = \frac{CTT_{E_iSC_j}}{CTT_{E_i}} \tag{5.2}$$

$$NWTT_{E_iSC_j} = \frac{WTT_{E_iSC_j}}{WTT_{E_i}} \tag{5.3}$$

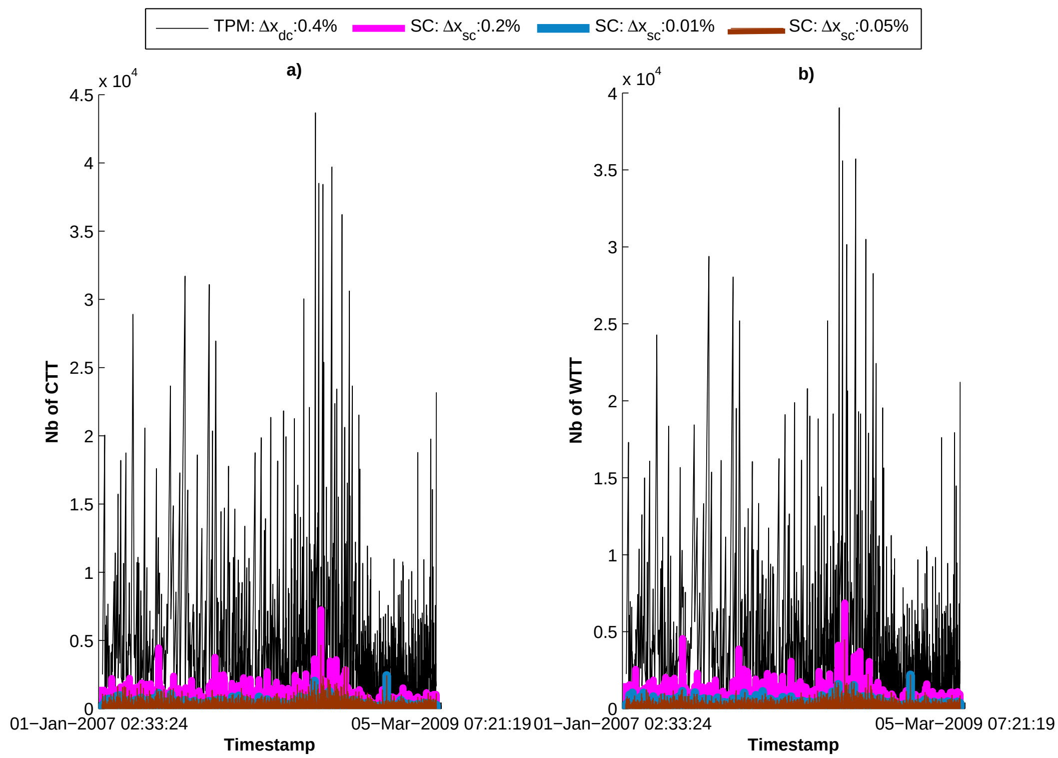

Figure 5.6: Number of a) counter and b) with trend trades over the 0.4% main price movement events and its relevant sub-clocks (0.2%, 0.1%, 0.05%) for the 2.25 year data set period. The time scale represents the physical timestamps at which each price movement (TPM) or sub-clock (SC) event occurred.

Figure 5.7 is used here only for illustration purposes. It shows the pattern of normalised counter trend trades (NCTT) along the 0.2% sub-clocks of the price movement events $\triangle x_{dc} = 0.4\%$. Although, we base our discretisation analysis on a much smaller sub-clock than the threshold used in Figure 5.7, yet we report here general properties that apply on all price movement thresholds using whatever sub-clock, indicating new features of FX market dynamics.

The most remarkable observation is the deminishing market liquidity along the various 0.4% price movement events depicted by Figure 5.7. For all captured events, we can clearly observe a decaying activity in the overshoot period, an important finding that has not previously been reported in the literature.

However, it is important to note that even though each price movement is composed of a *dc* section and an *os* section, it does not necessarily have sub-clock events of a specified threshold $\triangle x_{sc}$ during its overshoot period. This is due to the fact that each price movement event can have a various number of sub-clocks. As mentioned earlier in Table 5.1, a 2.25 year time span of EUR/USD has 1,762 directional change events of $\triangle x_{dc} = 0.4\%$ that can be discretised into 6,430 sub-clocks of $\triangle x_{sc} = 0.2\%$. However, these sub-clocks are not evenly distributed among the different price movement events. As shown in Figure 5.7, each price movement event can have from 1 to 27 sub-clocks. So, in some cases, we have events with no sub-clocks during the overshoot period of the main price movement event, i.e., a price move of 0.2% has not occurred along the *os* section of the 0.4% main price movement.

Naturally, for those events with no sub-clocks during the overshoot period, we cannot observe any diminishing market liquidity. To overcome this, we need to choose a very small threshold size of for $\triangle x_{sc}$, so that more sub-clock events can take place during both the *dc* and *os* sections of the main total price movement event TPM.

Therefore, from this point onwards, we base our analysis on very small fixed size sub-clock to account for all possible microscopic market activity patterns. We use a sub-clock of $\triangle x_{sc} = 0.05\%$ to discretise price movements of different thresholds $\triangle x_{dc}$(0.6%, 0.4%, 0.2%). This will be further illustrated and discussed in the following sections.

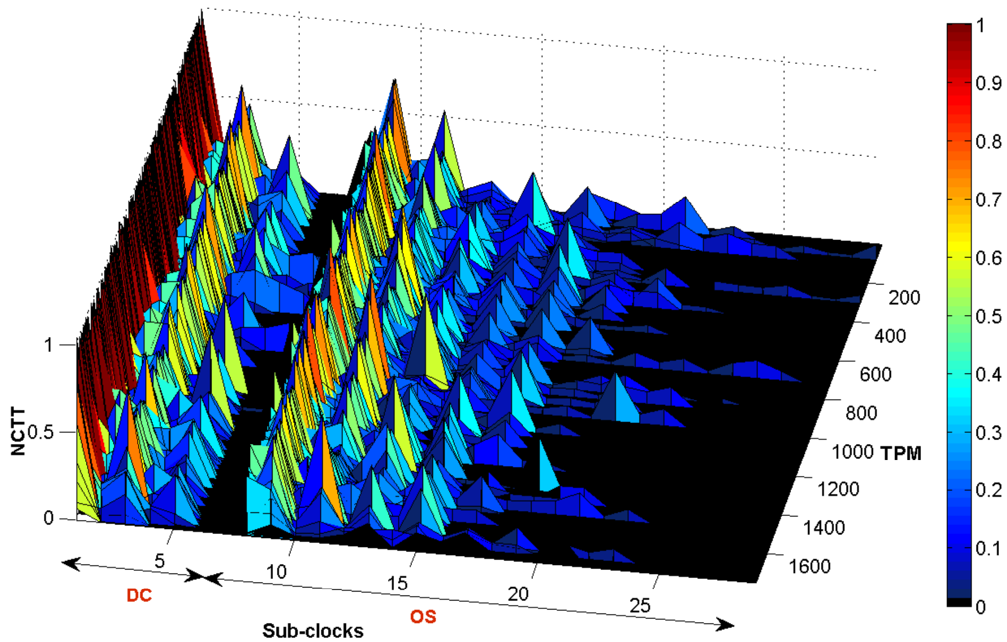

Figure 5.7: Normalised number of counter trend trades (NCTT) over the 0.2% sub-clocks of price movement events of threshold $\triangle x_{dc} = 0.4\%$ for the whole data set period (2.25 years). As each event has a variable number of sub-clocks during the dc and os periods, they have been aligned on a fixed size window composed of two sections. The size of the first window section is the max nb of sub-clocks during the dc period. The size of the second window section is the max nb of sub-clocks during the os period. The activity generally decays of the end of the os period.

5.6.2 DiminishingMarket Liquidity During Overshoot Periods

In this section, we undertake an in-depth analysis of the diminishing market liquidity for the various currency pairs supplied by OANDA; EUR/USD, USD/JPY, AUD/USD and GBP/USD. For each currency pair, we decipher the market activity along different intrinsic time scales of $\triangle x_{dc}$ (0.6%, 0.4%, 0.2%), where each scale is discretised into sub-clocks of $\triangle x_{sc} = 0.05\%$. Table 5.2 shows the total number of price movement events

TPM of the different thresholds and the associated total number of 0.05% sub-clock events for each currency pair.

Table 5.2: Total number of price movement events and the associated sub-clocks of over the whole data set period (2.25 years) for the currency pairs : EUR/USD, USD/JPY, AUD/USD, GBP/USD. For the price movement (TPM) events, we use different threshold sizes $\triangle x_{dc}$(0.6%, 0.4%, 0.2%). All price movements are discritzed into sub-clocks of $\triangle x_{sc} = 0.05\%$.

		$\triangle x_{dc} = 0.6\%$	$\triangle x_{dc} = 0.4\%$	$\triangle x_{dc} = 0.2\%$
EUR/USD	Total nb of TPM	780	1,762	6,472
	Total nb of 0.05% sub-clocks	69,256	69,284	69,264
USD/JPY	Total nb of TPM	1,093	2,420	8,944
	Total nb of 0.05% sub-clocks	91,654	91,650	91,604
AUD/USD	Total nb of TPM	2,150	4,548	14,866
	Total nb of 0.05% sub-clocks	124,491	124,452	124,227
GBP/USD	Total nb of TPM	907	1,956	6,965
	Total nb of 0.05% sub-clocks	64,426	64,426	64,380

As illustrated in Table 5.2, using a small sub-clock threshold of 0.05% generates tens of thousands of sub-clock events, making its visualisation the same way as is done in Figure 5.7 very complex. Therefore, we use 2D graphs to depict the pattern flow of market activity (NCTT/NWTT) over the different sub-clocks. However, since, each price movement event has a variable number of sub-clocks, we normalise the x-axis sub-clock ticks, so that all price moves have the same starting and ending points, i.e. the same number of sub-clocks. We apply Algorithm 5.2 for the sub-clock ticks (x-axis) normalisation.

Algorithm 5.2 Normalisation of Sub-clock Ticks

Input:

$X = (sc_1, sc_2, ..., sc_n)$ *//Vector of sub-clocks activities for one event of* $\triangle x_{dc}$

$maxCount$

/ maxCount is the desired number of sub-clock ticks along which to normalise the market activity. We set it as the maximum nb of the 0.05% sub-clocks that occurred for all events of a sepecific* $\triangle x_{dc}$ **/*

Output:

S *//Normalized vector of sub-clocks activities of one event of* $\triangle x_{dc}$*of length maxCount*

Begin

$count \leftarrow length(X)$; *//number of sub-clocks in vector X*

$f \leftarrow count/maxCount$;

$r \leftarrow 0$; $S \leftarrow []$; *//initialize variables*

For $j \leftarrow 1 : maxCount$

$v \leftarrow X_j + r$;

$s \leftarrow f * v$; *//compute the fractional activity of the relevant subclock*

$S \leftarrow [S; s]$;

$r \leftarrow v - s$; *//remaining activity that will be added to following sub-clock*

End

End

Normalising the sub-clock ticks and depicting the FX activity flow over price movement events spanning 2.25 years, we observe a consistent behaviour for all directional change thresholds across the different currency pairs. Both counter and with trend traders show high trading intensity at the beginning of the price move. They then decrease their activity gradually as the price move continues, whereas the minimum activity is at the end of the overshoot period. Figures 5.8- 5.11 show the diminishing market liquidity (NCTT/NWTT) over the sub-clocks of the different directional change events for EUR/USD, USD/JPY, AUD/USD and GBP/USD respectively.

By this analysis, we provide empirical evidence of new FX market dynamics during price overshoots that are valid across all magnitudes of the price movement events of the different currency pairs.

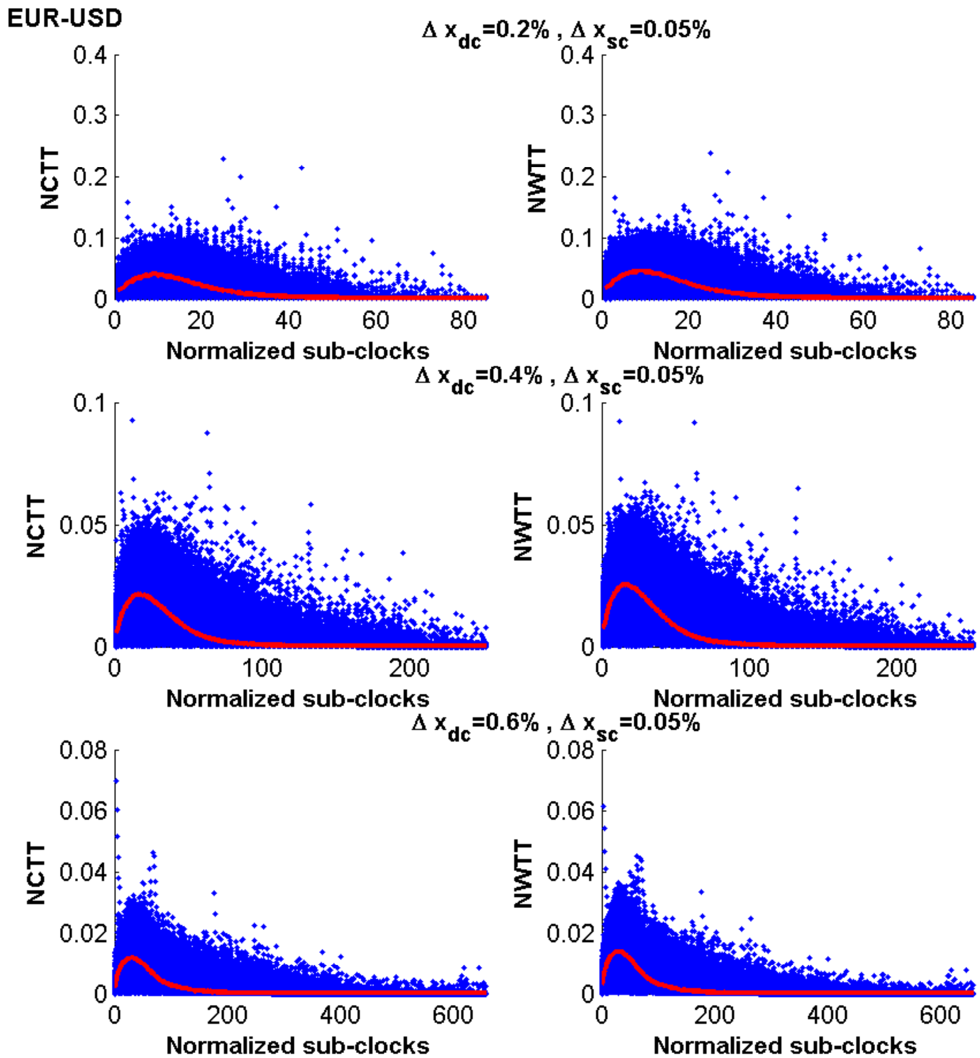

Figure 5.8: EUR/USD normalised counter and with trend trades over the 0.05% sub-clocks of different price movement events (0.2%, 0.4% and 0.6%) spanning 2.25 years. The blue dots represent the normalised activity of each sub-clock in each price movement event. Refer to Table 5.2 for the number of events of each threshold $\triangle x_{dc}$ and the associated number of sub-clocks. The solid red line is a simple average of the activity in each sub-clock tick of the x-axis. A deminishing market liquidity is observable for both NCTT and NWTT during the overshoot period of all price movements.

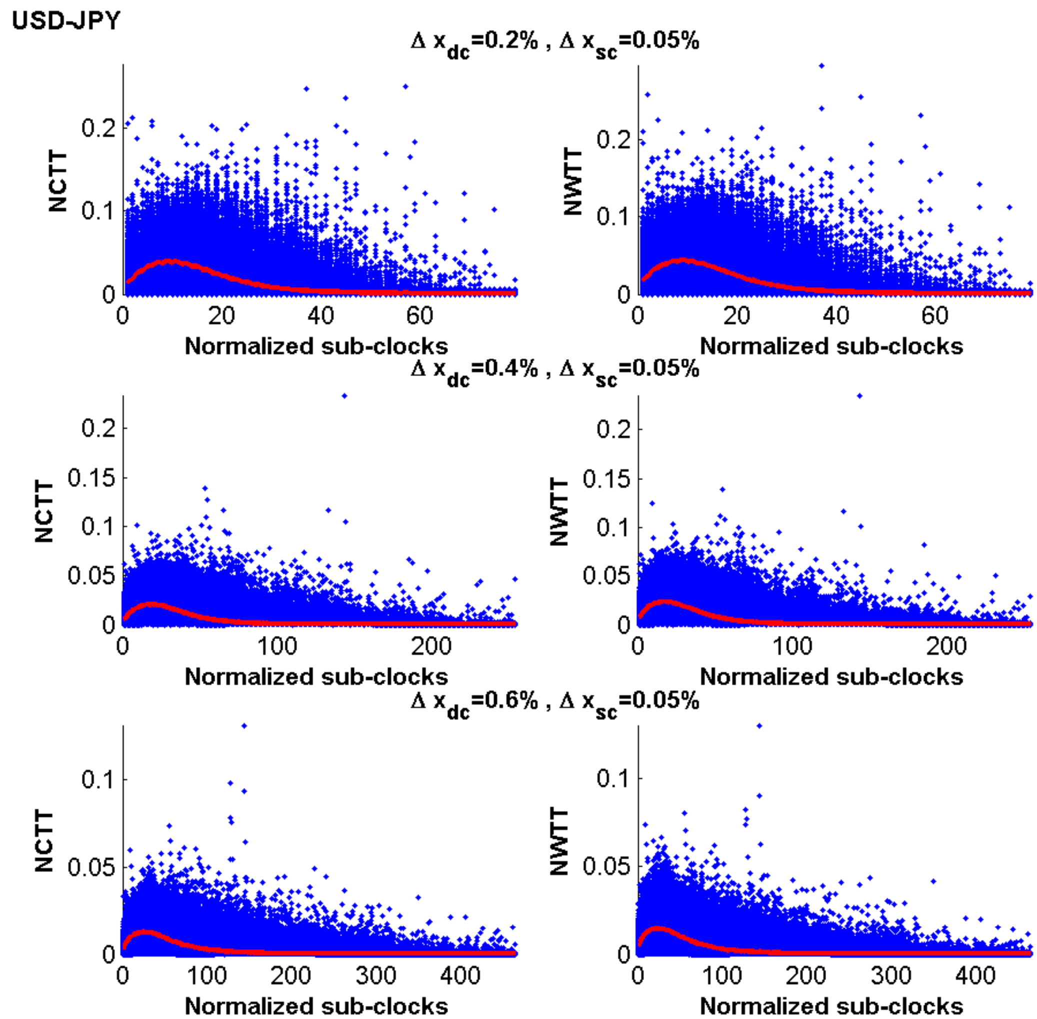

Figure 5.9: USD/JPY normalised counter and with trend trades over the 0.05% sub-clocks of different price movement events (0.2%, 0.4% and 0.6%) spanning 2.25 years. The blue dots represent the normalised activity of each sub-clock in each price movement event. Refer to Table 5.2 for the number of events of each threshold $\triangle x_{dc}$ and the associated number of sub-clocks. The solid red line is a simple average of the activity in each sub-clock tick of the x-axis. A diminishing market liquidity is observable for both NCTT and NWTT during the overshoot period of all price movements.

