Electronic Trading in the Foreign Exchange Spot Market

Martin Gould

Somerville College

University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Michaelmas Term 2013

Abstract

During the past 30 years, the proliferation of electronic trading has catalysed profound structural change in the global foreign exchange (FX) spot market. Today, more than 60% of the market's volume occurs via electronic trading platforms [13], which provide traders with round-the-clock market access from anywhere in the world. Such platforms offer several practical benefits that have encouraged market participation from a broad new class of financial institutions and have thereby spurred market growth.

The most widely used electronic trading platforms in the FX spot market incorporate several features that differentiate them from those used in other financial markets. These features raise many important questions about order flow, market state, price formation, trader behaviour, and volatility. Despite the enormous trade volumes that such platforms facilitate, these questions have received almost no attention to date.

In this thesis, we study a recent, high-quality data set from a large electronic trading platform in the FX spot market in order to investigate several aspects of trading via this mechanism. We calculate a wide range of statistics regarding order flow and market state, and we highlight how our findings contrast to those reported by empirical studies of electronic trading platforms in other markets. We study the autocorrelation properties of returns, absolute returns, and order flow, and we investigate the extent to which the market's organization impacts price formation. We also introduce a model designed to reproduce the most important properties of trading via such a platform. We derive several results regarding the model's temporal evolution, and we simulate the model to investigate how the interactions between individual traders influence volatility. We conclude that electronic trading platforms in the FX spot market retain many desirable features of centralized markets while providing traders with explicit control over their personal trading partnerships.

Acknowledgements

I would like to express my sincere thanks to several people who have made this project possible. I would like to thank my academic supervisors, Mason Porter and Sam Howison, for their guidance, wisdom, insight, and support throughout my time at Oxford. I would like to thank my industrial supervisors, Stacy Williams, Mark McDonald, and Dan Fenn, for their knowledge and practical expertise, and for enabling me to immerse myself with hands-on market experience. I would also like to thank both my academic and industrial supervisors for their detailed comments on many manuscripts during the last four years, which have improved my writing skills and helped to make my arguments more precise.

I am grateful to Jean-Philippe Bouchaud, J. Doyne Farmer, and Austin Gerig for the many discussions that have shaped and challenged my understanding of the rapidly evolving field of market microstructure. I would like to thank Nikolaus Hautsch for his insight and for hosting me as a visitor at the Humboldt Universität zu Berlin during Michaelmas Term 2012. I am thankful to Lajos Gyurko, Jonas Haase, Terry Lyons, Justin Sharp, and Rich Plummer-Powell for their technical support regarding the computationally intensive elements of this project. I gratefully acknowledge academic support and feedback from Stephen Roberts, Thaleia Zariphopoulou, and Ben Hambly. I would also like to thank Cosma Shalizi, Gabriele La Spada, D. Eric Smith, Samuel Cohen, and Neil Shephard for many useful discussions.

I would particularly like to thank my family for their continued support and encouragement throughout the many years of my education and for equipping me with the skills and determination required for an undertaking such as this.

Finally, I would like to thank HSBC Bank, EPSRC, and the Oxford-Man Institute for Quantitative Finance for funding this research.

Papers

This thesis draws on the work described in the following papers and working papers.

- M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, S. Howison, *"Limit order books"*, Quantitative Finance, 13 (2013), pp 1709–1742.
- M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, S. Howison, *"Statistical properties of limit order books in the foreign exchange spot market"*, in preparation, 2014.
- M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, S. Howison, *"Long memory in the foreign exchange spot market"*, in preparation, 2014.
- M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, S. Howison, *"Price formation on a multi-institution trading platform in the foreign exchange spot market"*, in preparation, 2014.
- M. D. Gould, N. Hautsch, S. Williams, M. McDonald, D. J. Fenn, S. Howison, M. A. Porter, *"An agent-based model of trade on a quasi-centralized trading platform"*, in preparation, 2014.

Statement of Originality

The research in this thesis is a result of collaboration between my coauthors and me. My collaborators have helped develop the ideas described in this thesis, but I have performed all of the analysis and computations leading to the results that I present.

Contents

Chapter 1

Introduction

According to the 2013 Bank for International Settlements (BIS) Triennial Central Bank Survey [13], the average daily turnover of the global foreign exchange (FX) spot market[1] exceeds US $2 trillion. This enormous volume of trade surpasses the average daily turnover of the New York Stock Exchange (NYSE) by a factor of more than 50 [172], outstrips the daily global gross domestic product by a factor of more than 5 [83], and constitutes a more than five-fold increase in market activity since 2001 [11]. The relatively poor growth of the global economy during the same period [83] suggests that this sustained expansion in FX volumes is not a consequence of increased global trade [125]. To the contrary, several recent surveys [13, 14, 50] have concluded that the key factor driving FX growth in recent years is the proliferation of automated systems called *electronic trading platforms,* which enable traders to interact and conduct trades electronically via global telecommunications networks.

Electronic trading platforms provide traders with real-time, round-the-clock access to the FX market from anywhere in the world. They also offer traders a wide range of practical benefits, such as lower execution costs [14], higher execution speeds [70], faster transaction processing [50], and greater market transparency [185]. Electronic trading has encouraged market participation from a broad new class of financial institutions, which, in turn, has spurred market growth [125].

As we discuss throughout this chapter, there exist several different types of electronic trading platforms in the FX spot market. In this thesis, we seek to address several questions regarding the market's most widely used [13] type of electronic trading platform, which we call "multi-institution trading platforms" (see Section

[1]The FX market consists of several constituent parts, including the spot, forwards, options, and swaps markets. *Spot* transactions are contracts for the exchange of two currencies within two business days. The 2013 BIS Triennial Central Bank Survey [13] estimates that the spot market accounts for approximately 38% of the FX market's total volume.

1.3). Multi-institution trading platforms incorporate several features that differentiate them from the trading platforms used in other financial markets. Given the enormous trade volumes that these platforms facilitate [13], the use of these features raises many important questions for research. Although several authors have discussed these features in detail [14, 43, 75, 85, 185, 189], both empirical and theoretical study of their effects has remained extremely limited due to the relatively recent uptake of electronic trading [185] and the difficulty of obtaining adequate data [70, 179].

Thanks to our collaboration with HSBC bank, we have been granted access to a recent, high-quality data set from a large multi-institution trading platform called Hotspot FX [128, 129, 130]. Empirical study of this data, complemented by a model of trade via such a platform, enables us to investigate several important questions regarding order flow, market state, price formation, trader behaviour, and volatility. Therefore, we hope that our results are useful for practitioners, regulators, and scholars alike. Moreover, trading in a financial market is the outcome of a complex social process where human decision-making is recorded within a highly structured environment, and in which complicated global phenomena emerge as a result of the local interactions between many heterogeneous agents. This makes a financial market an example of a *complex system* [169]. The unusually rich, detailed, and high-quality historic data from electronic trading platforms provides an excellent testing ground for ideas regarding universality, scaling, and emergence, which are popular topics in the field of complex systems.

In the remainder of this chapter, we briefly introduce some of the common themes that run throughout this thesis and provide a brief outline of the work that we present in the the subsequent chapters.

1.1 A Brief History of the FX Market

As World War II drew to a close, delegates from 44 allied nations attended the United Nations Monetary and Financial Conference in Bretton Woods, New Hampshire. Following what was perceived as a disastrous period of floating FX rates during the 1930s [175], the primary aim of the conference was to establish an international monetary constitution based on stable exchange rates that promoted post-war prosperity via the growth of international trade [9, 31]. After weeks of discussions, all participating nations signed the so-called Bretton Woods Agreement, which stated that their central banks would implement monetary policy to maintain a "par-value" system of

fixed exchange rates between their domestic currencies and the US dollar. Similarly, the US government agreed to fix the price of gold at $35 per ounce, in order to assure the participating nations that the value of the US dollar was dependable.

Although it took more than a decade for all 44 nations to perform the necessary economic adjustments to join the par-value system, by the late 1950s the fixed-rate convertibility set out in the Bretton Woods Agreement became a practical reality. Consequently, the US dollar quickly became the de facto currency of international trade [86]. According to several economic indicators, this era of full convertibility between the member currencies was the most stable economic period in the last century [31, 73]. However, weaknesses of the system gradually appeared, and within less than a decade it became apparent that the Bretton Woods Agreement was unsustainable [9]. Amid mounting political and financial pressure, President Nixon announced on 15 August 1971 that the US government would no longer honour its agreed fixed price on gold. Without this crucial link, the par-value system quickly fell into turmoil.

On 1 March 1973, the collapsing Bretton Woods Agreement was replaced by a regime of free-floating FX rates [31]. Throughout the remainder of the 1970s, trading volumes in the global FX market grew rapidly [185]. At this time, market participants fell into one of two categories: *customers,* who sought to convert one type of currency into another (often for the purpose of conducting retail transactions overseas), and *dealers,* who acted as market makers for customers. Trades in which one counterparty is a customer and the other is a dealer are often deemed to belong to the *customer market.* Dealers also traded among themselves in a closed *inter-dealer market,* in which spreads were typically much tighter than those accessible to customers [125, 185].

Throughout the 1970s, both the customer and dealer markets relied almost exclusively on phone-based technologies [70, 185]. When customers wished to obtain a quote for an FX transaction, they did so by telephoning one of more dealers. Similarly, dealers who sought to arrange inter-dealer trades telephoned either another dealer or a so-called *voice broker*, who (in exchange for a fee) searched for trading interest among other dealers and matched suitable counterparties to transact [85]. Because market conditions changed rapidly, it was necessary for dealers to make very frequent calls throughout each trading day to stay up-to-date with current market prices [14].

1.2 The Dawn of the Electronic Era

The history of electronic trading in the FX spot market dates back to the launch of an electronic bulletin board called the Reuters Market Data Service (RMDS) in 1981 [185]. RMDS enabled dealers to post messages to convey their trading interest in the inter-dealer market. If a dealer browsing RMDS identified a potential trading partner, he/she contacted the relevant counterparty by telephone to discuss and arrange the trade.

In 1987, Reuters replaced RMDS with a new electronic platform called D2000-1 [85]. D2000-1 was a closed network for bilateral communication that enabled dealers to engage in instant-message conversations to arrange trades in the inter-dealer market.[2] In contrast to RMDS, which ultimately required dealers to finalize trades via telephone, D2000-1 facilitated all stages of the trade-agreement process electronically.

In 1992, Reuters launched a new electronic trading platform called D2000-2 [85]. D2000-2 enabled dealers to submit orders that stated their desire to buy or sell a given quantity of a given currency pair at a given price. Whenever a dealer submitted a new order, the D2000-2 servers checked whether it was possible to perform a matching to an outstanding order in the system. If so, the owners of the respective orders performed the relevant trade at the agreed price; if not, the incoming order was added to the list of outstanding orders. Due to its automatic order-matching functionality, the D2000-2 system was an early example of an electronic limit order book (LOB) [98]. We provide a detailed discussion of LOBs in Chapter 2.

A year later, a consortium of banks launched a rival inter-dealer platform called Electronic Broking Services (EBS), which had similar functionality to D2000-2 [189]. The popularity of D2000-2 and EBS grew throughout the 1990s, and by the end of the decade these platforms dominated transaction volumes in the inter-dealer market [14]. Due to their similarities with the functions of a voice broker (i.e., matching outstanding trading interest among dealers), D2000-2 and EBS came to be known as *electronic broking systems*.

1.3 A Changing Market Landscape

Throughout the 1990s, access to the D2000-2 and EBS platforms — and, therefore, to their many practical benefits — was restricted to traders in the inter-dealer market. Towards the end of the decade, however, the increasing capabilities of the internet

[2] Evans [75] provides a detailed description of the process via which dealers agreed trades on D2000-1.

facilitated the development of two new types of trading platforms that broadened the accessibility of electronic trading to a much wider class of market participants [50].

The first new type of electronic trading platforms were proprietary systems called *single-institution trading platforms* [14]. The exact functionality of early single-institution trading platforms varied somewhat, but ultimately their purpose was to enable customers to arrange trades with dealers via a computer, rather than via telephone [50]. The second new type of electronic trading platforms were *electronic communication networks (ECNs)* [185]. ECNs were designed to mimic the functionality of electronic broking systems, but without the requirement that their users had access to closed inter-dealer networks [43, 50, 125]. Because of their similarities, we use the term *multi-institution trading platforms* to describe both electronic broking systems and ECNs.

The widespread availability of ECNs triggered an important change in the FX spot market [14]: participation from a wide new class of financial institutions that did not fit the traditional "customer/dealer" partitioning. Examples of these institutions — which the BIS calls *other financial institutions* [12] — include smaller commercial banks, investment banks, securities houses, mutual funds, pension funds, hedge funds, currency funds, money market funds, building societies, leasing companies, insurance companies, financial subsidiaries of corporate firms, and central banks. As participation on ECNs widened, liquidity on the platforms increased, and today many ECNs offer similar spreads to those in the inter-dealer market [43, 53, 85].

Another major consequence of the widespread use of ECNs was an increased automation of trading tasks. Today, many traders use computer algorithms that monitor and assess the state of the market from direct data feeds and submit or cancel orders without human intervention [44]. There are several possible uses for such algorithms — including automatic hedging of inventory risk, automatic market making, and capitalizing on statistical arbitrage opportunities [14, 125]. According to the 2010 BIS Triennial Central Bank Survey [12], the rise in popularity of algorithmic trading was a key factor in the surge of FX spot volumes towards the end of the last decade. We return to the topic of algorithmic trading at several points throughout the subsequent chapters.

1.4 Electronic Trading Today

Electronic trading now accounts for more than 60% of turnover in the FX spot market [13]. Although Reuters have replaced D2000-2 with a new platform called Dealing

3000 [43], the new platform inherits most of its functionality directly from the old D2000-2 system. The EBS platform continues to operate, and a wide selection of new ECNs and single-institution trading platforms have emerged during the last decade [176]. Most multi-institution trading platforms offer similar core functionality [43], but electronic broking systems remain primarily aimed at the inter-dealer market whereas ECNs welcome participation from a much wider class of financial institutions.

1.5 A Quasi-Centralized Marketplace

Multi-institution trading platforms use electronic LOBs to facilitate trade. In contrast to those used in many other financial markets, however, the LOBs used by multi-institution trading platforms in the FX spot market do not allow all traders to trade with all others. Although the prices of all trades on such platforms are announced publicly, each trader is only able to access the trading opportunities that are offered by his/her trading partners. We present a detailed discussion of this trading mechanism — which we call a *quasi-centralized* LOB — in Chapter 2.

An important consequence of the quasi-centralized structure of trade in the FX spot market is that different traders on the same multi-institution trading platform face different trading opportunities at the same time. This raises several questions for research. Do the statistical properties of order flow and LOB state vary from those in other markets? Do arbitrage opportunities arise for some traders? What impact do trading partnerships have on price formation? Does this market structure influence volatility? How do traders assess their local information about market state when deciding how to act? Despite the widespread use of quasi-centralized LOBs in the FX spot market [43, 189], these questions have received almost no attention to date. Therefore, detailed study of this trading mechanism is a very important task.

1.6 Outline

This thesis is organized into seven additional chapters. In each chapter that contains new research, we discuss our motivations and provide an overview of the literature most relevant to our work. In Chapter 2, we discuss the microstructure of multi-institution trading platforms in the FX spot market and highlight several important features that distinguish these platforms from those used in other markets. We also list many important definitions that we use throughout the thesis. In Chapter 3, we introduce the Hotspot FX trading platform, provide a detailed description of the data

that we study in the subsequent chapters, and discuss our data pre-processing and cleaning methodology.

In Chapter 4, we perform an empirical investigation of order flow and LOB state on Hotspot FX. We calculate several statistics regarding order arrival and cancellation rates, the mean LOB depth at given prices, and the distributions of limit order arrivals and cancellations. We identify several properties of order flow and LOB state that vary across trading days. We also uncover a simple, linear rescaling that causes each day's order-flow distributions to collapse onto a single curve.

In Chapter 5, we employ a wide range of estimators and techniques to perform a scientifically rigorous investigation of the the long-memory properties of returns, absolute returns, and order flow on Hotspot FX. Due to the extremely high levels of activity on the platform, we are able to detect long memory with high levels of statistical significance, even in cases when it is relatively weak. We find that autocorrelations in the return series disappear within less than 1 second, but we find strong, statistically significant evidence for long memory in absolute returns and order flow. We also uncover a low-order, negative autocorrelation in order flow that (to our knowledge) has not been reported by studies of other markets.

In Chapter 6, we investigate how heterogeneity in traders' access to trading opportunities impacts price formation on Hotspot FX. Specifically, we calculate the difference between the price at which each trade occurred and the price at which it would have occurred if the owner of the relevant market order had access to all trading opportunities on the platform at its time of its submission. We call this difference the *bypassing price* of a trade. Although we identify a small number of trades that have a large bypassing price, we find that the vast majority of market orders match at a price that is close to the best quotes at their time of execution, and we find that changes in bypassing price provide much weaker explanatory power for changes in trade price than do changes in the best quotes.

In Chapter 7, we introduce an agent-based model of trade on a quasi-centralized trading platform. We design our model to be sufficiently simple to enable us to derive several analytical results about its temporal evolution. We simulate our model using different trade-partnership networks and calculate the volatility of the resulting trade-price series in each case. We find that volatility is lowest when the trade-partnership network is complete, and that reducing the total number of edges in the trade-partnership network often — but not always — leads to an increase in volatility.

Finally, in Chapter 8, we offer some conclusions and discuss several possible avenues for future research.

Chapter 2

The Microstructure of Electronic Trading in the FX Spot Market

In this chapter, we discuss several important aspects of multi-institution trading platforms (see Section 1.3) in the FX spot market. The work in this chapter is also described in the jointly authored papers [98, 100].

2.1 Introduction

During the past 30 years, a growing body of research has concluded that macroeconomic activity fails to explain several important features of price formation in financial markets [79, 81, 82, 121, 149, 166, 195]. The poor performance of models that portray price movements as reflections of changes in exogenous fundamentals suggests that many important factors are extraneous to the variables that they consider [81, 82]. These findings provide strong motivation to explore alternative approaches.

One such alternative is the *market-microstructure approach* [79, 82, 87, 149, 150, 195]. This approach emphasizes that price formation is a direct consequence of traders' actions, and thereby highlights the need to study in detail the processes by which traders' latent trading intentions translate into trades. In recent years, microstructure models have provided alternative explanations for many important phenomena. For example, several authors have highlighted that volatility can emerge as a natural consequence of trader heterogeneity [32, 34, 42, 58, 204], and have therefore brought into question the notion that increased volatility necessarily signifies fundamental economic revaluations.

In this chapter, we provide a detailed description of the microstructure of multi-institution trading platforms in the FX spot market. We highlight several important features that distinguish these platforms from the electronic trading platforms used

in other markets. We also list several important definitions that we use throughout the subsequent chapters.

The remainder of this chapter is organized as follows. In Section 2.2, we introduce a general framework for LOB trading. In Section 2.3, we discuss several features of the LOBs used by multi-institution trading platforms in the FX spot market that distinguish them from those used in other financial markets. In Section 2.4, we discuss a network representation for the trading partnerships between traders on a multi-institution trading platform in the FX spot market. In Section 2.5, we present a brief summary of the main points from the chapter.

2.2 Limit Order Books

More than half of the world's financial markets use an electronic limit order book (LOB) mechanism to facilitate trade [187]. In contrast to quote-driven systems, in which buy and sell prices are set by designated market makers, price formation in an LOB is a self-organized process that is driven by traders' submissions and cancellations of orders. LOBs record comprehensive digital transcriptions of this process and thereby illuminate the complex interplay between order flow, liquidity, and price formation [57]. In this section, we provide a brief introduction to the main features of LOBs. For a more comprehensive discussion, see our survey article [98].

2.2.1 The Mechanics of LOBs

In an LOB, traders interact by submitting *orders.*

Definition. *An* order $x = (p_x, \omega_x, t_x)$ *submitted at time* t_x *with price* p_x *and size* $\omega_x > 0$ *(respectively,* $\omega_x < 0$*) is a commitment to sell (respectively, buy) up to* $|\omega_x|$ *units of the traded asset at a price no less than (respectively, no greater than)* p_x*.*

For a given LOB, the units of order size and price are set as follows.

Definition. *The* lot size σ *of an LOB is the smallest amount of the asset that can be traded within it. All orders must arrive with a size* $\omega_x \in \{\pm k\sigma | k = 1, 2, \ldots\}$*.*

Definition. *The* tick size π *of an LOB is the smallest permissible price interval between different orders within it. All orders must arrive with a price that is specified to the accuracy of* π*.*

Whenever a trader submits a buy (respectively, sell) order x, an LOB's trade-matching algorithm checks whether it is possible for x to *match to* a sell (respectively, buy) order y such that $p_y \leq p_x$ (respectively, $p_y \geq p_x$). If so, the matching occurs immediately and the relevant traders perform a trade for the agreed amount at the agreed price. If $|\omega_x| > |\omega_y|$, any residue size of x is then considered for matching to other active sell (respectively, buy) orders, and so on until either x becomes fully matched or there are no further active sell (respectively, buy) orders with a price that makes them eligible for matching to x.

Any portion of x that does not match instead becomes *active* at the price p_x, and it remains active until either it matches to an incoming sell (respectively, buy) order or it is *cancelled*. Cancellation usually occurs because the owner of an order no longer wishes to offer a trade at the stated price. Many platforms allow traders to specify that any portion of an arriving order that does not match immediately is automatically cancelled upon arrival. The rules governing specific trading platforms can also lead to the cancellation of active orders. For example, Hotspot FX automatically cancels all active orders at 22:00:00 GMT each day, to prevent a large accumulation of stale orders [130].

Orders that result in an immediate matching upon arrival are called *market orders*. Orders that do not, instead becoming active orders, are called *limit orders*. Some trading platforms allow traders to submit a wide variety of other order types (such as fill-or-kill, stop-loss, or peg orders [33, 130]), but it is always possible to decompose the corresponding order flow into limit and/or market orders. Therefore, we study order flow in terms of these simple building blocks.

Definition. *On a given electronic trading platform, the* LOB $\mathcal{L}(t)$ *for a given asset is the set of all active orders for the asset at time* t*.*

2.2.2 LOB State

At any time t, the set of active orders can be partitioned into the set of active buy orders $\mathcal{B}(t)$, for which $\omega_x < 0$, and the set of active sell orders $\mathcal{A}(t)$, for which $\omega_x > 0$.

Definition. *The* bid price at time t *is the highest price among active buy orders at time* t*:*

$$b(t) = \max_{x \in \mathcal{B}(t)} p_x. \tag{2.1}$$

The ask price at time t *is the lowest price among active sell orders at time* t*:*

$$a(t) = \min_{x \in \mathcal{A}(t)} p_x. \tag{2.2}$$

Definition. *The* bid-ask spread at time t *is* $s(t) = a(t) - b(t)$.

Definition. *The* mid price at time t *is* $m(t) = [b(t) + a(t)]/2$.

The evolution of $\mathcal{L}(t)$ is a càdlàg process.[1] It is often useful to consider the state of $\mathcal{L}(t)$ immediately before a specified time, so we introduce the underline notation

$$\underline{\mathcal{L}}(t) = \lim_{t' \uparrow t} \mathcal{L}(t') \tag{2.3}$$

to denote the state of $\mathcal{L}(t)$ immediately prior to time t. For example, if a limit order x becomes active upon arrival, then $x \in \mathcal{L}(t_x)$ but $x \notin \underline{\mathcal{L}}(t_x)$.

Each of $b(t)$, $a(t)$, $s(t)$, and $m(t)$ are also càdlàg processes. We use the same underline notation as in Equation (2.3) to denote the values of these processes immediately before time t. For example,

$$\underline{b}(t) = \lim_{t' \uparrow t} b(t').$$

2.2.3 Priority

Often, LOBs contain several active orders that are eligible for matching to an incoming market order. Different LOBs employ different priority systems to rank active orders in these situations. By far, the most common priority mechanism currently used is *price-time*, for which priority is first given to the active orders with the best (i.e. highest buy or lowest sell) price, and ties are broken by selecting the active order with the earliest submission time t_x among those at the best price. An alternative priority mechanism is *price-size*, for which priority is also first given to the active orders with the best price, but ties are broken by selecting the active order with the largest size among those at the best price.

2.3 Quasi-Centralized LOBs

Let $\Theta = \{\theta_1, \theta_2, \ldots, \theta_M\}$ denote the set of M traders in a given LOB. In some LOBs, every arriving buy (respectively, sell) market order x matches to the highest-priority active sell (respectively, buy) order $y \in \underline{\mathcal{L}}(t_x)$, irrespective of which traders own x and y. We call such LOBs *centralized LOBs*. In a centralized LOB, each trader $\theta_i \in \Theta$ faces the same trading opportunities at the same time because all traders are able to trade with all others.

[1] A sample function on a well-ordered set is *càdlàg* if it is continuous from the right and has left limits everywhere. A stochastic process is càdlàg if almost all of its sample paths are càdlàg.

Multi-institution trading platforms in the FX spot market adopt an alternative LOB configuration, which we call a quasi-centralized LOB.

Definition. *A* quasi-centralized LOB *is an LOB in which each trader's orders can only match to those of a specified set of trading partners but in which the price of each trade is announced publicly.*

We choose the name "quasi-centralized LOB" because these LOBs share some, but not all, features of centralized LOBs. At some points throughout the thesis, and particularly in Chapter 7, we also use the more general term *quasi-centralized trading platform* to refer to any electronic trading platform (LOB or otherwise) on which each trader can only trade with a specified set of trading partners but in which the price of each trade is announced publicly.

2.3.1 Trade Partnerships

On a quasi-centralized trading platform, each trader $\theta_i \in \Theta$ provides the exchange with a list of the other traders with whom they are willing to trade.

Definition. *Trader* $\theta_i \in \Theta$ accepts *trader* $\theta_j \in \Theta$ *if* θ_i *declares that he/she is willing to trade with* θ_j*. Otherwise,* θ_i blocks θ_j*.*

If θ_i accepts θ_j, we write $\theta_j \to \theta_i$. If θ_i blocks θ_j, we write $\theta_j \nrightarrow \theta_i$. On a quasi-centralized trading platform, each trader can only access the trading opportunities that are offered by the subset of traders with whom they posses a mutual trading agreement.

Definition. *If* $\theta_i \to \theta_j$ *and* $\theta_j \to \theta_i$*, then* θ_i *and* θ_j *are* trading partners*.*

If θ_i and θ_j are trading partners, we write $\theta_i \leftrightarrow \theta_j$. Otherwise, we write $\theta_i \nleftrightarrow \theta_j$.

In a quasi-centralized LOB, when a trader $\theta_i \in \Theta$ submits a buy (respectively, sell) market order x, the order matches to the highest-priority active sell (respectively, buy) order $y \in \underline{\mathcal{L}}(t_x)$ that is owned by a trader $\theta_j \in \Theta$ such that $\theta_i \leftrightarrow \theta_j$. Therefore, the temporal evolution of a quasi-centralized LOB depends not only on order flow, but also on which traders own which orders. Where necessary, we write x_i to denote an order x that is owned by the trader θ_i.

2.3.2 Blocking

If a trader θ_i chooses to block another trader θ_j (or vice-versa), then θ_i and θ_j are unable to trade with each other. We suggest several possible reasons why this might occur. First, if θ_i believes θ_j to possess superior private information or faster manoeuvrability after the release of relevant news, then θ_i might fear adverse selection from θ_j. For example, a trader who typically leaves orders active in the LOB might be unwilling to trade with aggressive algorithmic traders [44, 115]. Second, θ_i might block counterparties that only submit small orders. For example, a large bank that typically performs very large trades might not wish to trade with small financial institutions. Third, θ_i might block other traders whose activity clogs the trading screen. For example, the strategies adopted by many high-frequency traders involve the submission and rapid cancellation of limit orders over very short time intervals [22, 44, 115, 126]. Other traders might find it difficult to assess the state of the market if a huge number of orders rapidly appear and disappear from their trading screens.

2.3.3 Local LOBs

In a quasi-centralized LOB, traders are not in general able to view the state of the whole LOB $\mathcal{L}(t)$. Instead, each trader $\theta_i \in \Theta$ views only the active orders that are owned by his/her trading partners.

Definition. *Trader θ_i's* local LOB *is the subset of active orders in $\mathcal{L}(t)$ that are owned by θ_i's trading partners:*

$$\mathcal{L}_i(t) = \{x_j \in \mathcal{L}(t) |\ \theta_i \leftrightarrow \theta_j\}\,.$$

Definition. *Trader θ_i's* local bid price at time t *is the highest stated price among active buy orders at time t that are owned by θ_i's trading partners:*

$$b_i(t) = \max_{\{x_j \in \mathcal{B}(t) | \theta_i \leftrightarrow \theta_j\}} p_x. \tag{2.4}$$

Trader θ_i's local ask price at time t *is the lowest stated price among active sell orders at time t that are owned by θ_i's trading partners:*

$$a_i(t) = \min_{\{x_j \in \mathcal{A}(t) | \theta_i \leftrightarrow \theta_j\}} p_x. \tag{2.5}$$

Definition. *Trader θ_i's* local bid-ask spread at time t *is $s_i(t) = a_i(t) - b_i(t)$.*

Definition. *Trader θ_i's* local mid price at time t *is $m_i(t) = [b_i(t) + a_i(t)]/2$.*

To highlight the differences between the definitions in Section 2.2 (which correspond to the set of all active orders) and those in this section, we use the term *global LOB* to describe $\mathcal{L}(t)$, *global bid price* to describe $b(t)$, *global ask price* to describe $a(t)$, and so on. Figure 2.1 shows a schematic of a global LOB and a corresponding local LOB at some instant in time.

2.3.4 The Trade-Data Stream

Because each trader's local LOB only includes the active orders that are owned by his/her trading partners, different traders in a quasi-centralized LOB do not in general face the same trading opportunities at the same time. However, quasi-centralized LOBs also publicly disseminate a trade-data stream that lists the price, time, and direction[2] of every trade that occurs. Each trader is able to observe these details in real time, irrespective of whether he/she is trading partners with the traders involved in the trade.

Definition. *The* traded bid-price at time t*, denoted* $B(t)$*, is the price at which the last seller-initiated trade occurred at or before time* t*. The* traded ask price at time t*, denoted* $A(t)$*, is the price at which the last buyer-initiated trade occurred at or before time* t*.*

At any time t, all traders in a quasi-centralized LOB are able to calculate $B(t)$ and $A(t)$ from their trade-data stream. In contrast to $b_i(t)$ and $a_i(t)$, whose values vary across traders, the values of $B(t)$ and $A(t)$ are common to all traders.

2.4 Trade-Partnership Networks

In this section, we propose a simple representation of the trading partnerships on a quasi-centralized trading platform. We call this representation a *trade-partnership network.* Consider a network with M nodes, where node i represents trader θ_i, for $i = 1, \ldots, M$. For each pair of traders $\theta_i \neq \theta_j$, connect node i to node j with an edge if and only if $\theta_i \leftrightarrow \theta_j$. Figure 2.2 shows a trade-partnership network for a quasi-centralized trading platform that is populated by the traders $\Theta = \{\theta_1, \theta_2, \theta_3\}$ and in which $\theta_1 \leftrightarrow \theta_2$ and $\theta_1 \leftrightarrow \theta_3$ but $\theta_2 \nleftrightarrow \theta_3$.

[2]The *direction* of a trade describes whether it is buyer-initiated or seller-initiated. A *buyer-initiated* (respectively, *seller-initiated*) trade occurs when a buy market order matches to an active sell order (respectively, sell market order matches to an active buy order).

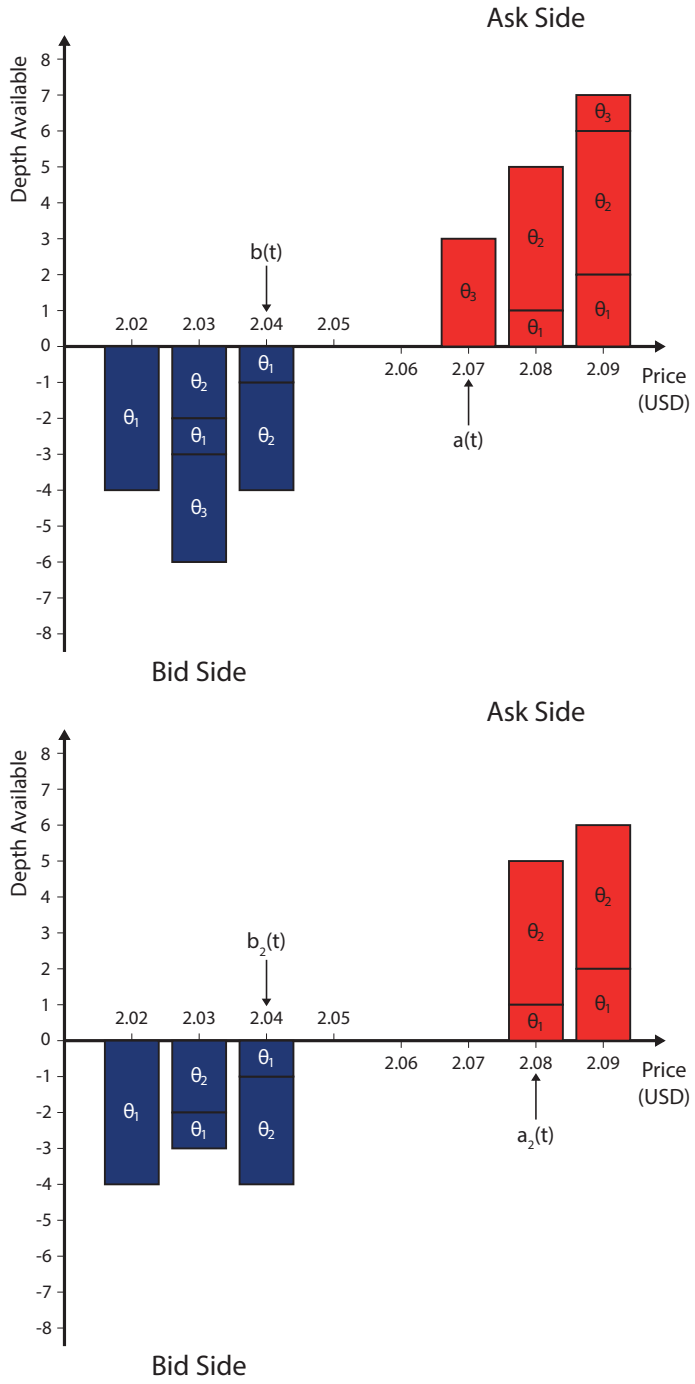

Figure 2.1: Schematic of (top) global LOB $\mathcal{L}(t)$ and (bottom) θ_2's local LOB $\mathcal{L}_2(t)$ for a quasi-centralized LOB populated by traders $\Theta = \{\theta_1, \theta_2, \theta_3\}$ with $\theta_1 \leftrightarrow \theta_2$ and $\theta_1 \leftrightarrow \theta_3$ but $\theta_2 \nleftrightarrow \theta_3$. To illustrate the role of the trading partnerships, each active order in the figure is labelled according to its owner. However, in order to protect traders' identities, electronic trading platforms do not disseminate this information to traders in real time.

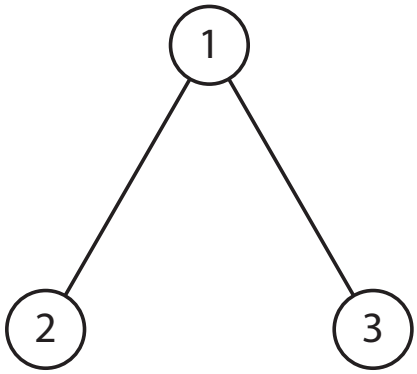

Figure 2.2: Network representation of the trading partnerships between traders $\Theta = \{\theta_1, \theta_2, \theta_3\}$ on a quasi-centralized trading platform where $\theta_1 \leftrightarrow \theta_2$ and $\theta_1 \leftrightarrow \theta_3$ but $\theta_2 \nleftrightarrow \theta_3$.

Throughout this thesis, we do not consider the directed nature of traders accepting or blocking other traders. However, we note that this network representation can be extended to incorporate directed relationships by using directed edges instead of undirected edges.

2.5 Summary

In this chapter, we have discussed several important features of trading via multi-institution trading platforms in the FX spot market. We also introduced a network representation of the trading partnerships between the traders on a quasi-centralized trading platform. Many of the topics that we have discussed in this chapter form the basis of our investigations throughout the remainder of the thesis.

Chapter 3

Hotspot FX

In this chapter, we introduce the Hotspot FX trading platform, provide a detailed description of the data that we study in the subsequent chapters, and discuss the methods that we use to pre-process and clean the data. The work in this chapter is also described in the jointly authored paper [100].

3.1 The Hotspot FX Platform

Since its launch in 2002, Hotspot FX [128, 129, 130] has grown to be one of the most widely used multi-institution trading platforms in the FX spot market [176]. Hotspot FX provides direct market access to a broad range of trading professionals, including banks, financial institutions, hedge funds, high-frequency traders, corporations, and commodity trading advisors [128]. More than 60 different currency pairs are traded simultaneously on Hotspot FX, and the platform is open for trades 24 hours a day, 5 days a week.

3.1.1 The Hotspot FX Trading Screen

Figure 3.1 shows a screenshot of the Hotspot FX trading screen. The right panel of the screen contains the trade-data stream that is common to all traders. For a given trader $\theta_i \in \Theta$, the left panel of the screen contains the local LOBs for a selection of different currency pairs. A price for the currency pair XXX/YYY denotes[1] how many units of the *counter currency* YYY are exchanged per unit of the *base currency* XXX. For example, a GBP/USD trade of size 1 at price 1.52342 corresponds to an exchange of 1.52342 US dollars for 1 pound sterling. Upon matching, the owner of a buy order

[1]To avoid potential confusion with fractions (e.g., x/y), some practitioners instead label currency pairs as XXXYYY or XXX-YYY.

$x = (p_x, \omega_x, t_x)$ gives $|\omega_x|\, p_x$ units of the counter currency and receives $|\omega_x|$ units of the base currency. Similarly, upon matching, the owner of a sell order $y = (p_y, \omega_y, t_y)$ gives ω_y units of the base currency and receives $\omega_y p_y$ units of the counter currency. Each currency pair has its own lot size σ and tick size π [130].

3.1.2 Trading Partnerships

Hotspot FX provides traders two ways to manage their trading partnerships on the platform: by specifying simple accept/block choices (see Section 2.3.1) or by specifying *credit limits*. If a trader θ_i specifies a credit limit for another trader θ_j, then $\theta_i \leftrightarrow \theta_j$ if and only if the total size of trades between θ_i and θ_j that have been agreed but not yet settled is less than a specified threshold. If the size of such trades ever reaches the threshold, θ_i temporarily blocks θ_j until either sufficiently many trades have been settled to reduce the total size of outstanding trades back below the threshold or θ_i increases the credit limit for θ_j.

3.1.3 Hidden Liquidity

Hotspot FX provides traders with an option that helps to hide the extent of their intentions to trade. An *iceberg order* is an order that specifies not only a total size and price but also a *visible size*. Other traders see only the visible size, which appears (to them) as a standard active order. The hidden component of an iceberg order has lower priority than the visible component of the same order, but it can still be matched if a suitably large market order arrives. Whenever any of the visible portion of an iceberg order is matched, the visible portion is automatically replenished by reducing the invisible portion by the appropriate amount. An iceberg order with a visible size of 0 is called a *hidden order*. For further discussion of iceberg orders, see [98].

3.2 The Hotspot FX Data

In Chapters 4–6, we use data from Hotspot FX to investigate several features of order flow, LOB state, and price formation in quasi-centralized LOBs. The data that we study describes all limit order arrivals and cancellations and all trades for the EUR/USD (Euro/U.S. dollar), GBP/USD (Pounds sterling/U.S. dollar), and EUR/GBP (Euro/Pounds sterling) currency pairs[2] on the platform during May 2010

[2]According to the 2013 BIS Triennial Central Bank Survey [13], these currency pairs constitute about 24%, 9%, and 2%, respectively, of the total turnover of the FX market.

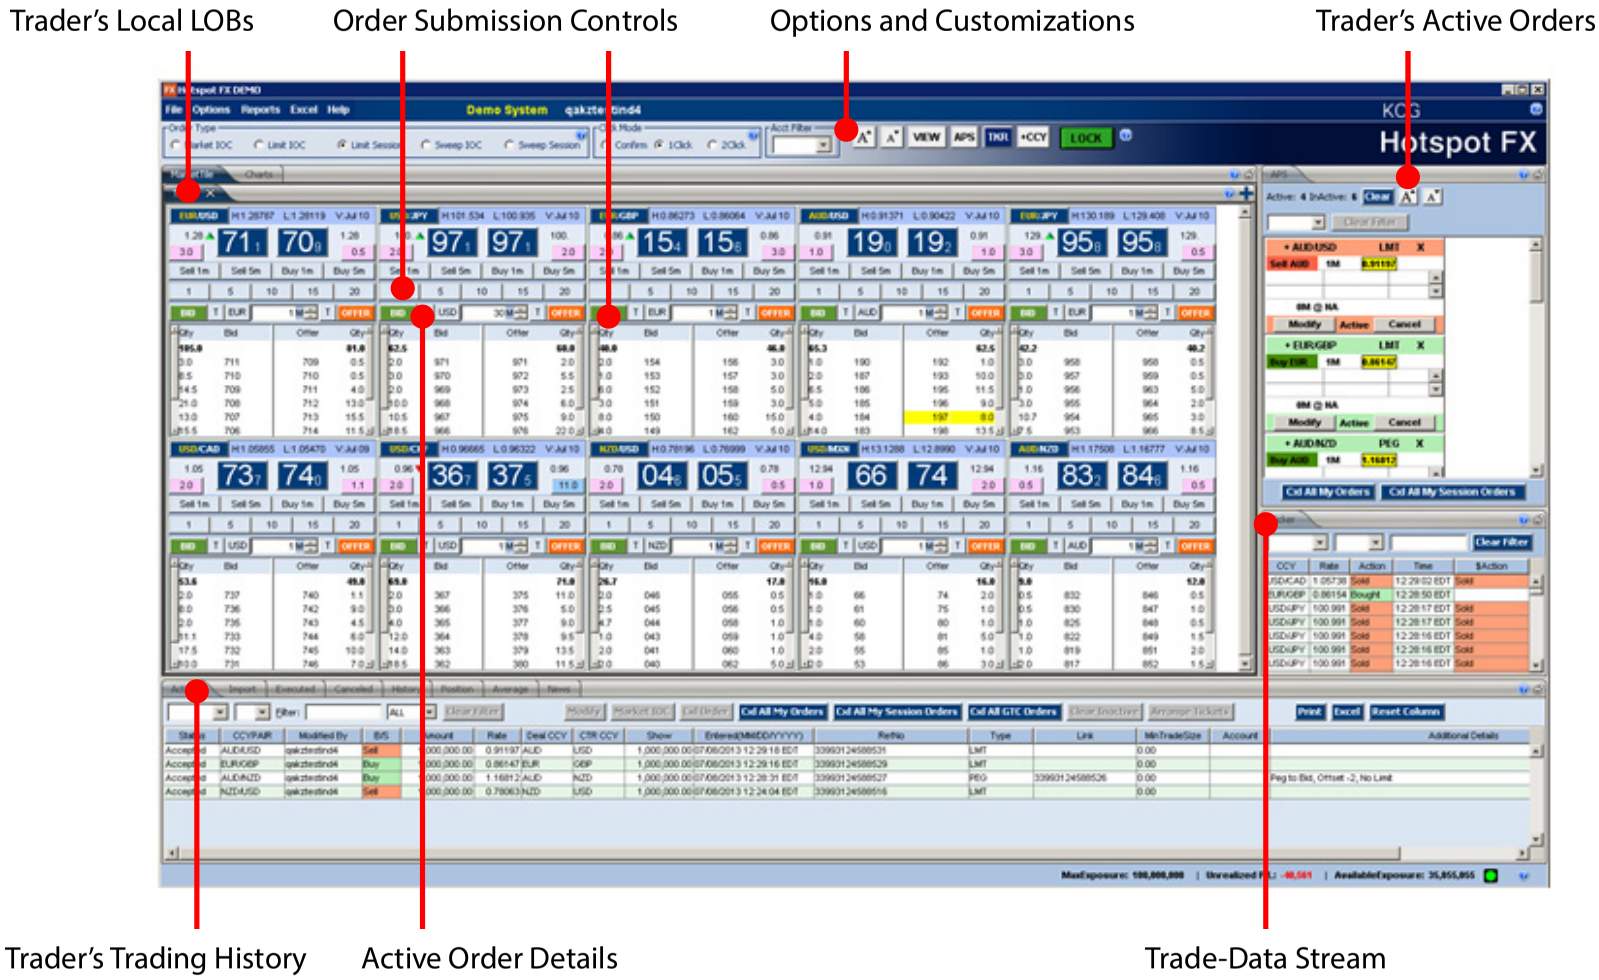

Figure 3.1: The Hotspot FX trading screen. Low-resolution screenshot reproduced with permission from Hotspot FX.

and June 2010. On Hotspot FX, each of the three currency pairs has a lot size of $\sigma = 0.01$ units of the base currency and a tick size of $\pi = 0.00001$ units of the counter currency.

Due to the extremely high levels of activity on the platform, the Hotspot FX data is extremely large. After cleaning (see Section 3.3), the sizes of the EUR/USD, GBP/USD, and EUR/GBP data files are approximately 13GB, 9GB, and 6GB, respectively.

3.2.1 File Types

For each currency pair, the Hotspot FX data consists of two types of files: the *trade-data file*, which describes all trades that occur, and the *tick-data file*, which lists all updates to $\mathcal{L}(t)$.

3.2.1.1 The Trade-Data File

The trade-data file contains a transcription of the trade-data stream that traders on Hotspot FX access in real time (see Section 2.3.4). Additionally, the trade-data file lists the size of each trade. A short extract from a trade-data file is:

```
2010/05/03,00:11:28.547,S,1.32141,100000.00
2010/05/03,00:14:14.449,S,1.32148,100000.00
2010/05/03,00:14:19.608,B,1.32157,10000.00
2010/05/03,00:14:30.543,S,1.32141,700000.00
2010/05/03,00:14:30.741,S,1.32141,2300000.00
2010/05/03,00:15:06.838,B,1.32140,100000.00
```

From left to right, the fields of the trade-data file are date, time, direction [buyer initiated (`B`) or seller initiated (`S`)], price, and size.

3.2.1.2 The Tick-Data file

The tick-data file consists of a separate data block for each minute during the trading day. Each block begins with an up-to-date snapshot of the global LOB $\mathcal{L}(t)$. After the snapshot, each block contains a time series of all changes that occurred to $\mathcal{L}(t)$ during the given minute. A short extract from a tick-data file is:

```
00:11:28.518,657730,N,2,1.32162,8000000
00:11:28.519,657694,C,,,
```

```
00:11:28.519,657731,N,1,1.30819,200000
00:11:28.519,657724,M,,,900000
00:11:28.520,657732,N,2,1.32185,1000000
00:11:28.520,657617,C,,,
```

From left to right, the fields are time, order ID, event type [new-order event (`N`), modify-order event (`M`), or cancel-order event (`C`)], buy order (`1`) or sell order (`2`), price, and size. For modify order and cancel order events, some fields are left blank; these can be identified by looking up the relevant order ID entries earlier in the file. For iceberg orders, only the visible portion is described in the tick-data file. Entirely hidden orders do not appear in the tick-data file.

New-order events indicate that a limit order arrived and became active. Modify-order events indicate that part — but not all — of an active order matched to an incoming market order, leaving some of the order still active.[3] The size field of a modify-order event indicates the remaining size of the active order after the matching occurred. If the active order involved in the matching was an iceberg order, then the size field of the relevant modify-order event indicates the size of the visible component after the automatic replenishment by the hidden component (see Section 3.1.3). Cancel-order events indicate that an active order departed from the LOB. This can occur for two reasons:

- an active order was cancelled;
- an active order fully matched to an incoming market order.

The tick-data file does not detail whether each cancel-order event relates to an active order cancellation or a full matching. We discuss this issue further in the next section.

3.2.2 Understanding Active Order Departures

Whenever a market order matches to a visible active order,[4] Hotspot FX records the activity in both the trade-data file (as a trade) and the tick-data file (as either a cancel-order event or a modify-order event, depending on whether or not the corresponding active order matches fully). Although both files record all events with millisecond time stamps, the trade-data and tick-data files are recorded by different servers (as

[3] If a trader wishes to change the size or price of one of their active orders, then he/she must first cancel the order and then submit a new one. Such an action is not recorded as a modify-order event in the tick-data file.

[4] If an incoming market order matches to an entirely hidden order, the trade is reported in the trade-data file but the tick-data file contains no record of the matching.

is common with LOB data [114]), so the time stamps for corresponding events in the two files often differ by a small, variable amount. Therefore, we are not able to deduce with certainty whether cancel-order events in the tick-data file correspond to active order cancellations or full matchings.

In Chapters 4 and 5, we study all active order departures together, irrespective of whether they correspond to active order cancellations or full matchings. For our calculations in these chapters, we study all cancel-order events as listed in the tick-data file. In Chapter 6, we study only trades. For our calculations in this chapter, it is necessary for us to infer which trades (as listed in the trade-data file) correspond to which cancel-order events in the tick-data file. We perform such inference via Algorithm 6.1, which we describe in Section 6.3.1.

3.2.3 Constructing $\mathcal{L}(t)$

By first constructing the initial LOB state (according to the snapshot) and then processing each subsequent line in the tick-data file, it is possible to reconstruct the full state of $\mathcal{L}(t)$ at any given time t. Figure 3.2 shows the temporal evolution of $\mathcal{L}(t)$, $b(t)$, $a(t)$, $B(t)$, and $A(t)$ for the EUR/USD LOB on Hotspot FX between 11:00:00 and 12:00:00 GMT on 3 May 2010.

Because trade for each currency pair on Hotspot FX occurs via a quasi-centralized LOB (see Chapter 2), the trading opportunities that a given trader can access depends on his/her trading partnerships. However, Hotspot FX do not disclose any information regarding the trade-partnership network on the platform, and the data does not contain any information about traders' identities, so it is not possible to reconstruct the local LOB $\mathcal{L}_i(t)$ for a specific trader $\theta_i \in \Theta$. We return to this discussion at several points throughout the subsequent chapters.

3.3 Data Pre-Processing and Cleaning

Before we perform our empirical calculations in the subsequent chapters, we pre-process and clean the Hotspot FX data in the following ways.

3.3.1 Excluded Days

For each of the three currency pairs, the mean inter-arrival time for buy and sell limit orders is less than 0.03 seconds. However, the data occasionally contains long periods lasting up to several seconds that include no order arrivals or cancellations. Such

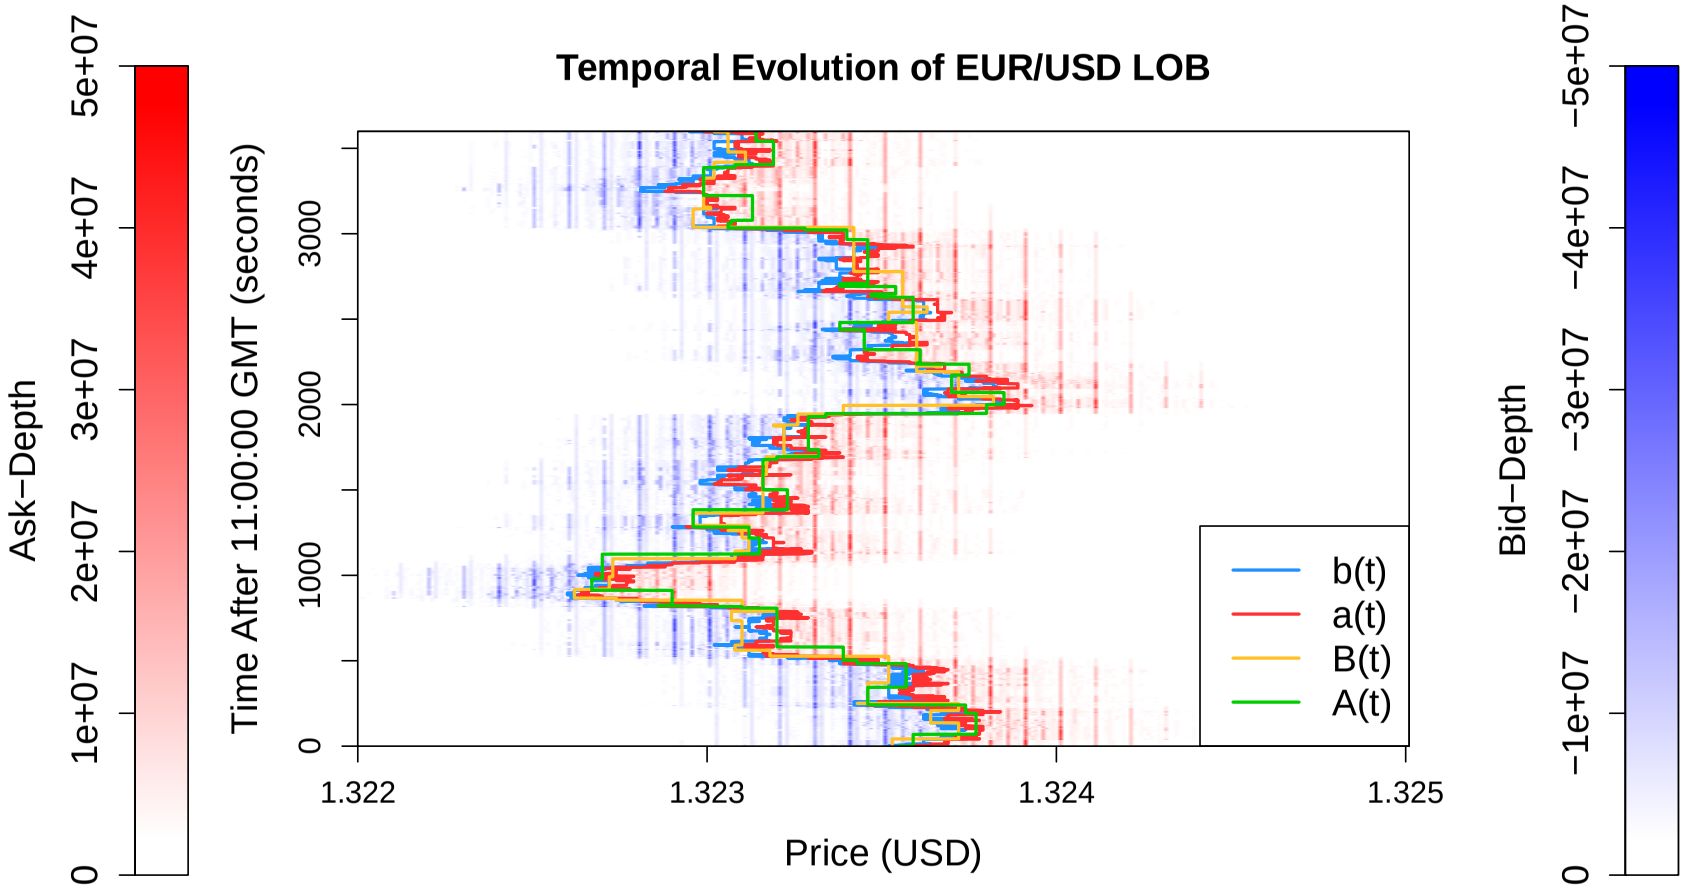

Figure 3.2: Temporal evolution of EUR/USD LOB on Hotspot FX between 11:00:00 and 12:00:00 GMT on 3 May 2010. The blue, red, green, and yellow curves denote $b(t)$, $a(t)$, $B(t)$, and $A(t)$, respectively. Blue and red dots denote active buy and sell orders at a given price.

gaps in recording are caused by a temporary loss of connectivity between the Hotspot FX servers. Although such periods are rare, they cause substantial difficulties in calculating the desired statistical properties from the data. Therefore, we exclude from our study any days with gaps in recording of 30 seconds or more.[5] We also exclude 3 May 2010 (May Bank Holiday in the UK) and 31 May 2010 (Spring Bank Holiday in the UK/Memorial Day in the US) because market activity on these days is extremely low. After these exclusions, a sample of 30 trading days remains.

3.3.2 Peak Trading Hours

On Hotspot FX, each trading day begins at 22:00:00 GMT and runs continuously until 21:59:59 GMT, at which point the platform automatically cancels all active orders [130]. However, far more trades occur during 08:00:00 – 17:00:00 GMT than outside these times (see Figure 3.3). To ensure that our results reflect LOB activity during these *peak trading hours*, we restrict our attention to the period 08:00:00 – 17:00:00 GMT each day.

3.3.3 Inferring Market Orders

If a market order matches to several different active orders (either across several prices, which is often called *walking up the book*,[6] or at a single price), then the trade-data file reports each partial matching as a separate event. Entries that are generated by the same market order appear as consecutive lines in the data, and the time stamps of each such pair of consecutive lines differ by at most 1 millisecond. We use Algorithm 3.1 to infer which lines in the trade-data file correspond to the same market orders.

It is possible that a pair of traders could submit market orders with the same sign that arrive within 0.001 seconds of each other. In this case, Algorithm 3.1 will incorrectly identify the two lines in the trade-data file as corresponding to a single market order. However, the mean time difference between consecutive lines in the trade-data file is of the order of several seconds for each of the three currency pairs, so we deem it extremely unlikely that this will occur often. In the absence of explicit details regarding which line corresponds to which market order, we perform inference

[5] We verified that the choice of 30 seconds was not too short by repeating all of our calculations but excluding any days with gaps in recording of 5 seconds or more. The results were identical because the list of days to exclude was the same in both cases.

[6] Some practitioners use the term *drilling the book* to describe this behaviour, but this terminology is far less common in the published literature.