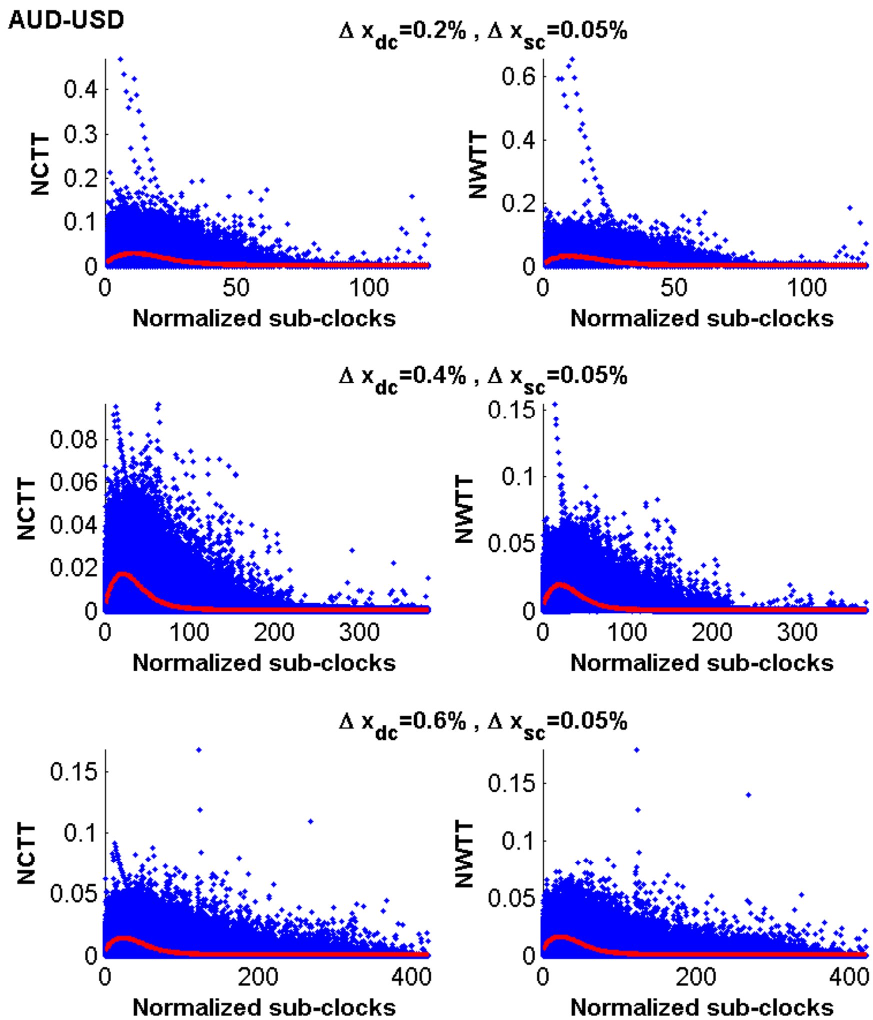

Figure 5.10: AUD_USD normalised counter and with trend trades over the 0.05% sub-clocks of different price movement events (0.2%, 0.4% and 0.6%) spanning 2.25 years. The blue dots represent the normalised activity of each sub-clock in each price movement event. Refer to Table 5.2 for the number of events of each threshold $\triangle x_{dc}$ and the associated number of sub-clocks. The solid red line is a simple average of the activity in each sub-clock tick of the x-axis. A diminishing market liquidity is observable for both NCTT and NWTT during the overshoot period of all price movements.

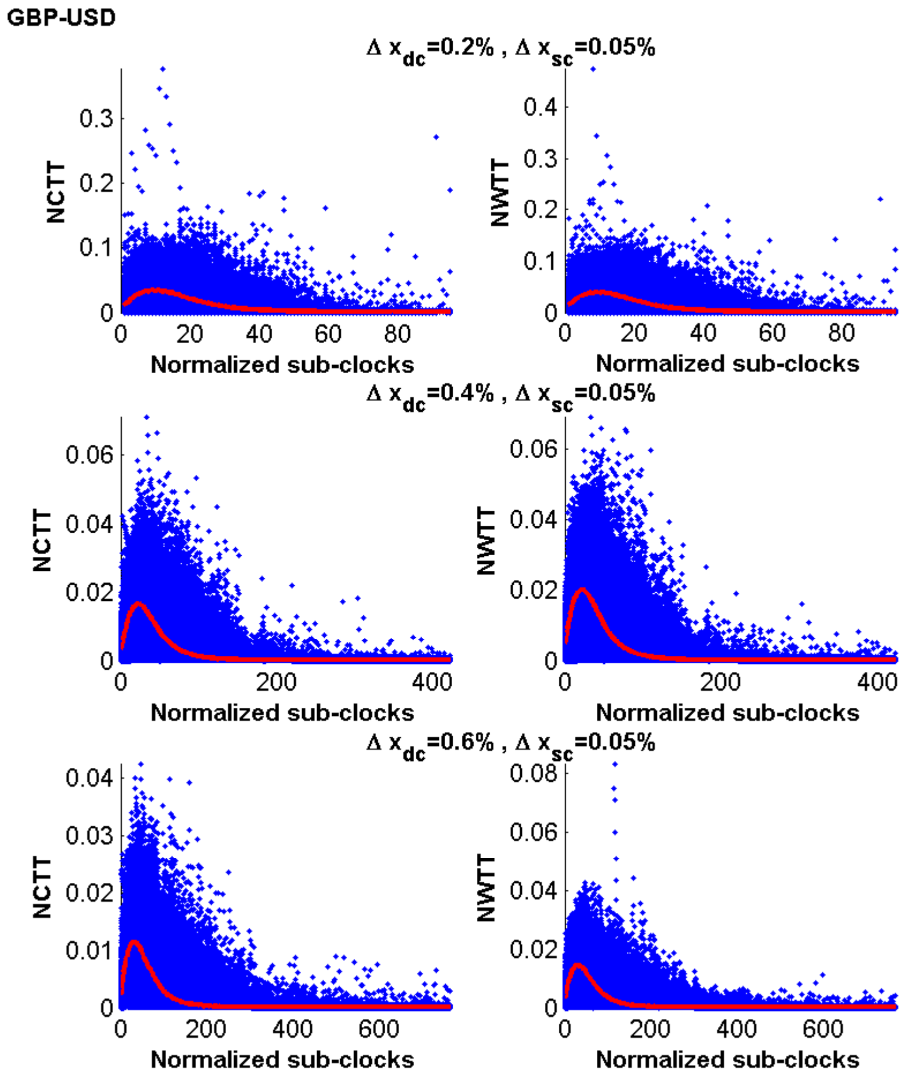

Figure 5.11: GBP/USD normalised counter and with trend trades over the 0.05% sub-clocks of different price movement events (0.2%, 0.4% and 0.6%) spanning 2.25 years. The blue dots represent the normalised activity of each sub-clock in each price movement event. Refer to Table 5.2 for the number of events of each threshold $\triangle x_{dc}$ and the associated number of sub-clocks. The solid red line is a simple average of the activity in each sub-clock tick of the x-axis. A diminishing market liquidity is observable for both NCTT and NWTT during the overshoot period of all price movements.

5.6.3 Further Deciphering of the Overshoot Period

In this section, we further decipher the overshoot period in attempt to explain the reason behind its stoppage. As shown in the previous section, both NCTT and NWTT behave similarly over the discretised price movement events and reach their minimum during the overshoot period. Despite this very similar pattern flow, we find that the intensity of NCTT and NWTT differs along the sub-clocks of each price movement event. This is revealed by analysing the cumulative difference between NCTT and NWTT over the different sub-clocks (eq. 5.4). For each event sub-clock E_iSC_j, we cumulate the difference between NCTT and NWTT of all sub-clocks preceding sub-clock SC_j of the same event E_i.

The results of the cumulative difference between NCTT and NWTT for all studied currency pairs over the sub-clocks of the different price movements are depicted in Figure 5.12. We find that with trend trades usually exceed counter trend ones at the beginning of the price movement. This behaviour then reverses, and counter trend trades become much more intense relative to with trend trades. As the price movement continues, both NCTT and NWTT decrease, as shown in the previous section, but still with the superiority of NCTT over NWTT. This finding is consistent across all scales for all currency pairs, indicating that price overshoots stop due to more participants placing counter trend trades. This means that, with this diminishing market liquidity, a small with trend order at the end of the price move can stimulate the underlying price trend.

Given this important information about FX traders' behaviour during overshoot periods, we do not only visualise the general trend of the FX market movement but also the small events (sub-clocks) leading to this movement.

$$\text{CumDiff}_{E_iSC_j} = \sum_{k=0}^{j} NCTT_{E_iSC_k} - NWTT_{E_iSC_k} \tag{5.4}$$

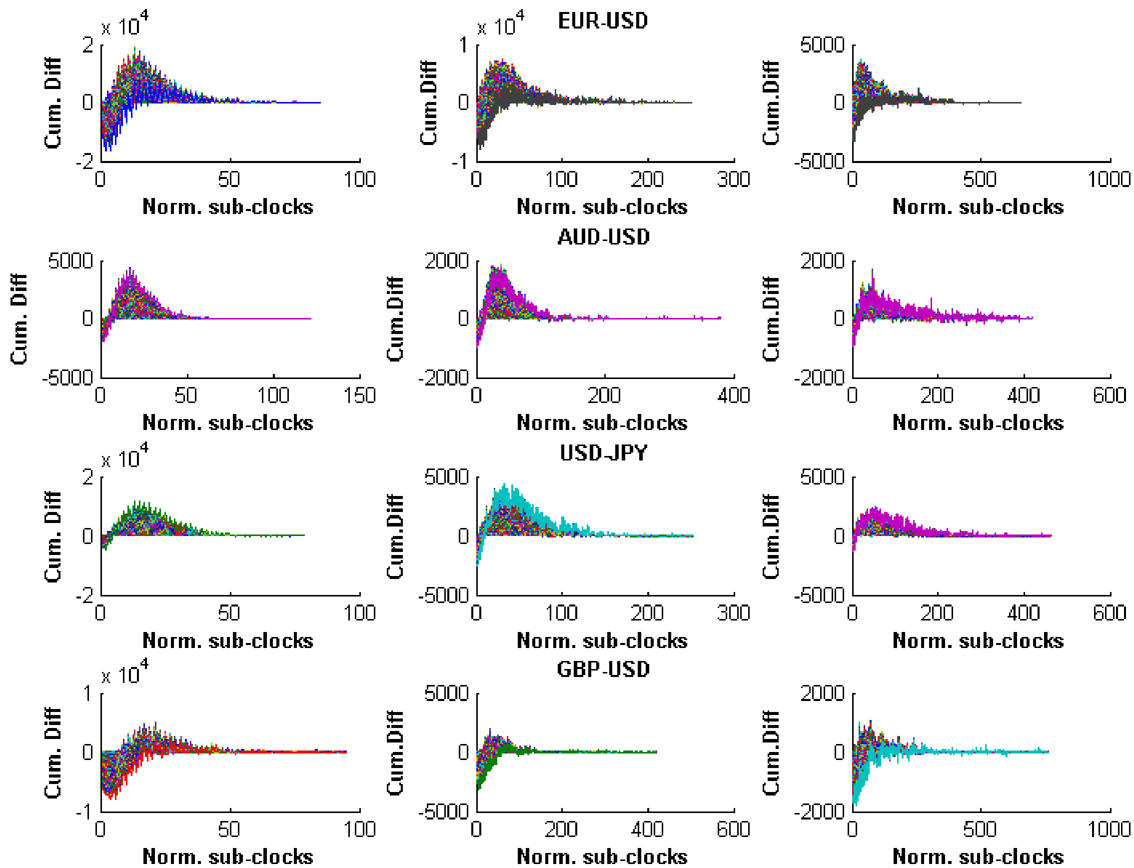

Figure 5.12: Cumulative difference between NCTT and NWTT over the normalised 0.05% sub-clocks of different price movement events spanning 2.25 years. Each row represents the cumulative difference between NCTT and NWTT for a certain currency pair over the following price movement events thresholds: (1^{st}col: $\triangle x_{dc} = 0.2\%$, 2^{nd}col: $\triangle x_{dc} = 0.4\%$, 3^{rd}col: $\triangle x_{dc} = 0.6\%$).

5.7 Market Liquidity and Price Tick Patterns

One of the well-known stylized facts with regard to financial time series is the heavy tailed empirical distribution of price changes in both stock and foreign exchange markets. A great deal of supporting evidence can be found in [49, 91, 100, 127, 135, 147], where all researchers agree that price returns are fat tailed over short time horizons, with a higher probability of extreme events than in a Gaussian random walk model.

Several theories have been developed for explaining this statistical regularity. The most widely accepted theory is the strong role of trading volume fluctuations in determining asset price changes [36, 47, 84]. However, this theory has been criticized in a comprehensive empirical study by [74] showing that trading volume is not the key factor. According to [74], heavy tailed returns can be observed on a microscopic time scale and the size of price changes is not affected by the volume of orders. In [74], price changes are associated with liquidity fluctuations, in other words, the market's ability to fulfil new market orders. This view is also supported by [183], where an agent based model of a double-auction market is set up to simulate liquidity fluctuations through the emergence of order book gaps. They show that liquidity changes are the main cause for large price changes or heavy tailed return distributions.

In this section, we check the impact of market liquidity on price changes, but on a micrsocopic level. We measure price changes by price ticks moves pattern over the different sub-clocks. A price tick move is generally calculated as the difference between the mid price from one tick and the next. A tick is a bid/ask pair, whereas the mid price m is the average price $((bid + ask)/2)$. Since a sub-clock can have n price ticks, then a sub-clock price move SPV represents the total number of price moves between the different n ticks in the sub-clock (eq. 5.5).

$$SPV = \sum_{k=0}^{n} |m_k - m_{k+1}| \tag{5.5}$$

Analysing the SPV pattern over the 0.05% sub-clocks of the different intrinsic time scales ($\triangle x_{dc}$: 0.2%, 0.4%, 0.6%), we find a consistent pattern of price changes, even on this microscopic level. As depicted by Figures 5.13- 5.16, we find that the SPV pattern over the different sub-clocks follow the same decaying pattern of market liquidity for all underlying currency pairs. The SPV decreases, the more we approach the end of the overshoot period of the price movement event. This pattern suggests a strong association between microscopic market liquidity and microscopic exchange rate price changes in the FX market. This is in line with the stock market findings by [36, 74, 183], where liquidity fluctuations are found to be the prime reason for price changes.

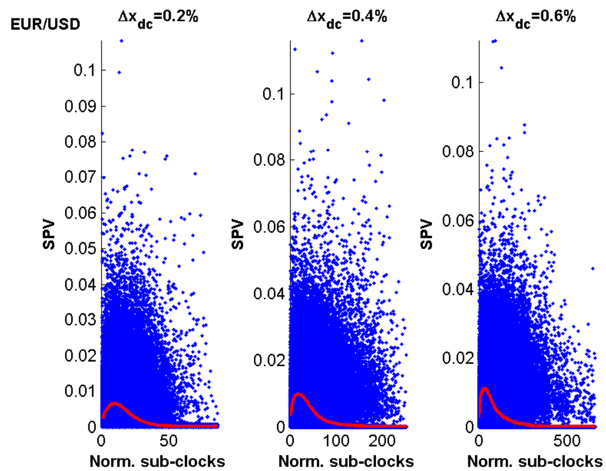

Figure 5.13: EUR/USD price ticks moves (SPV) over the normalised 0.05% sub-clocks of different price movement events (0.2%, 0.4% and 0.6%) spanning 2.25 years. The solid red line is a simple average of SPV in each sub-clock tick of the x-axis. SPV decreases the more it approaches the overshoot period.

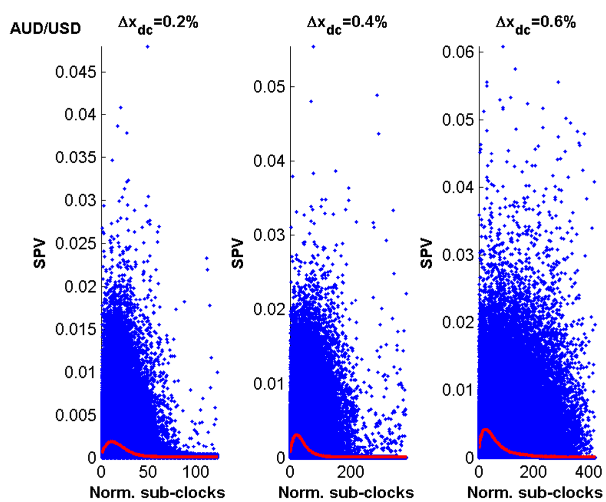

Figure 5.14: AUD/USD price ticks moves (SPV) over the normalised 0.05% sub-clocks of different price movement events (0.2%, 0.4% and 0.6%) spanning 2.25 years. The solid red line is a simple average of SPV in each sub-clock tick of the x-axis. SPV decreases the more it approaches the overshoot period.

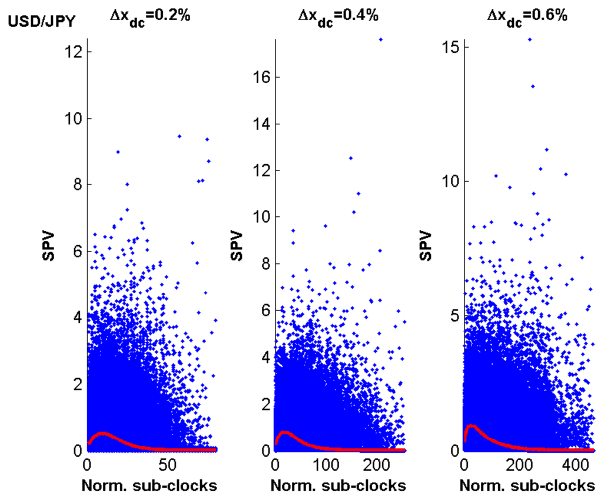

Figure 5.15: USD/JPY price ticks moves (SPV) over the normalised 0.05% sub-clocks of different price movement events (0.2%, 0.4% and 0.6%) spanning 2.25 years. The solid red line is a simple average of SPV in each sub-clock tick of the x-axis. SPV decreases the more it approaches the overshoot period.