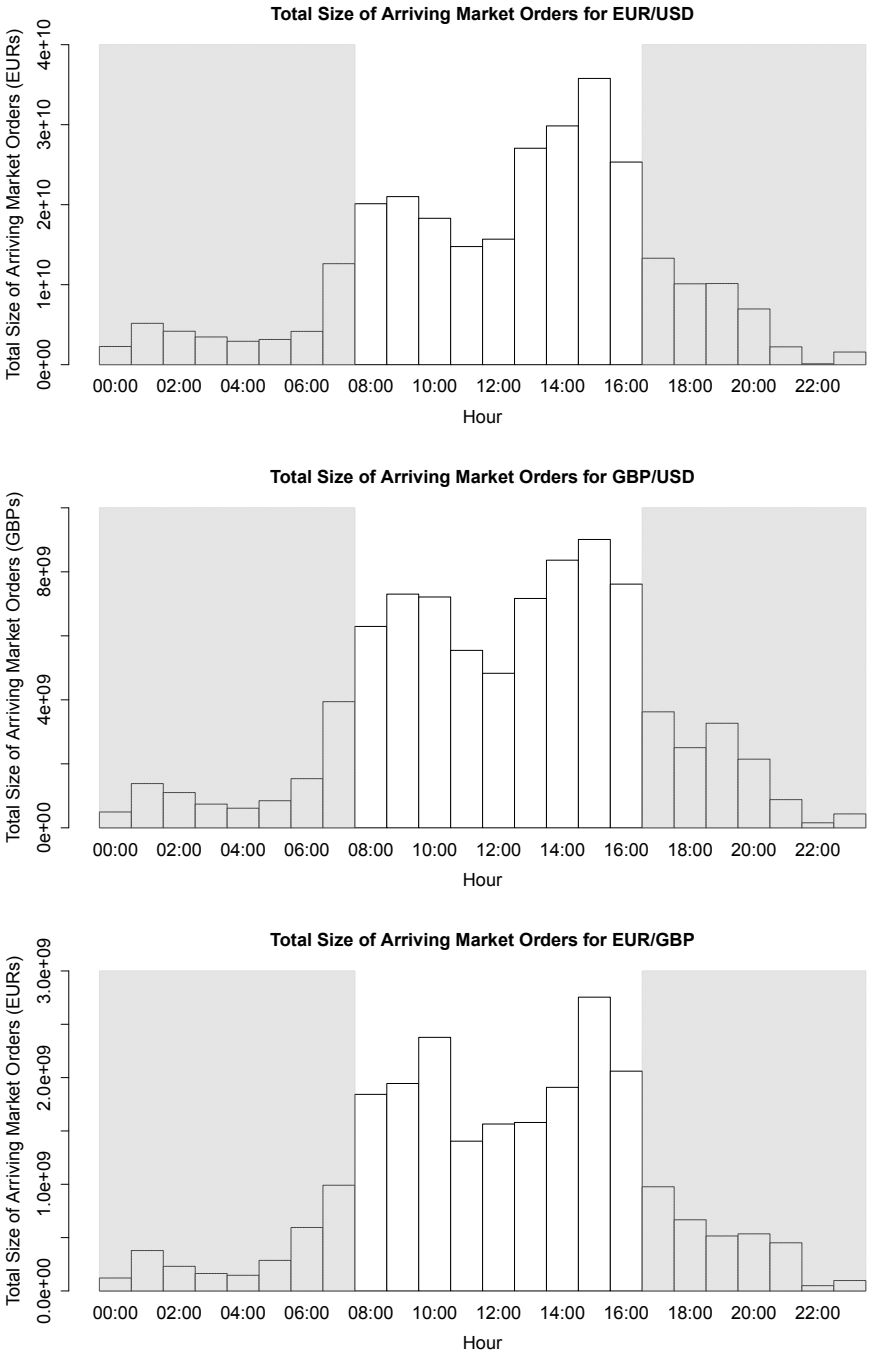

Figure 3.3: Total size of market orders each hour for (top) EUR/USD, (middle) GBP/USD, and (bottom) EUR/GBP. The percentage of total trading volume that occurs between the hours 08:00:00 – 17:00:00 GMT during the data that we study is 71.6%, 72.8%, and 73.7% for EUR/USD, GBP/USD, and EUR/GBP, respectively. The grey regions indicate the hours outside of this time interval.

```
Let j ↤ 0.
Let t_j ↤ −1.
Let p_j ↤ 0.
Let ω_j ↤ 1.
for each line of the trade-data file do
    Denote the time by T.
    Denote the price by P.
    Denote the size by S.
    if T − t_j ≤ 0.001 and sign(S) = sign(ω_j) then
        Let t_j ↤ T.
        Let ω_j ↤ ω_j + S.
        if S < 0 then
            Let p_j ↤ max(p_j, P)
        end
        else
            Let p_j ↤ min(p_j, P)
        end
    end
    else
        if j > 0 then
            Record that the j^th market order arrived at time t_j with price
            p_j and size ω_j.
        end
        Let j ↤ j + 1.
        Let t_j ↤ T.
        Let p_j ↤ P.
        Let ω_j ↤ S.
    end
end
```

Algorithm 3.1: Algorithm to infer whether consecutive lines in the trade-data file correspond to the same market order arrival.

according to Algorithm 3.1 and regard any incorrectly grouped market orders as a source of noise in the data.

The data does not provide a reliable way to perform inference about incoming orders for which part matches immediately and part becomes active. For such orders, we treat the matched part as a market order and the unmatched part as a separate limit order.

3.4 Summary

In this chapter, we have discussed several important features of the Hotspot FX platform and provided a detailed description of the Hotspot FX data. This data — which, after cleaning, describes all limit order arrivals and cancellations and all trades for the EUR/USD, GBP/USD, and EUR/GBP currency pairs during the peak trading hours of 08:00:00 – 17:00:00 GMT for 30 trading days during May 2010 and June 2010 — forms the basis of our empirical investigations throughout Chapters 4–6.

Chapter 4

Statistical Properties of Order Flow and LOB State on Hotspot FX

In this chapter, we perform a statistical analysis of several properties of order flow and LOB state on Hotspot FX. The work in this chapter is also described in the jointly authored paper [100].

4.1 Introduction

During the past 20 years, empirical study of LOB data has provided new insights into many long-standing questions about financial markets — such as the origins and nature of large price changes, market efficiency, volatility, market impact, and the determinants of the bid-ask spread [33, 34, 136, 204]. Although the breadth of such work is substantial (see, e.g., [23, 36, 45, 47, 60, 107, 108, 167, 181, 207]), almost all empirical investigations of LOBs to date have studied data from centralized LOBs. The quasi-centralized LOBs used by multi-institution trading platforms in the FX spot market distinguish them from centralized LOBs in several important ways (see Chapter 2) and thereby present several alternative questions for empirical research.

In this chapter, we perform a statistical analysis of historical data that describes all LOB activity for three liquid currency pairs on Hotspot FX during May 2010 and June 2010 (see Chapter 3). We study several important features of order flow and LOB state, including order arrival and cancellation rates, the mean LOB depth at given prices, and the distributions of limit order arrivals and cancellations. To our knowledge, all empirical studies of order flow and LOB state to date have measured prices relative to the best quotes (see [98] for our recent survey). In a quasi-centralized

LOB, however, traders do not in general know the values of the best quotes in real time. Therefore, we also study an alternative coordinate frame in which we measure prices relative to the most recent trades. We uncover several interesting statistical regularities in each of these two coordinate frames that together provide insight into how traders assess order flow and LOB state.

To date, most empirical studies of LOBs have examined data sets that aggregate LOB activity from many different trading days [23, 36, 45, 60, 78, 107, 108, 167, 181, 207]. We call such data sets *aggregate data*. Although these studies provide useful insight into long-run market behaviour [47, 78, 167], it is unclear whether the statistical properties of aggregate data provide an accurate reflection of LOB activity on shorter time scales. Due to the extremely high levels of market activity on Hotspot FX, we can analyze market behaviour on each trading day separately. We call such data *daily data*. We uncover several interesting differences between the statistical properties of the daily data and those of the aggregate data. In both quote-relative and trade-relative coordinates, we find that the distributions of limit order arrivals and cancellations vary substantially across different trading days. We also uncover a remarkably simple empirical universality: rescaling the data to account for daily differences in the first two moments causes the distributions from each day to collapse onto a single curve. We quantify the importance of daily differences in market activity by comparing two statistical models of the distributions of limit order arrivals and cancellations on a single trading day, and we find that the statistical properties of the aggregate data poorly reflect market activity on a single trading day.

The results that we present in this chapter are important for several reasons. First, they provide a statistical overview of recent activity in a highly liquid market. Second, they illustrate the usefulness of quote-relative coordinates for studying and modelling quasi-centralized LOBs. Third, they illuminate how the microstructure of multi-institution trading platforms in the FX spot market affects order flow and LOB state. Fourth, they provide a new statistical model for the distributions of limit order arrivals and cancellations on a single trading day. Fifth, they demonstrate that long-run statistical properties can poorly reflect market activity on shorter time horizons. This finding is particularly important because many traders implement their investment decisions and trading strategies during a single trading day [7, 25, 188]. Our results suggest that using long-run statistical properties to forecast daily market activity could produce misleading results and lead to poor management of risk exposure.

The remainder of this chapter is organized as follows. In Section 4.2, we survey several statistical analyses of other markets. In Section, 4.3 we describe the methods

that we use to perform our calculations. We present our results for the aggregate data in Section 4.4 and for the daily data in Section 4.5. In Section 4.6, we compare our two statistical models of order flow to assess the importance of daily variations in market activity. We discuss our findings in Section 4.7. In Section 4.8, we summarize our findings and discuss several possible avenues for future research.

4.2 Related Literature

Several authors have conducted statistical analyses of data from centralized LOBs. For a comprehensive discussion of these studies, see our recent survey article [98]. In this section, we review a small selection of empirical results most relevant to our work.

The distribution of limit order arrival rates was reported to exhibit its global maximum at the best quotes $b(t)$ and $a(t)$ in data from the London Stock Exchange (LSE) during 2000–2002 [167], the Shenzhen Stock Exchange during 2003 [107], and the Paris Bourse during 1991 [23] and 2001 [36]. The distribution of limit order arrivals was reported to exhibit a power-law tail [51] in LOB data from the Paris Bourse during 2001 [36], NASDAQ during 2002 [181], the Shenzhen Stock Exchange during 2003 [107], and the LSE during 1998–2000 [207] and 2004 [164], although the reported power-law exponent varies across the different markets. Bouchaud *et al.* [36] conjectured that the heavy tails of the distribution of limit order arrivals suggest that some traders hold optimistic beliefs that large price swings could occur.

Several empirical studies of centralized LOBs for a wide range of different assets (including equities on the Island ECN during 1999 [47, 112], an exchange-traded fund that tracks the NASDAQ 100 during 2002 [181], and S&P 500 futures contracts during 2010–2012 [16]) concluded that cancellations play an important role in the evolution of $\mathcal{L}(t)$. Among these studies, the reported percentage of cancelled active orders ranges from approximately 65% (on the Island ECN in 1999 [112]) to approximately 80% (also on the Island ECN but with no date range specified [47]). The cancellation rates were reported to be highest among active orders close to $b(t)$ and $a(t)$ on NASDAQ during 2002 [181] and on the Tokyo Stock Exchange (no date range specified) [60]. Several authors conjectured that the low cancellation rates among active orders far from $b(t)$ and $a(t)$ provide evidence that traders who place such orders aim to profit from large price swings on long time horizons and must therefore be willing to leave such orders active for long periods of time [47, 181, 207].

The mean LOB state was reported to exhibit a similar "hump" shape in a wide range of centralized LOBs and over a wide range of dates, including the Paris Bourse during 2001 [36], NASDAQ during 2002 [181], the Stockholm Stock Exchange during 1992 [116], and the Shenzhen Stock Exchange during 2003 [108]. Rosu [187] conjectured that such a hump represents a trade-off between the optimism that limit orders placed far from the spread might eventually result in a significant profit and the pessimism that such orders might never match.

Although the breadth of empirical work on centralized LOBs is substantial, the overwhelming picture painted by the existing literature is that the data studied is old (see Figure 4.1). In recent years, the speed and frequency of trading has increased dramatically [57], volumes have soared [13], and traders' strategies have evolved to incorporate new technologies [50]. Therefore, it is unlikely that old data provides an accurate reflection of today's markets.

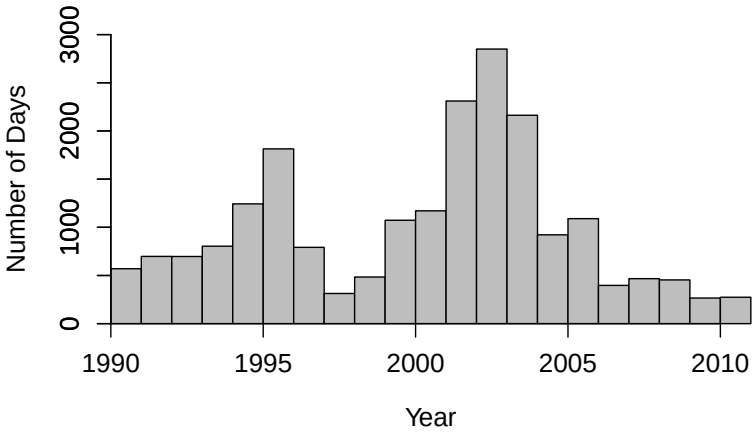

Figure 4.1: Approximate total number of days' data per year that was examined by the empirical studies discussed in our recent survey article [98].

4.3 Methodology

4.3.1 Time Scales and Rates

Throughout this chapter, we perform all of our calculations in *event time,* whereby we advance the clock by 1 unit whenever a limit order arrives. Event time helps to remove nonstationarities in the statistical properties of order flow and LOB state that occur in calendar time due to irregular bursts of trading activity [45, 102, 162, 193, 201].

The number of cancel-order events (which denote either an active order cancellation or a full matching to an arriving market order — see Section 3.2.1.2) in each unit of time varies. We reset the clock at the start of each trading day so that the first limit order each day has $t_x = 1$.

4.3.2 Coordinate Frames

Almost all existing publications on LOBs measure order prices relative to $b(t)$ and $a(t)$ (see, e.g., [23, 36, 45, 47, 60, 107, 108, 167, 181, 207]). We call this coordinate frame *quote-relative coordinates.*

Definition. *The* quote-relative price of an order x at time t *is*

$$\delta(p_x, t) := \begin{cases} b(t) - p_x, & \text{if } x \text{ is a buy order,} \\ p_x - a(t), & \text{if } x \text{ is a sell order.} \end{cases} \tag{4.1}$$

The use of quote-relative pricing in centralized LOBs is motivated by the notion that traders monitor $b(t)$ and $a(t)$ when deciding how to act. Because of the quasi-centralized structure of the platform, each trader θ_i on Hotspot FX is able to view their local bid and ask prices $b_i(t)$ and $a_i(t)$ in real time (see Section 2.3.3) but does not know the values of the global bid and ask prices $b(t)$ and $a(t)$. We are unable to study the Hotspot FX data using a coordinate frame that measures prices relative to $b_i(t)$ and $a_i(t)$ because we do not have access to the values of $b_i(t)$ and $a_i(t)$ for any $\theta_i \in \Theta$ (see Section 3.2.3). However, all traders on Hotspot FX are able to calculate the time series of previous trade prices $B(t)$ and $A(t)$ from their trade-data stream (see Section 2.3.4), and we are able to construct these series from the trade-data files (see Section 3.2.1.1). Therefore, we also study the Hotspot FX data using *trade-relative coordinates.*

Definition. *The* trade-relative price of an order x at time t *is*

$$\Delta(p_x, t) := \begin{cases} B(t) - p_x, & \text{if } x \text{ is a buy order,} \\ p_x - A(t), & \text{if } x \text{ is a sell order.} \end{cases} \tag{4.2}$$

4.3.3 The Trade-Data Stream

To prevent traders from exploiting trade information in real time, Hotspot FX lags dissemination of the trade-data stream by 1 second. For example, if a trade occurred at 08:34:26.346, the time stamp in the trade-data file would be 08:34:26.346 but the information would not be disseminated in the trade-data stream until 08:34:27.346. In order to align the activity reported in the trade-data file (which is a transcription

of the trade-data stream that traders on Hotspot FX access in real time) with the time that the relevant information arrived on traders' screens, we add 1 second to the times reported in the trade-data file when calculating the $B(t)$ and $A(t)$ series.

4.3.4 Trading Days

In Section 4.5, we study the statistical properties of each trading day separately. We choose a single trading day as our longitudinal unit for three reasons. First, a trading day on Hotspot FX represents a structural cycle because the platform automatically cancels all active orders at the end of each day (see Section 3.3.1). Second, a single trading day provides a compromise between including sufficiently many data points in each longitudinal unit to ensure statistically stable estimation and including sufficiently many longitudinal units to perform useful cross-sectional comparisons. Third, several empirical studies have reported that most traders implement their investment decisions and trading strategies during a single trading day [7, 25, 188]. The most useful statistics for such traders are consequently those that describe market behaviour on a single trading day.

4.4 Aggregate Market Activity

In this section, we present the results of our calculations for the aggregate data (see Section 4.1).

4.4.1 Summary Statistics

Table 4.1 lists several summary statistics regarding order flow and the state of the global LOB $\mathcal{L}(t)$. Although the total size and total number of limit order arrivals and cancellations are similar for each of the three currency pairs, the corresponding statistics for market orders vary across the currency pairs by more than a factor of 10. Similarly, the mean total size of active orders is about 10^8 units of the base currency in all cases, but the ratio of the mean total size of active orders to the total size of market orders[1] varies between roughly 10^{-2} for EUR/USD to 10^{-1} for EUR/GBP.

The Hotspot FX data does not provide a way to reconstruct the local LOB $\mathcal{L}_i(t)$ for any specific trader $\theta_i \in \Theta$ (see Section 3.2.1.2), so it is unclear what fraction of the total size of active orders each trader is able to access. Only a small percentage of market orders walk up the book, which suggests that traders assess their local LOBs

[1] This ratio was proposed by Wyart *et al.* [204] as a simple measure of liquidity.

		EUR/USD		GBP/USD		EUR/GBP	
		Buy	Sell	Buy	Sell	Buy	Sell
Total Size (units of base currency)	Limit Orders	1.50×10^{14}	1.51×10^{14}	1.18×10^{14}	1.18×10^{14}	9.24×10^{13}	9.22×10^{13}
	Market Orders	6.68×10^{10}	7.06×10^{10}	2.32×10^{10}	2.26×10^{10}	5.49×10^{09}	6.51×10^{9}
	Cancellations	1.50×10^{14}	1.51×10^{14}	1.18×10^{14}	1.18×10^{14}	9.24×10^{13}	9.22×10^{13}
Total Number (orders)	Limit Orders	6.80×10^{7}	6.80×10^{7}	6.54×10^{7}	6.57×10^{7}	4.41×10^{7}	4.39×10^{7}
	Market Orders	8.32×10^{4}	8.48×10^{4}	4.31×10^{4}	4.44×10^{4}	7.22×10^{3}	8.23×10^{3}
	Cancellations	6.80×10^{7}	6.80×10^{7}	6.54×10^{7}	6.57×10^{7}	4.41×10^{7}	4.39×10^{7}
Mean Size (units of base currency)	Limit Orders	2.21×10^{6}	2.22×10^{6}	1.80×10^{6}	1.80×10^{6}	2.10×10^{6}	2.10×10^{6}
	Market Orders	8.03×10^{5}	8.33×10^{5}	5.39×10^{5}	5.09×10^{5}	7.61×10^{5}	7.91×10^{5}
	Cancellation	2.21×10^{6}	2.22×10^{6}	1.80×10^{6}	1.80×10^{6}	2.10×10^{6}	2.10×10^{6}
Mode Size (units of base currency)	Limit Orders	1.00×10^{6}	1.00×10^{6}	1.00×10^{6}	1.00×10^{6}	1.00×10^{6}	1.00×10^{6}
	Market Orders	1.00×10^{6}	1.00×10^{6}	1.00×10^{6}	1.00×10^{6}	1.00×10^{6}	1.00×10^{6}
	Cancellations	1.00×10^{6}	1.00×10^{6}	1.00×10^{6}	1.00×10^{6}	1.00×10^{6}	1.00×10^{6}
Mean Inter-Arrival Time (seconds)	Limit Orders	0.0143	0.0143	0.0149	0.0148	0.0220	0.0221
	Market Orders	11.7	11.5	22.6	21.9	135	118
	Cancellations	0.0143	0.0143	0.0149	0.0148	0.0221	0.0221
Percentage of Market Orders	in Arriving Order Flow	0.0444%	0.0469%	0.0197%	0.0191%	0.00595%	0.00706%
	that Walk Up the LOB	8.27%	8.54%	6.70%	5.90%	3.93%	4.52%
Mean Total Size of Active Orders (units of base currency)		2.83×10^{8}	2.96×10^{8}	1.64×10^{8}	1.66×10^{8}	9.40×10^{7}	9.46×10^{7}

Table 4.1: Summary statistics for aggregate market activity in EUR/USD, GBP/USD, and EUR/GBP LOBs.

$\mathcal{L}_i(t)$ and rarely submit market orders that exceed the depth at $b_i(t)$ or $a_i(t)$. The mean size of market orders is smaller than that of limit orders (see Table 4.1), which further suggests that traders are reluctant to submit large market orders. These observations are consistent with a strategy called "order splitting" [34, 33, 137], by which traders who wish to conduct large trades submit several small market orders to minimize the impact of their actions [3, 4].

For each of the three currency pairs, the modal size of limit order arrivals, market order arrivals, and cancellations is 1 million units of the base currency. This indicates that even though the lot size for EUR/USD, GBP/USD, and EUR/GBP on Hotspot FX is $\sigma = 0.01$ units of the base currency (see Section 3.2), many traders choose much larger sizes for their orders. We discuss several possible reasons that traders may favour "round-number" order sizes in Section 4.7.

Market orders constitute less than 0.05% of the total arriving order flow for each of the three currency pairs. This percentage is much smaller than those reported for other markets (which range from about 10% to about 30% [47, 89, 112, 143, 181]). This implies that market orders play a much less prominent role in the evolution of $\mathcal{L}(t)$ on Hotspot FX than has been reported elsewhere in the literature.

4.4.2 Bid-Ask Spread

The mean global spreads (i.e., the mean of $s(t) = a(t) - b(t)$) for the EUR/USD, GBP/USD, and EUR/GBP LOBs are 3.62 ticks, 9.54 ticks, and 10.1 ticks, respectively. Figure 4.2 shows the empirical cumulative distribution functions (ECDFs) for each currency pair's global spread. The global spread tends to be smaller for EUR/USD than for the other two currency pairs. In a centralized LOB, a smaller value of $s(t)$ is often regarded as a sign of greater liquidity [70], because $s(t)$ determines the cost of conducting a round-trip trade (i.e., buying a single unit at $a(t)$ and selling a single unit at $b(t)$ using a pair of simultaneous market orders). In a quasi-centralized LOB, however, $s(t)$ does not have a clear interpretation because the liquidity available to each trader depends on his/her local LOB $\mathcal{L}_i(t)$.

If $s(t) < 0$, then $b(t) < a(t)$, so an arbitrage opportunity exists for any trader θ_k who is trading partners with both the owner θ_i of an active order at $b(t)$ and the owner θ_j of an active order at $a(t)$. The percentage of time that this occurs in the EUR/USD, GBP/USD, and EUR/GBP LOBs is 9.99%, 4.08%, and 0.23%, respectively. This suggests that among the traders who place limit orders closest to the best quotes, there exist many pairs of traders $\theta_i \nleftrightarrow \theta_j$ with no mutual trading partners.

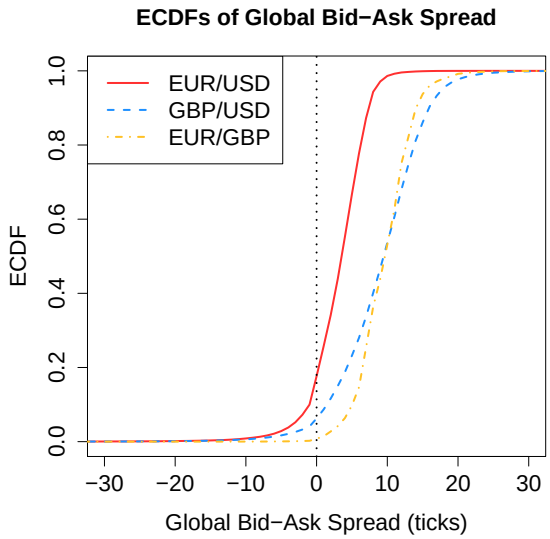

Figure 4.2: ECDFs of global bid-ask spread $s(t) = a(t) - b(t)$ for EUR/USD, GBP/USD, and EUR/GBP. The dotted black line denotes $s(t) = 0$.

4.4.3 Limit Order Arrivals and Cancellations

In this section, we study the distributions of limit order arrivals and cancellations[2] in quote-relative and trade-relative coordinates.

Definition. *Let L denote the set of buy limit order arrivals. The* quote-relative distribution of buy limit order arrivals *is given by*

$$g(p) = \frac{\sum_{\{x \in L | \delta(p_x, t_x) = p\}} \omega_x}{\sum_{x \in L} \omega_x}. \tag{4.3}$$

The corresponding definitions for sell limit orders are similar. The corresponding definitions in trade-relative coordinates are similar, using the trade-relative price $\Delta(p_x, t_x)$ instead of the quote-relative price $\delta(p_x, t_x)$. The corresponding definitions for cancellations are similar, using the set of active order cancellations and corresponding cancellation times instead of the set of active order arrivals and corresponding arrival times.

Figure 4.3 shows the distributions of limit order arrivals for each of the three currency pairs. Limit order arrivals for EUR/USD and EUR/GBP tend to occur closer to the best quotes than do those for GBP/USD, and limit order arrivals for

[2]Recall from Section 3.2.2 that "cancellations" refer to cancel-order events as listed in the Hotspot FX tick-data files, which occur when an active order is either cancelled by its owner or fully matched to an incoming market order.

EUR/USD tend to occur closer to the most recently traded price than do those for GBP/USD and EUR/GBP.

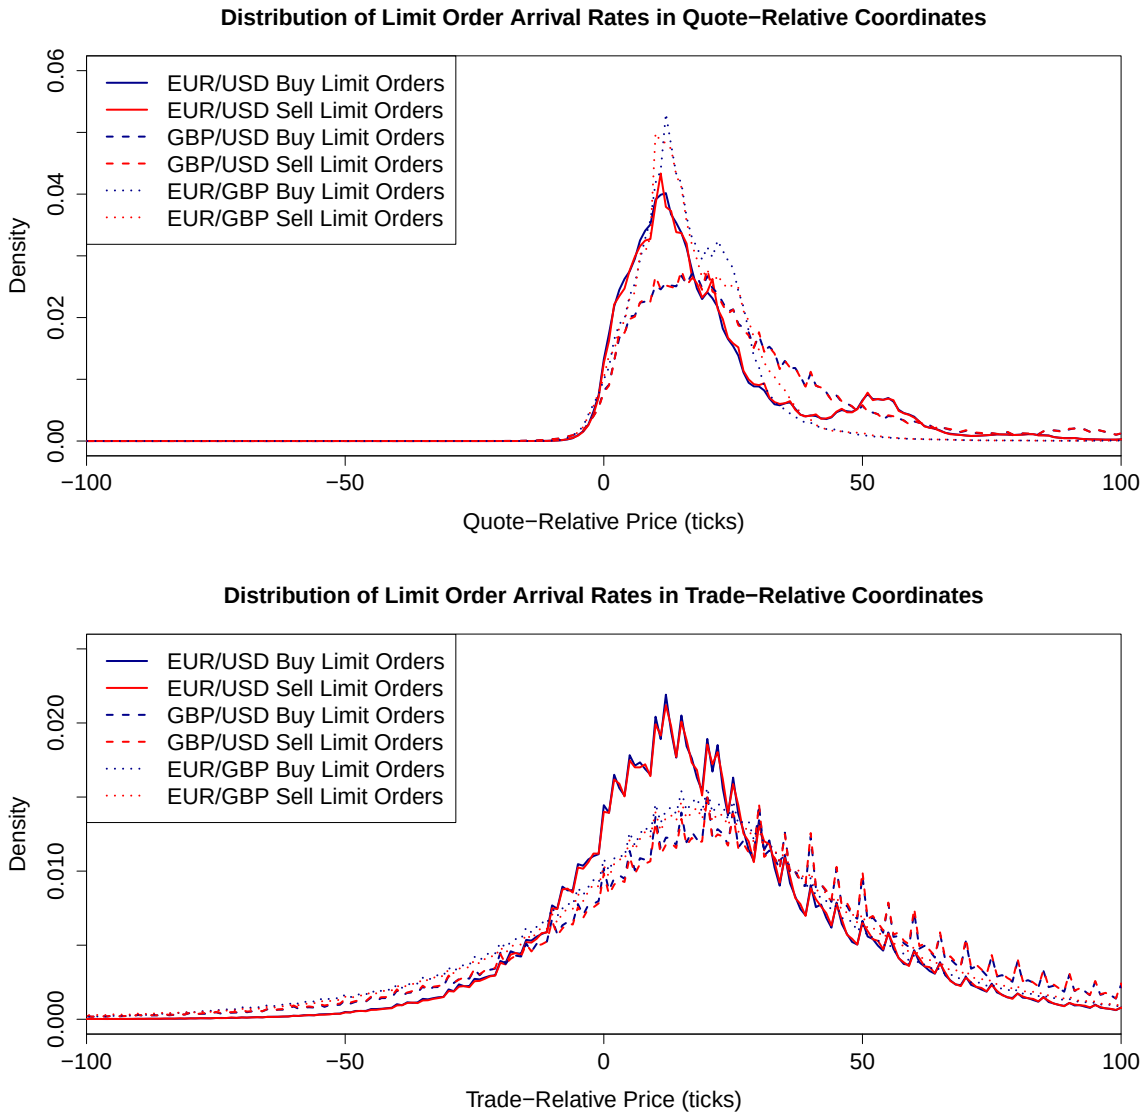

Figure 4.3: Distributions of limit order arrivals in (top) quote-relative and (bottom) trade-relative coordinates for EUR/USD, GBP/USD, and EUR/GBP.

In both quote-relative and trade-relative coordinates, each of three currency pairs' maximum limit order arrival rates occur at strictly positive relative prices. This contrasts to the findings of several empirical studies of other LOBs [23, 36, 116, 167, 181], which report that the maximum limit order arrival rate occurs at the best quotes. We propose two possible reasons for why this is not the case on Hotspot FX. First, traders on Hotspot FX cannot view the values of $b(t)$ and $a(t)$, but instead

view $b_i(t) \leq b(t)$ and $a_i(t) \geq a(t)$. If each trader θ_i bases his/her trading decisions on $b_i(t)$ and $a_i(t)$, and if $b_i(t)$ and $a_i(t)$ both typically reside at a strictly positive quote-relative price, then the maximum arrival rate of the aggregate limit order flow generated by all traders will similarly occur at a strictly positive quote-relative price. Second, the tick size for each of the three currency pairs on Hotspot FX is smaller than the tick sizes for the LOBs studied in Refs [23, 36, 116, 167, 181]. Therefore, traders on Hotspot FX are able to specify order prices with greater accuracy than was the case for the LOBs in the other empirical studies. Given this extra choice, traders might prefer to place limit orders slightly behind the best quotes.

Figure 4.4 shows the distributions of cancellations for each of the three currency pairs.[3] In contrast to limit order arrivals, it is not possible for cancel-order events to occur at negative quote-relative prices, because the lowest possible quote-relative price of an active order is 0 (which occurs for orders at $b(t)$ or $a(t)$).

Each of three currency pairs' quote-relative cancellation distributions have a local maximum at 0, but cancellations for EUR/USD tend to occur closer to the most recently traded price than do those for the other two currency pairs. Although some cancel-order events are caused by market order arrivals, the total size of market orders constitutes less than 0.05% of order flow for each of the three currency pairs (see Table 4.1). Moreover, market orders in a quasi-centralized LOB do not necessarily match with a quote-relative price of zero, because a sell (respectively, buy) market order x submitted by a trader $\theta_i \in \Theta$ matches to the highest-priority active buy (respectively, sell) order that is owned by a trader θ_j such that $\theta_i \leftrightarrow \theta_j$, but it bypasses any higher-priority active buy (respectively, sell) orders that are owned by traders θ_k such that $\theta_k \nleftrightarrow \theta_i$ (we study this topic in detail in Chapter 6). Therefore, the large number of cancellations at a quote-relative price of 0 are unlikely to be strongly influenced by market order arrivals. This suggests that traders cancel active orders at the best quotes particularly often. For strictly positive quote-relative prices, the cancellation distributions have qualitatively similar shapes to those of limit order arrivals (see Figure 4.3).

For each of the three currency pairs, the trade-relative distributions of limit order arrivals have qualitatively similar shapes to those of cancellations. In all cases, limit order arrivals and cancellations tend to occur particularly frequently at trade-relative prices that are integer multiples of 10. This suggests that many traders calculate the trade-relative prices of their limit order submissions and cancellations, which in turn

[3]Observe that these distributions count the total size of cancel-order events at a given relative price, not the cancellation rate per active order (which we calculate in Section 4.4.5).

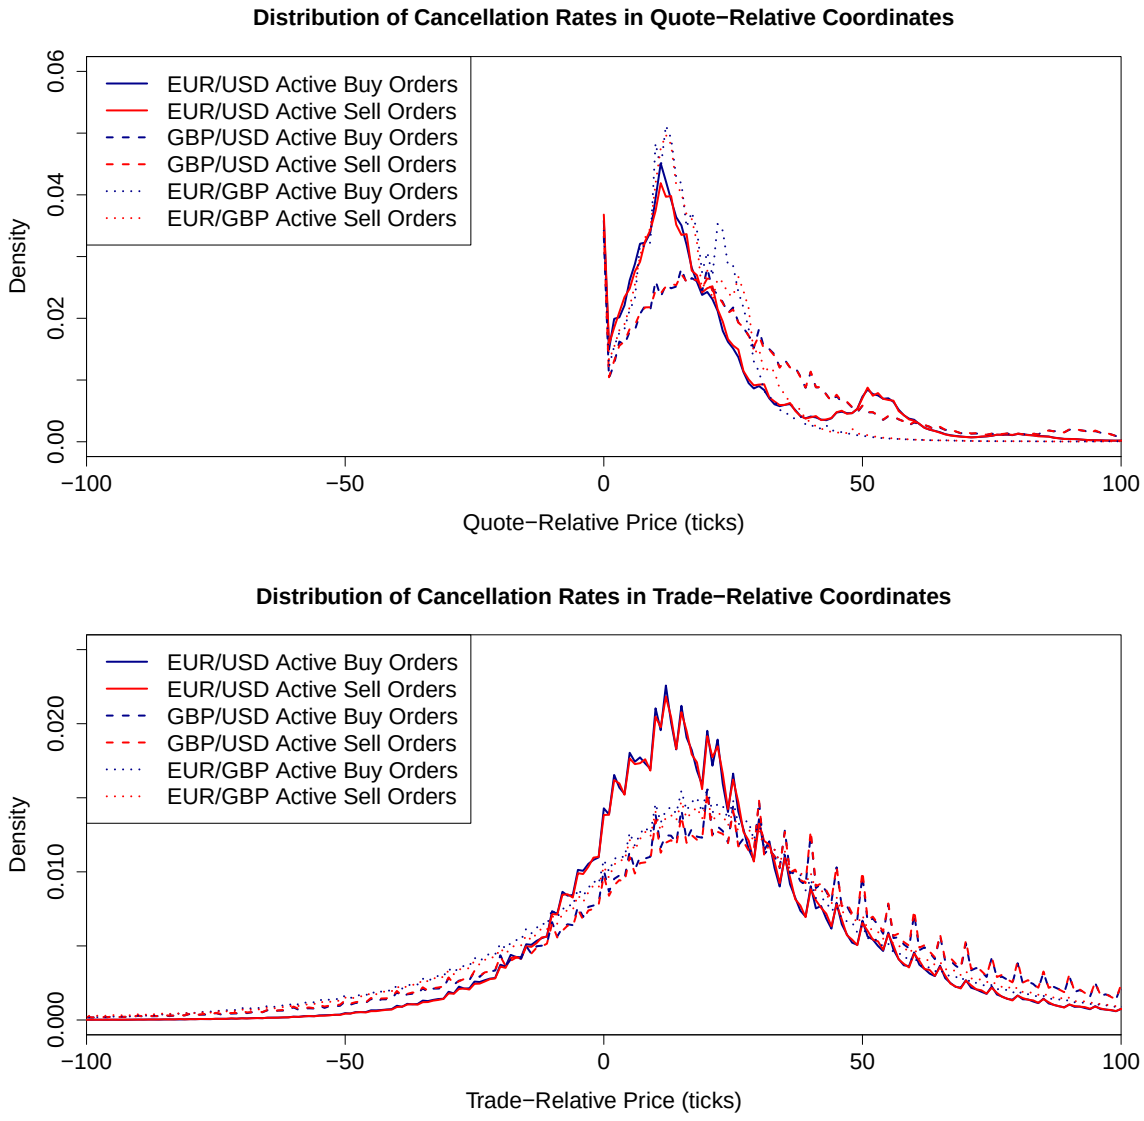

Figure 4.4: Distributions of active order cancellations in (top) quote-relative and (bottom) trade-relative coordinates for EUR/USD, GBP/USD, and EUR/GBP.

suggests that traders regard the trade-price series $B(t)$ and $A(t)$ to contain important information when deciding how to act.

4.4.4 Mean Relative Depth Profiles

In this section, we study the mean LOB depths at given quote-relative and trade-relative prices. Recall from Section 2.2.2 that $\mathcal{B}(t)$ and $\mathcal{A}(t)$ denote, respectively, the set of active buy orders and sell orders at time t.

Definition. *The* bid-depth at quote-relative price p and time t *is*

$$n(p,t) = \sum_{\{x \in \mathcal{B}(t) | \delta(p_x,t)=p\}} |\omega_x| \,. \tag{4.4}$$

Definition. *The* mean bid-depth at quote-relative price p during the time interval $[t_0, t_1]$ *is*

$$\overline{n}(p,t_0,t_1) = \frac{1}{t_1 - t_0} \int_{t_0}^{t_1} n(p,t) \mathrm{d}t. \tag{4.5}$$

Definition. *The* mean bid-side quote-relative depth profile during the time interval $[t_0, t_1]$ *is the set of all ordered pairs* $(\overline{n}(p,t_0,t_1), p)$.

The corresponding definitions for the ask side are similar, using $\mathcal{A}(t)$ instead of $\mathcal{B}(t)$. The corresponding definitions in trade-relative coordinates are similar, using the trade-relative price $\Delta(p_x, t)$ instead of the quote-relative price $\delta(p_x, t)$.

Figure 4.5 shows the mean relative depth profiles for each of the three currency pairs. The mean total depth of active orders with a trade-relative price close to 0 is larger for EUR/USD than it is for GBP/USD and EUR/GBP. Similarly, the mean total depth of active orders near the best quotes is larger for EUR/USD than it is for the other two currency pairs.

In quote-relative coordinates, all of the mean relative depth profiles have a local maximum at 0. This implies that, on average, larger depths are available at $b(t)$ and $a(t)$ than at neighbouring prices. In both quote-relative and trade-relative coordinates, all three currency pairs' mean relative depth profiles exhibit the "hump" shape that has been reported widely in the literature [36, 108, 116, 181]. For the LOBs examined by the other empirical studies, however, market orders played a central role in the evolution of $\mathcal{L}(t)$ near the best quotes. On Hotspot FX, by contrast, market orders play a much less prominent role (see Table 4.1). Therefore, the hump shape of the mean relative depth profile is primarily maintained by limit order arrivals and active order cancellations.

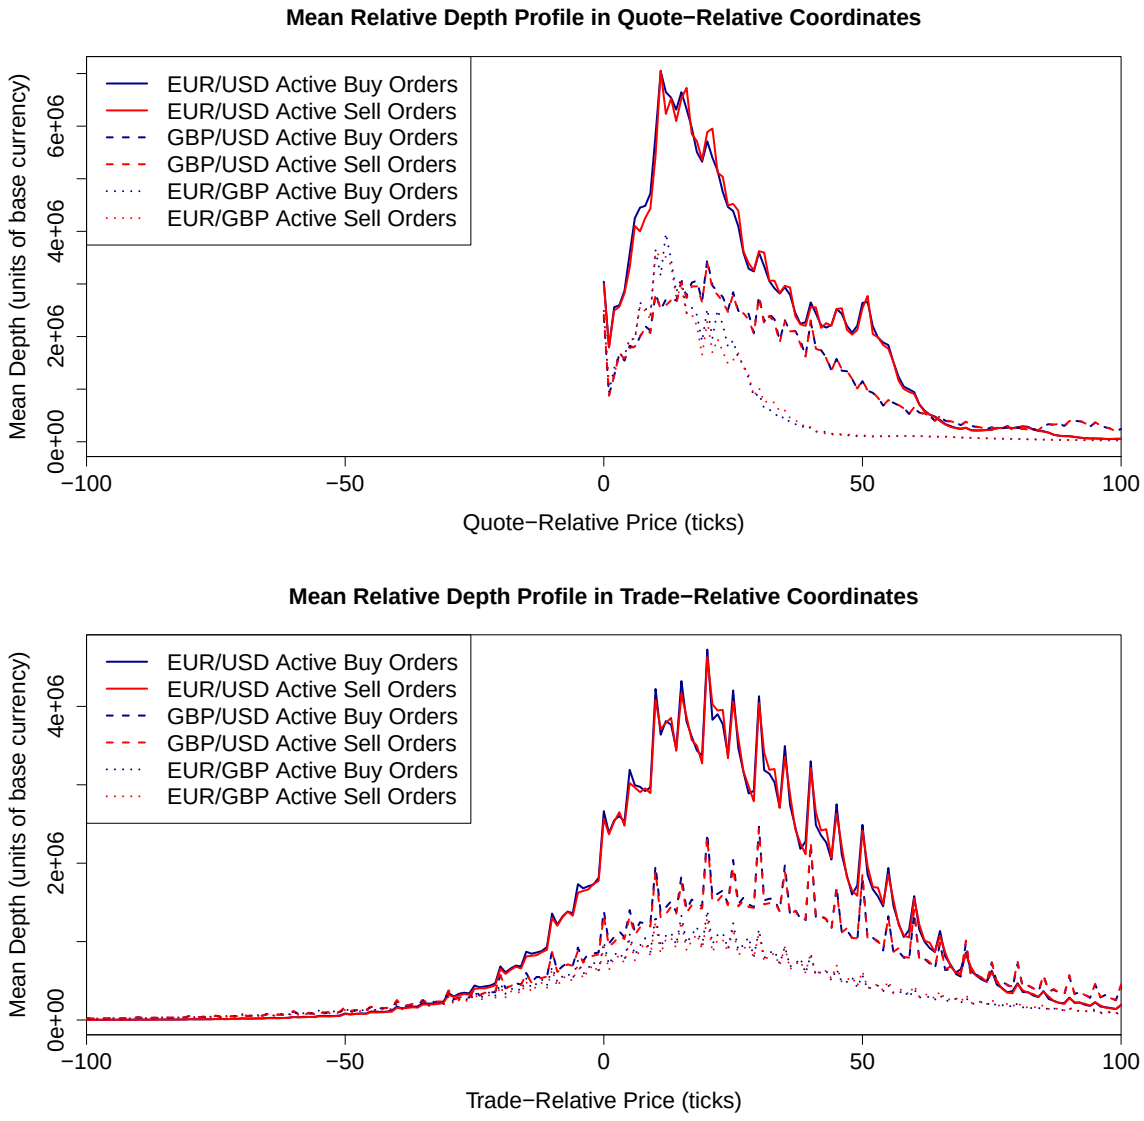

Figure 4.5: Buy-side and sell-side mean relative depth profiles in (top) quote-relative and (bottom) trade-relative coordinates for EUR/USD, GBP/USD, and EUR/GBP.

Similarly to the distributions of limit order arrivals and cancellations (see Section 4.4.3), the mean depths at trade-relative prices that are integer multiples of 10 tend to be larger than those at neighbouring relative prices. We discuss this observation further in Section 4.7.

4.4.5 Cancellation Profiles

In this section, we study the cancellation rate per active order at given quote-relative and trade-relative prices.

Definition. *Let J denote the set of active buy orders that were cancelled during the time interval $[t_0, t_1]$ and let $c_1, c_2, \ldots, c_N$ denote the cancellation times. The* mean buy-order cancellation rate at quote-relative price p during $[t_0, t_1]$ *is*

$$h(p, t_0, t_1) = \frac{\sum_{\{x \in J | \delta(p_x, c_x) = p\}} |\omega_x|}{(t_1 - t_0)\overline{n}(p, t_0, t_1)}. \tag{4.6}$$

The function h describes the mean fraction of active orders at trade-relative price p that are cancelled per arriving order during time interval $[t_0, t_1]$.

Definition. *The* buy-order quote-relative cancellation profile during $[t_0, t_1]$ *is the set of all ordered pairs* $(h(p, t_0, t_1), p)$.

The corresponding definitions for sell orders are similar. The corresponding definitions in trade-relative coordinates are similar, using the trade-relative price $\Delta(p_x, c_x)$ instead of the quote-relative price $\delta(p_x, c_x)$.

Figure 4.6 shows the cancellation profiles for each of the three currency pairs. In quote-relative coordinates, the cancellation rates vary considerably with relative price. This is unsurprising, however, given that traders on Hotspot FX are not able to calculate the quote-relative prices of their orders in real time and therefore cannot use such information when deciding whether to cancel their active orders.

A different picture emerges in trade-relative coordinates. For each of the three currency pairs, the cancellation rate is lower at trade-relative prices that are integer multiples of 5 than it is at neighbouring trade-relative prices. This implies that on average, traders leave orders active for longer at these "round-number" (see Section 4.7) trade-relative prices than they do at other trade-relative prices. Aside from these round-number effects, each currency pair's cancellation rate is approximately constant among orders with negative trade-relative price, then decreases for positive trade-relative prices (although this decrease occurs for larger trade-relative prices for EUR/GPB than it does for EUR/USD and GBP/USD). This finding supports the

hypothesis that traders who place orders with larger relative prices must be willing to leave such orders active for longer times [47, 181, 207].

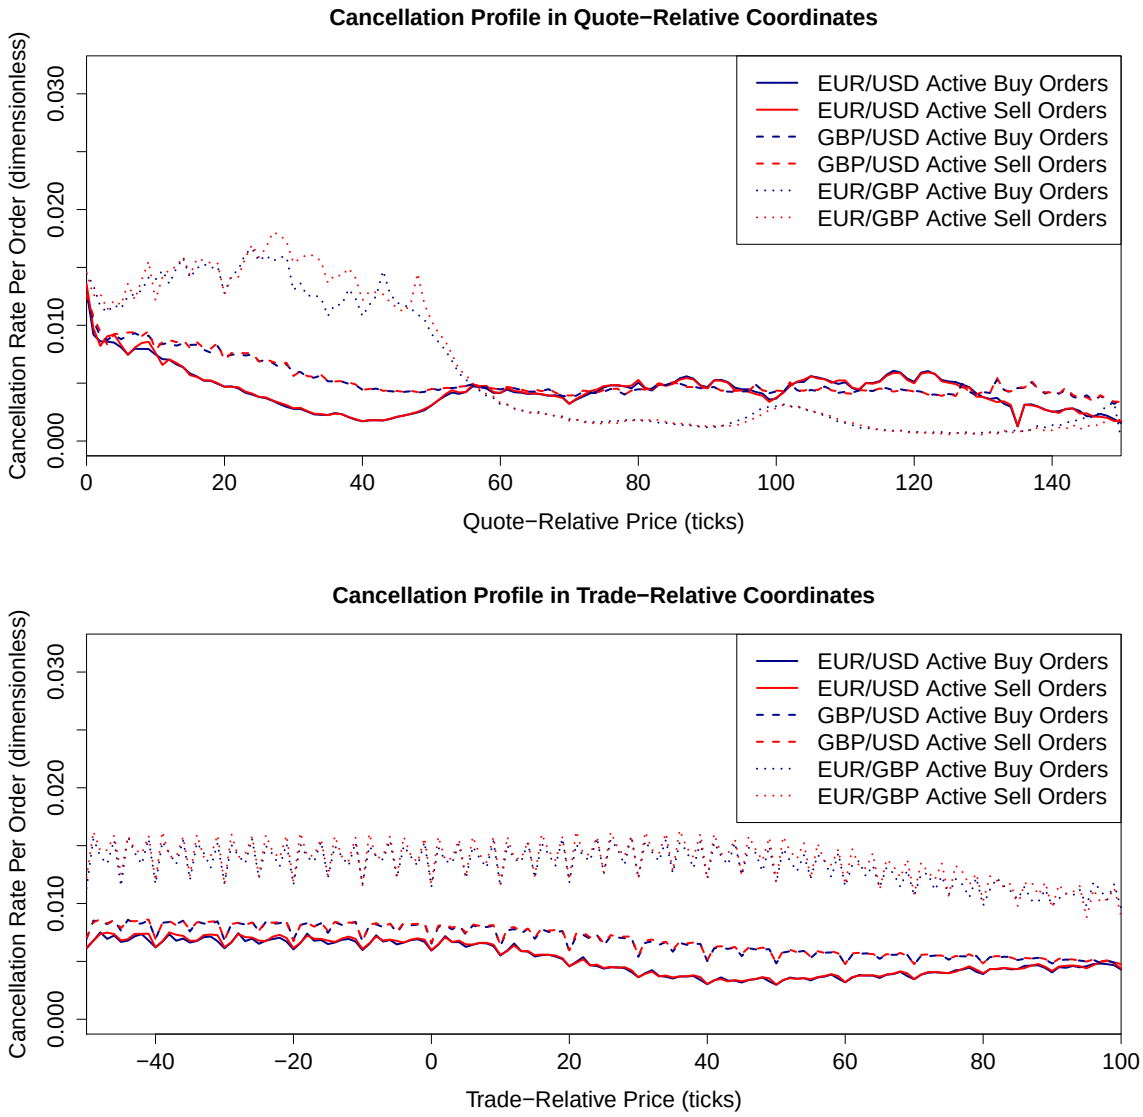

Figure 4.6: Buy-order and sell-order cancellation profiles in (top) quote-relative and (bottom) trade-relative coordinates for EUR/USD, GBP/USD, and EUR/GBP.

4.5 Daily Market Activity

In this section, we present the results of our calculations for the daily data (see Section 4.1).

4.5.1 Summary Statistics

Figure 4.7 shows the total size of limit orders, market orders, and cancelled active orders each day. Figure 4.8 shows similar plots for the mean size. The total size of limit orders each day exhibits substantial longitudinal variation, but the longitudinal variation in mean limit order size is relatively small. Similar results hold for market orders and cancellations. Therefore, differences in order flow each day are primarily caused by differences in the total number of order arrivals and cancellations rather than by differences in their mean size.

We calculate the following imbalance factors to assess the imbalance between buy and sell order flow.

Definition. *Let β_i and α_i denote, respectively, the total size of arriving buy limit orders and sell limit orders for a given currency pair on day i. The* limit order imbalance factor *is*

$$\frac{1}{30}\sum_{i=1}^{30}\frac{|\beta_i - \alpha_i|}{(\beta_i + \alpha_i)/2}. \tag{4.7}$$

The corresponding definition for market orders is similar, using the total size of arriving buy market orders and sell market orders. The corresponding definition for cancellations is similar, using the total size of cancelled active buy orders and active sell orders.

Table 4.2 lists the limit order, market order, and cancellation imbalance factors for each of the three currency pairs.

	EUR/USD	GBP/USD	EUR/GBP
Limit Orders	0.008	0.008	0.005
Market Orders	0.109	0.101	0.249
Cancellations	0.008	0.008	0.005

Table 4.2: Limit order, market order, and cancellation imbalance factors for EUR/USD, GBP/USD, and EUR/GBP.

For each of the three currency pairs, the limit order and cancellation imbalance factors are small, but the market order imbalance factors exceed 10% in all cases. This

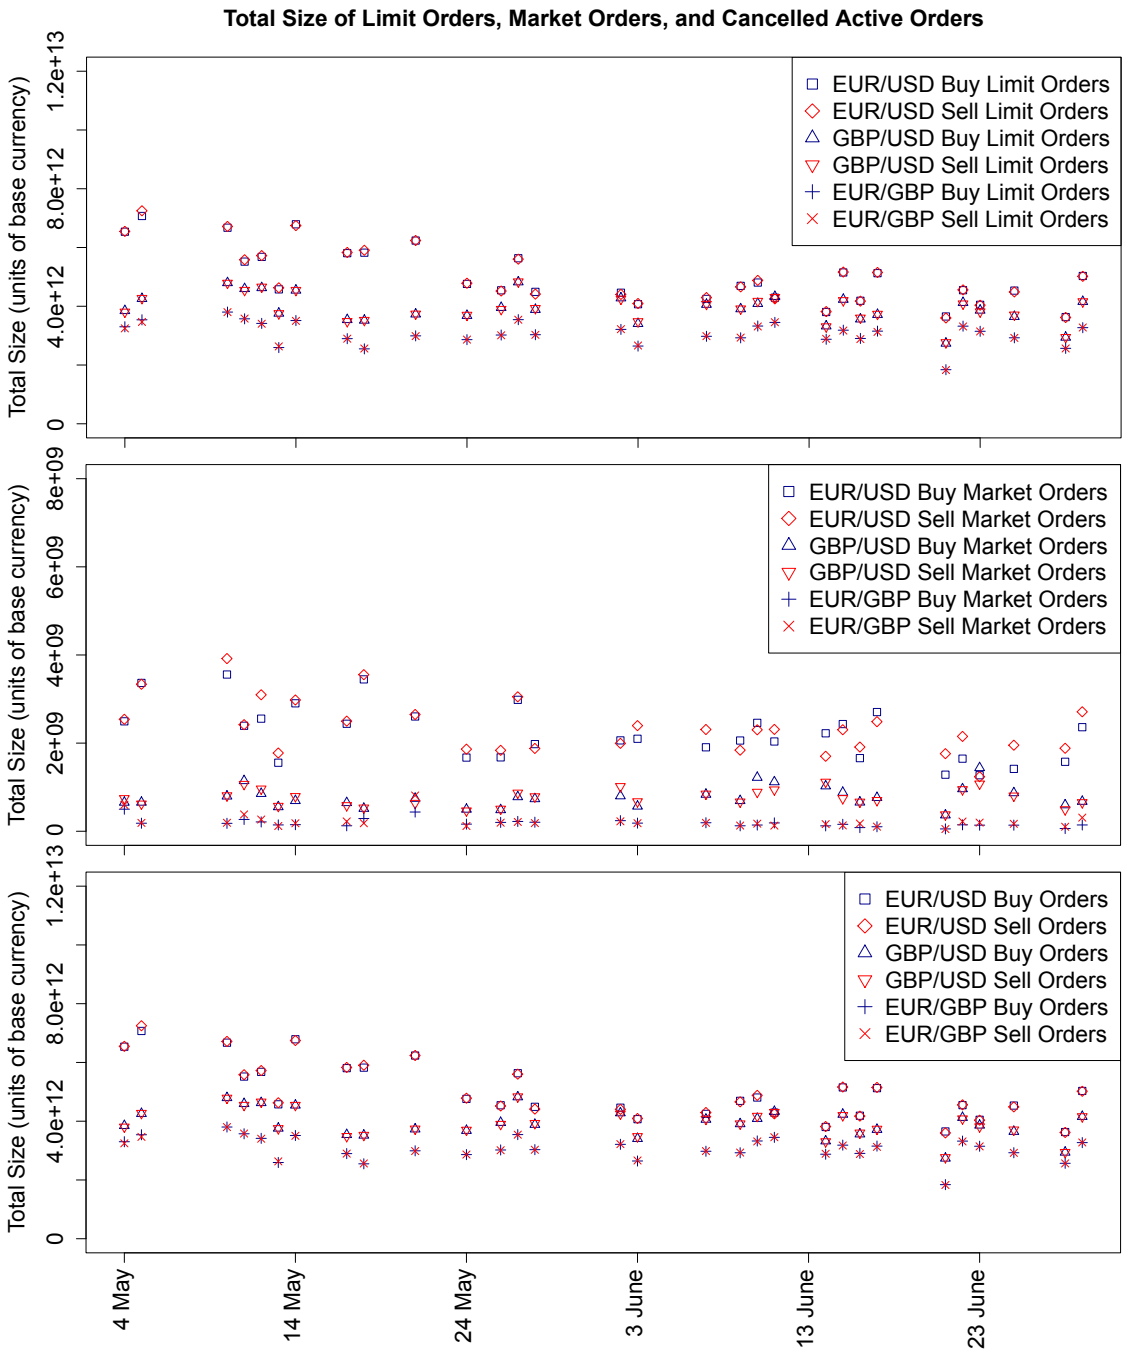

Figure 4.7: Total size of (top panel) limit orders, (middle panel) market orders, and (bottom panel) cancelled active orders each day for EUR/USD, GBP/USD, and EUR/GBP.

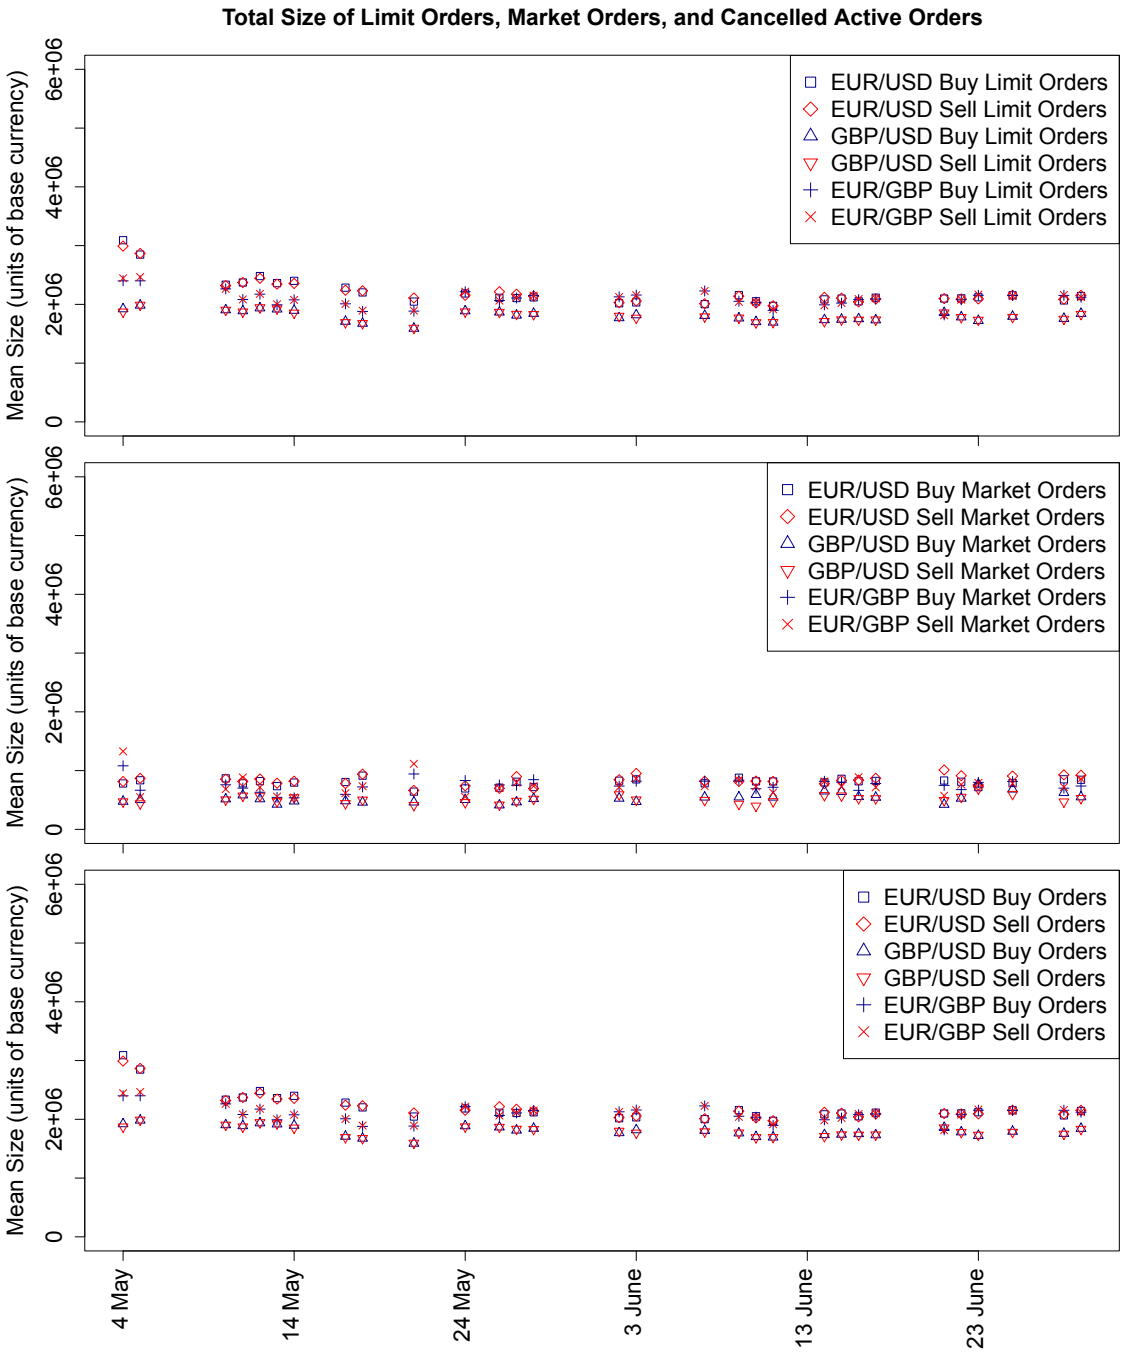

Figure 4.8: Mean size of (top) limit orders, (middle) market orders, and (bottom) cancelled active orders each day for EUR/USD, GBP/USD, and EUR/GBP.

implies that the total size of buy market orders on a given day often substantially exceeds that of sell market orders (or vice-versa).