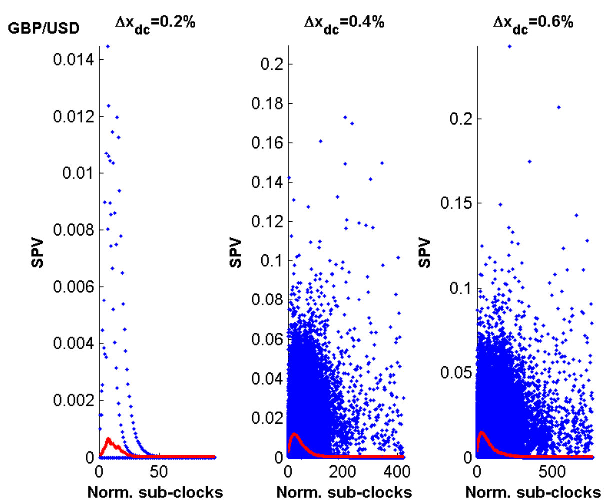

Figure 5.16: GBP/USD price ticks moves (SPV) over the normalised 0.05% sub-clocks of different price movement events (0.2%, 0.4% and 0.6%) spanning 2.25 years. The solid red line is a simple average of SPV in each sub-clock tick of the x-axis. SPV decreases the more it approaches the overshoot period.

5.8 Conclusion

In this chapter, we carry out the first study of its kind that synchronizes and deciphers FX market activity along price movement events, and in particular the price overshoot period. A price overshoot is of special importance as it measures the excess price move of a given threshold. The longer the overshoot, the higher the imbalance in the market for the specified threshold.

The results of this study have been only enabled by utilizing the intrinsic time scale. The concept of intrinsic time samples the tick price data at points, where there is a price action (movement) of a given threshold, and zooms into the market behaviour between these focal points. The approach yields valuable microscopic insight in how FX market activity changes as the price movement progresses. We discover that market activity declines as the price overshoots. This diminishing market liquidity has in turn a strong impact on price ticks changes, which follow the same decaying pattern during price overshoots.

We find that a price overshoot stops due to more market participants placing counter trend trades acting against the market force. This discovery indicates that already small imbalances of market activity in large overshoots can alter the price trajectory. Having this valuable information, we can also identify when the market would be vulnerable to the placement of large orders, and the impact of opening counter trend or with trend positions on price overshoots, and hence on the global FX market price move.

Chapter 6

Formalisation of Market Dynamics

To understand financial markets and prevent crisis we need to analyse market microstructure. This chapter introduces market-calculus, as a new approach, to formalise market dynamics. We define calculus to analyse market processes and feedback loops like cascading margin calls, which exist in both FX and stock markets, in the context of a simple double auction market model. The objective to get a better understanding of risk scenarios, not to forecast exogenous order flow. The price trajectory is determined by the present market state and new orders arriving in the market. By studying the market microstructure, we can compute the impact of orders of any size, or how big a sell order has to be to cause the market to fall by a certain percentage. Using this formalism reduces ambiguity and enables rigorous reasoning. An algorithm for risk assessment is proposed and implemented in a double auction demo simulator that is based on the same calculus logic. The demo simulator demonstrates the potentiality of the market-calculus approach in understanding market microstructure. Real markets are more complex than the models presented here, but this chapter is a step towards building a solid foundation for studying market models.

6.1 Introduction

For many years, researchers have adopted various approaches, trying to understand the market behaviour and its dynamics with an ultimate goal of avoiding or reducing losses, enhancing market performance, and most important providing early warning signals for crashes, especially after the last world crisis [41]. According to [165], this can be only achieved by digging deeper into financial markets and tracking the tick-by-tick data generated from the interaction of market participants. This means looking at and analysing every single price movement and every single observable action (placing an order) by each trader in relation to his/her open positions and trading margin [195]. A margin is the minimum amount required to be deposited and held by the investor relative to the total value of his/her trades. Given that this collateral is maintained, a margin trader can open positions larger than his/her actual balance. In contrast to other financial markets, the FX market can increase its traders' purchasing power up to 50 times the underlying capital or even more [156]. This is always attractive as long as the trader makes profits. However, once the account balance depresses below the specified margin requirement, a margin call is triggered, which closes out the trader's open position by force. This in turn can have a great impact on the market movement, since many traders may be trapped in the same situation.

Therefore, in this chapter we study the impact of orders on the market state and on price movement, when associated with margin trading, an important aspect that has not been discussed in the literature.

As mentioned earlier in chapter 3, OANDA tries to achieve an open transparent market by publishing the underlying market open short and long positions to help traders judge the depth of the market for a given currency pair. A 24-hour summary of open positions held by OANDA's clients on any given day, shows that many traders incur losses by getting trapped on the wrong side of the market move. This is illustrated in Figure 6.1, which depicts the summary of OANDA's traders' open positions on December 13, 2012. We can clearly observe that traders get locked into counter trend positions, and herding behaviour comes about. When the market price moves up, a large number of traders get locked into a short position against the current trend, and when the market price moves down, a large number of traders are long against the falling price, hence end up making losses as shown in Figure 6.1. With such an imbalance, any small 'random price spike', for whatever reason, can trigger margin calls to close out losing positions of many traders. For instance, if a trader was short and the price moves up to an extent that a margin call is triggered, then the trader's short position is closed by buying the underlying traded currency, which increases the price imbalance. Thus, with many traders locked into losses as depicted in Figure 6.1, several market buy (sell) orders will take place to close out the counter trend losing positions. This in turn will amplify the price move, triggering further margin calls and close-outs [167]. According to [165, 167] cascading margin calls can cause dangerous price moves.

To understand these market dynamics and prevent possible crisis, we study the market microstructure using a new microscopic event-based approach. In this chapter, we introduce market-calculus as a new approach for defining and formalising market mechanisms in a scientific way. It is based on the observation that every mechanism in a market is by design [195]. Consequently, the market clearance process of new orders placed into the market can be written down and formalised. The usage of formal logic reduces ambiguity and enables rigorous reasoning and studying of micro-behaviour.

To illustrate the potential of the market-calculus approach, we use a simplified version of the popular double-auction market model. The FX market maker model cannot be used in this scope due to its complexity and the lack of knowledge of this model in the literature. Yet, the concept of cascading margin calls and its impact on market price is the same. By formalising the market mechanism using market-calculus, we are able to compute the impact of any placed order and the percentage by which the market would fall, when placing a sell order of a specific size. We also propose an algorithm for assessing market risk. Having this definite formalisation, we then implement a double auction demo simulator that demonstrates market clearance dynamics, feedback loops and market risk assessment.

Real markets are more complex than the market models presented in this chapter. Yet, we demonstrate how a new approach like the market-calculus can build a solid foundation for studying and formalising complex market models. The rest of this chapter is organized as follows: Section 6.2 describes the market-calculus approach. Section 6.3 formalises the dynamics of a simple double auction market model. Using this formal logic, we propose an algorithm for assessing market risk in section 6.3. The algorithm, in addition to the formalised market clearance model are then simulated in section 6.5. Conclusions are given in section 6.6.

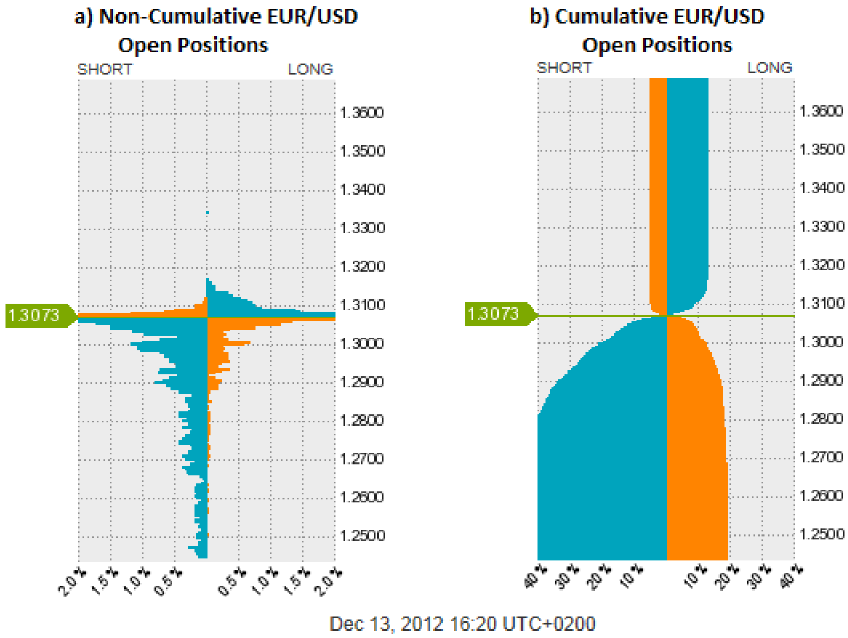

Figure 6.1: OANDA FX Open Positions: A 24-hour summary of a) non-cumulative and b) cumulative open positions held by OANDA's clients on Dec 13, 2012 16:20 UTC. Each graph quadrant shows whether positions are long or short and whether entry price points are below or above the market price (horizontal green line) at the time the snapshot was taken. Number of positions at each price level is shown as a percentage of the maximum number of all open positions during the past 24 hours. The orange colour in the top left and bottom right quadrants indicates traders with profitable open positions. The blue colour in the top right and bottom left quadrants indicates traders with losing open positions. Source: www.fxtrade.OANDA.com

6.2 The Market-Calculus Approach

The fact that classical economics have not captured many of the complex financial markets dynamics (e.g. see [164, 189]), has opened the gate for other various research approaches. One novel approach to market studies is based on actual observations of how market participants behave [189, 190], with the attempt to model the micro-behaviour of markets and discover general dynamics [1, 87].

Following the same methodology, we introduce a new event-based approach, the market calculus, in which we closely look at every single action by the trader and study its effect on both the trader and the whole market. This approach is still in its infancy. We look at simple market models, and attempt to define the market dynamics formally. The intended contribution of this chapter is not in modelling micro-behaviour, but in formalising such models and analysing their properties, to examine what can be usefully inferred from market information.

In the defined calculus, the world is described by states, which are changed by events. In this chapter, we limit our attention to buy and sell events initiated by market participants. Even though the behaviour of market participants may in general be unpredictable, certain inferences can be made. Given a set of buy and sell orders, the calculus can define state transitions. We can make an analogy with weather forecasts, where we may not know the long term weather changes, but we can predict the immediate future given the current state; e.g. air flows from high pressure to low pressure regions.

Event calculus is useful for reasoning [45, 114, 145]. Shanahan states: *"The event calculus is a logical mechanism that infers what's true when given what happens when and what actions do"* [187]. Although we have not adopted conventional event calculus, this chapter formalises the components relevant to the calculus of market transitions. It highlights the fact that the consequences of an order, which are the only events considered in this chapter, can be complex: the consequences are dependent on the positions and margins held by market participants. With this analysis, one can determine, for example, how big orders need to be to cause market crashes.

This chapter formalises the obvious. But it is better to state the obvious with mathematical rigor rather than ambiguity, which needs repeated clarification later in our research. Besides, what is obvious to some may not be obvious to others. Stating the obvious through a formal description enables us to study micro-behaviour rigorously.

6.3 Market Clearance Models

In this section, we formalise two models for market clearance in a double auction market. The first model is a simplistic one and describes the calculus for fulfilling new orders entered into the market. The second model extends the first one and describes the calculus of the impact of any order on the market price with respect to traders' positions and margin constraints. Finally, we propose a general form for the calculus of any market maker model. The latter is not implemented in the demo simulator in section 6.5.

6.3.1 Model 1: Simple Market Clearance

This model is defined under a double auction market.

$$State + Orders \rightarrow State$$

Where:

State = Queue_Profile = (Bid_Queue, Offer_Queue)

Bid_Queue = $((order_1,\ price_1,\ volume_1),\ (order_2,\ price_2,\ volume_2),$
$\ldots,(order_n,\ price_{bq},\ volume_{bq}))$
where $price_1 > price_1 > \ldots\ price_{bq}$

Offer_Queue = $((order_1,\ price_1,\ volume_1),\ (order_2,\ price_2,\ volume_2),$
$\ldots,(order_n,\ price_{bq},\ volume_{bq}))$
where $price_1 < price_2 < \ldots\ price_{bq}$

The Bid_Queue comprises the bids to buy. The Offer_Queue comprises the offers to sell. Buy (sell) orders having the same price are not merged.

Orders refer to a sequence of orders, where each order is either a bid or an offer, together with its volume.

Orders = $(Order_1,\ Order_2,\ \ldots,\ Order_n)$

We assume that the orders are processed in sequence:

State+ $(Order_1,\ Order_2,\ \ldots,\ Order_n) \rightarrow (\text{State} + Order_1) + (Order_2,\ \ldots,\ Order_n)$

For simplicity, we assume only two types of orders. A market order is to buy or sell at the market price. A limit order is to buy a certain volume up to a price specified, or to sell a certain volume above a price specified. For notional convenience, we write a market buy order as a limit buy order with the price set at infinity; a market sell order sets its price to minus infinity.

Order = (Order_No, Order_Type, Price, Volume)

Order_Type = bid | offer

Order_No = O_i

We define a symbol Inf, which stands for both infinity and minus infinity. We write a market buy order as (buy, Inf, Volume), a market sell order as (sell, Inf, Volume). The calculus for clearance of a limit sell order can be defined below.

Let

Bid_Queue1 = $((O_1, P_1, V_1), (O_2, P_2, V_2) \ldots)$

Offer_Queue1 = $((O_3, P_3, V_3), (O_4 , P_4, V_4), \ldots)$

Limit_Order = $(O_n$, sell, $P_n, V_n)$.

The calculus for a limit order is very simple. If the price of the sell order is less than or equal to at least the bid order at the head of the bid queue, the limit order can be fully or partially fulfilled. The sell order of volume V_n removes from the head of the Bid_Queue (P_1, V_1) the minimum of V_n or V_1. If V_n is greater than V_1, then the head of the Bid_Queue is removed. If the limit price is reached, clearing stops and the remaining unfulfilled sell order joins the offer queue. If the limit price is not yet reached, clearing continues with the remaining Bid_Queue until V_n is reduced to 0. If the price of the limit sell order is larger than the first bid order in the bid queue, then the sell limit order joins the offer queue. This can be formalised as follows.

$$(((O_1, P_1, V_1), (O_2, P_2, V_2), \ldots), ((O_3, P_3, V_3), (O_3, P_3, V_3), \ldots))$$
$$+(O_n, \text{sell}, P_n, V_n) \rightarrow (((O_1, P_1, V_1), (O_2, P_2, V_2), \ldots),$$
$$((O_3, P_3, V_3), (O_4 ,P_4, V_4), \ldots) \oplus (O_n, \text{sell}, P_n, V_n))$$
$$\text{if } P_1 < P_n$$
$$(((O_1, P_1, V_1 -(\min(V_1, V_n)), (O_2, P_2, V_2), \ldots), ((O_3, P_3, V_3), (O_4, P_4, V_4), \ldots))$$
$$+ (O_n, \text{sell}, P_n, V_n -(\min(V_1, V_n))); \text{ Transaction Price } (TP) = P_n \text{ is defined}$$
$$\text{if } P_1 \geq P_n$$

The + operation is recursive when$P_1 \geq P_n$, in which case transaction takes place; it stops when $P_1 < P_n$ or V_n is reduced to 0. Here $\oplus$ is the queue joining operator which simply put the orders in ascending order according to their prices[1]. Cleared orders are removed from the bid queue:

[1]In functional programming convention, $\oplus$ is defined below:
((P1, V1), (P2, V2), ...) $\oplus$ (sell, P, V) $\rightarrow$
((P, V), (P1, V1), (P2, V2), ...) if P < P1
((P1, V1), ((P2, V2), ...)$\oplus$ (sell, P, V))) if P $\geq$ P1

$((O_1,\ P_1,\ 0),\ (O_2,\ P_2,\ V_2),\ldots) \rightarrow ((O_2,\ P_2,\ V_2),\ \ldots)$

In the above rule, we highlight Transaction Price (TP) at the point where it is defined. We shall refer to it later.
Limit buy orders are handled symmetrically.

In the calculus above, the clearing of a market order is exactly the same as the limit order, except that market orders do not have limit prices and hence are always completely fulfilled as long as there are buyers (sellers). They do not join the bid or offer queues. Generally, the handling of unmatched large market orders depends on the order book configuration of the trading system.

6.3.1.1 Example 1 for Model 1

With Model 1, the calculus for computing state transition is straight-forward. This example shows the state change for a given market order.

State1.1 = (Bid_Queue1.1, Offer_Queue1.1)
Bid_Queue1.1 = (($O1$, 1.60, 2500), ($O2$, 1.59, 2000), ($O3$, 1.58, 2500),
($O4$, 1.57, 1500), ($O5$, 1.56, 4000))
Offer_Queue1.1 = (($O6$, 1.61, 3000), ($O7$, 1.62, 2000), ($O8$, 1.63, 1500))

Let

Order1.1 = ($Order_9$, $Order_{10}$, $Order_{11}$), where
$Order_9$= (O_9, sell, Inf, 5000)
$Order_{10}$= (O_{10}, buy, 1.57, 1000)
$Order_{11}$= (O_{11}, buy, 1.62, 6000)

With $Order_9$, which is a market order, the following transactions ensue:

2500 will be transacted at 1.60

This will result in the Bid_Queue being reduced to:

(O_2, 1.59, 2000), (O_3, 1.58, 2500), (O_4, 1.57, 1500), (O_5, 1.56, 4000))

Next, the following two transactions will take place:

2000 will be transacted at 1.59
500 will be transacted at 1.58

The resulting state is:

State1.2 = (Bid_Queue1.2, Offer_Queue1.2)
Bid_Queue1.2 = ((O_3, 1.58, 2000), (O_4, 1.57, 1500), (O_5, 1.56, 4000))
Offer_Queue1.2 = Offer_Queue1.1

With Limit_$Order_{10}$, the offer queue is not changed as the price of the buy limit order is less than the price of the head of the offer queue. Since Limit_$Order_{10}$ is not matched; it is added to the bid queue.
The resulting state is:

State1.3 = (Bid_Queue1.3, Offer_Queue1.3)
Bid_Queue1.3 = ((O_3, 1.58, 2000), (O_4, 1.57, 1500), (O_{10}, 1.57, 1000), (O_5, 1.56, 4000))
Offer_Queue1.3 = Offer_Queue1.2

With Limit_$Order_{11}$ (to buy 6000 with limit price 1.62), the offer queue is changed. Since the price 1.62 is greater than or equal to the first two orders in the offer queue, the following transactions will take place:

3000 will be transacted at 1.61
2000 will be transacted at 1.62

The remaining 1000 units will join the bid queue. Therefore, the resulting state is:

State1.4 = (Bid_Queue1.4, Offer_Queue1.4)
Bid_Queue1.4 = ((O_{11}, 1.62, 1000), (O_3, 1.58, 2000), (O_4, 1.57, 1500), (O_{10}, 1.57, 1000), (O_5, 1.56, 4000))
Offer_Queue1.4 = ((O_8, 1.63, 1500))

6.3.2 Model 2: Market Clearance when Positions and Margins are considered

The market dynamics will change when traders trade with margins. A trader with margin m, where 0<m≤1, will pay up only proportion m of the value that it trades. We make the following assumptions in our analysis:

Assumption 2.1. *For a trader with a short (long) position with margin m, its position is closed automatically when the price rises (falls) by more than m.*

For example, a trader who trades with a margin of 4% will have its short position closed automatically when the price rises by 4% or more.