This imbalance between buy and sell market orders is important in the context of price formation. Most existing LOB models assume that the buy-side and sell-side rate parameters are equal [47, 60, 78, 167, 192], but there can exist strong imbalances between the arrivals of buy and sell market orders on a single trading day. Our results suggest that LOB models of price formation on short time scales should be sufficiently flexible to incorporate short-term asymmetries in order flow.

4.5.2 Daily ECDFs

In this section, we compare the ECDFs of limit order arrivals and cancellations each day. Figure 4.9 shows the ECDFs for EUR/USD limit order arrivals. The corresponding plots for cancellations and for the other currency pairs are qualitatively similar. In both quote-relative and trade-relative coordinates, the ECDFs vary substantially across the trading days. On some days, most market activity occurs with a quote-relative (respectively, trade-relative) price that is close to 0; on other days, market activity is spread over a much wider range of quote-relative (respectively, trade-relative) prices.

Figure 4.10 shows the same ECDFs for EUR/USD limit order arrivals, after first rescaling each day's data by subtracting the day's sample mean and dividing by the day's sample standard deviation. In both quote-relative and trade-relative coordinates, the daily variation between the rescaled ECDFs is much smaller than that between the ECDFs for the raw data. We quantify the strength of this effect in the next section.

4.6 Model Comparison

In this section, we quantify the effectiveness of the rescalings described in Section 4.5.2 for reducing the differences between the order-flow ECDFs on different trading days. Specifically, we choose a single trading day i and calculate the ECDF F_i of the relevant order-flow property (e.g., EUR/USD buy limit order arrivals in quote-relative coordinates) on this day. We then compare F_i to the outputs of two statistical models. For the first model, we aggregate the data from the other 29 days to construct a data set Z^1_{-i} (where the subscript $-i$ denotes that we exclude day i), then calculate the ECDF F^1_{-i} of Z^1_{-i}. For the second model, we rescale the data from the other 29 days to account for the daily differences in the first two sample moments, aggregate this

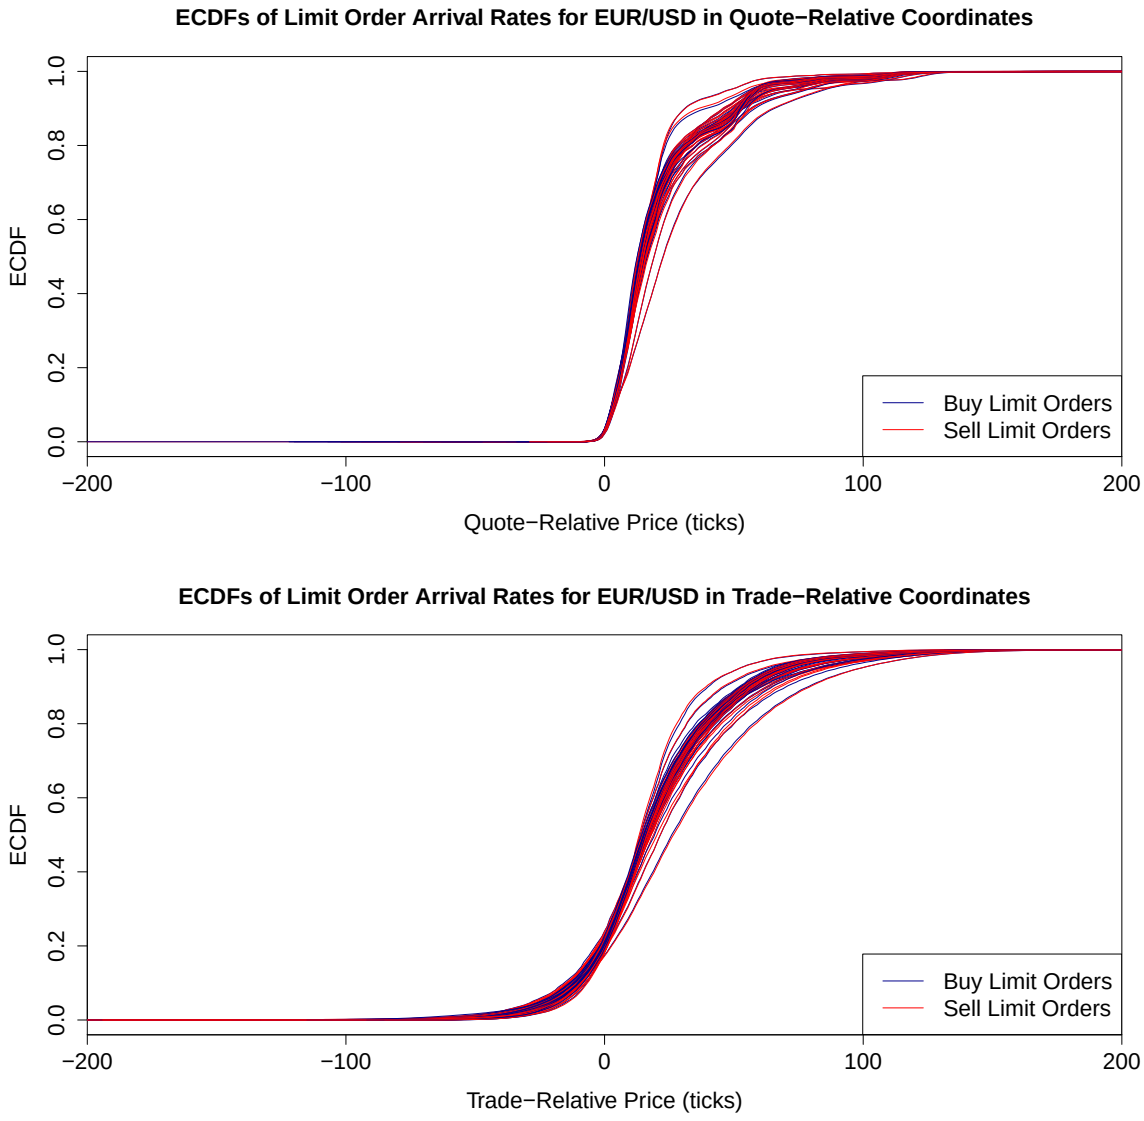

Figure 4.9: ECDFs for EUR/USD limit order arrivals each day in (top) quote-relative and (bottom) trade-relative coordinates.

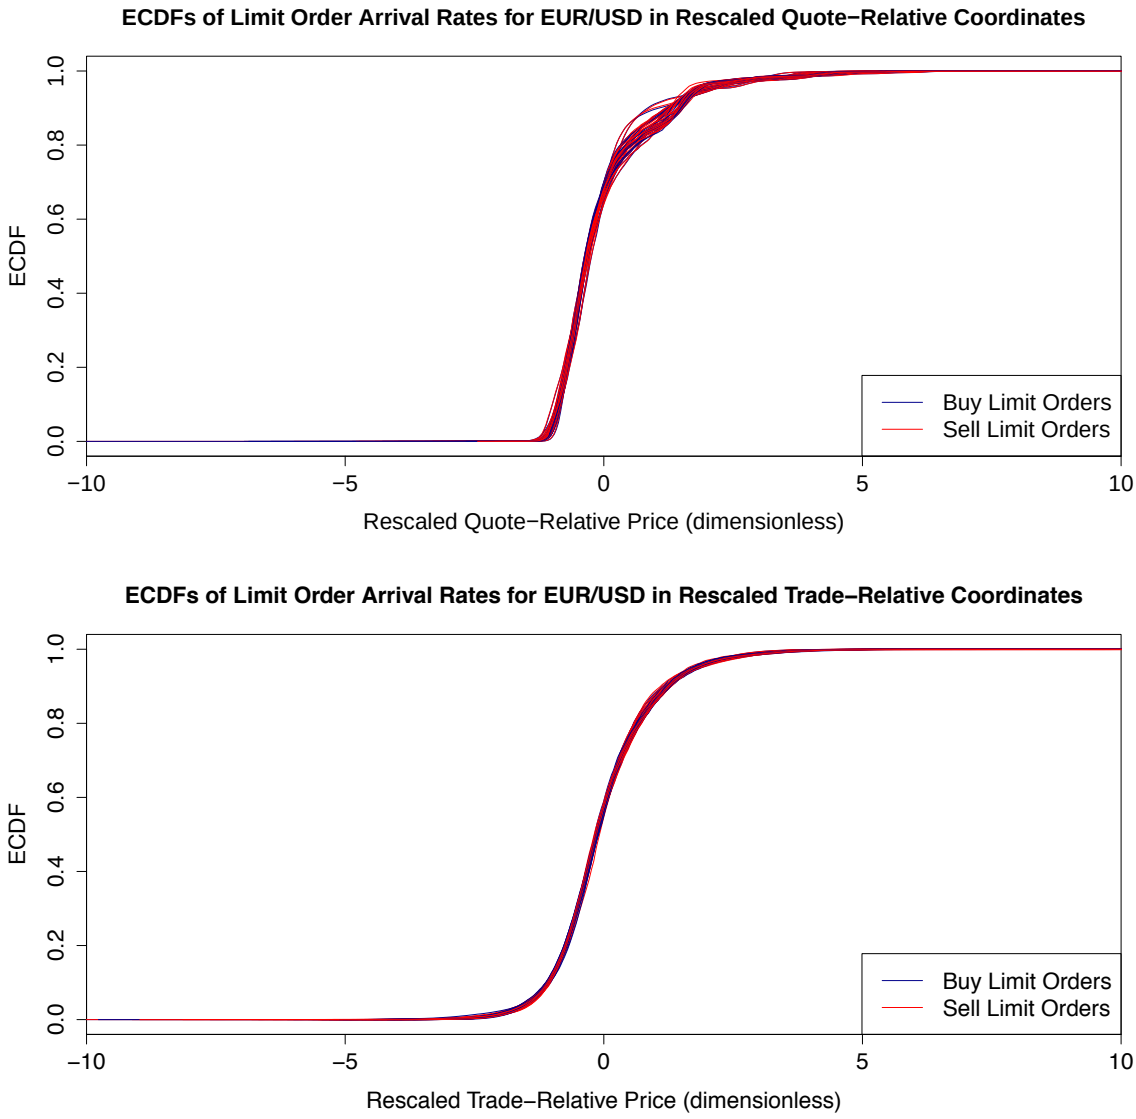

Figure 4.10: ECDFs for EUR/USD limit order arrivals each day in (top) quote-relative and (bottom) trade-relative coordinates after first rescaling each day's data by subtracting the day's mean and dividing by the day's standard deviation.

rescaled data to construct a data set Z^2_{-i}, then invert the rescaling (by multiplying the data in Z^2_{-i} by day i's sample standard deviation and adding day i's sample mean) and calculate the ECDF F^2_{-i} of the resulting data.

We compare the performance of each model by calculating the *Kolmogorov–Smirnov (KS)* distances [191]

$$K_i^j = \sup_p \left| F^j_{-i}(p) - F_i(p) \right|, \qquad i = 1, 2, \ldots, 30, \ j = 1, 2 \tag{4.8}$$

and the *Cramér–von Mises (CvM)* distances [61]

$$C_i^j = \int_p \left(F^j_{-i}(p) - F_i(p) \right)^2 \mathrm{d}k, \qquad i = 1, 2, \ldots, 30, \ j = 1, 2 \tag{4.9}$$

between the F^j_{-i}s and the true ECDFs.[4] Our model comparisons are based on the mean distances

$$K^j = \frac{1}{30} \sum_{i=1}^{30} K_i^j, \qquad C^j = \frac{1}{30} \sum_{i=1}^{30} C_i^j, \qquad j = 1, 2. \tag{4.10}$$

Smaller values of K^j and C^j correspond to smaller distances between the daily order-flow ECDFs.

Table 4.3 lists the values of K^1, K^2, C^1, and C^2 for each of the three currency pairs. In all cases, the values of K^2 and C^2 are smaller than the corresponding values of K^1 and C^1. This implies that rescaling each day's data to account for differences in the first two moments reduces the KS and CvM distances between the order-flow ECDFs on different days.

For EUR/USD and GBP/USD, all values of K^2 and C^2 for trade-relative coordinates are smaller than the corresponding values for quote-relative coordinates. For EUR/GBP, there are some instances for which the converse of this statement is true. We provide two possible reasons for why this could be the case. First, trades for EUR/GBP occur much less frequently for EUR/GBP than they do for EUR/USD and GBP/USD (see Table 4.1). Therefore, updates to the $B(t)$ and $A(t)$ series occur less frequently for EUR/GBP, which could cause traders to regard these series as less informative than the corresponding series for EUR/USD and GBP/USD. Second, heterogeneity in traders' access to trading opportunities — and, in particular, differences between different traders' local bid-prices $b_i(t)$ and local ask-prices $a_i(t)$

[4]The KS distance measures the largest difference between the two ECDFs. The CvM distance measures the squared L_2 norm between the two ECDFs. There are many other possible distance measures [67], but we choose these because they are widely used, easy to interpret, and fast to compute.

Distance	Order Flow	Coordinates	EUR/USD		GBP/USD		EUR/GBP	
K^1	Limit Orders	Quote Relative	0.055	0.052	0.039	0.038	0.073	0.072
K^1	Limit Orders	Trade Relative	0.038	0.040	0.040	0.044	0.062	0.059
K^1	Cancellations	Quote Relative	0.056	0.052	0.040	0.038	0.075	0.075
K^1	Cancellations	Trade Relative	0.039	0.041	0.040	0.044	0.063	0.059
C^1	Limit Orders	Quote Relative	0.099	0.096	0.091	0.089	0.107	0.108
C^1	Limit Orders	Trade Relative	0.106	0.115	0.160	0.173	0.270	0.273
C^1	Cancellations	Quote Relative	0.100	0.097	0.091	0.089	0.109	0.110
C^1	Cancellations	Trade Relative	0.106	0.114	0.160	0.173	0.269	0.273

Distance	Order Flow	Coordinates	EUR/USD		GBP/USD		EUR/GBP	
K^2	Limit Orders	Quote Relative	0.033	0.033	0.029	0.028	0.050	0.044
K^2	Limit Orders	Trade Relative	0.016	0.017	0.016	0.017	0.027	0.032
K^2	Cancellations	Quote Relative	0.036	0.035	0.030	0.030	0.053	0.045
K^2	Cancellations	Trade Relative	0.016	0.017	0.016	0.017	0.027	0.032
C^2	Limit Orders	Quote Relative	0.024	0.023	0.024	0.023	0.035	0.027
C^2	Limit Orders	Trade Relative	0.006	0.007	0.012	0.017	0.042	0.044
C^2	Cancellations	Quote Relative	0.026	0.024	0.025	0.024	0.038	0.029
C^2	Cancellations	Trade Relative	0.006	0.007	0.012	0.018	0.042	0.045

Table 4.3: Mean Kolmogorov–Smirnov (KS) and Cramér-von Mises (CvM) distances for limit order arrivals and cancellations in quote-relative and trade-relative coordinates for the (top) first and (bottom) second statistical models described in Section 4.6.

— could vary across the three currency pairs. If traders consider the values of $b_i(t)$ and $a_i(t)$ when deciding how to act, and if the variation across $b_i(t)$ and $a_i(t)$ is smaller for EUR/GBP than it is for EUR/USD and GBP/USD, then measuring the distributions in quote-relative coordinates could produce stronger curve collapse for EUR/GBP than it does for the other two currency pairs.

To illustrate how closely the output of our second model fits the data, we construct quantile-quantile (Q-Q) plots of F^2_{-i} against F^i (see Figure 4.11 for EUR/USD limit order arrivals on 29 June 2010; results for the other currency pairs, for cancellations, and for other days are qualitatively similar). Although random fluctuations prevent perfect curve collapse, our results demonstrate that the first two moments of the order-flow distributions account for most of the daily variation in the ECDFs, and therefore suggests that these two moments provide significant explanatory power for daily order flow.

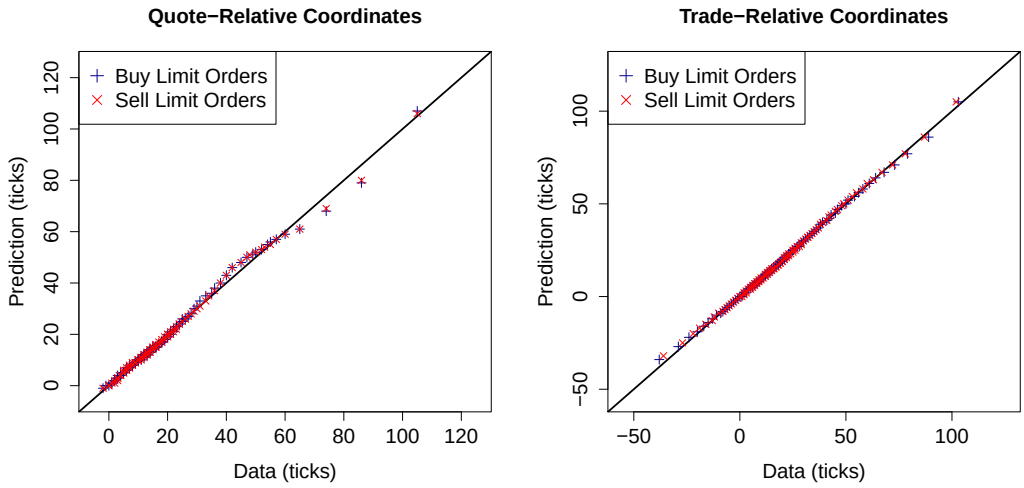

Figure 4.11: Quantile-quantile (Q-Q) plots for F^2_{-i} against F_{-i} for EUR/USD limit order arrivals in (left) quote-relative and (right) trade-relative coordinates. The points denote the $0.01, 0.02, \ldots, 0.99$ quantiles of distributions. The solid black lines denote the diagonal.

4.6.1 Aggregate Rescaled Data

In order to compare aggregate market market activity for each of the three currency pairs, we also calculate the ECDFs of limit order arrivals and cancellations across all 30 days in the sample, after first separately rescaling each day's data by subtracting the day's sample mean and dividing by the day's sample standard deviation (see Figure 4.12 for limit order arrivals; results for cancellations are qualitatively similar).

The results reveal that the tails of the rescaled distributions of quote-relative prices are heavier for EUR/GBP than for EUR/USD and GBP, but that otherwise the distributions in both quote-relative and trade-relative coordinates are extremely similar for each of the three currency pairs.

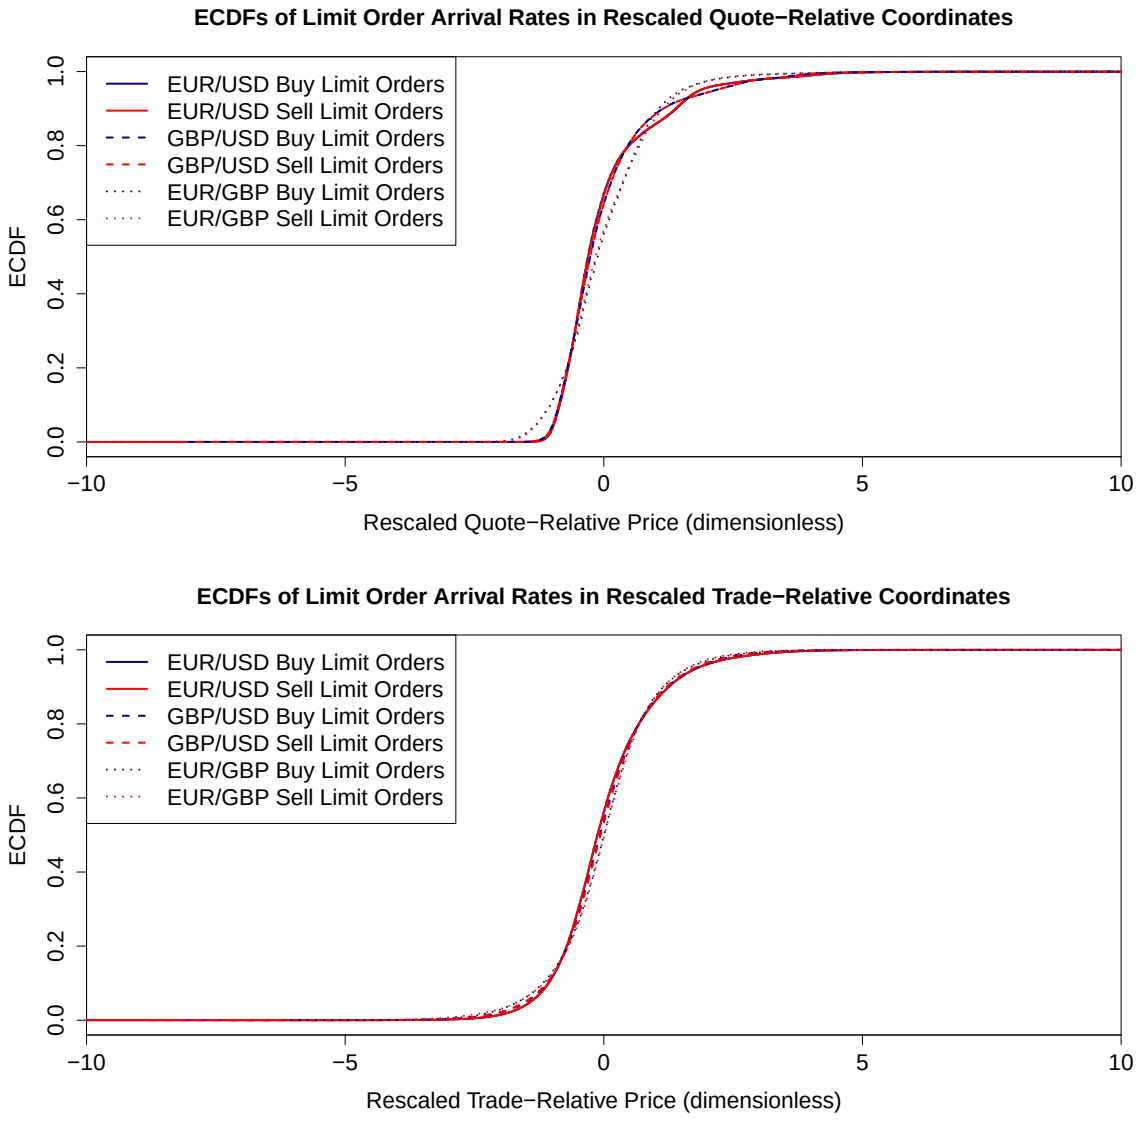

Figure 4.12: ECDFs for limit order arrivals for all 30 days in (top) quote-relative and (bottom) trade-relative coordinates after first rescaling each day's data by subtracting the day's mean and dividing by the day's standard deviation.

4.7 Discussion

In this section, we discuss several of our empirical results in more detail and suggest possible explanations for our observations.

4.7.1 Market Orders

A particularly surprising feature of the Hotspot FX data is the extremely low percentage of market orders in the arriving order flow. This finding raises the question of why traders submit so many limit orders, given that so few of them result in trades. We propose three possible explanations. First, some traders might submit and immediately cancel active orders to search for hidden liquidity. Second, some traders might place limit orders with the same price on several different trading platforms simultaneously, to increase their chance of receiving a matching. If one of their orders match, such traders can immediately cancel all duplicate orders on other platforms at no cost. Third, some traders might submit and cancel limit orders to compete for priority with other traders. For example, so-called "high-frequency traders" [22, 44, 115, 126] use electronic trading algorithms that monitor order flow and LOB state and rapidly submit and cancel large numbers of orders near the best quotes. Given the surge in popularity of algorithmic trading towards the end of the last decade [12], we believe that the third hypothesis provides the most plausible explanation for the behaviour that we observe on Hotspot FX.

4.7.2 Time Scales

Many existing LOB models assume that order flow is governed by stochastic processes with fixed rate parameters [47, 60, 78, 167, 192, 203]. To date, the empirical verification of such models has typically consisted of using large data sets to estimate their parameters and then comparing the long-run statistical properties of their output to those of the data. Our results demonstrate that LOB statistics on short time horizons (such as a single trading day) can differ substantially from their long-time-horizon averages. Even though several models have been reported to mimic certain long-time-horizon statistical averages from real markets, our results demonstrate that their output over shorter time horizons could be misleading.

4.7.3 Round Numbers

Similarly to several other empirical studies of other markets [47, 108, 170, 206], our results reveal several examples in which traders demonstrate a preference for round numbers, such as trade-relative prices that are integer multiples of 5 or 10 ticks and order sizes of 1 million units of the base currency. We propose three possible reasons to explain why traders might behave in this way. First, the tick sizes and lot sizes on many other electronic trading platforms in the FX spot market are larger than those on Hotspot FX. If traders who access several different trading platforms choose to use the same order sizes (respectively, prices) on all platforms, then these sizes (respectively, prices) will be integer multiples of the lot size (respectively, tick size) on Hotspot FX. Second, many older trading mechanisms that predate the widespread uptake of electronic trading used larger tick sizes and lot sizes than those on Hotspot FX, so traders may still regard these numbers as important quantities. Third, traders may simply find the sizes or prices that they use to be appealing or convenient.

4.7.4 Trade-Relative Coordinates

Trade-relative coordinates provide two important benefits when studying the Hotspot FX data: they illuminate several interesting properties of order flow and LOB state that are not apparent in quote-relative coordinates (such as the shape of the cancellation profiles in Figure 4.6) and they provide a simple and efficient way to quantify daily differences between the order-flow distributions using just two scalar parameters (see Section 4.6). Although our use of trade-relative coordinates was motivated by the quasi-centralized structure of Hotspot FX, we conjecture that this coordinate frame could also provide useful insight in centralized LOBs. In many markets, the rise in popularity of electronic trading algorithms has led to a sharp increase in the frequency of order arrivals in recent years [22, 44, 57, 115, 126]. Due to the high levels of market activity near the best quotes, the values of $b(t)$ and $a(t)$ — and therefore the quote-relative prices of all orders — often fluctuate rapidly. By contrast, $B(t)$ and $A(t)$ change only when a trade occurs, so measuring market activity relative to these trade-price series avoids the excessive update frequency of quote-relative coordinates.

4.8 Summary and Future Work

In this chapter, we have uncovered several robust statistical properties of order flow and LOB state in a recent, high-quality data set from Hotspot FX, which is a large

multi-institution trading platform in the FX spot market. We have highlighted several important ways in which the quasi-centralized structure of the platform causes our results to differ from those reported for other LOBs, but we have also discussed how other factors — such as the high levels of liquidity and the recent date range of our data — could also play an important role in our results. Given that the substantial literature on LOBs is based on a relatively small set of empirical studies of centralized LOBs [98], we hope that these findings underline the urgent need for detailed investigations of other recent, high-quality data sets.

4.8.1 Future Work

Our results suggest many possible avenues for future research.

- **Investigating different time periods:** Several publications have identified intra-day seasonalities in market behaviour (see, e.g., [89, 102, 122, 143, 162]) that could cause order flow to vary within a trading day. The extremely high levels of market activity on Hotspot FX would enable detailed study of these effects. It would be interesting to see whether the in-sample and out-of-sample performance of the statistical models in Section 4.6 improved when using shorter time windows, such as an hour, rather than a whole trading day.

- **Predicting scaling parameters:** The statistical models that we describe in Section 4.6 require prior knowledge of the sample mean and standard deviation to rescale each trading day's data. Therefore, the methods that we describe do not provide a way to predict future order flow. A model that predicts the values of the first two moments of the order-flow distributions, given a recent history, would facilitate such predictions.

- **Nonlinear rescalings:** The rescalings that we used throughout this chapter are simple linear transformations based on the first two sample moments of the data. It would be interesting to formulate and solve a nonlinear optimization problem to identify the rescalings that cause the optimal curve collapse.

- **Round-number preferences:** We have identified several different examples of traders' preferences for round numbers. In our second model in Section 4.6, the process of rescaling and aggregating order-flow data from different trading days has the undesired consequence of smoothing the peaks at round-number relative prices. Given how strongly such peaks appear in the aggregate data (see Figure 4.3), it seems likely that incorporating an additional step to redistribute

the density and restore the peaks would improve the performance of our second model. One possible approach would be to perturb our models' output ECDFs to slightly increase the probabilities of events at round-number trade-relative prices and slightly decrease the probabilities of other events.

- **Analyzing other FX spot platforms:** Both the Reuters Dealing 3000 [200] and EBS [72] platforms operate quasi-centralized LOBs with very similar trading rules to Hotspot FX, but they also disseminate the values of $b(t)$ and $a(t)$ to all traders in real time [43]. Empirical study of data from these platforms would help address the question of whether providing traders with this additional information causes detectable differences in the order-flow distributions.