Assumption 2.2. *All consequences of an automatic position closure take place before any new event occurs.*

Today, market orders are cleared by computer programs, which will typically handle one order at a time. A program must clearly specify how orders are processed even if they reach the computer simultaneously with parallel hardware. A calculus can be written down for every clearly defined clearing mechanism. Without loss of generality, we assume in this chapter that the market clearing process cannot be interrupted. We assume that the recursive application of the rule will not be interrupted before the clearing mechanism handles new orders.

Assumption 2.3[2]**.** *We assume that a position cannot be adjusted and is only opened by a market or limit order. Position closure takes place automatically through margin calls. The relaxation of this assumption does not affect the generality of the results shown in our chapter.*

Assumption 2.4. *We assume that the orders, positions and margins are available.*

Under this model, the description of a state must include traders' position profiles:

$$\text{State} = (\text{Queue\_Profile, Position\_Profile})$$

Where:

Queue_Profile = (Bid_Queue, Offer_Queue)
Position_Profile = {Position | Position = (Position_Code, Position_Type, Volume, Value, Price, Margin)}
Position_No= $P(O_i)$, where O_i is the Order_No of the order opening the position, given Assumption 2.3
Position_Type = long | short
Value = the value of the order(s) against which the opening position order has been matched. Given:

[2] In a real market, a position is constructed via a set of orders. It can be opened, adjusted and closed by market and limit orders. Position closure takes place as a result of either a margin call or the trader's decision.

Bid_Queue = ((O_1, P_1, V_1), (O_2, P_2, V_2), .., (O_{n-1}, P_{n-1}, V_{n-1}))
Order= (On, sell, Inf, Vn)
$P(O_n)$Value = ($P_1\times$ min(V_1,V_n)) + (P_2 $\times$ min(V_2, (V_n - min(V_n,V_2)))+...
+ ($P_{n-1}\times$ min(V_{n-1},V_n - min($\ldots$)))
Price= Unit Price= Value/Volume

The clearance calculus is exactly the same as in Model 1, except that new events, namely new orders, can be triggered by state transitions.

The last transaction price (TP) is defined by the order clearing rule described in Model 1. TP may trigger margin calls, which force some positions to be closed. The margin-triggered set of new orders is NO:

NO = {(Oi, buy, Inf, V) | ($P(Oi)$, short, Vol, Val, P, m) ϵ Position_Profile
such that
P$\times$ (1+m) $<$ TP} $\cup$ {(Oi, sell, Inf, V) | ($P(Oi)$, long, Vol, Val P, m)ϵ
Position_Profile such that P$\times$ (1$-m$) $>$ TP}
Orders = Orders + NO

Here we make no assumption on how the set of new orders (NO) join the Orders queue; i.e. the "+" operator between orders is yet to be defined. This is left to future refinement of the model.

6.3.2.1 Example 2 for Model 2: The effect of margin constraints

The following example shows the state transitions and how new events (which are limited to market orders in this model) are triggered.

Let

State2.1 = ((Bid_Queue 2.1, Offer_Queue 2.1), Positions 2.1)

BidQueue2.1 = ((O_4, 1.60, 2500), (O_5, 1.59, 2000), (O_6, 1.58, 2500),
(O_7, 1.57, 1500), (O_8, 1.56, 4000))

Offer_Queue2.1= ((O_9, 1.61, 3000), (O_{10}, 1.62, 2000), (O_{11}, 1.63, 1500))

Positions2.1 = (($P(O_1)$, long, 4000, 6600, 1.65, 4%),
($P(O_2)$, long, 2000, 3280, 1.64, 4%),
($P(O_3)$, long, 2000, 3280, 1.64, 5%))

For illustration purposes let us assume the following:

1. The position profile (Positions2.1) represents the current positions in the market created from previous orders.

2. Any new position in the market has a margin of 4%

3. Only one market order in the queue: Order2.1 = ((O_{12}, sell, Inf, 5000))

This is the same order that we used in Example 1. When it is cleared, as explained above, the bid queue will be changed. The state will be changed to:

State2.2 = (Bid_Queue2.2, Offer_Queue2.2, Positions2.2)
Bid_Queue2.2 = ((O_6, 1.58, 2000), (O_7, 1.57, 1500), (O_8, 1.56, 4000))
Offer_Quene2.2 = Offer_Queue2.1
Positions2.2 = (($P(O_1)$, long, 4000, 6600, 1.65 4%),
($P(O_2)$, long, 2000, 3280, 1.64, 4%),
($P(O_3)$, long, 2000, 3280, 1.64, 5%),
($P(O_4)$, long, 2500, 4000, 1.60, 4%),
($P(O_5)$, long, 2000, 3180, 1.59, 4%),
($P(O_6)$, long, 500, 790, 1.58, 4%),
($P(O_{12})$, short, 5000, 7970, 1.594, 4%))
Where:
$P(O_{12})$ Value = (1.6×2500)+(1.59×2000)+(1.58×500)= 7970
$LastTP$= 1.58 (the price of the last matched order in the Queue_Profile)

At this point, the bid queue and the position $P(O_1)$ together will trigger a new market order. This is because 1.65×(1−4%) = 1.584, which is above the last transaction price, which was 1.580. Therefore, the margin is exceeded, and this position must be closed (Assumption 2.1). That means the order queue will be changed to:

Order 2.2 = ((O_{13}, sell, Inf, 4000))

The following transactions take place:

2000 will be transacted at 1.58
1500 will be transacted at 1.57
500 will be transacted at 1.56

This will change the state to:

State 2.3 = ((Bid_Queue2.3, Offer_Queue2.3), Positions2.3)

Bid_Queue2.3 = ((O8, 1.56, 3500))

Offer_Quene2.3 = Offer_Queue2.2

Positions2.3 = ((P(O2), long, 2000, 3280, 1.64, 4%),
(P(O3), long, 2000, 3280, 1.64, 5%),
(P(O4), long, 2500, 4000, 1.60, 4%),
(P(O5), long, 2000, 3180, 1.59, 4%),
(P(O6), long, 2500, 3950, 1.58, 4%),
(P(O12), short, 5000, 7970, 1.594, 4%),
(P(O7), long, 1500, 2355, 1.57, 4%),
(P(O8), long, 500, 780, 1.56, 4%))

Where:

$LastTP$= 1.56

Note that order O_6 has opened a new position $P(O_6)$ in State2.2. However, it was only partially matched. In State2.3, O_6 is fully matched. Thus, we do not open a new position but we update the already opened position $P(O_6)$.
The long position $P(O_2)$ must be closed when the last transaction price (1.56 in this case) falls below its margin, which is $1.64 \times (1-4\%) = 1.574$. This means the order queue will be updated by the new market order:

Order2.3 = (O_{14}, sell, Inf, 2000)

When the order (sell, Inf, 2000) is matched, 2000 will be transacted at 1.56. This will reduce the state to:

State 2.4 = ((Bid_Queue2.4, Offer_Queue2.4), Positions2.4)

Bid_Queue2.4 = ((O8, 1.56, 1500))

Offer_Queue2.4 = Offer_Queue2.3

Positions2.4 = ((P(O3), long, 2000, 3280, 1.64, 5%),
(P(O4), long, 2500, 4000, 1.60, 4%),
(P(O5), long, 2000, 3180, 1.59, 4%),
(P(O6), long, 2500, 3950, 1.58, 4%),
(P(O12),short, 5000, 7970, 1.594, 4%),
(P(O7), long, 1500, 2355, 1.57, 4%),

(P(O8), long, 2500, 3900, 1.56, 4%))

Where:

$LastTP$= 1.56

Note that order O_8 has opened a new position $P(O_8)$ in State2.3. However, O_8 was only partially matched. In State2.4, O_8 is fully cleared. Thus, we update the already opened position $P(O_8)$. The position $P(O_3)$ will only be closed when the last transaction price falls below $1.64 \times (1-5\%) = 1.558$.

To summarize, a single market order of 5000 units led to the closure of two positions, which led to a total clearance of 11000 units, and a drop of 2.5% (from ≥1.60 to 1.56) in the market. It should be useful to compute, given a particular state of the market, how big an order is needed to drop the price by, say, 10%.

Besides, what would happen if the (P(O3), long, 2000, 3280, 1.64, 5%) position has a 4% margin, instead of 5%? This will mean that this position has to be closed, but only 1500 of the 2000 will be bought (by the last bid in the queue); the remaining 500 units will not be cleared. The analysis of these properties goes beyond the scope of this simple calculus

6.3.3 Market Making

The market maker is an aggregator who nets the flow of buyers and sellers. His profit is a reward for managing the uncertainty of this process. He manages the flow by dynamically skewing bid and ask prices. The market maker sets the "bid" and "ask" price on a tick-by-tick basis. The bid price is the price at which the market maker offers to buy; the ask price is the price at which the market maker offers to sell.

State = (Bid_price, Ask_price, MaxVol, Queue_Profile, Position_Profile)

Where:

Bid_price and Ask_price are the bid and ask prices quoted by the market maker;

MaxVol is the maximum volume that the market maker is willing to deal per order;

Queue_Profile and Position_Profile are the same as those defined in Model 2.

Here we assume that the clearing mechanism is completely automated. The key to the clearing mechanism is in the way that the market maker updates its bid and ask prices.

In this research scope, we make no assumption on f, which could vary from market maker to market maker; f should be a complex function.

Let Bid_price and Ask_price be the bid and ask prices in the current state, and Bid_price' and Ask_price' be the bid and ask prices in the next state. We generalize that the market maker sets the Bid_price' and Ask_price' with a function f, without specifying exactly what f is. f is a function that involves Bid_price, Ask_price, Queue_Profile, Position_Profile and many other factors, which may include the market maker's own position, bid and ask prices by the other market makers, the balance of payment between countries, interest rates, news and other economic indicators of the countries involved.

$$
\begin{array}{ll}
\multicolumn{2}{l}{(\text{Bid\_price, Ask\_price, MaxVol, (Bid\_Queue, Offer\_Queue), Positions})} \\
\quad +(\text{sell, P, V})\rightarrow & \\
\multicolumn{2}{l}{(\text{Bid\_price', Ask\_price', MaxVol, (Bid\_Queue, Offer\_Queue), Positions})} \\
 & \text{if } P \leq \text{Bid\_price} \ \& \ V \leq \text{MaxVol} \\
\multicolumn{2}{l}{(\text{Bid\_price', Ask\_price', MaxVol, (Bid\_Queue, Offer\_Queue), Positions})} \\
\quad + (\text{sell, P, V}-\text{MaxVol}) & \\
 & \text{if } P \leq \text{Bid\_price} \ \& \ V > \text{MaxVol} \\
\multicolumn{2}{l}{(\text{Bid\_price', Ask\_price', MaxVol, (Bid\_Queue, Offer\_Queue} \oplus} \\
\quad \text{(sell, P, V)), Positions)} & \\
 & \text{if } P > \text{Bid\_price}
\end{array}
$$

The queue joining operator $\oplus$ is defined in the Model 1 section. For any well specified f, we should be able to formalise market making.

6.4 Assessing Market Risk

In this section, we propose two methods for assessing market risk and stability given the formal logic described above. The first method evaluates market fragility at any point of time by computing the consequential closure and the resulting price drop of any potential sell order placed into the underlying market state. The second method assesses market stability by measuring market liquidity, in other words, the extent to which the market can absorb a new order.

6.4.1 Consequential Closure

One can compute the consequential closure with respect to margin constraints. By doing so, one can evaluate the final state of any given event. For example, one would be able to say that "a market order to sell 6 million will lead to a price drop of 4%". One may also compute the condition for minimum price changes, e.g. "What is the minimum size of a market sell order to lead to a price drop of r%?" Answering questions like this would help to assess the stability of the market and value at risk. It could provide early warnings.

An algorithm as outlined below returns the volume of a market sell order that would lead the price to drop to or below price P_{drop}. This function traverses the bid queue and examines the effect of hypothetical market sell orders on the underlying market state, with respect to traders' positions and their margin constraints. The function takes three inputs; the Queue_Profile and the Positions_Profile of the underlying market state and the desired P_{drop}. In each iteration of the function, a new market sell order is placed to walk through the bid queue until P_{drop} is reached.

The market sell order is fed into to the procedure closure (Queue_Profile, Position_Profile, Order). The only variable input to closure is the market order, as it has a different volume in each iteration. The procedure computes the resulting Queue_Profile' after consequential closure is maintained using the calculus shown in the Model 2 Section[3]. This involves matching the market sell order with the bid queue; updating the market positions profile; updating and sorting the queue profile; checking for margin calls and its consequential forced positions closures, while keeping record of the last transaction price.

[3] Strictly speaking, the termination condition P1' $\leq P_{drop}$ should be replaced by LTP $\leq P_{drop}$, where LTP is the Last transaction price which could be returned by the closure function. This is simplified for clarity. When the head of the queue in Queue_Profile' is below P_{drop}, any market order to sell will drop the price below P_{drop}. Therefore, the Volume returned is correct, which is our justification for the compromise.

Algorithm 6.1 Market Risk Assessment

Function MinDrop(Queue_Profile, Position_Profile, Pdrop)

Begin

/ Let Queue_Profile = (Bid_Queue, Offer_Queue)*

*If Bid_Queue is not empty, let it be ((P_1, V_1), (P_2, V_2), . . . , (P_{bq}, V_{bq})) */*

i ← 1; Volume ← 0;

Bid_Queue' ← Bid_Queue;

// Bid_Queue' is a working structure; if it is not empty, then let its head be (P_1', V_1')

While P_1' > P_{drop} and Bid_Queue' is not empty

If ($Vi \leq V_1$') *// V_i is the volume at index i of Bid_Queue*
Then {Volume ← Volume + V_i; i ← i + 1}
Else Volume ← Volume + V_1'; *//See if incrementing Volume by V_1' makes any difference*
Queue_Profile' ← closure(Queue_Profile, Position_Profile, (offer, Inf, Volume));
(P_1', V_1') ← Head of the bid queue in Queue_Profile'

End While

If P_1' > P_{drop} Then report that P_{drop} cannot be reached in this market as Bid_Queue is exhausted

Return Volume;

End

The Queue_Profile' is a working structure, which is discarded on exit. It is used to define the potential price P_1' (head of the bid queue in Queue_Profile'), the market would reach after executing the market sell order. If $P_1'>P_{drop}$, the colure procedure is called again to evaluate the impact of a bigger market sell order. The algorithm increments i (which has the effect of increasing volume) until enough volume is accumulated to see the price drop to P_{drop}. The function will terminate when P_{drop} is reached or when the bid queue is completely cleared. Once terminated, the function returns the Volume required to reach P_{drop}, giving a preview of the potential multiplied effect on the underlying market once a market sell order of a specific volume is placed.

If the market does not have enough depth, all the buy orders will be exhausted before P_{drop} is reached. Otherwise, there exists a minimum k such that, for all the orders (P_i, V_i) at the front of the Bid_Queue, $P_{drop} \leq P_i$ and Volume $\leq V_1+V_2+...+V_k$. In the worst case, Function MinDrop has to go through all such (P_i, V_i)s[4] . Volume increases monotonically in Function MinDrop. Therefore this function must terminate.

Let M be the list of positions in the Position_Profile which margin calls are above P_{drop}. In the worst case, the procedure has to go through all of them. So each cycle of the "While" loop will have complexity of $|M|$. Each "Then" part in each cycle of the "While" loop would increase Volume to include one (Pi, Vi) pair. It is more complex to analyse the number of times that the "Else" part could be entered. In the worst case, each of the positions could bring the loop into the "Else" part through a margin call. Therefore, the complexity of the algorithm is bounded by $O(k \times |M|^2)$.

6.4.2 Market Liquidity

Coherent measures of risk have been proposed by [13]. This was scrutinized by [1], for not taking full consideration of liquidity risk. [1] introduced the marginal supply-demand curves (MSDCs), which defines at any time instance the available prices of a given asset in the market. The attractiveness of their formalism is that liquidity risk is measured by market data; no assumptions are required. Figure 6.2 shows the MSDC in State2.1. After clearing of Order2.1, the market loses a certain amount of liquidity in State2.2. This is shown by MSDC in Figure 6.3. Like [1], we are looking at the microstructure of illiquid markets, and free from hypotheses on the dynamics of the market.

[4] This is an upper-bound because any margin calls that might be triggered will absorb some of the volume.

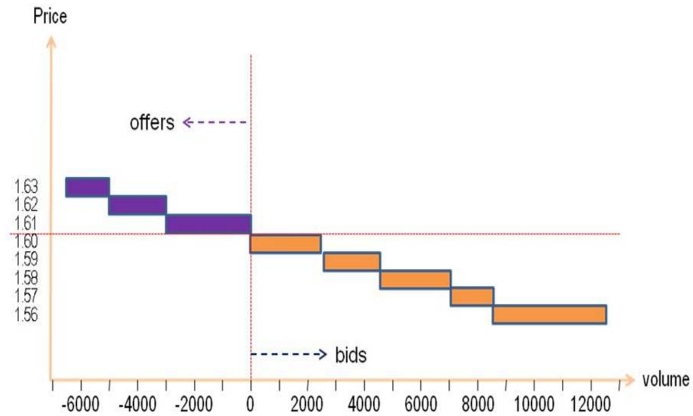

Figure 6.2: The Marginal Supply-Demand Curve defined by the Queue_Profile at State2.1

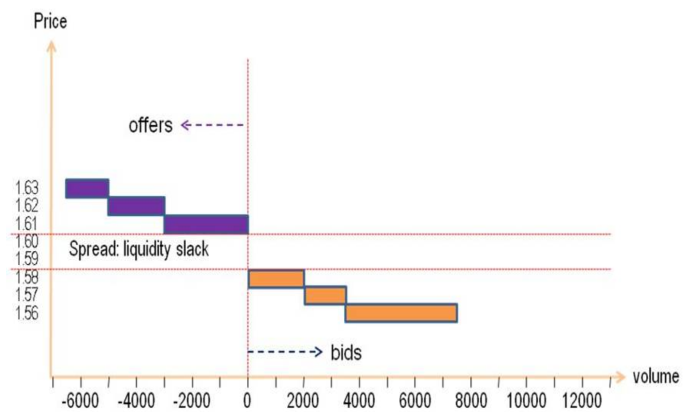

Figure 6.3: The Marginal Supply-Demand Curve defined by the Queue_Profile at State2.2

The work by [1] is based on the concept of mark-to-market. When position and margin information are not considered (Model 1 in section 6.3.1), the shape of the MSDC curve depends on queue profiles alone. When position and margin information are available, the mark-to-market values are changed. In fact, the shape of the MSDC could be changed by the orders processing procedure of Model 2 explained in section 6.3.2. Therefore, this study complements the work done by [1].

The queue profile defines how liquid an asset is at any given time. Liquidity of an asset is therefore determined by how steep one ascends or descends in the MSDC. Following the above example, suppose at State2.1, two traders bid 1.60 for 500 shares, and 1.59 for another 500 shares. Although the highest bid price is still 1.60, the new MSDC is actually steeper than the one shown in Figure 6.2. This means, to sell over 1000 shares in this market (as opposed to the market shown in Figure 6.2) , the seller must be prepared to accept lower bids.

Unfortunately, the ordinary investors/traders who have no access to order books have no means of fully assessing their liquidity risk [5]. Therefore, market making provides investors/traders with market liquidity up to a certain limit (MaxVol in section 6.3.3 of the market making model). It also offers transparency in market liquidity. The MSDC under market making is shown in Figure 6.4.

It is worth noting the obvious that, as a Queue Profile does not have to be symmetric, an asset could be highly liquid when one wants to buy, but illiquid when one wants to sell (and vice versa).

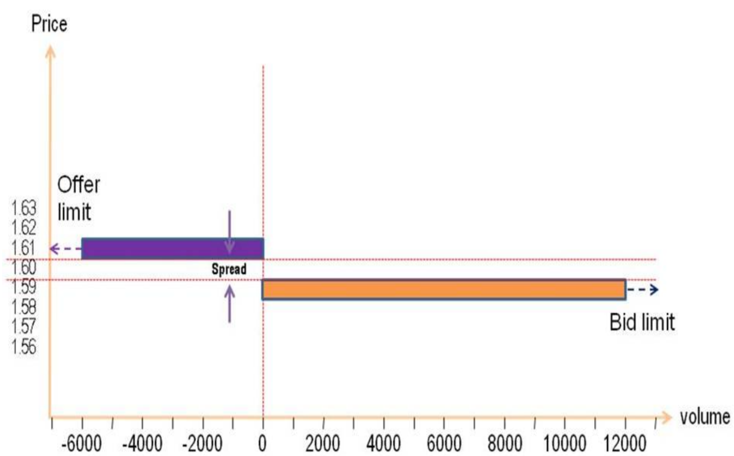

Figure 6.4: The Marginal Supply-Demand Curve defined by the Queue Profile at State2.3

[5]OANDA provides information on trader positions. This could help conjecturing (with low confidence) marginal supply and demands (because eventually those in long positions have to sell, and those in short positions have to buy).

6.5 Simulating Clearance Models and Market Risk Assessment

In this section, we build a simulated market that is an exact implementation of the calculus described earlier. We briefly introduce the simulator and one simple comprehensive experiment to illustrate its features. The simulator is named CDAM (Calculus Double Auction Market) and is available online[6].

6.5.1 CDAM Description

CDAM is a laboratory for demonstrating the clearance mechanism in a double auction market; monitoring and analysing the different market state transitions at any point of time; explaining the consequences of new orders coming into the market; and assessing market risk and fragility through examining the consequences of potential orders of any size. The purpose is to demonstrate the potentiality of this new market-calculus approach in understanding the market microstructure. Figure 6.5 depicts a snapshot of this demo simulator.

CDAM provides all relevant information of the formalised market clearance mechanism presented in the previous sections. This includes the basic components describing the market state:

- The order book (Queue_Profile) which is composed of the bid and offer queues. The setup of the order book is discussed further in section 6.5.2.
- The market positions (Position_Profile) which include information about the position code (number), type, volume, value, unit price and trading margin.
- A graph depicting the transaction prices, at which the orders are executed, to monitor the market price movement. The graph also depicts the different bid and offer (ask) prices of the order book.

[6]CDAM URL: http://cswww.essex.ac.uk/CSP/demos/eventcalcsim/

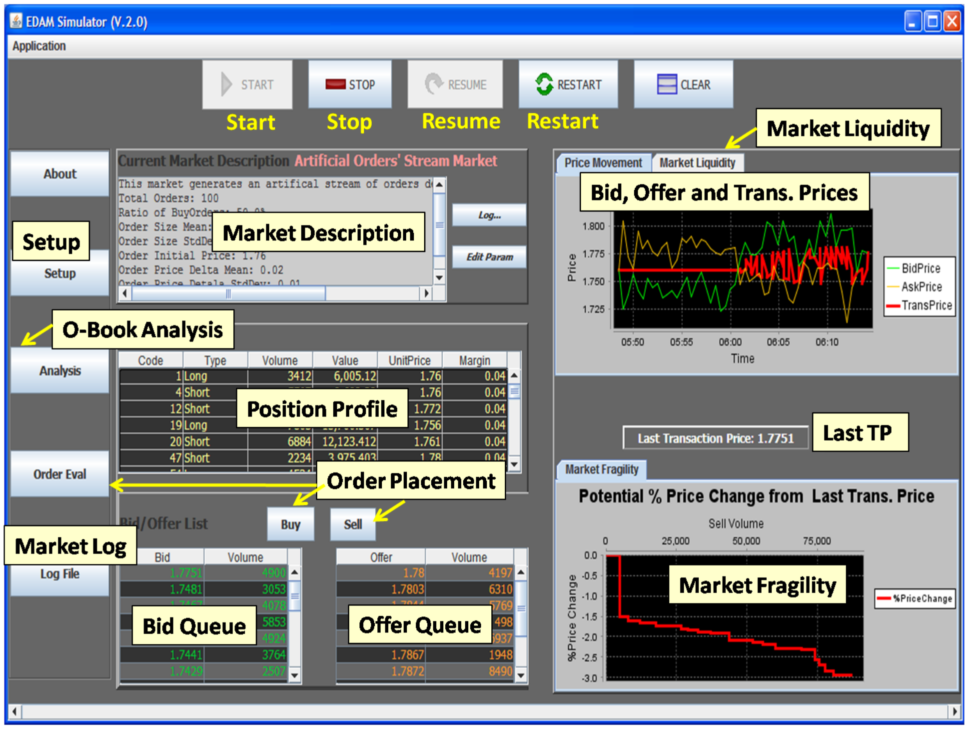

Figure 6.5: Screenshot of the double-auction market simulator CDAM. Available online at: http://cswww.essex.ac.uk/CSP/demos/eventcalcsim/

The user can monitor and analyse the different market state transitions at any point of time by running (starting/restarting), stopping and resuming the simulator. Stopping and resuming the simulator market clearance allows for understanding step-by-step the transition of the market from one state to another. When stopping the simulation run, the user can also observe what would happen to the underlying market state under hypothetical large orders. A step-by-step description of the market transition from one state to another be viewed in the "Market Log" of CDAM.

The user can evaluate the market risk through the fragility and liquidity graphs described in section 6.5.3. The user can also query the minimum size of a market sell order that would lead to a price drop of r%, which is further discussed in section 6.5.3.

6.5.2 CDAM Order Book Setup

Before running the simulator, the user has to setup first the market order book, by selecting one of the following three simulation modes available in CDAM:

- The built-in sample order book of CDAM.
- An order book built by an artificial stream of orders.
- A custom-defined order book, where a batch of orders are entered by the user as a script file.

Each of these simulation modes requires some input parameters to make this demo-market as customizable as possible. It is important to note that in parallel to these simulation modes, the user can manually place limit and market orders at any point of time.

As mentioned in the calculus above, we assume that a trader can only open and close but not adjust his/her position (Assumption 2.3 in section 6.3.2). Hence, each order entered into the market is handled as a new trader. If the order is executed, a position is opened for this trader and is coded by the opening order number. Following Assumption 2.1 and Assumption 2.2 in section 6.3.2, positions closure take place through margin calls and all consequences of an automatic position closure take place before any new event occurs.

1. Parameters of Built-In Market:

   This market includes four built-in scenarios to be run on the market.

   The "Default Scenario" represents a simple list of limit orders (buys and sells) to be executed on the market. Matched buys and sells do not join the order book and create market positions. Unmatched (full or partially) limit orders join the order book, which is sorted according to a price priority basis as explained in section 6.3.

   The three other scenarios analyse the effect of placing a large order on the market dynamics and fragility.

   Only two parameters are needed for this market; the scenario to run and the margin of traders entered into the market.

2. Parameters of Artificial Orders' Stream Market:

   In this mode, the simulator generates a stream of artificial limit orders (buys and sells) to build the order book. This stream of limit orders is generated using the following parameters:

   - Number of Orders/Traders: Unlike many other market models, we do not fix the number of traders. In each simulation run, the user can choose any value for the total number of traders that will be automatically entered into the market. These traders will place limit orders, buys (bids) and sells (offers), to constitute the market order book. As for market orders, they can be either manually entered by the user or automatically forced through margin calls.
   - Probability of Buy Orders: This determines if with equal or different probabilities a new trader is a buyer or a seller. With probability q_{buy} the trader places a limit order to buy and with probability 1- q_{buy} he places a sell limit order.
   - Size of Order: The size of order is generated from a probability distribution of two input parameters mean μ_{size} and standard deviation σ_{size}.
   - Price of Limit Order: Each order price is determined by simply offsetting the last transaction price (LTP) by a random number Δ. This positive random number is drawn each time from the same probability distribution $P(\Delta)$. To make sure that the market is liquid enough, we do not allow for the price ranges of limit buy and sell orders to overlap for 50% of the total orders. After that, we allow for overlapping to observe some executed transactions on the market. To avoid overlapping in the beginning, we set the limit order (buy/sell) price as follows:

     A new sell order (offer) is placed above the current market price $(LTP + \Delta)$, while a new buy order (bid) is below the market price by $(LTP - \Delta)$. This way there is always a positive gap between the highest bid at the head of bid queue and the lowest offer at the head of the offer queue. In the case of overlapping, we make a slight modification by signing Δ. Each time a new order is generated, the sign of Δ is offset and the price of both bid and offer orders is computed by $(LTP + \Delta)$.

Input parameters entered by the user to determine the price are: $initialPrice$ of first bid/offer entered into the market; and mean of Δ $\mu_{\triangle}$ and standard deviation$\sigma_{\triangle}$of the probability distribution $P(\Delta)$.

- Margin: Each order entered into the market represents a new trader, having a specific margin. This parameter sets a default margin, unless changed by user, for any new trader entering the market.

3. Parameters of Custom Market:

   The custom market is mainly designed for developing market scenarios that the user would like to analyse. Using a very simple script, the user can write a text file containing a batch of limit and market orders that can be executed on the market. To run the custom market the following parameters (script commands) are required: type of order (limit/market), price of limit order, size of order and margin of trader. The user can also add a “STOP” command at any point of the file to force the simulator to pause there. The simulator can then be resumed from the application interface to execute the remaining orders in the file. For future work, given that we have access to real market data, we can enhance this custom market module to rebuild a real market order book.

6.5.3 CDAM Features

In this section, we describe the main features of CDAM and demonstrate how the simulation of the market-calculus model facilitates the understanding of the dynamic market micro-behaviour.

1. ***Simulation of Market Clearance and State Transitions***

   The clearance in CDAM applies the same rules of the market-calculus model presented in section 6.2. Each time a new limit or market order is placed, a sequence of updates of the order book (Queue_Profile), the market positions (Position_Profile), the last market transaction price (LTP), the market liquidity and the market fragility is carried out. The latter two are further explained in the coming section. This update process can be repeated over and over in case margin calls are triggered, which would force further market orders to close out open market positions. The following code fragment shows the different functions called whenever a new limit order LO is entered into the market.

```
procedure addLimitOrder(LO)  \\LO = (Order_No,
Order_Type, Price, Volume)
begin
    if (Order_Type=="sell")
        /*traverses bid queue to match the limit sell order and
        to open the relevant market positions. The last
        transaction price LTP is returned .*/
        LTP=matchLimitSellToBidQueue(LO);
    else if (Order_Type=="buy")
        /*traverses offer queue to match the limit buy order
        and to open the relevant market positions. The last
        transaction price LTP is returned .*/
        LTP=matchLimitBuyToOfferQueue(LO);
    updateBids(); //sorts Bids
    updateOffers(); //sorts Offers
    updateLiquidity();
    updateMarketFragility()
    checkMarginCalls(LTP);
end
```

Once this update process is executed and margin constraints are considered with every new order, a new market state transition takes place. The consequences of the new order on the order book, market price and traders' positions can be clearly visualised on CDAM .

In case of entering many orders into the market through running one of the simulation modes discussed above, the different market state transitions can be traced through stopping and resuming the simulator. For instance, assume a simulation run of 1000 traders in

the artificial orders' stream mode. The user can choose to stop and resume this run at any point of time to analyse the market transition from one state to another. To analyse the extent to which the underlying market state (stopping point) would change under certain conditions, the user can manually enter one or more orders and watch their effect on the underlying market state. The original simulation run of the artificial orders' stream can be then resumed, which will be appended to the last manually entered order. The visualization and the ability of analysing each of the different market state transitions under different conditions enables precise studying and understanding of the financial market micro-behaviour.

2. ***Simulation of Market Risk Assessment***

   - ***Market Liquidity***

As mentioned earlier in section 6.4.2 , [1] introduced a recent interesting quantitative approach for assessing market liquidity, that requires no assumptions. They use marginal supply-demand curves (MSDCs) to summarize all market information (bids and offers prices) of an asset at any given time. The same idea is used here, since the Queue_Profile (bid and offer queues) can define the liquidity of an asset at any point of time.

In real life, investors have no access to the order book and hence no means for assessing their liquidity risk. However, the extent of market transparency, the ability of market participants to observe information about the trading process, has been always one of the strong debatable issues. Since in CDAM, all data is available, market liquidity is constantly shown. It is continuously updated after processing every new order placed into the market. As shown in Figure 6.8 liquidity curves summarize the total available units at each different price level for both sides, bids and offers. Similar to the order book, bid prices are ordered ascending-wise, whereas prices on the offers side have a descending order. Having this information available on the spot, allows for analysing market stability, in terms of the extent to which the market can absorb more orders at any point of time.

   - ***Market Fragility***

With CDAM we can examine the effects of hypothetical orders at any point of the simulation run, in order to evaluate how fragile the underlying market state is. Market fragility helps assessing the stability and the risk of the current market with respect to the current order book, traders' positions and their margin constraints. In CDAM, we examine market fragility using the "Fragility Curve" feature (Figures 6.7 and 6.8). It can answer questions like "What is the minimum size of a market sell order to lead to a price drop of r%?". The price drop is the potential change of the market price relative to the last transaction price (LTP).

To generate the "Fragility Curve", we clone the underlying current market state (Position_Profile, Queue_Profile, LTP) into a new one. The analysis of the market fragility is then performed on this cloned state. To assess the market fragility, we traverse the order book (Queue _Profile) of the cloned market version by executing "potential" market sell orders. Each "potential" market sell order placed, has the same size as the underlying head of the bid queue of the cloned market.

The clearing mechanism is exactly the same. Once a market sell order is placed, we process the cloned market state transition by updating its market positions (Positions_Profile), order book (Queue_Profile), execution price of sell order (LTP), and check for margin calls. All consequences of clearing the head of the bid queue of the cloned market take place before any potential market sell order is placed to clear the new head of the bid queue of the cloned market.

The transaction price at which each "potential" market sell order has been executed in the cloned market is then compared with the actual market last transaction price (LTP) to compute the percentage change (drop). The "Fragility Curve" is then depicted by plotting each percentage price drop resulting from executing each "potential" market sell order.

It is important to note that once the actual market state changes as a result of any new event, the "Fragility Curve" is re-computed and re-plotted on the spot by cloning the new actual market state and by repeating the clearance of the new BID_Queue of the cloned market.

In the following section, we carry out an experiment illustrating the usage of the "Fragility Curve". We show that market fragility can be used as an early warning signal to prevent possible financial crisis. Even though the market-calculus model used here is simple, however its precise and mathematical definition of the market states has allowed for developing a mean for assessing market risk.

6.5.4 Experiment

In this section, we present one inclusive experiment examining the consequences of entering new orders into the market, the resulting state transitions and the associated market risk and stability.

For illustration purposes, we use the simple built-in market of CDAM. As mentioned earlier, this market has built-in scenarios for orders placement. We use the "Default Scenario", where every half a second a limit order (buy or sell) is entered into market. For all traders we use a margin of 4%. The different market state transitions can be viewed at any point of time by stopping and resuming the market clearance. Table 6.1 lists all the limit orders entered into the market. They are sorted according to their codes, indicating the sequence in which each sell/buy order has been entered into the market.

Table 6.1: Limit Orders of the built-in market of CDAM

Limit Sell Orders			**Limit Buy Orders**		
Code	**Price**	**Units**	**Code**	**Price**	**Units**
1	1.8	2000	**2**	1.69	2000
5	1.78	2500	**8**	1.68	3000
12	1.76	8000	**16**	1.7	1000
21	1.79	2000	**26**	1.71	2000
32	1.82	3000	**38**	1.72	3000
45	1.83	2000	**52**	1.73	2000
60	1.86	3000	**68**	1.74	4000
77	1.783	2000	**86**	1.67	4000
96	1.75	5000	**106**	1.712	1000
113	1.73	3000	**128**	1.75	6000
141	1.77	1000	**154**	1.76	5000
167	1.84	2000	**180**	1.66	2000
193	1.85	1000	**206**	1.65	3000

Once each limit order is entered into the market, it joins the order book (Queue_Profile), in case not matched with any previously entered order. The order book is then sorted according to a price time priority basis, where the lowest offer is at the head of the Offer_Queue, and the highest bid is at the head of the Bid_Queue. Once an order is fully or partially fulfilled, the market order book, positions, price graph, fragility and liquidity are updated.

In this case here, all limit orders placed in the "Default Scenario" join the order book and are not executed up till order code: O(106). In this experiment, we shall present the main market states since entering the buy order O(106). The details of the market states can be also viewed in the CDAM "Market Log", when running the simulator.

1. **State1.1**

   In this state, the market is stopped right after entering offer order *O(106)*. At this point, no orders have been executed (fulfilled) yet. As a result, there are no open positions on the market and the current market state shows the following Queue_Profile (order book) along with the bid/offer (ask) price curves (Figure 6.6).

   Offer_Queue1.1:
   [**96** 1.75 5000 **12** 1.76 3000, **5** 1.78 2500, 77 1.783 2000,
   21 1.79 2000, **1** 1.8 2000, **32** 1.82 3000, **45** 1.83 2000,
   60 1.86 3000]

   Bid_Queue1.1:
   [**68** 1.74 4000 **52** 1.73 2000, **38** 1.72 3000**,** **106** 1.712 1000**,**
   26 1.71 2000, **16** 1.7 1000, **2** 1.69 2000, **8** 1.68 3000,
   86 1.67 4000]

Figure 6.6: Bid and ask (offer) prices of the order book at State1.1. Since both price curves do not overlap, no orders are executed.