- **Trade-relative analysis of centralized LOBs:** To our knowledge, all existing empirical studies of order flow and LOB state in centralized LOBs have used quote-relative coordinates to perform their analysis. It would be interesting to repeat these calculations in trade-relative coordinates to compare the results with our findings for Hotspot FX.

- **Short-horizon LOB modelling:** Our results suggest that the statistical properties of order flow and LOB state on a single trading day can differ substantially from those of aggregate data. Given suitable parameter estimates from a single trading day, it would be interesting to study the performance of several existing LOB models that were originally designed to mimic market behaviour over longer periods.

Chapter 5

Long Memory on Hotspot FX

In this chapter, we study the long-range autocorrelation properties of returns, absolute returns, and order flow on Hotspot FX. The work in this chapter is also described in the jointly authored paper [99].

5.1 Introduction

The long-memory properties of financial markets have been the subject of statistical debate for several decades [56]. During this time, empirical studies have reported long memory in a wide range of financial time series [29, 30, 34, 45, 59, 106, 136, 162], but many authors have argued that these results are inconclusive for reasons such as inadequacy of statistical techniques [141], inappropriateness of modelling assumptions [7, 168], and lack of suitable data [103].

Understanding the long-memory properties of financial markets is an important task for several reasons. First, the present values of a long-memory process are correlated with values in the distant past [18]. Therefore, quantifying the strength of long memory is useful for forecasting. Second, understanding the various interacting types of long memory in financial markets helps to illuminate the complex interplay between order flow, liquidity, and price formation [33, 34, 136, 204]. Third, identifying the sources of long memory in order flow may provide insight into traders' strategic decision-making processes [34, 137]. Fourth, several models suggest that long-range autocorrelations may hold the key to understanding some of the so-called "stylized facts" of financial markets, such as volatility clustering and the heavy-tailed distribution of returns [55, 167, 203].

In this chapter, we employ a wide range of estimators and techniques to perform a formal and scientifically rigorous investigation of the the long-memory properties of returns, absolute returns, and order flow on Hotspot FX. Due to the extremely high

levels of activity on the platform, we are able to study the long-memory properties of intra-day time series without needing to aggregate data from different trading days or different currency pairs. Moreover, we are able to detect long memory with high levels of statistical significance, even in cases when the effect is relatively weak. Therefore, our results address several criticisms of previous empirical studies of long memory [7, 103, 141, 168]. We find no evidence for long memory in returns, but we find strong, statistically significant evidence for long memory in absolute returns and order flow. We also uncover a low-order, negative autocorrelation in order flow that (to our knowledge) has not been reported elsewhere in the literature.

To investigate the impact of aggregating data from different trading days, we also repeat our calculations using data that spans daily boundaries. We find that the ensuing results are very similar to those for intra-day data. We test and reject the hypothesis that the long-range autocorrelations in these series are an artifact caused by a structural break at the daily boundaries, in favour of the alternative hypothesis of true long memory. We conclude that on Hotspot FX, long memory is a robust statistical property of absolute returns and order flow that persists across daily boundaries.

The remainder of this chapter is organized as follows. In Section 5.2, we present a technical discussion of long memory. In Section 5.3, we discuss several statistical techniques for assessing the long-memory properties of an empirically observed time series. In Section 5.4, we review the empirical findings of several other studies of long memory in financial markets. In Section 5.5, we discuss our methods for analyzing the Hotspot FX data. We present our empirical results in Section 5.6 and discuss these results in Section 5.7. In Section 5.8, we summarize our findings and discuss several possible avenues for future research.

5.2 Long-Memory Processes

Let

$$\{W_t\} = W_1, W_2, \ldots \tag{5.1}$$

denote a real-valued, second-order stationary[1] time series with mean

$$\mathbb{E}(W_t) = \mu. \tag{5.2}$$

[1] A time series $\{W_t\}$ is *second-order stationary* if its first and second moments are finite and do not vary with time [48, 196].

Definition. *For a time series* $\{W_t\}$*, the* autocovariance function $\gamma : \mathbb{Z} \to \mathbb{R}$ *is given by*

$$\gamma(k) = \operatorname{cov}\left(W_t, W_{t+|k|}\right). \tag{5.3}$$

Definition. *For a time series* $\{W_t\}$*, the* autocorrelation function (ACF) $\rho : \mathbb{Z} \to [-1, 1]$ *is given by*

$$\rho(k) = \frac{\gamma(k)}{\gamma(0)}. \tag{5.4}$$

5.2.1 Short Memory

Definition. *A time series* $\{W_t\}$ *is said to exhibit* short memory *if*

$$\lim_{N\to\infty} \sum_{k=-N}^{N} |\rho(k)| < \infty. \tag{5.5}$$

For example, if $\{W_t\}$ is an invertible autoregressive moving average (ARMA) process [8], then there exist constants $z_1 > 0$ and $z_2 \in (0, 1)$ such that

$$\lim_{k\to\infty} |\rho(k)| \leq z_1 z_2^{-k}. \tag{5.6}$$

Therefore, by the convergence properties of geometric series,

$$\lim_{N\to\infty} \sum_{k=-N}^{N} |\rho(k)| < \infty, \tag{5.7}$$

so $\{W_t\}$ is a short-memory process.

5.2.2 Long Memory

Definition. *A time series* $\{W_t\}$ *is said to exhibit* long memory *if*

$$\lim_{N\to\infty} \sum_{k=-N}^{N} |\rho(k)| = \infty. \tag{5.8}$$

One way in which a time series can exhibit long memory is if there exists some $\alpha \in (0, 1)$ such that $\rho(k)$ decays asymptotically as a power of k:

$$\rho(k) \sim k^{-\alpha} L(k), \quad k \to \infty, \tag{5.9}$$

where L is a slowly varying function[2] [56, 136, 137]. Smaller values of α correspond to slower decay of the long-range autocorrelations in $\{W_t\}$ [33, 136].

[2] A function L is *slowly varying* if $\lim_{k\to\infty} \frac{L(zk)}{L(k)} = 1$ for all $z > 0$.

5.2.3 Spectral Density

The autocorrelation properties of a time series $\{W_t\}$ can also be characterized by the behaviour of its *spectral density* [18]

$$f(\lambda) = \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} \gamma(k) e^{-ik\lambda}. \tag{5.10}$$

If

$$f(\lambda) \to \text{constant as } \lambda \to 0, \tag{5.11}$$

then $\{W_t\}$ is a short-memory process. If

$$f(\lambda) \sim \lambda^{-\beta} \text{ as } \lambda \to 0 \tag{5.12}$$

for some constant $\beta \in (0, 1)$, then $\{W_t\}$ is a long-memory process. Larger values of β correspond to slower decay of the long-range autocorrelations in $\{W_t\}$ [18].

5.2.4 The Hurst Exponent

The Hurst exponent [18, 120, 157, 158, 159, 160] of a time series $\{W_t\}$ provides an alternative way to characterize its long-range autocorrelation properties.[3]

Definition. *Let*

$$\overline{W}_k = \frac{1}{k} \sum_{j=t+1}^{t+k} W_j. \tag{5.13}$$

The rescaled-range statistic *[159] is the ratio*

$$Q(t, k) = \frac{R(t, k)}{S(t, k)}, \tag{5.14}$$

where

$$R(t, k) = \max_{1 \le i \le k} \left[\sum_{j=t+1}^{t+i} \left(W_j - \overline{W}_k \right) \right] - \min_{1 \le i \le k} \left[\sum_{j=t+1}^{t+i} \left(W_j - \overline{W}_k \right) \right] \tag{5.15}$$

[3] Study of the Hurst exponent dates back to the 1950s, when Harold Edwin Hurst [120] considered how to build reservoirs to ensure steady flow from rivers whose natural behaviour is characterized by long periods of dryness followed by long periods of returning floods [18]. Hurst plotted the rescaled-range statistic (5.14) versus k for several different hydrological records (including, most famously, flow data from the river Nile) on doubly logarithmic axes. He noted that for large values of k, each hydrological record's data points were scattered around a straight line with slope $H > 1/2$, whereas similar plots for the number of successes from independent repetitions of a Bernoulli trial yielded data points that were scattered around a straight line with slope $H \approx 1/2$. The steeper slope of the plots for the hydrological records came to be known as the *Hurst effect* [18].

and

$$S^2(t,k) = \frac{1}{k}\sum_{j=t+1}^{t+k}\left(W_j - \overline{W}_k\right)^2. \tag{5.16}$$

The rescaled-range statistic Q measures the range of partial sums of deviations of the time series $\{W_1, W_2, \ldots\}$ from its mean, rescaled by an estimate of its standard deviation [141]. The following theorem by Mandelbrot [154] provides a formal relationship between a time series' long-range autocorrelations and its Hurst exponent H.

Theorem 5.1. *If $\{W_t\}$ is a second-order stationary process such that W_t^2 is ergodic and $t^{-H}\sum_{i=1}^{t} W_i$ converges weakly to a fractional Brownian motion[4] with parameter H as $t \to \infty$, then*

$$k^{-H}Q(t,k) \to_d \eta \text{ as } k \to \infty, \tag{5.17}$$

where η is a nondegenerate random variable and $\to_d$ denotes convergence in distribution.

The constant H is called the *Hurst exponent* of $\{W_t\}$. A time series with a Hurst exponent $H = 1/2$ is a short-memory process. For a long-memory process that satisfies the conditions of Theorem 5.1, H is related to α in Equation (5.9) by

$$H = 1 - \frac{\alpha}{2} \tag{5.18}$$

and to β in Equation (5.12) by

$$H = \frac{\beta + 1}{2}. \tag{5.19}$$

5.3 Empirical Assessment of Long Memory

In many empirical applications, it is common to observe only a single, finite-length realization $\{w_1, w_2, \ldots, w_N\}$ of $\{W_t\}$. If the statistical properties of $\{W_t\}$ are unknown, then estimating the long-memory properties of $\{W_t\}$ from $\{w_1, w_2, \ldots, w_N\}$ entails considerable challenges [17, 18, 161].

For an empirically observed time series $\{w_1, w_2, \ldots, w_N\}$, let

$$\overline{w} = \frac{1}{N}\sum_{i=1}^{N} w_i \tag{5.20}$$

[4] A *fractional Brownian motion* [156] is a Gaussian process $B_H(t)$ with zero drift that satisfies $B_H(0) = 0$ and $\mathbb{E}\left[B_H(t)B_H(s)\right] = \frac{1}{2}\left(|t|^{2H} + |s|^{2H} - |t-s|^{2H}\right)$ for some $H \in (0,1)$.

denote the *sample mean* of $\{w_1, w_2, \ldots, w_N\}$ and let

$$K = \{-(N-1), -(N-2), \ldots, N-2, N-1\}. \tag{5.21}$$

Definition. *For an empirically observed time series* $\{w_1, w_2, \ldots, w_N\}$*, the* sample autocovariance function $\hat{\gamma} : K \to \mathbb{R}$ *is given by*

$$\hat{\gamma}(k) = \frac{1}{N} \sum_{i=1}^{N-|k|} \left(w_{i+|k|} - \overline{w}\right)(w_i - \overline{w}), \ k \in K. \tag{5.22}$$

Definition. *For an empirically observed time series* $\{w_1, w_2, \ldots, w_N\}$*, the* sample ACF $\hat{\rho} : K \to [-1, 1]$ *is given by*

$$\hat{\rho}(k) = \frac{\hat{\gamma}(k)}{\hat{\gamma}(0)}, \ k \in K. \tag{5.23}$$

It is very difficult to estimate the large-k decay of $\rho(k)$ from $\hat{\rho}$, so direct estimation of the long-memory properties of $\{W_t\}$ from $\hat{\gamma}$ often produces very poor results [136]. Instead, most empirical studies employ heuristic methods for this task. The performance of such techniques on empirically observed series varies considerably, so it is common for empirical studies to evaluate the output of several heuristic methods rather than relying on a single estimator [194]. In this section, we describe a selection of such tests, which we use throughout the remainder of the chapter. For a more detailed discussion of these and other techniques, see [18].

5.3.1 Rescaled-Range Plots

For a given *block number* $B \in \mathbb{N}$, let

$$G(k) = \left\{ t = \frac{N(i-1)}{B} + 1 \,\middle|\, i = 1, \ldots, B; \ t + k \leq N \right\} \tag{5.24}$$

and let $|G(k)|$ denote the number of elements in $G(k)$.

Definition. *A* rescaled-range plot *(also known as a* pox plot*) [157, 159, 198] is a plot of*

$$\overline{R}(k) = \frac{1}{|G(k)|} \sum_{t \in G(k)} Q(t, k) \tag{5.25}$$

versus k *on doubly logarithmic axes.*

The slope of a rescaled-range plot for large values of k provides a rough estimate for H [159].

5.3.2 Lo's Modified Rescaled-Range Statistic

Lo [141] noted that if $\{W_t\}$ is subject to short-range autocorrelations, then the denominator $S(t,k)$ of the rescaled-range statistic $Q(t,k)$ is not a consistent estimator for the standard deviation of $\{W_t\}$. Therefore, an important problem with using the rescaled-range statistic to assess the long-memory properties of an empirically observed time series $\{w_1, w_2, \ldots, w_N\}$ is that the finite-sample properties of Q are not invariant to short-range dependence. Lo proposed replacing the denominator of Q with the so-called *Newey-West estimator* [173],[5] which discounts short-range dependence in $\{W_t\}$ up to a specified lag $q < N$.

Definition. *For an empirically observed time series* $\{w_1, w_2, \ldots, w_N\}$*,* Lo's modified rescaled-range statistic *[141] is*

$$\tilde{Q}(q) = \frac{R(1,N)}{\hat{\sigma}(q)}, \tag{5.26}$$

where

$$\hat{\sigma}^2(q) = \begin{cases} S^2(1,N), & \text{for } q = 0, \\ S^2(1,N) + 2\sum_{i=1}^{q}\left(1 - \frac{i}{q+1}\right)\hat{\gamma}(i), & \text{otherwise.} \end{cases} \tag{5.27}$$

5.3.2.1 Lo's Modified Rescaled-Range Test

The statistic

$$V(q) = \frac{\tilde{Q}(q)}{\sqrt{N}} \tag{5.28}$$

can be used as a test statistic for the hypothesis test

$$H_0 : \{W_t\} \text{ is a short-memory process,}$$
$$H_1 : \{W_t\} \text{ is a long-memory process.}$$

This hypothesis test is called *Lo's modified rescaled-range test* [141]. In the limit $N \to \infty$, the asymptotic critical region for the test at the 5% significance level is $[0.809, 1.862]$.

Teverovsky *et al.* [198] remarked that although Lo's modified rescaled-range statistic is a significant improvement over the original rescaled-range statistic, Lo's test can fail to reject H_0 for time series with weak long memory. Moreover, they noted that both the size and the power of the test depend on q.

[5] Newey and West [173] showed that for a suitable choice of q (see Section 5.3.2.2), $\hat{\sigma}(q)$ is a consistent estimator for the standard deviation of $\{W_t\}$, even if $\{W_t\}$ is subject to short-range autocorrelations.

5.3.2.2 Choosing q

The optimal choice of q for $\tilde{Q}(q)$ and $V(q)$ depends on the behaviour of the spectral density f of $\{W_t\}$ [5]. If f is unknown (as is often the case for empirically observed series), then there exists no universal rule via which to choose q. In empirical applications, it is therefore customary to calculate $\tilde{Q}(q)$ and $V(q)$ using several different choices of q and/or to calculate a so-called *plug-in estimator* $\hat{q}$ [5, 7, 141] by assuming that f is equal to the spectral density of a specified parametric process. Andrews [5] derived plug-in estimators for several different parametric processes, including autoregressive, moving-average, and ARMA models.

Throughout this chapter, we use the plug-in estimator for an AR(1) process

$$W_t = \phi W_{t-1} + \varepsilon_t, \tag{5.29}$$

where $\phi \in \mathbb{R}$ is the autocorrelation parameter and ε_t is uncorrelated Gaussian noise. This estimator is given by [5, 141]

$$\hat{q} = \left\lfloor \left(\frac{3N}{2}\right)^{1/3} \left(\frac{2\hat{\phi}}{1-\hat{\phi}^2}\right)^{2/3} \right\rfloor, \tag{5.30}$$

where $\hat{\phi}$ is the maximum likelihood estimate of ϕ given $\{w_1, w_2, \ldots, w_N\}$ and $\lfloor x \rfloor$ denotes the greatest integer less than or equal to x.

5.3.3 Detrended Fluctuation Analysis

Detrended fluctuation analysis (DFA) [178] is a technique for estimating the Hurst exponent from an empirically observed series $\{w_1, w_2, \ldots, w_N\}$. The technique is performed as follows. Let

$$w_i^* = \sum_{j=1}^{i} w_i, \ i = 1, 2, \ldots, N. \tag{5.31}$$

For a given window length $m \in \mathbb{N}$ such that $m \leq N$, divide $\{w_1^*, w_2^*, \ldots, w_N^*\}$ into non-overlapping windows of length m. For each $j = 1, 2, \ldots, N/m$, label the data points in window j as $y_{1,j}, y_{2,j}, \ldots, y_{m,j}$, perform an ordinary least-squares regression to fit a straight line to the m data points in the window, and let $\hat{y}_{i,j}$ denote the value of the regression line at the point $y_{i,j}$ for each $i = 1, 2, \ldots, m$. For each window, calculate the *detrended standard deviation*

$$F_j(m) = \sqrt{\frac{1}{m} \sum_{i=1}^{m} (y_{i,j} - \hat{y}_{i,j})^2}, \tag{5.32}$$

and then calculate the *length-m mean detrended standard deviation*

$$F(m) = \frac{1}{k/m} \sum_{j=1}^{k/m} F_j(m). \tag{5.33}$$

Repeat this process for several logarithmically spaced choices of window length m, and plot $F(m)$ versus m using doubly logarithmic axes. Identify the smallest value $m_{\min}$ such that the plot is approximately straight for all $m \geq m_{\min}$. The DFA estimate of H is given by the slope of the best-fit line for $m \geq m_{\min}$.

5.3.4 Log-Periodogram Regression

Recall from Section 5.2.3 that the long-memory properties of a time series $\{W_t\}$ can be characterized by the behaviour of its spectral density function $f(\lambda)$ as $\lambda \to 0$. Therefore, estimating the behaviour of $f(\lambda)$ close to 0 provides an alternative approach to estimating the Hurst exponent H. Let

$$\lambda_{k,N} = \frac{2\pi k}{N}, \qquad k = 1, 2, \ldots, \left\lfloor \frac{N-1}{2} \right\rfloor \tag{5.34}$$

denote the *Fourier frequencies* of $\{w_1, w_2, \ldots, w_N\}$, and let

$$I(\lambda_{k,N}) = \frac{1}{2\pi N} \left| \sum_{t=1}^{N} (w_t - \overline{w})\, e^{-it\lambda_{k,N}} \right|^2 \tag{5.35}$$

denote the corresponding *periodogram*. The slope of an ordinary least-squares regression of $\log\left(I(\lambda_{k,N})\right)$ onto $\log\left(\lambda_{k,N}\right)$ for small $\lambda_{k,N}$ is an estimator for $-\beta$, and therefore an estimator for H [18, 136, 194].

Despite the attractiveness of its computational simplicity, log-periodogram regression suffers from a substantial practical drawback [18]: there is no universal rule for choosing the number of Fourier frequencies c with which to perform the regression. The scaling in Equations (5.11) and (5.12) only holds for $\lambda \to 0$, so choosing c too large leads to large bias but choosing c too small leads to high variance. Robinson [186] derived an expression for the optimal choice of c to minimize the mean squared error of the estimated cumulative spectral distribution function, but the optimal choice turns out to depend on the unknown value of H. Therefore, Robinson's expression does not provide a method for choosing c for an empirically observed time series. Instead, most empirical studies use one of two rules of thumb for this choice, namely $c = 0.1 \times (N/2)$ [194] or $c = \sqrt{N}$ [91]. Despite their widespread use, it is important to note that neither of these rules are based on rigorous derivation or optimization.

5.3.5 Nonstationarities

Another difficulty with assessing the long-memory properties of empirically observed time series is that the heuristic methods that we have discussed in this section can produce similar output for a series with nonstationarities (e.g., a monotonic trend [20] or change in mean [92, 105]) as they do for a stationary series with long memory [183, 194, 205]. Several authors have thereby argued that the apparent long memory reported by other empirical studies is actually an artifact caused by nonstationarities in the underlying series [7, 19, 139, 168]. Assessing whether or not this is the case is important for two reasons. First, observing a long history of a long-memory process significantly improves forecasts [18], whereas including too many previous values can instead *harm* forecasts for nonstationary series [19]. Second, understanding the structure of a time series can help to illuminate the data-generating mechanism. For example, estimating the times at which a series undergoes a change in mean can help to identify important events in the underlying process [105].

Disentangling the statistical properties of a time series with long memory and a time series with nonstationarities is a difficult task, and the choice of whether to model such time series using a long-memory model or a nonstationary model often depends on the desired application. Long-memory models are appealing for several reasons: they are parsimonious and straightforward to simulate, and there are many techniques that require only mild assumptions to estimate their parameters from data [18]. Nonstationary models can illuminate important features of an underlying series (such as the locations and frequency of structural breaks) that are not addressed by long-memory models, but they typically require the inclusion of either a large number of parameters (which can lead to over-fitting) or latent parameters (which can be difficult to estimate from data).

Many standard tests for nonstationarities in an empirically observed series have low power in the presence of long memory [68, 113], and there is no universal test that is able to determine whether the apparent long-memory properties of an empirically observed time series are an artifact caused by some unknown form of nonstationarity [105]. However, Berkes *et al.* [19] developed a hypothesis test that can help distinguish between long memory and a specific type of nonstationarity, which they call a *structural break*.[6]

[6]We use Berkes' Change-Point Test to examine whether apparent long memory in the Hotspot FX data is caused by a structural break between consecutive pairs of trading days, and we therefore only test for a single change point. A similar framework can be used to test for any bounded number of structural breaks [19].

Definition. *The time series* $\{Z_t\}$ *is a* short-range dependent series with a structural break at time k^* *if there exists a real-valued, second-order stationary, short-memory process* $\{\zeta_t\}$ *and a constant* $\mu^* \neq 0$ *such that*

$$Z_t = \begin{cases} \zeta_t, & t \leq k^*, \\ \mu^* + \zeta_t, & t > k^*. \end{cases} \tag{5.36}$$

Given a finite-length empirical observation $\{z_1, z_2, \ldots, z_N\}$ of a time series $\{Z_t\}$, *Berkes' change-point test* is a hypothesis test

$$H_0 : \{Z_t\} \text{ is a short-memory process with a single structural break,}$$
$$H_1 : \{Z_t\} \text{ is a long-memory process.}$$

The test makes use of the so-called *cumulative-sum change-point estimator* [19]

$$\hat{k}^* = \min \left\{ k : \max_{1 \leq j \leq N} \left| \sum_{i=1}^{j} z_i - \frac{j}{N} \sum_{i=1}^{N} z_i \right| = \left| \sum_{i=1}^{k} z_i - \frac{k}{N} \sum_{i=1}^{N} z_i \right| \right\}. \tag{5.37}$$

For a given $q < N$, let

$$T^{(1)} = \frac{1}{\hat{\sigma}^2_{(1)}(q)\sqrt{\hat{k}^*}} \max_{1 \leq i \leq \hat{k}^*} \left| \sum_{j=1}^{i} z_j - \frac{i}{\hat{k}^*} \sum_{j=1}^{\hat{k}^*} z_j \right|,$$
$$T^{(2)} = \frac{1}{\hat{\sigma}^2_{(2)}(q)\sqrt{N - \hat{k}^*}} \max_{\hat{k}^* \leq i \leq N} \left| \sum_{j=\hat{k}^*}^{N} z_j - \frac{i - \hat{k}^*}{N - \hat{k}^*} \sum_{j=\hat{k}^*}^{N} z_j \right|,$$

where $\hat{\sigma}_{(1)}(q)$ and $\hat{\sigma}_{(2)}(q)$ are the values of $\hat{\sigma}(q)$ from Equation (5.27) for the $\{z_1, z_2, \ldots, z_{\hat{k}^*}\}$ and $\left\{z_{\hat{k}^*+1}, z_{\hat{k}^*+2}, \ldots, z_N\right\}$ series, respectively. The test statistic for Berkes' test is

$$M = \max\left(T^{(1)}, T^{(2)}\right). \tag{5.38}$$

In the limit $N \to \infty$, the asymptotic critical value of M at the 5% significance level is 1.48 [19].

5.4 Long Memory in Financial Markets

5.4.1 Returns

The idea that financial time series could exhibit long-range autocorrelations dates back to at least the 1970s, when Mandelbrot [153] proposed that returns followed a long-memory process.

Definition. *Given a price series* $\{P_1, P_2, \ldots\}$*, the* return series *is* $\{R_1, R_2, \ldots\}$*, where*

$$R_i = \ln\left(\frac{P_{i+1}}{P_i}\right), \quad i = 1, 2, \ldots. \tag{5.39}$$

Several early empirical studies [29, 30, 106, 155] used the rescaled-range statistic (5.14) to investigate this conjecture on time series of daily returns (i.e., where P_i denotes the closing price of a given asset on day i) and concluded in favour of Mandelbrot's hypothesis in a wide range of different markets.

Several years later, Lo [141] showed that the rescaled-range statistic used by these studies is not robust to short-range dependence and thereby cast doubt upon their findings. Lo proposed a new statistic (see Section 5.3.2), which he claimed could overcome these problems, and applied it to daily returns for US stocks. Based on his calculations, Lo rejected the hypothesis of long memory. However, Teverovsky *et al.* [198] later showed that Lo's test is too severe by giving examples of synthetic time series with weak long memory for which Lo's method does not reject the hypothesis of short-range dependence. They thereby argued that Lo's results are inconclusive.

More recently, the availability of high-frequency data has led to interesting insight into the long-memory properties of returns on shorter time scales [162]. Many empirical studies [2, 34, 37, 41, 45, 59, 95, 136, 206] covering a wide range of different markets — including the New York Stock Exchange (NYSE), the London Stock Exchange (LSE), the Paris Bourse, the Deutsche Bourse, Euronext, FX markets, the S&P 500 index, German interest rates futures, and crude oil futures — have all reported that intra-day return series exhibit no significant autocorrelations, except on short time scales. Short-term autocorrelations seem to disappear more quickly in recent data than in older data: they were reported to persist for about 30 minutes for the S&P 500 index during 1984–1987 [162], about 20–30 minutes in FX markets during 1991–1995 [37], about 10 minutes for the S&P 500 index during 1991–2001 [95], and only a few minutes on Euronext during 2007–2008 [45]. We return to this interesting phenomenon in Section 5.7.1.

5.4.2 Absolute Returns

The absence of autocorrelations in returns does not imply that returns are temporally independent [55]. Several empirical studies have reported that *absolute-return series* $\{|R_1|, |R_2|, \ldots\}$ exhibit long memory in a wide range of different markets, including the S&P 500 index, Euronext, the Tokyo Stock Exchange, the Paris Bourse, the

NYSE, and FX markets [45, 71, 104, 140]. The reported values of the Hurst exponent range from $H \approx 0.8$ to $H \approx 0.9$.

The long memory of absolute returns implies that large absolute price changes tend to follow other large absolute price changes and that small absolute price changes tend to follow other small absolute price changes. This important stylized fact was noticed by Mandelbrot more than 50 years ago [152] and is widely known as *volatility clustering* [56]. The existence of volatility clustering has stimulated the design of several important classes of models (such as GARCH [197] and its extensions) that seek to mimic this effect [56]. Volatility clustering eliminates the simple random walk as a model for the temporal evolution of asset prices, because large price changes are more likely to follow other large price changes than they are to occur unconditionally [142].

5.4.3 Order Flow

In recent years, some researchers have been granted access to data that describes the temporal evolution of financial markets at the resolution of individual order arrivals and cancellations.

Definition. *The* limit order sign series *is* $\{L_1, L_2, \ldots\}$*, where*

$$L_i = \begin{cases} -1, & \textit{if the } i^{th} \textit{ limit order arrival is a buy limit order arrival,} \\ +1, & \textit{if the } i^{th} \textit{ limit order arrival is a sell limit order arrival,} \end{cases} \tag{5.40}$$

for $i = 1, 2, \ldots$*. The* cancellation sign series *and* market order sign series *are defined similarly, using active order cancellations and market order arrivals, respectively, instead of limit order arrivals.*

Lillo *et al.* [136] reported long memory in the limit order sign series, cancellation sign series, and market order sign series on the LSE and the NYSE, and Bouchaud *et al.* [34] reported long memory in the market order sign series on Euronext. The reported values of the Hurst exponent range from $H \approx 0.7$ to $H \approx 0.9$, which suggests that the strength of long memory in order flow varies substantially across assets.

Both Lillo *et al.* [136] and Bouchaud *et al.* [34] reported that the sample ACFs for order-sign series contain no discernible breaks corresponding to the length of a trading day and thereby argued that long memory in order flow persists across daily boundaries. However, Axioglou and Skouras [7] reported different results for the market order sign series on the LSE. They noted that although standard statistical tests strongly support the existence of long memory, Berkes' change-point test (see

Section 5.3.5) detects daily boundaries in these series with high accuracy. They thereby argued that the apparent long memory in market order signs on the LSE is mostly an artifact caused by daily structural breaks. We return to this discussion in Section 5.7.4.

5.5 Methodology

For each of the three currency pairs, we construct intra-day and cross-day time series of returns, absolute returns, limit order signs, and cancellation signs as follows.