2. **State1.2**

   Resuming the simulation run to allow for the remaining limit buy and sell orders (O(113)-O(206)) to enter and execute, State 1.2 comes to place. At this point, several orders have been executed resulting in several open market positions. However, no margin calls have been triggered. The market state after placing the last buy order O(206) is shown below:

   Bid_Queue1.2:

   [**128** 1.75 1000, **68** 1.74 1000, **52** 1.73 2000, **38** 1.72 3000,
   106 1.712 1000, **26** 1.71 2000, **16** 1.7 1000, **2** 1.69 2000,
   8 1.68 3000, **86** 1.67 4000 **180** 1.66 2000 **206** 1.65 3000]

   Offer_Queue1.2:

   [**12** 1.76 3000, 141 1.77 1000 **5** 1.78 2500, 77 1.783 2000,
   21 1.79 2000, **1** 1.8 2000, **32** 1.82 3000, **45** 1.83 2000,
   167 1.84 2000, **193** 1.85 1000 **60** 1.86 3000]

   Positions_Profile1.2:

   [**12** Short 5000 8800 1.76 0.04, **68** Long 3000 5220 1.74 0.04,
   96 Short 5000 8750 1.75 0.04, **113** Short 3000 5220 1.74 0.04,
   128 Long 5000 8750 1.75 0.04, **154** Long 5000 8800 1.76 0.04]

   Last Transaction Price (LTP): 1.76

The transaction prices, at which these orders have been executed, can be observed from the transaction price red curve in Figure 6.7(a). Examining the stability of the current state through the "Fragility Curve" (Figure 6.7(b)), we find that the market becomes very fragile and a sudden price drop can appear, once a sell order of size >**13000 units** is placed. To test this, we first place a market order of size 13000 units and observe the new state transition (State 1.3).

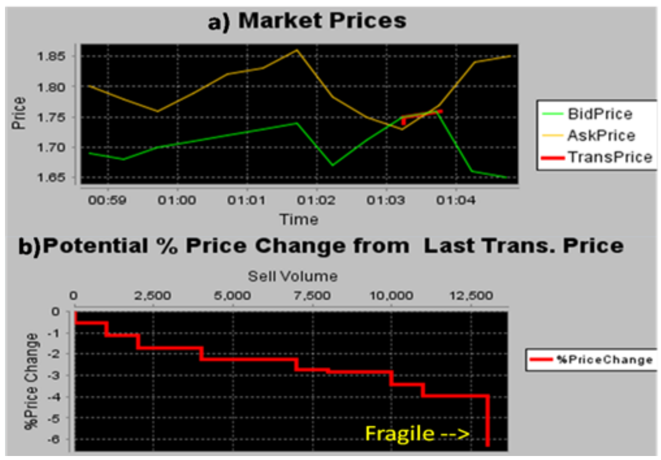

Figure 6.7: Market prices (a) and market fragility (b) at State1.2

1. **State1.3**

   State1.3 comes to place through the execution of the market sell order (221, sell, Inf, 13000). The current order book (Queue_Profile) and market positions (Position_Profile) are described below.

   Bid_Queue1.3:
   [**8** 1.68 3000, **86** 1.67 4000, **180** 1.66 2000, **206** 1.65 3000]

   Offer_Queue1.3:
   [**12** 1.76 3000, **141** 1.77 1000, **5** 1.78 2500, **77** 1.783 2000,
   21 1.79 2000, **1** 1.8 2000, **32** 1.82 3000, **45** 1.83 2000,
   167 1.84 2000, **193** 1.85 1000, **60** 1.86 3000]

Positions_Profile1.3:

[**2** Long 2000 3380 1.69 0.04, **12** Short 5000 8800 1.76 0.04, **16** Long 1000 1700 1.7 0.04, **26** Long 2000 3420 1.71 0.04, **38** Long 3000 5160 1.72 0.04, **52** Long 2000 3460 1.73 0.04, **68** Long 4000 6960 1.74 0.04, **96** Short 5000 8750 1.75 0.04, **106** Long 1000 1712 1.712 0.04, **113** Short 3000 5220 1.74 0.04, **128** Long 6000 10500 1.75 0.04, **154** Long 5000 8800 1.76 0.04, **221** Short 13000 22322 1.717 0.04]

LTP: 1.69

As shown in Figure 6.8(a), the market is now extremely unstable. After clearing the market sell order, the market has lost a great amount of liquidity. The market is now more liquid for buyers than for sellers. Also, the spread (0.08) between the best bid (1.68) and the best offer (1.76) is very high. This means, to execute more sell orders, the seller must accept lower bids.

As for the market fragility in Figure 6.8(b), placing an order of a minimum of one unit would cause a **price drop of >2%**. The only explanation for this potential behaviour is that a martingale of margin calls will occur. To test, if the market behaves as anticipated by the fragility curve, we place a market order of one unit to State1.3 and check for margin calls. This would lead to a new state transition (State1.4).

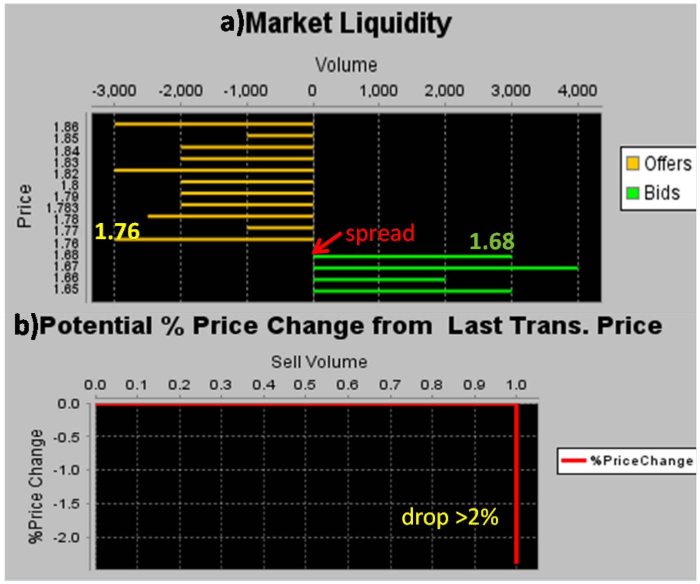

Figure 6.8: Market liquidity (a) and market fragility (b) State1.3

2. **State1.4**

   Placing the one unit market order (228, sell, Inf, 1) to State1.3 leads to the following state transition (State 1.4):

   Offer_Queue1.4= Offer_Queue1.3

   Bid_Queue1.4:

   [**8** 1.68 2999, **86** 1.67 4000, **180** 1.66 2000, **206** 1.65 3000]

   Positions_Profile1.4:

   [**2** Long 2000 3380 1.69 0.04, **12** Short 5000 8800 1.76 0.04,
   16 Long 1000 1700 1.7 0.04, **26** Long 2000 3420 1.71 0.04,
   38 Long 3000 5160 1.72 0.04, **52** Long 2000 3460 1.73 0.04,
   68 Long 4000 6960 1.74 0.04, **96** Short 5000 8750 1.75 0.04,
   106 Long 1000 1712 1.712 0.04, **113** Short 3000 5220 1.74 0.04,
   128 Long 6000 10500 1.75 0.04, **154** Long 5000 8800 1.76 0.04,
   221 Short 13000 22322 1.717 0.04 **8** Long 1 1.68 1.68 0.04,
   228 Short 1 1.68 1.68 0.04]

   LTP: 1.68

Checking the margin constraints of the different traders' positions, we find that a margin call is triggered to close position P(154), as $1.76 \times (1 - 0.04) = 1.689$ is greater than 1.68, the last transaction price in State1.4 (see margin rule in section 6.3.2). As a result, a market sell order is automatically placed to force closure of position P(154), leading to State1.5

3. **State1.5**

   To close P(154) a market sell order (235, sell, Inf, 5000) is automatically placed, which leads to the following state transition:

   Offer_Queue1.5= Offer_Queue1.4

   Bid_Queue1.5:
   [**86** 1.67 1999, **180** 1.66 2000, **206** 1.65 3000]

   Positions_Profile 1.5:
   [**2** Long 2000 3380 1.69 0.04, **12** Short 5000 8800 1.76 0.04,
   16 Long 1000 1700 1.7 0.04, **26** Long 2000 3420 1.71 0.04,
   38 Long 3000 5160 1.72 0.04, **52** Long 2000 3460 1.73 0.04,
   68 Long 4000 6960 1.74 0.04, **96** Short 5000 8750 1.75 0.04,
   106 Long 1000 1712 1.712 0.04, **113** Short 3000 5220 1.74 0.04,
   128 Long 6000 10500 1.75 0.04, **221** Short 13000 22322 1.717 0.04,
   8 Long 1 1.68 1.68 0.04, **228** Short 1 1.68 1.68 0.04]

   LTP 1.67

   Checking for margin calls, we find that two further positions, P(68) and P(128), have to be closed, as both positions violate the margin constraints. For P(68) $1.74 \times (1 - 0.04) = 1.6704$ is greater than LTP. For P(128) $1.75 \times (1 - 0.04) = 1.68$ is greater than LTP.

   This automatically forces the placement of two market sell orders of size 4000 and 6000 to close both positions (P(68) and P(128)) respectively. This in turn leads to a further price drop, which causes further margin calls and positions closure.

   In this experiment the martingale of margin calls and the resulting forced market orders clear the bid queue. However, real markets are much more liquid and will absorb the market orders. However, the consequences are the same; a dramatic price decline. The price drop can be clearly seen in Figure 6.9, which agrees with what has been spotted by the "Fragility Curve" in State1.3. Thus, the "Fragility Curve" can be a very useful mean to assess the market risk and stability or even provide early signals for potential market crisis.

Figure 6.9: Dramatic price drop resulting from margin calls after State1.3

6.6 Conclusion

In this chapter, we have introduced the first microscopic approach, a market-calculus, for formalising market dynamics. We have defined a calculus for mathematically describing state changes in a market as a consequence of new orders being posted. We base our analysis on simple market models. This is no attempt to predict what new exogenous orders will arrive. The aim is to formalise processes in the market so that we can study their dynamics and properties scientifically.

To illustrate the potential of this approach, we have implemented a demo market simulator (CDAM) that is an exact application of the defined market-calculus model. We show that even with the simple calculus defined, we can analyse feedback loops like cascading margin calls, and ask important questions such as "how big a sell order would push the price down by r%?". This research also supports Acerbi and Scandolo's call [1] to measure liquidity risks with market data.

We acknowledge the fact that state changes in real markets are far more complex than what is described in this chapter. It is up to the participants, including governing bodies, market makers and traders, to define the rules in an unambiguous mathematical and mechanical way. The aim is to create markets with properties that can be studied formally as well as extensively in simulations and validated empirically. Eliminating black boxes and laying the foundations for extensive scientific analysis may be the best way to ensure stability and prevent financial crises.

Part IV

Concluding Remarks

Chapter 7

Conclusions

This chapter provides a summary of the thesis, points out its contributions, and discusses possible future research work.

7.1 Summary

As implied by the thesis title, this research work conducts a microscopic analysis of FX market behaviour and dynamics using event-based approaches. Surveying the literature in chapter 2, reveals that researchers have adopted various approaches to understanding market behaviour and its dynamics. Different aspects of the FX market have been explained either through behavioural finance studies of traders' psychological biases, or micro-structural studies of order flow, or agent-based modelling and simulation of FX market stylized facts. However, many market phenomena have not yet been explained. We believe that the study of traders' tick-by-tick transactions and the analysis of their micro-behaviour would unveil many of the market dynamics and provide an insight into the collective macro-behaviour.

Having this convention, we carry out in this thesis three micro-behavioural studies aiming to understand some of the features of the FX market dynamics. In all studies an event-based analysis of traders' micro-behaviour is carried out and general inferences about market behaviour are made.

The first two studies, presented in chapters 4 and 5, are empirical research based on tracking and analysing a unique high frequency data set of FX traders' real transactions. It covers more than 140 million tick transactions carried out by more than 45,000 FX traders, on an account level, on the online OANDA FX trading platform for over two years. The uniqueness of the data set is owing to the fact that there is no platform except for OANDA that stores the details of traders' transactions and their positions' history over several years, representing the worldwide exposure of traders in FX.

The production process of the data set is described in detail in chapter 3. To the best of our knowledge, we have produced the biggest set of high frequency data ever. We have thoroughly explored the structure of our data set and the unique properties of the OANDA platform. Taking this into account, we have produced a high frequency transaction data set clean from any erroneous or misleading observations that could affect the research validity. We have also validated the quality of data by tracking individual traders' activities from opening to closing their positions, confirming the consistency of the accounts' trades and the reliability of any forthcoming results. Moreover, we have provided empirical evidence that the data set is representative of the whole market.

Having produced such an invaluable data set, the first empirical study is carried out in chapter 4, which revisits the FX market seasonal activity from a microscopic perspective. It examines whether the unique double U-shape signature of the FX market is due to the overlapping world trading sessions that allow FX traders to trade 24hrs/5days a week from any geographical location. In an attempt to answer this question, an event-based procedure is carried out to normalise the time zones of the different traders. The procedure is mainly based on tracking the intraday seasonal activity of each individual account in the data set and identifying its sleeping (inactive) hours during the day over the whole data set period and other time spans. Traders having the same number of sleeping hours are then grouped together and synchronized over one trading time window. The aggregation of the normalised seasonal activities of all traders produces a single U-shape pattern, akin to the ones observed in equity markets, where traders are more active at the beginning and at the end of the trading day. Using the findings of this empirical study, we have been also able to identify traders' preferences regarding their positions' management during the trading day.

Similar to chapter 4, chapter 5 also studies the micro-behaviour of traders by tracking their tick activity, but this time along price movements. Given the high frequency and discontinuous nature of FX tick data, we use intrinsic time, an event-based approach to synchronize market activity with price movements. To define an intrinsic time scale, we adopt the directional change event approach, which focuses on price turning points of a pre-specified threshold. In this chapter we define the price movement events for three thresholds for all currency pairs studied. To allow for a more in-depth analysis, we discretise the different price movement events into further smaller sub-clocks. We then synchronize traders' tick activity of the produced data set in chapter 3 along the sub-clocks of the different directional change events thresholds. Adopting this microscopic approach, we are able to decipher market activity during both sections, the directional change and the overshoot periods, of the different price movement events. We provide empirical evidence for diminishing market liquidity at the end of the overshoot period for all studied currency pairs. We find that the price overshoot stops due to more participants placing counter trend trades. We also provide empirical evidence that this market liquidity pattern has an impact on price ticks moves, which follow the same decaying pattern during overshoot periods. The findings in this study are valid across all magnitudes of price events.

The third microscopic event-based study is carried out in chapter 6. We define market-calculus to formalise market mechanisms and analyse market dynamics and feedback loops like cascading margin calls, which exist in both FX and equity markets. Due to the lack of knowledge of the FX market maker model in the literature, this chapter illustrates the potential of this calculus in the context of a simple double-auction market model. The price trajectory is determined by the present market state - including market order book, traders' positions and margin constraints - and new orders arriving in the market. We look closely at every single action by the trader and study its effect on the market. The objective of this exercise is to reduce ambiguity and get a better understanding of risk scenarios, not to forecast exogenous order flow. Given this formalism, we can compute the impact of orders of any size. We can also compute how big an order needs to be to cause a market crash. Finally, we build a demo simulator (CDAM), based on the same calculus logic, to demonstrate the potentiality of the market-calculus approach in understanding market microstructure. CDAM is a laboratory for demonstrating the clearance mechanism in a double auction market; monitoring and analysing the different market state transitions at any point of time; explaining the consequences of new orders coming into the market; and assessing market risk and fragility through examining the consequences of potential large orders.

7.2 Contributions

The major contributions of this thesis are as follows:

1. We have produced a data set in high frequency finance, storing the full transaction history of more than 45,000 traders on an account level on FX for 2.25 years. As far as we know, this is by far the largest data set of the kind available. Accounting for the complexity of the high frequency data and for the unique properties of the data set dictated by the specific environment of the OANDA platform, we have cleaned the data set from any observations that are not reliable in terms of actual market activity. We have also validated the quality of the data and have proved the representativeness of the data of the global FX market. The production of such a unique and invaluable HF data set is a major contribution for future research exploring the FX market microscopic behaviour.

2. We have demonstrated that the well-known double U-shape pattern is in fact composed of many single U-shape trading signatures. This has been achieved by developing an algorithm for normalising traders' time zones. To the best of our knowledge, no one has been able to empirically demonstrate that the double U-shape pattern is indeed an overlap of single U-shape patterns (characterising single participants) amid traders' behaviour in regulated exchanges. Given these results, we have also been able to identify how traders manage their positions during the day, with positions closures being the main cause for the trading intensity at the end of the trading day.

3. We have provided empirical evidence for decaying market liquidity and price ticks moves during price overshoots. We have also found that a price overshoot stops due to more market participants placing counter trend trades acting against the market force. This discovery indicates that already small imbalances of market activity in large overshoots can alter the price trajectory. To our knowledge, this is the first study that deciphers FX market activity during overshoot periods. This study is also the first of its kind that uses intrinsic time scale, defined by directional change events, for synchronizing market activity. We show that intrinsic time is a powerful tool that enables us to discover recurring market patterns, which otherwise are hidden if we use only physical time.

4. We have also introduced the first microscopic approach, a market-calculus, for formalising market dynamics. The aim is to formalise processes in the market so that we can study their dynamics and properties scientifically. In this calculus, we define the different market states mathematically and demonstrate the consequences of placing an order into the market. With this formalism, feedback loops like cascading margin calls are analysed and questions like "how big a sell order has to be to cause the market to fall by a certain percentage" can be answered. Despite the infancy of this approach and the simple models used, it eliminates black boxes and lays a solid foundation for extensive scientific analysis of complex market models. The calculus would also enable us to study market dynamics through simulation. In practice, it can help us to provide early warnings of market crashes.

7.3 Future Work

This thesis adopts a new approach for understanding the FX market behaviour. Akin to physics and other natural sciences, where findings are generalized from observations, we study market dynamics through observing and analysing traders' micro-behaviour.

Given a rich data set, like the one that we have created in this research, there is a wide range of microscopic studies that can be carried out for future work. For instance, many of the market dynamics can be explained by tracking traders' positions and their profitability. Defining the profit/loss thresholds at which traders choose to close their positions, and their trading reaction towards losing positions, can infer much information about the general market movement. Moreover, the positions' information, together with the price level at which the positions were opened, can indicate which groups of traders are most likely to run into losses and may be forced to close out their positions through margin calls. The danger of cascading margin calls in accelerating price moves have been already demonstrated in chapter 6. However, a further step would be creating a dynamic positions map, which would function like "weather maps" for the FX market. The map will indicate the market direction, given information about market price level, traders' positions and probability of cascading margin calls. The ultimate goal would be building a financial forecasting and market crisis warning system.

References

[1] Acerbi, C., Scandolo, G., 2008. Liquidity risk theory and coherent measures of risk. Quantitative Finance 8 (7), 681–692.

[2] Admati, A. R., Peiderer, P. A., 1988. Theory of intraday patterns: volume and price variability. Review of Financial Studies 1, 13–40.

[3] Aguirre, M. S., Saidi, R., March 1999. Feedback trading in exchange-rate markets, evidence from within and across economic blocks. Journal of Economics and Finance 23 (1), 1–14.