5.5.1 Returns and Absolute Returns

Let $P(t)$ denote the price of the most recent trade[7] that occurs at or before time t. For a given trading day D_i (see Section 3), we use the Hotspot FX trade-data file (see Section 3.2.1.1) to reconstruct $P(t)$ during the peak trading hours of 08:00:00 – 17:00:00 GMT (see Section 3.3.1). We then sample $P(t)$ once per second to construct the intra-day price series $\{p_1, p_2, \ldots, p_{32401}\}$, intra-day return series $\{r_1, r_2, \ldots, r_{32400}\}$, and intra-day absolute-return series $\{|r_1|, |r_2|, \ldots, |r_{32400}|\}$ on day D_i, according to (5.39). We repeat this process for each of the trading days in our sample, which we label $D_1, D_2, \ldots, D_{30}$.

For a given pair of consecutive trading days D_i and D_{i+1}, we construct each of the cross-day return series and absolute-return series by concatenating the second half of the relevant intra-day series from day D_i and the first half of the corresponding intra-day series from day D_{i+1}. Each intra-day return and absolute series contains 32400 entries, so the boundary between trading days D_i and D_{i+1} in the cross-day series occurs after the 16200^{th} entry.

Figure 5.1 shows a schematic to illustrate how we construct the intra-day and cross-day return series and absolute-return series.

5.5.2 Limit Order Signs and Cancellation Signs

For a given trading day D_i, we use the Hotspot FX tick-data file (see Section 3.2.1.2) to produce an ordered list of the limit order arrivals that occur during the peak trading

[7]Some authors study returns of the mid-price series $m(t) = [b(t) + a(t)]/2$. However, because traders on Hotspot FX do not have access to the values of $b(t)$ or $a(t)$ in real time due to the platform's quasi-centralized structure (see Chapter 3), we instead study returns of the trade-price series. This information is available to all traders on Hotspot FX in real time, via the trade-data stream (see Section 2.3.4). We return to this discussion in Section 5.7.

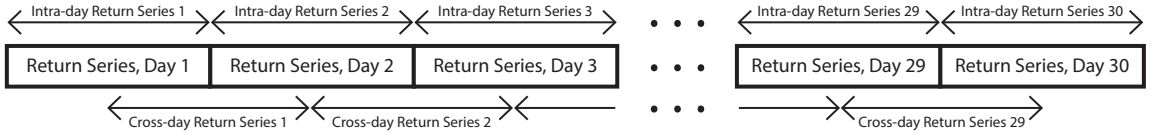

Figure 5.1: Schematic to illustrate construction of intra-day and cross-day return series. We construct intra-day and cross-day absolute-return series in the same way.

hours of 08:00:00 – 17:00:00 GMT. For each day, we use this list to construct the intra-day limit order sign series according to (5.40). We construct the intra-day cancellation sign series similarly, using an ordered list of the active order cancellations.[8] We repeat this process for each of the trading days $D_1, D_2, \ldots, D_{30}$. The lengths of the intra-day limit order sign and cancellation sign series vary according to the number of limit order arrivals and active order cancellations each day.

For a given pair of consecutive trading days D_i and D_{i+1}, we similarly construct each of the cross-day limit order sign series and cancellation sign series by concatenating the second half of the relevant intra-day series from day D_i and the first half of the corresponding intra-day series from day D_{i+1}. If an intra-day series has odd length, we round down to the next even number. In contrast to the returns and absolute returns, both the series lengths and the locations of the boundaries between different trading days vary according the number of limit order arrivals and active order cancellations. Therefore, the boundary between trading days D_i and D_{i+1} does not necessarily lie at the mid-point in the cross-day series.

Figure 5.2 shows a schematic to illustrate how we construct the intra-day and cross-day limit order signs series and cancellation sign series.

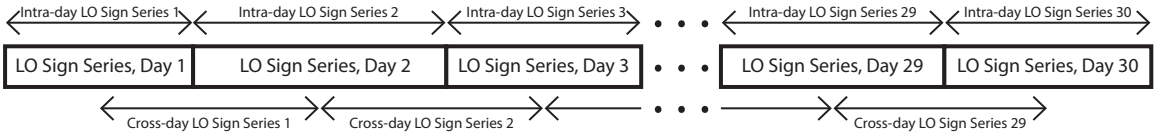

Figure 5.2: Schematic to illustrate construction of intra-day and cross-day limit order sign series. We construct intra-day and cross-day cancellation sign series in the same way.

[8]Recall from Section 3.2.2 that "cancellations" refer to cancel-order events in the Hotspot FX tick-data files. A cancel-order event occurs when an active order is either cancelled or fully matched to an incoming market order.

5.5.3 Normalizing Series Lengths

We introduce the following normalization to enable comparisons between the performance of the cumulative-sum change-point estimator (5.37) for the cross-day series on Hotspot FX, for which the locations of the daily boundaries differ for different series. Let $\{w_1, w_2, \ldots, w_N\}$ denote a cross-day series with daily boundary k^* that satisfies $1 < k^* < N$. The *normalized change-point estimator* $\tilde{k}$ is given by

$$\tilde{k} = \begin{cases} (\hat{k}^* - k^*)/k^*, & \text{if } \hat{k}^* \leq k^*, \\ (\hat{k}^* - k^*)/(k - k^*), & \text{otherwise.} \end{cases} \tag{5.41}$$

Observe that $\tilde{k} \in [-1, 1]$ and $\tilde{k} = 0$ if and only if $\hat{k}^* = k^*$.

5.6 Results

5.6.1 Intra-Day Series

5.6.1.1 Sample Autocorrelation Functions

For each of the three currency pairs, there is no evidence of autocorrelation in returns at any lag $k \geq 1$ (see the top row of Figure 5.3 for EUR/USD; other currency pairs are similar). By contrast, positive autocorrelations persist in the absolute-return series for several hundreds of seconds (see the left panel in the second row of Figure 5.3 for EUR/USD; other currency pairs are similar). The occasional negative values in the daily sample ACFs do not persist when averaging across all 30 days (see the right panel in the second row of Figure 5.4 for EUR/USD; other currency pairs are similar), which suggests that these negative values are statistical noise caused by estimating $\hat{\gamma}$ from a finite sample. It is difficult to discern the behaviour of the sample ACFs for lags longer than about 1000 seconds because the statistical errors associated with estimation are approximately constant for all k, but the signal strength decreases for larger values of k due to data sparsity [33]. Although we do not attempt to fit a curve to these plots (doing so often produces very poor results; see Section 5.3), the slow decay of the sample ACFs (see the top panel of Figure 5.4 for EUR/USD; other currency pairs are similar) suggests that it is plausible that absolute returns exhibit long memory. We use several other tests to address this question throughout the subsequent sections of this chapter.

Up to lags of about 25 events, the sample ACFs for both the limit order sign series and cancellation sign series fluctuate between positive and negative values (see the third and bottom rows of Figure 5.3 for EUR/USD; other currency pairs are similar).

This behaviour arises for all three currency pairs and persists when averaging the daily sample ACFs across all 30 days, which suggests that it is a robust statistical property of the data that is not caused by statistical noise. After these fluctuations subside, the mean sample ACFs remain positive for lags of several thousands of events (see the middle and bottom panels of Figure 5.4 for EUR/USD; other currency pairs are similar). Together, these observations suggest that the intra-day limit order sign series and cancellation sign series exhibit both short-range negative autocorrelations (which die out quickly but cause the sample ACFs to fluctuate between positive and negative values for small k) and long-range positive autocorrelations (which decay slowly over several thousands of lags).

5.6.1.2 Rescaled-Range Plots

For each of the three currency pairs, the slope of the rescaled-range plot (see Section 5.3.1) for each intra-day return series is close to 0.5 for all values of k (see the top left panel of Figure 5.5 for EUR/USD; other currency pairs are similar). This suggests that intra-day return series are short-memory processes. By contrast, the slope of the rescaled-range plot for each intra-day absolute-return series is greater than 0.5 for values of k larger than about 100 (see the top right panel of Figure 5.5 for EUR/USD; other currency pairs are similar) and the slope of the rescaled-range plots for each intra-day limit order sign and cancellation sign series is greater than 0.5 for values of k larger than about 10000 (see the bottom left and bottom right panels of Figure 5.5 for EUR/USD; other currency pairs are similar). This suggests that the intra-day absolute-return series and order-flow sign series are long-memory processes.

5.6.1.3 Lo's Modified Rescaled-Range Test

For most of the intra-day return series, Lo's modified rescaled-range test (see Section 5.3.2) does not reject the null hypothesis of short memory at the 5% significance level for a wide range of q (see Table 5.1 and the top panel of Figure 5.6 for EUR/USD; other currency pairs are similar). Therefore, Lo's test provides strong evidence that the return series are short-memory processes.

For the intra-day absolute-return series, the values of Lo's test statistic vary considerably across the three currency pairs (see Table 5.1 and Figure 5.7). For EUR/USD, the test rejects the null hypothesis of short memory at the 5% significance level on all days and at all choices of q that we study. For GBP/USD, the values of the test statistic tend to be smaller than the corresponding values for EUR/USD, but the test still rejects the null hypothesis of short memory on all but one or two

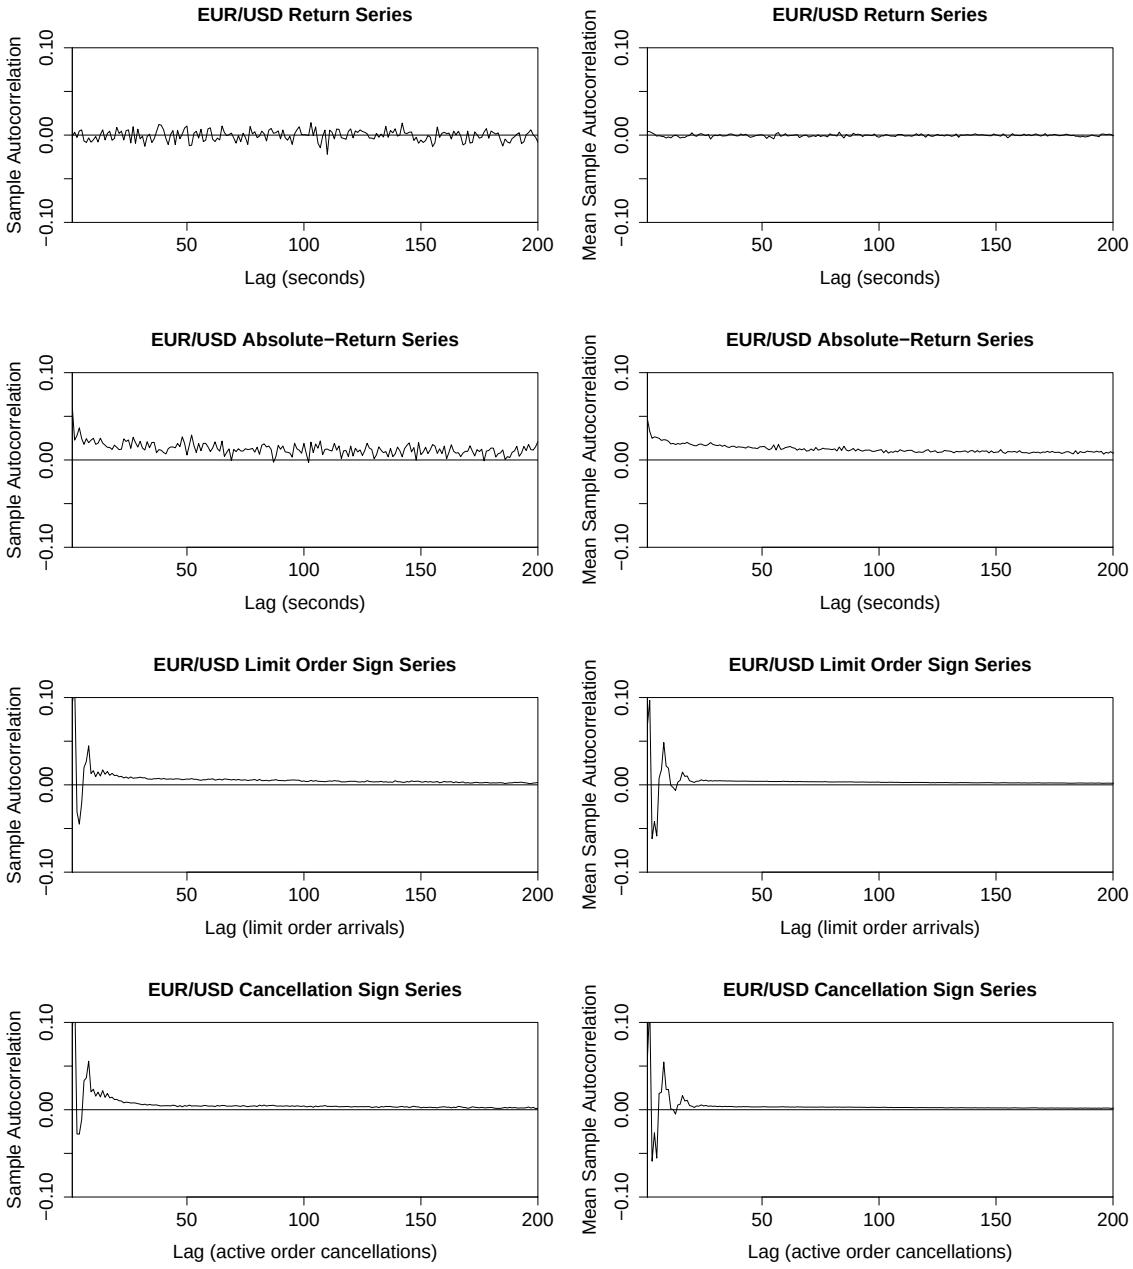

Figure 5.3: Sample autocorrelation functions (ACFs) for EUR/USD intra-day (top row) return series, (second row) absolute-return series, (third row) limit order sign series, and (bottom row) cancellation sign series. In each row, the left plot shows the sample ACF for 4 May 2010 (other days are similar) and the right plot shows the mean sample ACF across all 30 days.

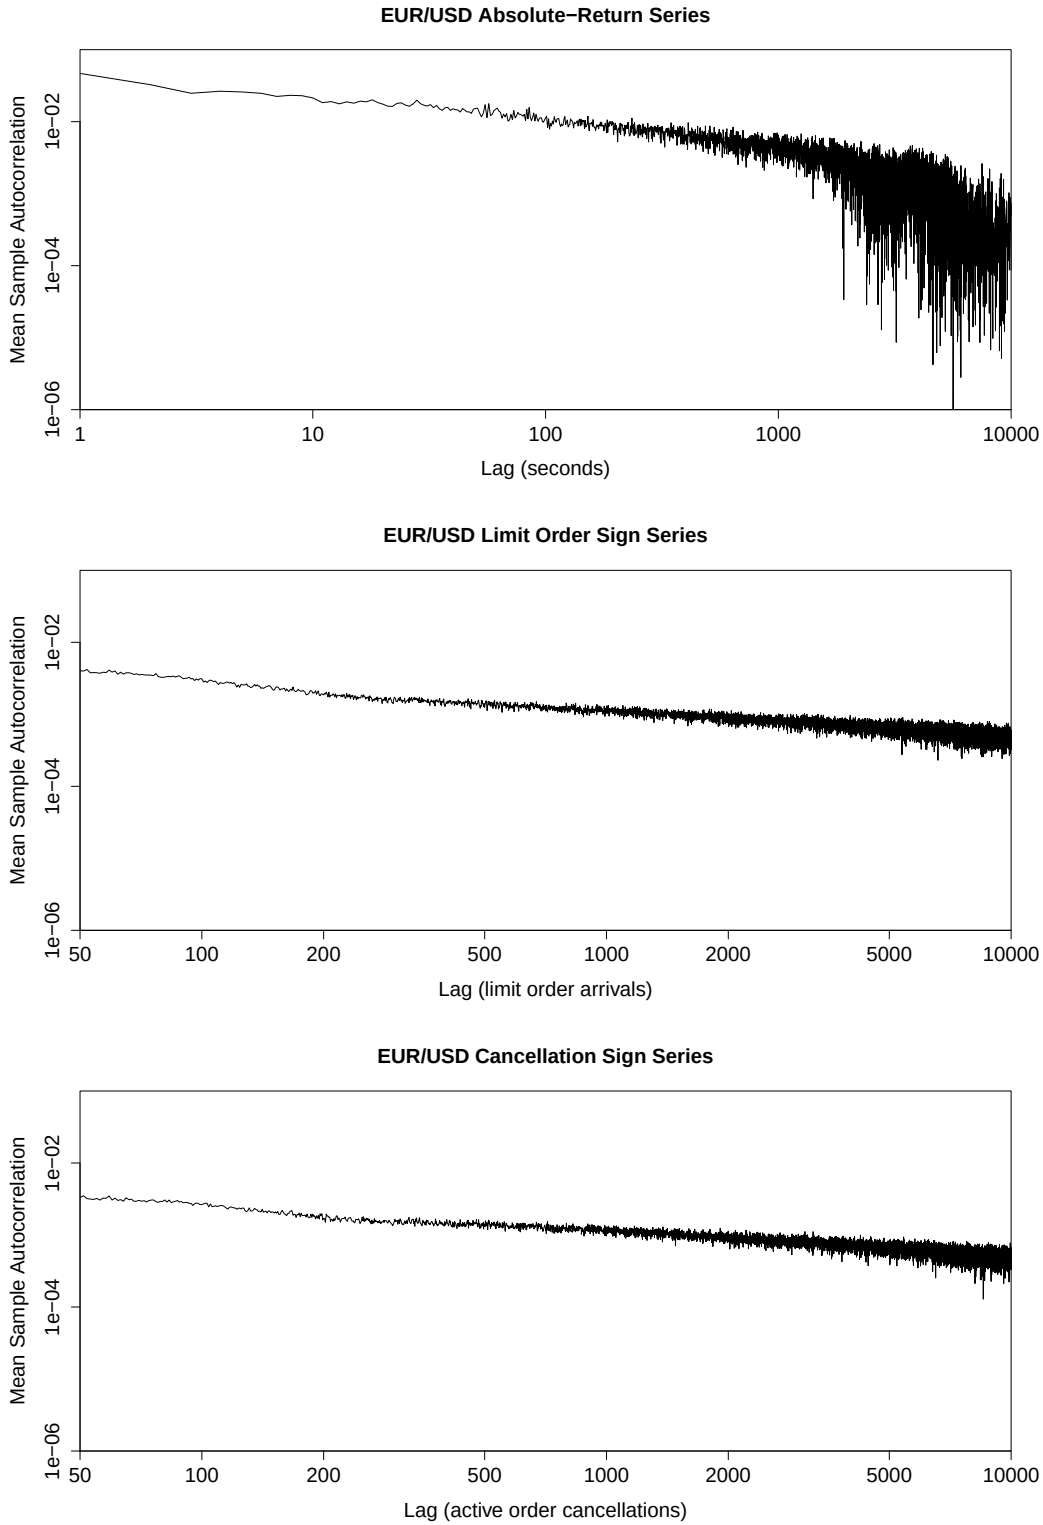

Figure 5.4: Mean sample ACFs averaged across all 30 days for EUR/USD intra-day (top) absolute-return series, (middle) limit order sign series, and (bottom) cancellation sign series.

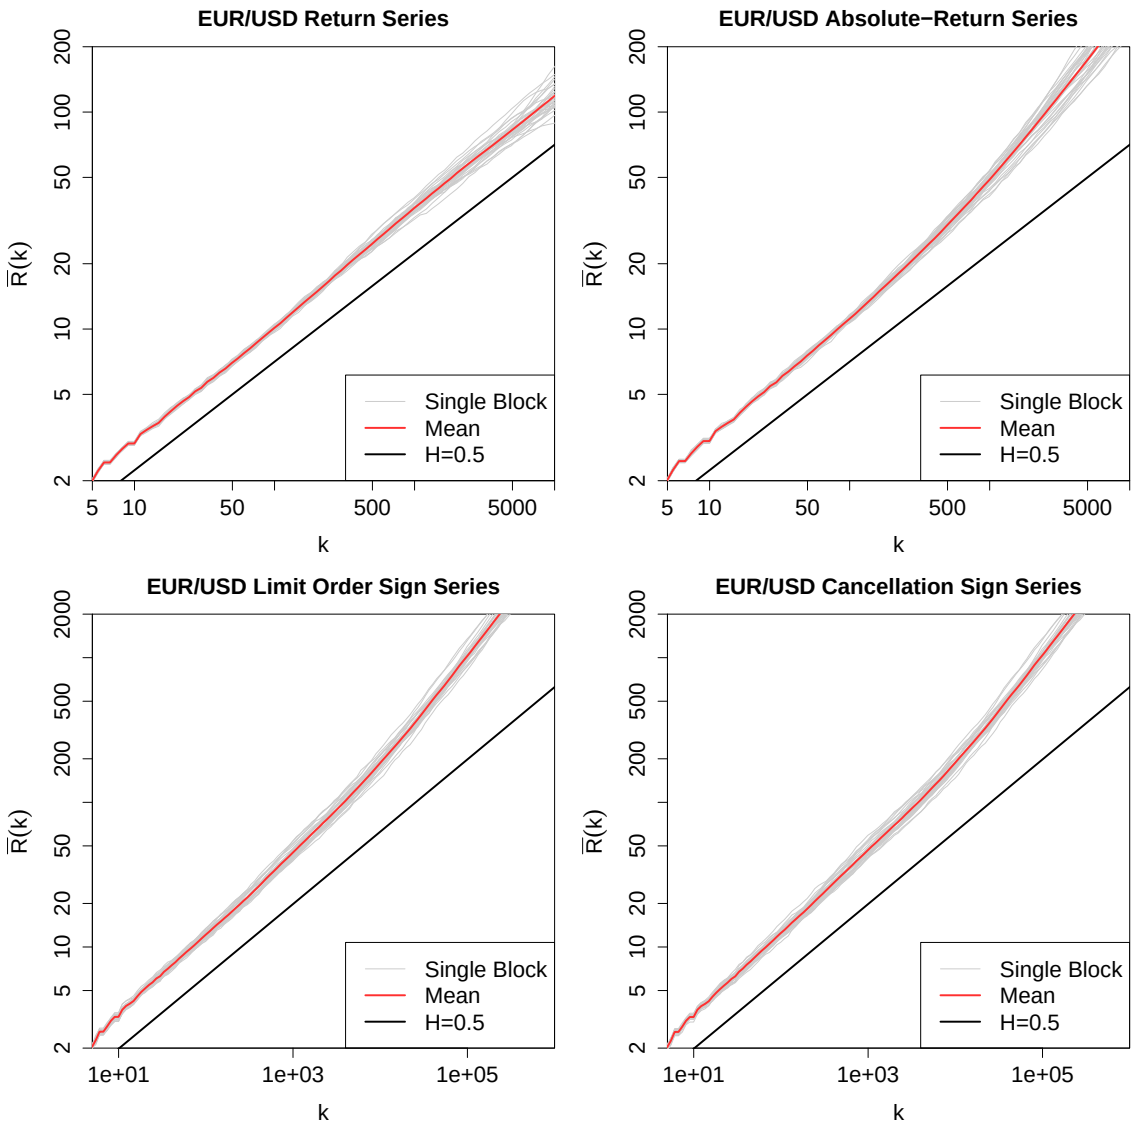

Figure 5.5: Rescaled-range plots for EUR/USD intra-day (top left) return series, (top right) absolute-return series, (bottom left) limit order sign series, and (bottom right) cancellation sign series. In each plot, the grey curves indicate the $\overline{R}(k)$ statistics for a single day and the red curves indicate the mean $\overline{R}(k)$ statistics across all 30 days. The solid black lines have gradient 0.5. These plots use $B = 100$ blocks (see Section 5.3.1). We produced similar plots for several different values of $B \in [10, 1000]$, and we found the results to be qualitatively the same for all B in this range.

days (depending on the choice of q). This provides strong evidence that the intra-day absolute-return series for EUR/USD and GBP/USD are long-memory processes. For EUR/GBP, however, the results are much less clear because the values of the test statistic cluster close to the critical threshold for Lo's test. When using Andrews' plug-in estimator $\hat{q}$ (see Table 5.1), Lo's test suggests that about half of the 30 intra-day series exhibit short memory, while the other days exhibit long memory. The results of Lo's test thus seem to be inconclusive for EUR/GBP.

For the intra-day limit order sign series and cancellation sign series, Lo's test rejects the null hypothesis of short memory at the 5% significance level for each currency pair, on all days, and at all choices of q that we study (see the middle and bottom panels of Figure 5.6 for EUR/USD; other currency pairs are similar). This provides strong evidence that the intra-day limit order sign series and cancellation sign series are long-memory processes.

Series	EUR/USD	GBP/USD	EUR/GBP
Returns	13.3%	6.7%	13.3%
Absolute Returns	100%	96.7%	53.3%
Limit Order Signs	100%	100%	100%
Cancellation Signs	100%	100%	100%

Table 5.1: Percentage of intra-day series for which Lo's modified rescaled-range test rejects the null hypothesis of short memory at the 5% significance level when using Andrews' [5] data-driven plug-in estimator $\hat{q}$ from Equation (5.30).

5.6.1.4 DFA

When calculating the DFA estimates (see Section 5.3.3) of H for the limit order sign series and cancellation sign series, the negative short-range autocorrelations (see Section 5.6.1.1) dominate the mean detrended standard deviations $F(m)$ for small window lengths m. Therefore, it is necessary to choose carefully the range of m over which to perform DFA estimates of H. For each intra-day series, we perform this task by plotting $F(m)$ versus m on doubly logarithmic axes, for logarithmically spaced values of $m \in [4, N]$.

For the intra-day return and absolute-return series, the plots follow an approximately straight line for all values of m in this range (see the top left and top right panels of Figure 5.8 for EUR/USD; other currency pairs are similar). For the intra-day limit order sign series and cancellation sign series, the plots follow an approximately straight line for values of m larger than about 80, but they differ for smaller values

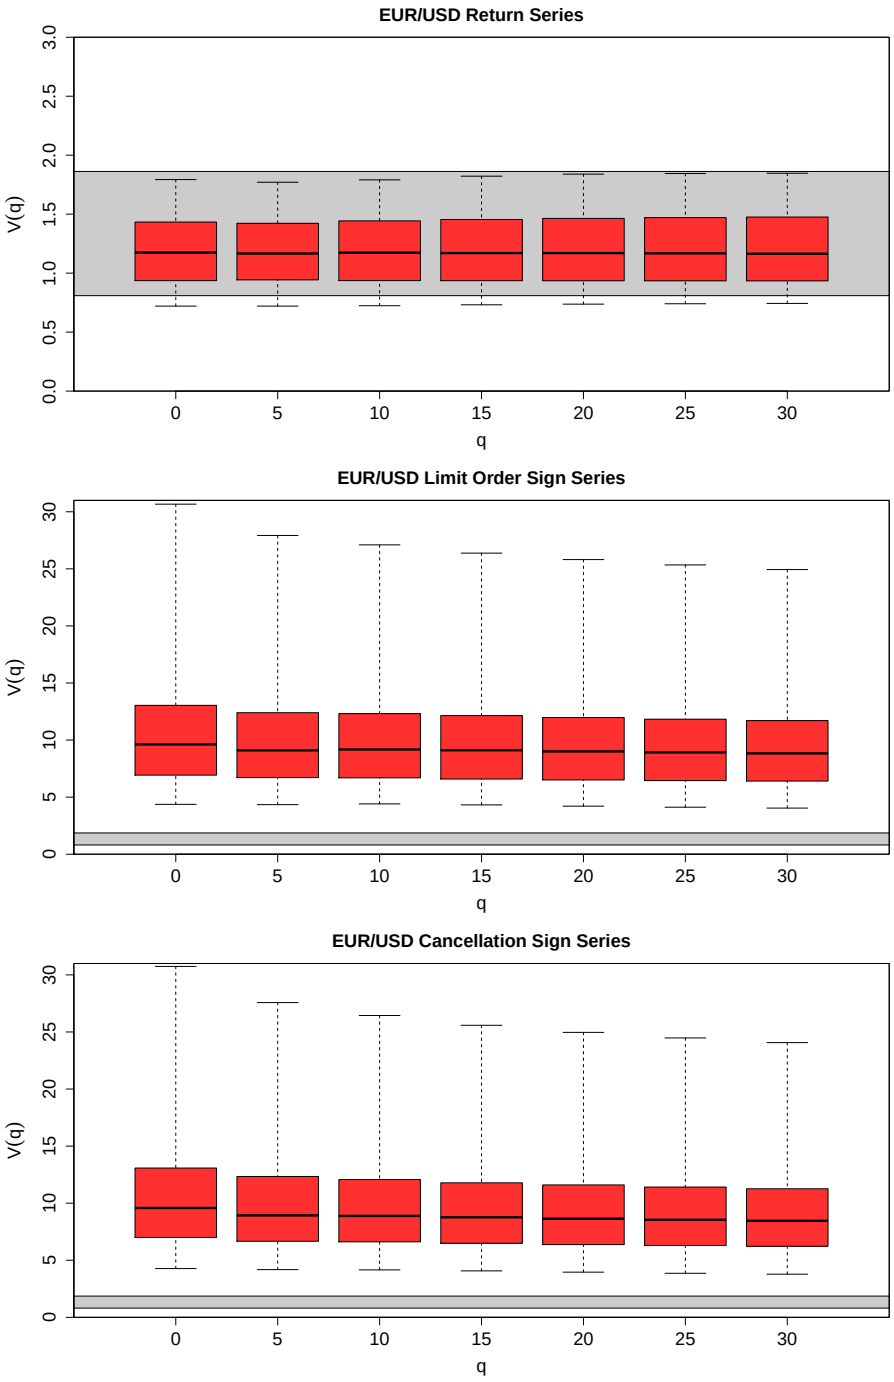

Figure 5.6: Box plots of Lo's modified rescaled-range test statistic $V(q)$ for the EUR/USD (top) return series, (middle) limit order sign series, and (bottom) cancellation sign series for several choices of q (see Section 5.3.2). The boxes indicate the lower quartile, median, and upper quartile of $V(q)$ across all 30 days at the given choice of q. The whiskers indicate the minimum and maximum of $V(q)$ across all 30 days at the given choice of q. The grey shading indicates the critical region for Lo's modified rescaled-range test at the 5% significance level.