[4] Ahrens, R., Reitz, S., 2005. Heterogeneous expectations in the foreign exchange market : Evidence from daily dm/us dollar exchange rates. Journal of Evolutionary Economics 15, 65 – 82.

[5] Aityan, S. K., 2011. The concept of event-driven time. Physics International 2 (2), 40–49.

[6] Akram, F., Rime, D., Sarno, L., 2008. Arbitrage in the foreign exchange market: Turning on the microscope. Journal of International Economics 76, 237–253.

[7] Almeida, A., Goodhart, C. A. E., Payne, R. G., 1998. The effects of macroeconomic news on high frequency exchange rate behavior. Journal of Financial and Quantitative Analysis 33 (3), 383–408.

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

[9] Andersen, T. G., Bollerslev, T., 1997. Intraday periodicity and volatility persistence in financial markets. Journal of Empirical Finance 4, 115–158.

[10] Andersen, T. G., Bollerslev, T., Diebold, F. X., Vega, C., 2003. American economic review. Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange 93, 38–62.

[11] Arifovic, J., March 2001. Evolutionary dynamics of currency substitution. Journal of Economic Dynamics and Control 25 (3-4), 395–417.
URL http://ideas.repec.org/a/eee/dyncon/v25y2001i3-4p395-417.html

[12] Arthur, W., Holland, J., LeBaron, B., Palmer, R., Taylor, P., 1997. Asset pricing under endogenous expectations in an artificial stock market. In: Arthur, W., Lane, D., Durlauf, S. (Eds.), The Economy as an evolving complex system II. Addison-Wesley.

[13] Artzner, P., Delbaen, F., Eber, J.-M., Heath, D., 1999. Coherent measures of risk. Mathematical Finance 9 (3), 203–228.

[14] Bacchetta, P., Wincoop, E., 2003. Can information dispersion explain the exchange rate disconnect puzzle? NBER Working Papers 9498, National Bureau of Economic Research, Inc.
URL http://www.nber.org/papers/w9498

[15] Bacchetta, P., Wincoop, E., 2006. Can information heterogeneity explain the exchange rate determination puzzle? American Economic Review 96, 552–576.

[16] Baillie, R. T., Bollerslev, T., 1989. The message in daily exchange rates: a conditional-variance tale. Journal of Business and Economic Statistics 7 (3), 297–305.

[17] Baillie, R. T., Bollerslev, T., 1990. Intra day and inter market volatility in foreign exchange rates. Review of Economic Studies 58, 565–585.

[18] Balduzzi, P., G., B., Foresi, S., 1995. Asset price dynamics and infrequent feedback trades. The Journal Of Finance 50 (5).

[19] Bank for International Settlements, B., June 2008. 78th annual report.
URL http://www.bis.org/publ/arpdf/ar2008e.htm

[20] Bank for International Settlements, B., December 2010. Triennial central bank survey of foreign exchange and derivatives market activity in 2010.
URL http://www.bis.org/publ/rpfxf10t.htm

[21] Barber, B., Odean, T., 2001. Boys will be boys: gender, overconfidence, and common stock investment. Quarterly Journal of Economics 141, 261–292.

[22] Barber, B. M., Odean, T., 2000. Trading is hazardous to your wealth: The common tock investment performance of individual investors. Journal of Finance 55 (2), 773–806.

[23] Barberis, N., Huang, M., 2001. Mental accounting, loss aversion and individual stock returns. Journal of Finance 56, 1247–1292.

[24] Barberis, N., Huang, M., Santos, T., 2001. Prospect theory and asset prices. Quarterly Journal of Economics 116, 1–53.

[25] Barberis, N., Shleifer, A., 2003. Style investing. Journal of Financial Economics 68, 161–199.

[26] Barberis, N., Thaler, R., 2003. A survey of behavioural finance. NBER Working Papers 9222, National Bureau of Economic Research, Inc.

[27] Bauer, C., De Grauwe, P., Reitz, S., 2008. Exchange rate dynamics in a target zone: A heterogeneous expectations approach. Journal of Economic Dynamics and Control 33 (2), 329–344.

[28] Benartzi, S., Thaler, R. H., February 1995. Myopic loss aversion and the equity premium puzzle. The Quarterly Journal of Economics 110 (1), 73–92.
URL http://ideas.repec.org/a/tpr/qjecon/v110y1995i1p73-92.html

[29] Berger, D., A., C., S., C., Howorka, E., Wright, J., 2006. Order flow and exchange rate dynamics in electronic brokerage system data. International Finance Discussion Papers 830, Board of Governors of the Federal Reserve System.

[30] Biais, B., Hilton, D., Mazurier, K., Pouget, S., 2005. Judgmental overconfidence, self-monitoring and trading performance in an experimental financial market. Review of Economic Studies 72, 287–312.

[31] Bisig, T., Dupuis, A., Impagliazzo, V.and Olsen, R., 2012. The scale of market quakes. Quantitative Finance 12 (4), 501–508.
URL URLhttp://www.tandfonline.com/doi/abs/10.1080/14697688.2011.609180

[32] Bjønnes, G. H., Rime, D., 2005. Dealer behaviour and trading systems in foreign exchange markets. Journal of Financial Economics 75, 571–605.

[33] Bjønnes, G. H., Rime, D., Haakon, O., 2005. Liquidity provision in the overnight foreign exchange market. Journal of International Money and Finance 24, 175–196.

[34] Bohl, M. T., Reitz, S., Winter-Spring 2006. Do positive feedback traders act in Germany's neuer markt? Quarterly Journal of Business and Economics 45 (1/2), 3–14.

[35] Bollerslev, T., Domowitz, I., 1993. Trading patterns and prices in the interbank foreign exchange market. The Journal of Finance 48, 1421–1443.

[36] Bouchaud, J.-P., Farmer, J. D., Lillo, F., Sep 2008. How markets slowly digest changes in supply and demand. Quantitative Finance Papers 0809.0822, arXiv.org.
URL http://arxiv.org/abs/0809.0822

[37] Breymann, W., Dias, A., Embrechts, P., 2003. Dependence structures for multivariate high-frequency data in finance. Quantitative Finance 3 (1), 1–14.

[38] Brown, P., Chappel, N., da Silva Rosa, R., Walter, T., 2002. The reach of the disposition effect: Large sample evidence across investor classes. Working paper, University of Western Australia.

[39] Brownlees, C. T., Gallo, G., December 2006. Financial econometric analysis at ultrahigh frequency: Data handling concerns. Computational Statistics Data & Analysis 51 (4), 2232–2245.

[40] Brueckner, U., Nolte, I., 2002. Intrinsic time. Tech. rep., University of Konstanz. URL http://www.olsen.ch/fileadmin/Publications/Archive//2002_konstanzIntrinTime

[41] Cajueiro, D. O., Tabak, B. M., Werneck, F. K., 2009. Can we predict crashes? the case of the brazilian stock market. Physica A: Statistical Mechanics and its Applications 388 (8), 1603–1609.
URL http://EconPapers.repec.org/RePEc:eee:phsmap:v:388:y:2009:i:8:p:1603-1609

[42] Calamia, A., 1999. Market microstructure: Theory and empirics. LEM Papers Series 1999/19, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa, Italy.
URL http://ideas.repec.org/p/ssa/lemwps/1999-19.html

[43] Carlson, J. A., Lo, M., November 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.
URL http://ideas.repec.org/a/eee/jimfin/v25y2006i7p1090-1102.html

[44] Chang, Y., Taylor, S. J., 2003. Information arrivals and intraday exchange rate volatility. International Financial Markets, Institutions, and Money 13, 85–112.

[45] Chen, C.-C., 2009. Complex event types for agent-based simulation. Ph.D. thesis, College of London University.

[46] Cheung, Y.-W., Chinn, M. D., 2001. Currency traders and exchange rate dynamics: A survey of the us market. Journal of International Money and Finance 20 (4), 439–471.

[47] Clark, P. K., 1973. Subordinated stochastic process model with finite variance for speculative prices. Econometrica 41 (1), 135–155.

[48] Coval, J. D., Shumway, T., 2005. Do behavioral biases affect prices? Journal of Finance 60 (1), 1–34.

[49] Dacorogna, M., Gencay, R., Muller, U., Olsen, R., Pictet, O., 2001. An Introduction to High Frequency Finance. Academic Press, London.

[50] Daniélsson, J., Love, R., 2006. Feedback trading. International Journal of Finance and Economics 11 (1), 35–53.

[51] Daniélsson, J., Payne, R., 2002. Real trading patterns and prices in spot foreign exchange markets. Journal of International Money and Finance 21, 203–222.

[52] De Grauwe, P., Dewachter, H., 1993. A chaotic model of the exchange rate: the role of fundamentalists and chartists. Open Economies Review 4, 351–79.

[53] De Grauwe, P., Grimaldi, M., 2006. Exchange rate puzzles: A tale of switching attractors. European Economic Review 50 (1), 1–33.

[54] De Long, J., Shleifer, A., Summers, L., Waldmann, R., 1990. Positive feedback investment strategies and destabilizing rational speculation. Journal of Finance 45 (2), 375–395.

[55] DeBondt, W. F., Thaler, R. H., 1995. Financial decision making in markets and firms: A behavioral perspective. NBER Working Papers 4777, National Bureau of Economic Research, Inc.
URL http://www.nber.org/papers/w4777

[56] Dhar, R., Zhu, N., 2002. Up close and personal: An individual level analysis of the disposition effect. Working Paper 02-20, Yale International Center for Finance.

[57] Dignum, S., Poli, R., 2004. Multi-agent foreign exchange market modelling via gp. In: Deb, K. (Ed.), Genetic and Evolutionary Computation– GECCO-2004. Vol. 3102 of Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 255–256.
URL http://dx.doi.org/10.1007/978-3-540-24854-5_24

[58] Dominguez, K., Panthaki, F., 2005. What defines news in foreign exchange markets? NBER Working Papers 11769, National Bureau of Economic Research, Inc.
URL http://www.nber.org/papers/w11769

[59] Engel, C., Kenneth, D., 2004. West exchange rates and fundamentals. NBER Working Papers 10723, National Bureau of Economic Research, Inc.
URL http:.//www.nber.org/papers/w10723

[60] Engle, R. F., January 2000. The econometrics of ultrahigh frequency data. Econometrica, Econometric Society 68 (1), 1–22.

[61] Evans, J. S., 1987. Beliefs and expectations as causes of judgmental bias. In: Wright, G., Ayton, P. (Eds.), Judgmental Forecasting. John Wiley & Sons, Chichester, pp. 31–47.

[62] Evans, M., Lyons, R., 2002. Information integration and fx trading. The Journal of International Money and Finance 21, 807–831.

[63] Evans, M., Lyons, R., 2002. Order flow and exchange rate dynamics. Journal of Political Economy 110, 170–180.

[64] Evans, M., Lyons, R., 2003. Inventory information. NBER Working Papers 9893, National Bureau of Economic Research, Inc.
URL http:.//www.nber.org/papers/w9893

[65] Evans, M., Lyons, R., 2006. Understanding order flow. International Journal of Finance and Economics 11 (1), 3–23.

[66] Evans, M., Lyons, R., 2007. Exchange-rate fundamentals and order flow. NBER Working Papers 13151, National Bureau of Economic Research, Inc.
URL http//www.nber.org/papers/w13151

[67] Evans, M., Lyons, R. K., 2004. Do currency markets absorb news quickly? Journal of International Money and Finance 24 (5), 197–217.

[68] Evans, M., Lyons, R. K., 2005. A new micro model of exchange rate dynamics. NBER Working Papers 10379, National Bureau of Economic Research, Inc.
URL http://www.nber.org/papers/w10379

[69] Evans, M., Lyons, R. K., 2006. How is macro news transmitted to exchange rates? NBER Working Papers 9433, National Bureau of Economic Research, Inc.
URL http//www.nber.org/papers/w9433

[70] Evans, M. D. D., December 2002. Fx trading and exchange rate dynamics. Journal of Finance 57 (6), 2405–2447.
URL http://ideas.repec.org/a/bla/jfinan/v57y2002i6p2405-2447.html

[71] Evans, M. D. D., 2008. Foreign exchange market microstructure. In: Durlauf, S. N., Blume, L. E. (Eds.), The New Palgrave Dictionary of Economics. Palgrave Macmillan, Basingstoke.

[72] Fan, M., Lyons, R., 2003. Customer trades and extreme events in foreign exchange. In: Mizen, P. (Ed.), Monetary History, Exchange Rates and Financial Markets: Essays in Honour of Charles Goodhart. Edward Elgar, Northampton, pp. 160–179.

[73] Farmer, J. D., Joshi, S., 2002. The price dynamics of common trading strategies. Journal of Economic Behavior & Organization 49, 149–171.

[74] Farmer, J. D., Lillo., F., 2004. On the origin of power laws in financial markets. Quantitative Finance 314, 7–10.

[75] Fenn, D. J., Howison, S. D., McDonald, M., Williams, S., Johnson, N. F., 2009. The mirage of triangular arbitrage in the spot foreign exchange market. International Journal of Theoretical and Applied Finance 12 (8), 1105–1123.

[76] Flood, R. P., Rose, A. K., 1999. Understanding exchange rate volatility without the contrivance of macroeconomics. Economic Journal 109 (459), 660–72.

[77] Flood, R. P., Taylor, M. P., 1996. Exchange rate economics: What's wrong with the conventional macro approach ? In: The Microstructure of Foreign Exchange Markets. NBER Chapters. National Bureau of Economic Research, Inc, pp. 261–302.

[78] Foster, F. D., Viswanathan, S., 1990. A theory of the interday variations in volume, variance, and trading costs in securities markets. Review of Financial Studies 3, 3–40.

[79] Frankel, J. A., Froot, K. A., 1986. Understanding the u.s. dollar in the eighties: The expectation of chartists and fundamentalists. The Economic Record 62, 24–38.

[80] Frankel, J. A., Froot, K. A., 1987. Using survey data to test standard propositions regarding exchange rate expectations. American Economic Review 77 (1), 133–153.

[81] Frankel, J. A., Froot, K. A., 1990. Chartists, fundamentalists, and trading in the foreign exchange market. American Economic Review 802, 181–185.

[82] Frankel, J. A., Froot, K. A., 1990. Exchange rate forecasting techniques survey: Data, and implications for the foreign exchange market. Working Paper WP/90, International Monetary Fund, Washington, DC.

[83] Froot, K. A., Ramadorai, T., 2005. Currency returns, intrinsic value and institutional-investor flow. Journal of Finance 60 (3), 1535–1566.

[84] Gabaix, X., Gopikrishnan, P., Plerou, V., Stanley., H. E., 2003. A theory of power-law distributions in financial market fluctuations. Nature 270, 423–267.

[85] Gerety, M., Mulherin, J., 1992. Trading halts and market activity: an analysis of volume at the open and the close. Journal of Finance 5, 1765–811.

[86] Glaser, M., Weber, M., 2007. Overconfidence and trading volume. The Geneva Risk and Insurance Review 32, 1–36.

[87] Glattfelder, J., Dupuis, A., Olsen, R., 2011. Patterns in high-frequency fx data: Discovery of 12 empirical scaling laws. Quantitative Finance 11 (4), 599 – 614.

[88] Goodhart, C., Hall, S., Henry, S., Pesaran, B., 1993. News effects in high-frequency model of the sterling-dollar exchange rate. Journal of Applied Econometrics 8, 1–13.

[89] Goodhart, C. A. E., 1989. 'news' and the foreign exchange market. In: Proceedings of the Manchester Statistical Society. pp. 1–79.

[90] Goodhart, C. A. E., Figliuoli, L., 1991. Every minute counts in financial markets. Journal of International Money and Finance 10, 23–52.

[91] Gopikrishnan, P., Meyer, M., Amaral, L., Stanley, H., 1998. Inverse cubic law for the probability distribution of stock price variations. European Physical Journal B. 3, 139–140.

[92] Grinblatt, M., Keloharju, M., 2001. What makes investors trade? Journal of Finance 56 (2), 589–616.

[93] Guillaume, D. M., Dacorogna, M. M., Davé, R. R., Müller, U. A., Olsen, R. B., Pictet, O. V., 1997. From the bird's eye to the microscope: a survey of new stylized facts of the intra-daily foreign exchange markets. Finance and Stochastics 1 (2), 95 –129.

[94] Guillaume, D. M., Pictet, O. V., Muller, U. A., Dacorogna, M. M., 1995. Unveiling non linearities through time scale transformations. Working Papers 1994-06-26, Olsen and Associates.
URL http://EconPapers.repec.org/RePEc:wop:olaswp:_002

[95] Hasbrouck, J., 2007. Empirical Market Microstructure. Oxford University Press.

[96] Hasbrouk, J., Hamao, Y., 1995. Securities trading in the absence of dealers: trades and quotes on the tokyo stock exchange. Review of Financial Studies 8, 849–878.

[97] Hau, H., Rey, H., 2002. Exchange rates, equity prices, and capital flows. NBER Working Papers 9398, National Bureau of Economic Research, Inc.
URL http://www.nber.org/papers/w9398

[98] Helen, Taylor, M. P., 1990. Charts, noise and fundamentals in the london foreign exchange market. Economic Journal 100 (400), 49–59.

[99] Hilton, D. J., 2001. The psychology of financial decision-making: Applications to trading, dealing, and investment analysis. Journal of Psychology and Financial Markets 2 (1), 37–53.

[100] Hols, M. C., De Vries, C. G., 1991. The limiting distribution of extremal exchange rate returns. Journal of Applied Econometrics 6, 287–302.

[101] Hommes, C., 2002. Modeling the stylized facts in finance through simple nonlinear adaptive systems. In: Proceedings of the National Academy of Sciences. Vol. 99. pp. 7221–7228.

[102] Hommes, C. H., 2006. Heterogeneous agent models in economics and finance. In: Tesfatsion, L., Judd, K. L. (Eds.), Handbook of Computational Economics. Vol. 2 of Handbook of Computational Economics. Elsevier, Ch. 23, pp. 1109–1186.
URL http://ideas.repec.org/h/eee/hecchp/2-23.html

[103] Hong, H., Stein, J., 1999. A unified theory of under-reaction, momentum trading, and overreaction in asset markets. Journal of Finance 54, 2143–2184.

[104] Hong, H., Stein, J., 2003. Differences of opinion, short-sale constraints and market crashes. Review of Financial Studies 16 (2), 487–525.

[105] Hsieh, D. A., 1988. The statistical properties of daily foreign exchangerates: 1974-1983. Journal of International Economics 24, 129–145.

[106] Ito, T., 1990. Foreign exchange rate expectations: Micro survey data. American Economic Review 80 (3), 434–449.

[107] Ito, T., Hashimoto, Y., 2004. Microstructure of the yen/dollar foreign exchange market: patterns of intraday activity revealed in the electronic broking system. NBER Working Paper 10856.

[108] Ito, T., Hashimoto, Y., December 2006. Intraday seasonality in activities of the foreign exchange markets: Evidence from the electronic broking system. Journal of the Japanese and International Economies 20 (4), 637–664.

[109] Jain, P., Joh, G., 1988. The dependence between hourly prices and trading volume. Journal of Financial and Quatitative Analysis 23, 269–83.

[110] Kim, K. A., Nofsinger, J. R., 2003. The behavior and performance of individual investors in japan. Working paper.

[111] King, M., Osler, C., Rime, D., 2012. The market microstructure approach to foreign exchange: Looking back and looking forward. Working Papers 54, Brandeis University, Department of Economics and International Businesss School.
URL http://EconPapers.repec.org/RePEc:brd:wpaper:54

[112] Kokar, M. M., Baclawski, K., 2000. Modeling combined time- and event-driven dynamic systems. Working paper, Northeastern University, Boston, MA 02115.

[113] Koutmos, G., Said, R., 2001. Positive feedback trading in emerging capital markets. Applied Financial Economics 11, 291–97.

[114] Kowalski, R., Sergot, M., 1986. A logic-based calculus of events. New Generation Computing 4 (1), 67–95.

[115] Kozhan, R., Salmon, M., 2009. Uncertainty aversion in a heterogeneous agent model of foreign exchange rate formation. Journal of Economic Dynamics and Control 33 (5), 1106–1122.

[116] Krugman, P. R., August 1991. Target zones and exchange rate dynamics. The Quarterly Journal of Economics 106 (3), 669–82.
URL http://ideas.repec.org/a/tpr/qjecon/v106y1991i3p669-82.html

[117] Laopodis, N. T., 2005. Feedback trading and autocorrelation interactions in the foreign exchange market: Further evidence. Journal of Economic Modelling 22 (5), 811–827.

[118] LeBaron, B., 2001. A builder's guide to agent-based financial markets. Quantitative Finance 1 (2), 254–261.
URL http://ideas.repec.org/a/taf/quantf/v1y2001i2p254-261.html

[119] LeBaron, B., 2006. Agent-based computational finance. In: Tesfatsion, L., Judd, K. L. (Eds.), Handbook of Computational Economics. Vol. 2 of Handbook of Computational Economics. Elsevier, Ch. 24, pp. 1187–1233.
URL http://ideas.repec.org/h/eee/hecchp/2-24.html

[120] Lechner, S., Nolte, I., 2009. Customer trading in the foreign exchange market: Empirical evidence from an internet trading platform. Working papers, Warwick Business School, Financial Econometrics Research Centre.
URL http://EconPapers.repec.org/RePEc:wbs:wpaper:wp09-01

[121] Levin, J. H., October 1997. Chartists, fundamentalists and exchange rate dynamics. International Journal of Finance & Economics 2 (4), 281–290.
URL http://ideas.repec.org/a/ijf/ijfiec/v2y1997i4p281-90.html

[122] Lo, A., Repin, D., 2002. The psychophysiology of real-time financial risk processing. Journal of Cognitive Neuroscience 14 (3), 323–39.

[123] Locke, P., Mann, S., 2005. Do professional traders exhibit loss realization aversion? Journal of Financial Economics 76, 401–444.

[124] Love, R., Payne, R., 2008. Macroeconomic news, order flows, and exchange rates. Journal of Financial and Quantitative Analysis 43 (2), 467–48.

[125] Lui, Y., Mole, D., 1999. The use of fundamental and technical analysis by foreign exchange dealers: Hong kong evidence. Journal of International Money and Finance 17, 535–545.

[126] Lui, Y.-H., Mole, D., June 1998. The use of fundamental and technical analyses by foreign exchange dealers: Hong kong evidence. Journal of International Money and Finance 17 (3), 535–545.
URL http://ideas.repec.org/a/eee/jimfin/v17y1998i3p535-545.html

[127] Lux, T., 1996. The stable paretian hypothesis and the frequency of large returns: an examination of major german stocks. Applied Financial Economics 6 (6), 463–475.

[128] Lux, T., Marchesi, M., 2000. Volatility clustering in financial markets: A microsimulation of interacting agents? International Journal of Theoretical and Applied Finance 4, 1–20.

[129] Lux, T., Schornstein, S., February 2005. Genetic learning as an explanation of stylized facts of foreign exchange markets. Journal of Mathematical Economics 41 (1-2), 169–196.
URL http://ideas.repec.org/a/eee/mateco/v41y2005i1-2p169-196.html

[130] Lyons, R., 1995. Tests of microstructural hypotheses in the foreign exchange market. Journal of Financial Economics 39, 321–351.

[131] Lyons, R., 2001. The Microstructure Approach to Exchange Rates. MIT Press.

[132] MacDonald, R., Marsh, I. W., 1996. Currency forecasters are heterogeneous: Confirmation and consequences. Journal of International Money and Finance 155, 665–685.

[133] Madhavan, A., 2000. Market microstructure: A survey. Journal of Financial Markets 8 (2), 217–264.

[134] Madhavan, A., Smidt, S., 1991. A bayesian model of intraday specialist pricing. Journal of Financial Economics 30, 99–134.

[135] Mandelbrot, B. B., Taylor, H. M., 1967. On distribution of stock price differences. Operations Research 15 (6), 1057–1062.

[136] Markets Committe Study Group, M., September 2011. High-frequency trading in the foreign exchange market. Bank for International Settlements.
URL http://www.bis.org/publ/mktc05.htm

[137] Marsh, I. W., O'Rourke, C., April 2005. Customer order flow and exchange rate movements: Is there really information content? Working paper, Cass Business School.
URL http://dx.doi.org/10.2139/ssrn.704944

[138] Masry, S., Dupuis, A., Olsen, R., Tsang, E. K., 2013. Time zone normalisation of fx seasonality. Quantitative Finance-in press.

[139] Meese, R., Kenneth, R., 1983. Empirical exchange rate models of the seventies. Journal of International Economics 14, 3–24.

[140] Menkhoff, L., 1997. Examining the use of technical currency analysis. International Journal of Finance and Economics 2 (4), 307–318.

[141] Menkhoff, L., 1998. The noise trading approach-questionnaire evidence from foreign exchange. Journal of International Money and Finance 17 (3), 547–564.

[142] Menkhoff, L., Rebitzky, R. R., Schröder, M., May 2009. Heterogeneity in exchange rate expectations: Evidence on the chartist-fundamentalist approach. Journal of Economic Behavior & Organization 70 (1-2), 241–252.
URL http://ideas.repec.org/a/eee/jeborg/v70y2009i1-2p241-252.html

[143] Menkhoff, L., Taylor, M. P., December 2007. The obstinate passion of foreign exchange professionals: Technical analysis. Journal of Economic Literature 45 (4), 936–972.
URL http://ideas.repec.org/a/aea/jeclit/v45y2007i4p936-972.html

[144] Miller, N. C., 2002. Balance of Payments and Exchange Rate Theory. Edward Elgar Publishing Limited.

[145] Mueller, E., 2009. Automated commonsense reasoning using the event calculus. Communications of the ACM 52 (1), 113–117.

[146] Mueller, U. A., Dacorogna, M. M., Dave, R. D., Pictet, O. V., Olsen, R. B., Ward, J., 1995. Fractals and intrinsic time - a challenge to econometricians. Working Papers 1993-08-16, Olsen and Associates.
URL http://ideas.repec.org/p/wop/olaswp/009.html

[147] Mueller, U. A., Dacorogna, M. M., Pictet, O. V., 1998. Heavy tails in high-frequency financial data. In: Adler, R. J., Feldman, R. E., Taqqu, M. S. T. S. (Eds.), A practical guide to heavy tails: Statistical Techniques for Analysing Heavy Tailed Distributions. Birkhaeuser, pp. 55–77.

[148] Neely, C. J., 1997. Technical analysis in the foreign exchange market: a layman's guide. Review (Sep), 23–38.
URL http://ideas.repec.org/a/fip/fedlrv/y1997isepp23-38.html

[149] Neely, C. J., 2005. An analysis of recent studies of the effect of foreign exchange intervention. Working Papers 2005-030, Federal Reserve Bank of St. Louis.
URL http://ideas.repec.org/p/fip/fedlwp/2005-030.html

[150] Neely, C. J., Weller, P. A., 2003. Intraday technical trading in the foreign exchange market. Journal of International Money and Finance 22 (2), 223–237.

[151] Ng, W. L., and, I. G. N. C., 2010. Periodicities of fx markets in intrinsic time. Working Paper WP2010-4, Essex Business School.
URL http://www.essex.ac.uk/ebs/research/working_papers/WP2010-4%20NConstantinous%20Periodicities.pdf

[152] Nolte, I., 2012. A detailed investigation of the disposition effect and individual trading behavior: a panel survival approach. The European Journal of Finance 18 (10), 885–919.
URL http://www.tandfonline.com/doi/abs/10.1080/1351847X.2011.601635

[153] Nolte, I., Lechner, S., Mar. 2007. Customer trading in the foreign exchange market: Empirical evidence from an internet trading platform. CoFE Discussion Paper 07-

03, Center of Finance and Econometrics, University of Konstanz.
URL http://ideas.repec.org/p/knz/cofedp/0703.html

[154] Nolte, I., Voev, V., Feb. 2007. Panel intensity models with latent factors: An application to the trading dynamics on the foreign exchange market. CoFE Discussion Paper 07-02, Center of Finance and Econometrics, University of Konstanz.
URL http://ideas.repec.org/p/knz/cofedp/0702.html

[155] O Connell, P. G., Teo, M., 2009. Institutional investors, past performance, and dynamic loss aversion. Journal of Financial and Quantitative Analysis, Cambridge University Press 44 (1), 155–188.

[156] OANDA, 1996-2012. Oanda ltd.
URL http://www.oanda.com/

[157] Oberlechner, T., 2001. Evaluation of currencies in the foreign exchange market: Attitudes and expectations of foreign exchange traders. Zeitschrift für Sozialpsychologie 32 (3), 180–188.

[158] Oberlechner, T., 2001. Importance of technical and fundamental analysis in the european foreign exchange market. International Journal of Finance and Economics 61, 81–93.

[159] Oberlechner, T., 2004. The psychology of the foreign exchange market. Wiley trading series. Chichester, Wiley.

[160] Oberlechner, T., Osler, C., 2008. Overconfidence in currency markets. Working Papers 02, Brandeis University, Department of Economics and International Businesss School.

[161] Obstfeld, M., Rogoff, K., Jul. 2000. The six major puzzles in international macroeconomics: Is there a common cause? NBER Working Papers 7777, National Bureau of Economic Research, Inc.
URL http://ideas.repec.org/p/nbr/nberwo/7777.html

[162] Odean, T., 1998. Are investors reluctant to realize their losses. Journal of Finance 53 (5), 1775–1798.

[163] O'Hara, M., 1995. Market Microstructure Theory. Blackwell, Oxford Press.

[164] Olsen, R., 2005. Classical economics: An emperor with no clothes. Wilmott Magazine 15, 84–85.

[165] Olsen, R., Cookson, C., 2009. How science can prevent the next bubble. Financial Times.
URL http://www.ft.com/intl/cms/s/0/c25e1008-f93e-11dd-90c1-000077b07658.html#axzz2D4StQLjZ

[166] Olsen, R. B., December 2008. The fallacy of the invisible hand. Olsenblog, Olsen Ltd.

[167] Olsen, R. B., January 2010. How to trade: Why butterflies cause cascading margin calls. Olsenblog, Olsen Ltd.
URL http://www.olsenblog.com/2009/10/scaling-laws-as-powerful-tools-of-econo

[168] Oomen, R. C., April 2006. Properties of realized variance under alternative sampling schemes. Journal of Business & Economic Statistics 24, 219–237.

[169] Osler, C. L., 2005. Stop-loss orders and price cascades in currency markets. Journal of International Money and Finance 24, 219–41.

[170] Osler, C. L., July 2008. Foreign exchange microstructure: A survey of the empirical literature.
URL http://homepages.rpi.edu/~tealj2/Osler2008.pdf

[171] Osler, C. L., Mende, A., Menkhoff, L., 2011. Price discovery in currency markets. Journal of International Money and Finance 30 (8), 1696–1718.
URL http://ideas.repec.org/a/eee/jimfin/v30y2011i8p1696-1718.html

[172] Payne, R., December 2003. Informed trade in spot foreign exchange markets: an empirical investigation. Journal of International Economics 61 (2), 307–329.
URL http://ideas.repec.org/a/eee/inecon/v61y2003i2p307-329.html

[173] Payne, R., Vitale, P., December 2003. A transaction level study of the effects of central bank intervention on exchange rates. Journal of International Economics

61 (2), 331–352.
URL http://ideas.repec.org/a/eee/inecon/v61y2003i2p331-352.html

[174] Ramadorai, T., 2008. What determines transaction costs in foreign exchange markets? International Journal of Finance & Economics 13 (1), 14–25.
URL http://ideas.repec.org/a/ijf/ijfiec/v13y2008i1p14-25.html

[175] Ranguelova, E., 2001. Disposition effect and firm size: New evidence on individual investor trading activity. Working paper, Goldman Sachs Group, Inc. - Equity Derivatives Research.
URL http://dx.doi.org/10.2139/ssrn.293618

[176] Reitz, S., Westerhoff, F., 2003. Nonlinearities and cyclical behavior: The role of chartists and fundamentalists. Studies in Nonlinear Dynamics & Econometrics 7, 4–3.

[177] Reitz, S., Westerhoff, F., Wieland, C., 2006. Target zone interventions and coordination of expectations. Journal of Optimization Theory and Applications 128 (2), 453–467.

[178] Rime, D., Sarno, L., Sojli, E., January 2010. Exchange rate forecasting, order flow and macroeconomic information. Journal of International Economics 80 (1), 72–88.
URL http://ideas.repec.org/a/eee/inecon/v80y2010i1p72-88.html

[179] Romeu, R., Jan. 2004. A puzzle of microstructure market maker models. IMF Working Papers 04/6, International Monetary Fund.
URL http://ideas.repec.org/p/imf/imfwpa/04-6.html

[180] Sager, M., Taylor, M. P., 06 2008. Commercially available order flow data and exchange rate movements: "caveat emptor". Journal of Money, Credit and Banking 40 (4), 583–625.
URL http://ideas.repec.org/a/mcb/jmoncb/v40y2008i4p583-625.html

[181] Sager, M. J., Taylor, M. P., 2006. Under the microscope: the structure of the foreign exchange market. International Journal of Finance & Economics 11 (1),

81–95.
URL http://ideas.repec.org/a/ijf/ijfiec/v11y2006i1p81-95.html

[182] Samanidou, E., Zschischang, E., Stauffer, D., Lux, T., 2007. Agent-based models of financial markets. Reports on Progress in Physics 70, 409–450.

[183] Schmitt, T. A., Schäfer, R., Münnix, M. C., Guhr, T., Jul 2012. Microscopic understanding of heavy-tailed return distributions in an agent-based model. Quantitative Finance 1207.2946, arXiv.org.
URL http://arxiv.org/abs/1207.2946

[184] Sentana, E., Wadhwani, S. B., March 1992. Feedback traders and stock return autocorrelations: Evidence from a century of daily data. Economic Journal 102 (411), 415–25.
URL http://ideas.repec.org/a/ecj/econjl/v102y1992i411p415-25.html

[185] Serbinenko, A., Rachev, S. T., 2009. Intraday spot foreign exchange market. analysis of efficiency,liquidity and volatility. Investment Management and Financial Innovations 6, 35–45.

[186] Sewell, M., February 2007. Behavioural finance.
URL http://www.behaviouralfinance.net/behavioural-finance.pdf

[187] Shanahan, M., 1999. The event calculus explained. In: M.J.Wooldridge, M.Veloso (Eds.), Artificial Intelligence Today. Vol. 1600. Springer Lecture Notes in Artificial Intelligence, pp. 409–430.

[188] Shefrin, H., Statman, M., July 1985. The disposition to sell winners too early and ride losers too long: Theory and evidence. Journal of Finance 40 (3), 777–90.
URL http://ideas.repec.org/a/bla/jfinan/v40y1985i3p777-90.html

[189] Shleifer, A., 2000. Inefficient Markets: An Introduction to Behavioral Finance. Oxford University Press.

[190] Solomon, S., Levy, H., Levy, M., 2000. Microscopic Simulation of Financial Markets from Investor Behaviour to Market Phenomena. Academic Press.

[191] Statman, M., Thorley, S., Vorkink, K., 2006. Investor overconfidence and trading volume. Review of Financial Studies 19 (4), 1531–1565.

[192] Steenberger, B. N., 2003. The Psychology of Trading: Tools and Techniques for Minding the Markets. John Wiley & Sons, Hoboken, New Jersey.

[193] Subrahmanyam, A., 2007. Behavioural finance: A review and synthesis. European Financial Management 14 (1), 12–29.

[194] Taylor, M., Allen, H., 1992. The use of technical analysis in the foreign exchange market. Journal of International Money and Finance 11 (3), 304–314.

[195] Tsang, E., October 2010. New ways to understand financial markets. Working Paper WP046-10, Centre for Computational Finance and Economic Agents (CCFEA), University of Essex.

[196] Tsang, E., Martinez-Jaramillo, S., 2004. Computational finance. IEEE Computational Intelligence Society Newsletter, 3–8.

[197] Vigfusson, R., 1997. Switching between chartists and fundamentalists: A marko regimeswitching approach. International Journal of Finance and Economics 24 (4), 291– 305.

[198] Vitale, P., May 2006. A market microstructure analysis of foreign exchange intervention. Working Paper Series 629, European Central Bank.
URL http://ideas.repec.org/p/ecb/ecbwps/20060629.html

[199] Vitale, P., December 2007. A guided tour of the market microstructure approach to exchange rate determination. Journal of Economic Surveys 21 (5), 903–934.
URL http://ideas.repec.org/a/bla/jecsur/v21y2007i5p903-934.html

[200] Wang, F., Y. K. H. S., Stanely, H., 2009. Statistical regularities of equity market activity. arXiv:0911.4258v1.

[201] Weber, M., Camerer, C. F., January 1998. The disposition effect in securities trading: an experimental analysis. Journal of Economic Behavior & Organization 33 (2), 167–184.
URL http://ideas.repec.org/a/eee/jeborg/v33y1998i2p167-184.html

[202] Westerhoff, F. H., July 2003. Expectations driven distortions in the foreign exchange market. Journal of Economic Behavior & Organization 51 (3), 389–412.
URL http://ideas.repec.org/a/eee/jeborg/v51y2003i3p389-412.html

[203] Westerhoff, F. H., Reitz, S., 2003. Nonlinearities and cyclical behavior: The role of chartists and fundamentalists. Studies in Nonlinear Dynamics & Econometrics 7 (4), 3.
URL http://ideas.repec.org/a/bpj/sndecm/v7y2003i4n3.html

[204] Yan, B., Zivot, E., Nov 2003. Analysis of highfrequency financial data with s-plus. Working Papers UWEC-2005-03, University of Washington, Department of Economics.
URL http://ideas.repec.org/p/udb/wpaper/uwec-2005-03.html