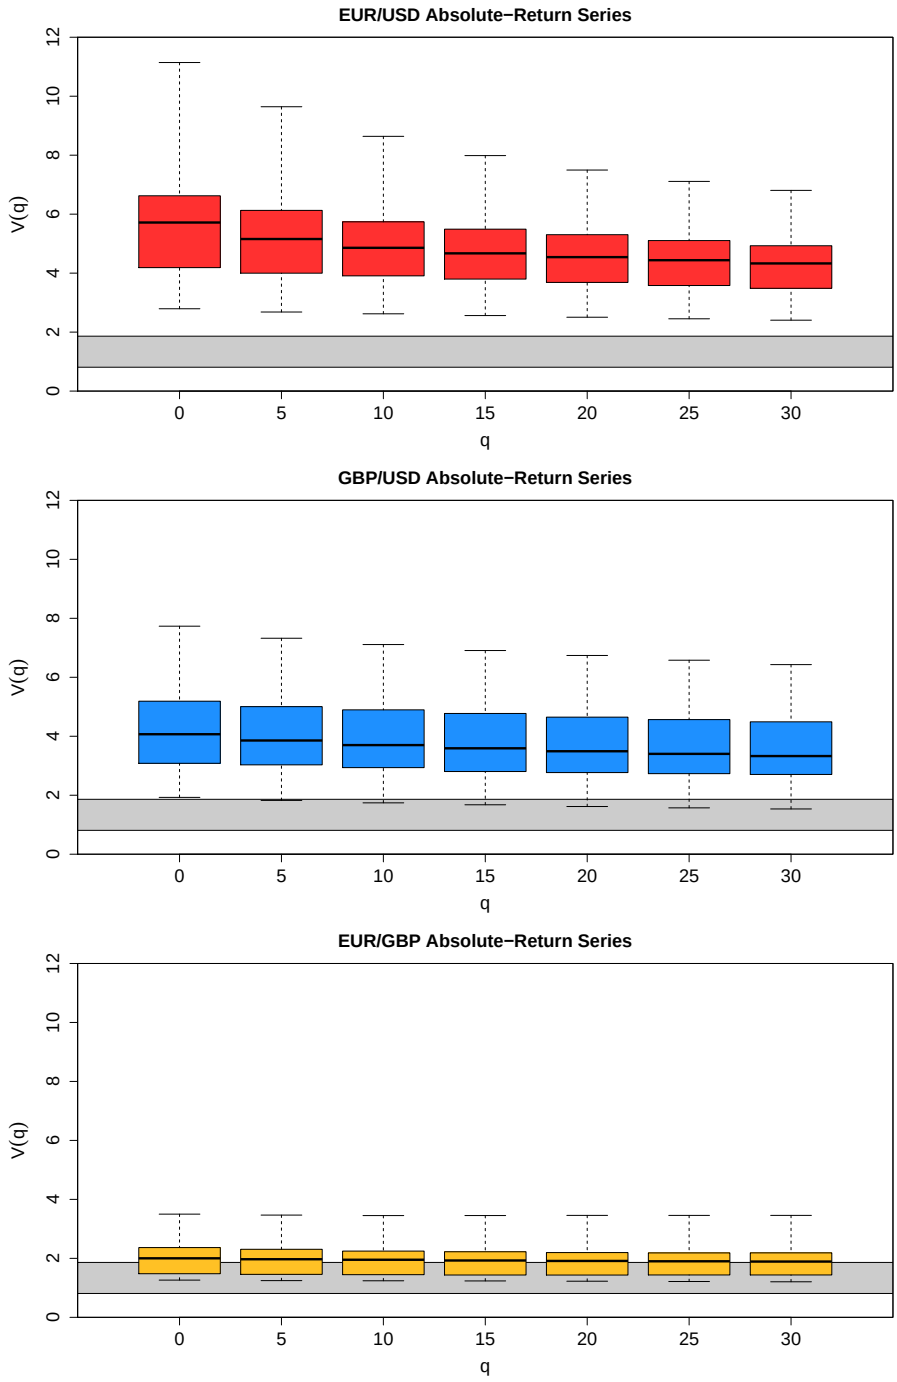

Figure 5.7: Box plots of Lo's modified rescaled-range test statistic $V(q)$ for the (top) EUR/USD, (middle) GBP/USD, and (bottom) EUR/GBP absolute-return series for several choices of q (see Section 5.3.2). The boxes indicate the lower quartile, median, and upper quartile of $V(q)$ across all 30 days at the given choice of q. The whiskers indicate the minimum and maximum of $V(q)$ across all 30 days at the given choice of q. The grey shading indicates the critical region for Lo's modified rescaled-range test at the 5% significance level.

of m due to the negative short-range autocorrelations in these series (see the bottom left and bottom right panels of Figure 5.8 for EUR/USD; other currency pairs are similar). We therefore perform our DFA estimates of H using the values $m_{\min} = 4$ for the intra-day return series and absolute return series and $m_{\min} = 100$ for the intra-day limit order sign series and cancellation sign series. We list our DFA estimates of H using these values of $m_{\min}$ in Section 5.6.1.6.

5.6.1.5 Log-Periodogram Regression

To investigate the sensitivity of the log-periodogram regression estimates of H to the number of Fourier frequencies c included in the regression (see Section 5.3.4), we plot the log-periodogram regression estimates of H for several different values of $c \in [100, 5000]$. For the intra-day return series, the estimates of H are similar for all choices of c larger than about 1500 (see the top left panel of Figure 5.9 for EUR/USD; other currency pairs are similar). For the intra-day absolute-return series, limit order sign series, and cancellation sign series, however, the estimates of H tend to decrease as c increases (see the top right, bottom left, and bottom right panels of Figure 5.9 for EUR/USD; other currency pairs are similar), and there is no obvious plateau over which the estimates of H are stable. Therefore, the log-periodogram regression estimates of H for these series depend heavily on the choice of c.

As we discussed in Section 5.3.4, there is no universal rule for choosing a value of c that captures only the small-λ behaviour of the periodogram. In the absence of an obvious choice for c, we try two popular rules of thumb, namely $c = \sqrt{N}$ [91] and $c = 0.1 \times (N/2)$ [194]. We stress that although these rules of thumb are used widely in the empirical literature, neither are based on a rigorous derivation or optimization.

For the intra-day return series and absolute-return series, $c = \sqrt{N}$ results in high variance between the estimates of H on different days, but $c = 0.1 \times (N/2)$ provides relatively stable estimates (see the top right and top let panels of Figure 5.9 for EUR/USD; other currency pairs are similar). Therefore, we use $c = 0.1 \times (N/2)$ to estimate H for these series. For the intra-day limit order sign series and cancellation sign series, both rules of thumb produce estimates with relatively low variance but $c = 0.1 \times (N/2)$ produces estimates of $H \approx 0.5$ due to the extremely large size of the series (and, therefore, the inclusion of too many Fourier frequencies far from $\lambda = 0$). Given the results of our other statistical tests in this chapter, we do not regard $H \approx 0.5$ to be sensible, so we use $c = \sqrt{N}$ to estimate H for the intra-day limit order sign series and cancellation sign series (see the bottom left and bottom right panels of Figure 5.9 for EUR/USD; other currency pairs are similar). We list our

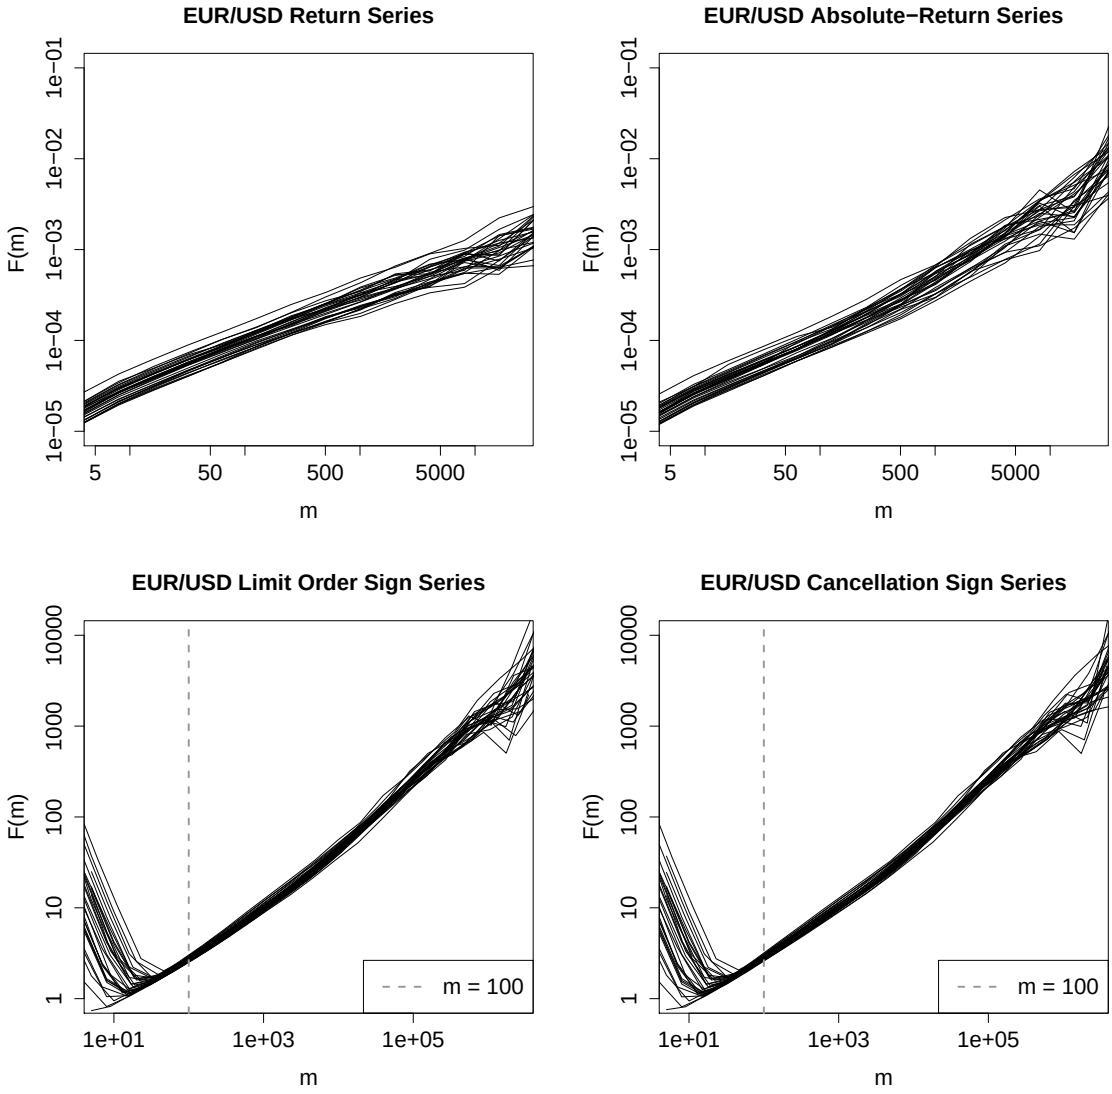

Figure 5.8: Doubly logarithmic plots of the length-m mean detrended standard deviation $F(m)$ for logarithmically spaced values of $m \in [4, N]$ for the EUR/USD intra-day (top left) return series, (top right) absolute-return series, (bottom left) limit order sign series, and (bottom right) cancellation sign series. For the intra-day return series and absolute-return series series, we perform our DFA estimates of H using logarithmically spaced values of $m \in [4, N]$. For the intra-day limit order sign series and cancellation sign series, we perform our DFA estimates of H using logarithmically spaced values of $m \in [100, N]$.

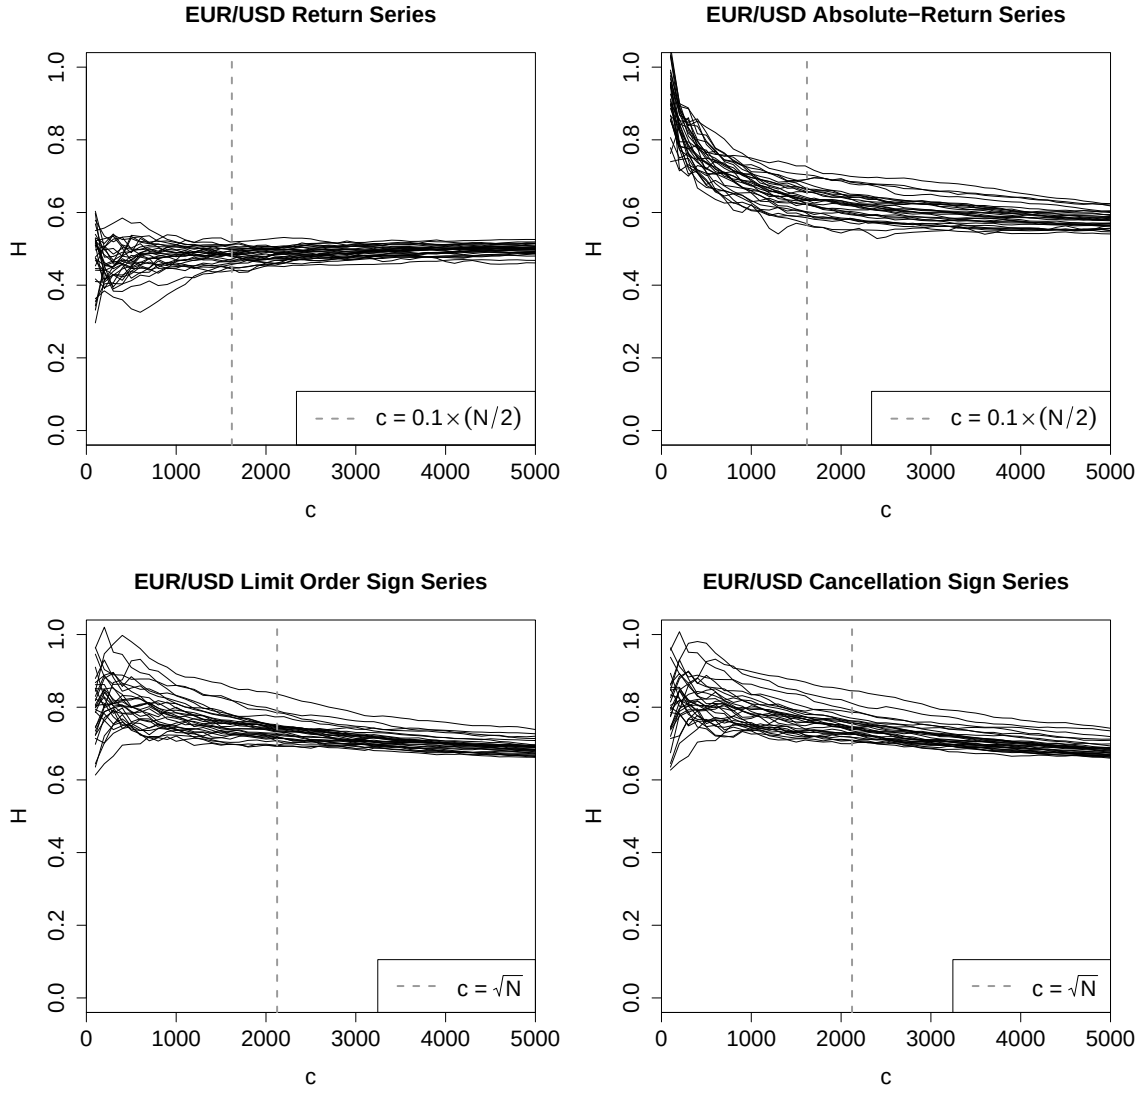

Figure 5.9: Log-periodogram regression estimates of H for given numbers of Fourier frequencies c for the EUR/USD intra-day (top left) return series, (top right) absolute-return series, (bottom left) limit order sign series, and (bottom right) cancellation sign series. For the intra-day return series and absolute-return series series, the dashed grey lines denote $c = 0.1 \times (N/2)$ [194]; for the intra-day limit order sign series and cancellation sign series,the dashed grey lines denote $c = \sqrt{N}$ [91] (see the discussion in the main text).

log-periodogram regression estimates of H using these values of c in Section 5.6.1.6. We stress that our results depend heavily on these choices, and that other values of c result in different estimates of H. The absence of a clear choice for c highlights a weakness of log-periodogram regression in the present application.

5.6.1.6 Comparing DFA and Log-Periodogram Regression

For each of the three currency pairs, both the DFA and log-periodogram regression estimates of H for the intra-day return series cluster around $H \approx 0.5$ (see the top left panel of Figure 5.10 and Table 5.2). This provides strong evidence that the intra-day return series are short-memory processes. For the intra-day absolute-return series, both the DFA and log-periodogram regression estimates of H vary substantially across the three currency pairs (see the top right panel of Figure 5.10 and Table 5.2), but in all cases suggest that the absolute-return series exhibit long memory. The DFA estimates suggest that the Hurst exponent is $H \approx 0.6$ to $H \approx 0.7$ for EUR/USD and GBP/USD, but that it is only slightly larger than 0.5 for EUR/GBP. The log-periodogram regression estimates of H tend to be slightly lower than the DFA estimates, but we stress that the log-periodogram estimates depend heavily on the choice of c (for which we use a rule of thumb). Therefore, we believe that DFA provides more robust estimates of H than log-periodogram regression.

For the intra-day limit order sign series and cancellation sign series, both DFA and log-periodogram regression strongly support the hypothesis of long memory. The DFA estimates suggest that $H \approx 0.7$ for each of the three currency pairs (see Table 5.2). Again, we believe that DFA provides more robust estimates of H because the log-periodogram regression estimates depend on the choice of c.

5.6.2 Cross-Day Series

To assess how aggregating data from different trading days impacts our results, we repeat all of our calculations using the cross-day series (see Section 5.5). In all cases, we find that the sample ACFs, rescaled-range plots, and results from Lo's modified rescaled-range test are qualitatively the same as those for the intra-day series (thus, we do not show these plots). Moreover, we find that the DFA and log-periodogram regression estimates of H for the cross-day series are qualitatively the same as those for the intra-day series (see Figure 5.11 and Table 5.2).

To provide benchmarks against which to compare the performance of the cumulative-sum change-point estimator (5.37) for the cross-day series, we simulate the series

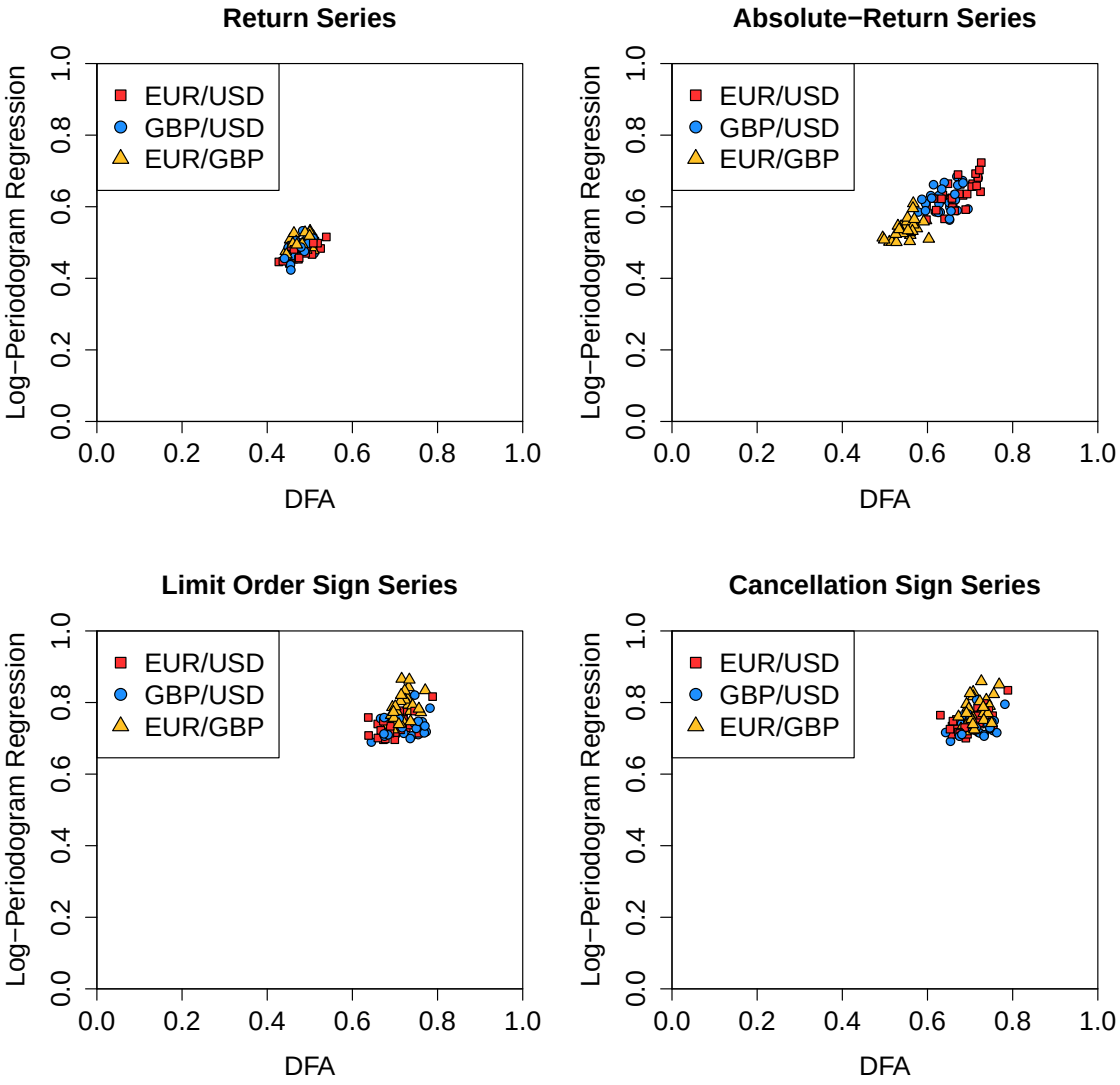

Figure 5.10: DFA and log-periodogram regression estimates of the Hurst exponent H for the EUR/USD, GBP/USD, and EUR/GBP (top left) return series, (top right) absolute-return series, (bottom left) limit order sign series, and (bottom right) cancellation sign series.

EUR/USD				
	Intra-Day		Cross-Day	
	DFA	LP	DFA	LP
Returns	0.48 (0.03)	0.48 (0.02)	0.48 (0.03)	0.48 (0.02)
Absolute Returns	0.67 (0.04)	0.64 (0.04)	0.68 (0.04)	0.64 (0.04)
Limit Order Signs	0.70 (0.04)	0.74 (0.03)	0.71 (0.03)	0.74 (0.03)
Cancellation Signs	0.70 (0.04)	0.74 (0.03)	0.70 (0.03)	0.74 (0.03)

GBP/USD				
	Intra-Day		Cross-Day	
	DFA	LP	DFA	LP
Returns	0.48 (0.02)	0.49 (0.02)	0.47 (0.02)	0.49 (0.02)
Absolute Returns	0.64 (0.03)	0.62 (0.04)	0.64 (0.04)	0.62 (0.04)
Limit Order Signs	0.72 (0.04)	0.74 (0.03)	0.73 (0.03)	0.74 (0.03)
Cancellation Signs	0.72 (0.03)	0.74 (0.03)	0.73 (0.03)	0.74 (0.03)

EUR/GBP				
	Intra-Day		Cross-Day	
	DFA	LP	DFA	LP
Returns	0.48 (0.02)	0.50 (0.02)	0.48 (0.02)	0.50 (0.01)
Absolute Returns	0.55 (0.03)	0.54 (0.03)	0.55 (0.04)	0.54 (0.03)
Limit Order Signs	0.72 (0.02)	0.79 (0.04)	0.72 (0.03)	0.79 (0.04)
Cancellation Signs	0.71 (0.02)	0.78 (0.04)	0.72 (0.03)	0.78 (0.04)

Table 5.2: DFA and log-periodogram regression estimates of the Hurst exponent H for the EUR/USD, GBP/USD, and EUR/GBP intra-day and cross-day return series, absolute-return series, limit order sign series, and cancellation sign series. The entries in each column indicate the mean of the estimates across all intra-day or cross-day series and the numbers in brackets indicate 1 standard deviation of the estimates across all intra-day or cross-day series.

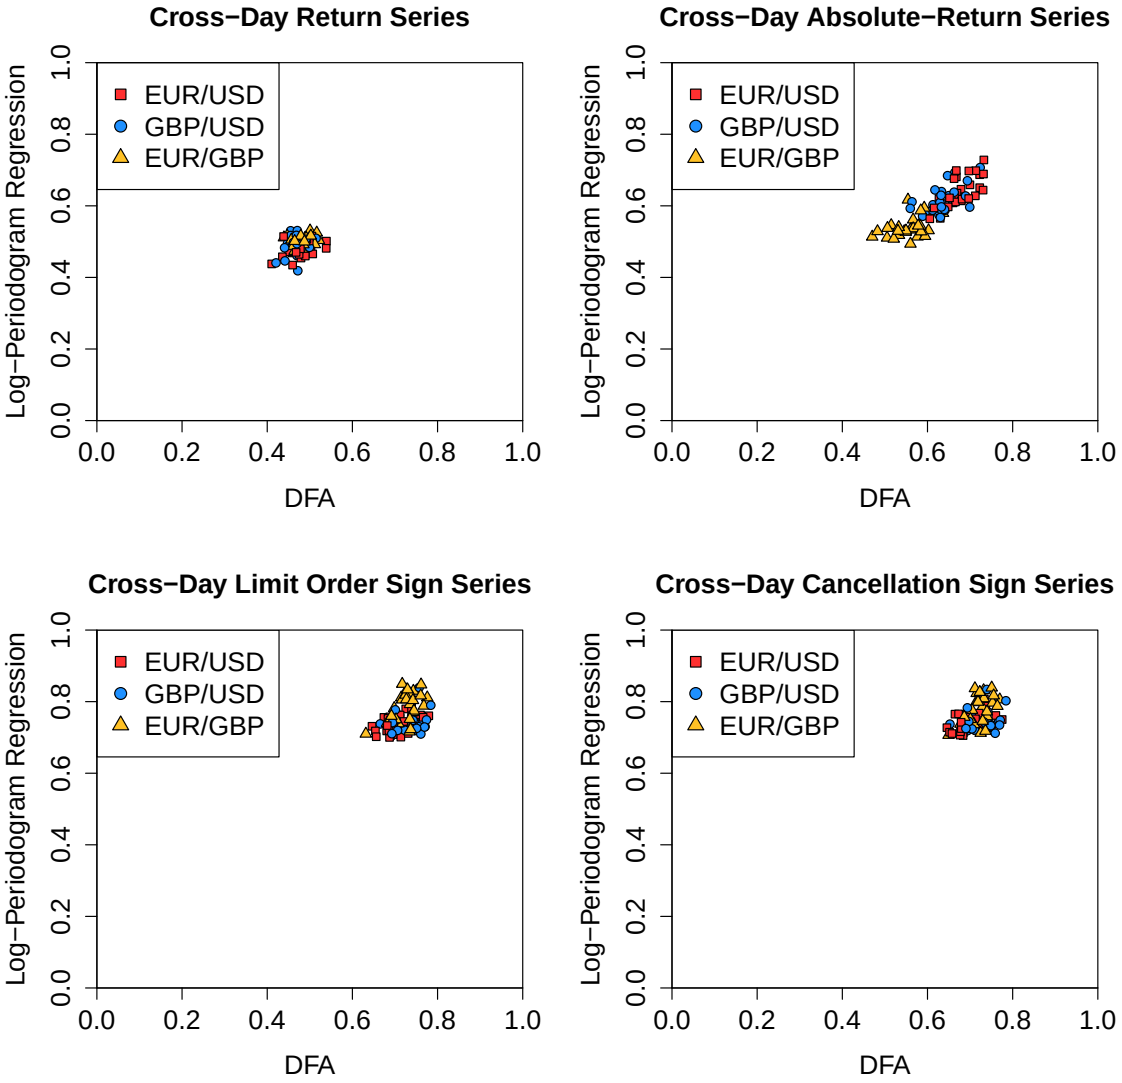

Figure 5.11: DFA and log-periodogram regression estimates of the Hurst exponent H for the EUR/USD, GBP/USD, and EUR/GBP cross-day (top left) return series, (top right) absolute-return series, (bottom left) limit order sign series, and (bottom right) cancellation sign series.

$\{X_1, X_2, \ldots, X_{32400}\}$ and $\{Y_1, Y_2, \ldots, Y_{32400}\}$, where

$$X_i \overset{i.i.d.}{\sim} \mathcal{N}(1,1), \text{ for } i = 1, 2, \ldots, 32400,$$
$$Y_i \overset{i.i.d.}{\sim} \begin{cases} \mathcal{N}(1,1), & \text{for } i = 1, 2, \ldots, 16200, \\ \mathcal{N}(-1,1), & \text{for } i = 16201, 16202, \ldots, 32400. \end{cases}$$

We calculate the normalized cumulative-sum change-point estimates (see Section 5.3.5) for $\{X_1, X_2, \ldots, X_{32400}\}$ and $\{Y_1, Y_2, \ldots, Y_{32400}\}$, and then repeat this process 100000 times and calculate the ECDFs of these estimates. The ECDF for the $\{X_1, X_2, \ldots, X_{32400}\}$ series provides a benchmark of the estimator's output for a second-order stationary series with no structural breaks, and the ECDF for the $\{Y_1, Y_2, \ldots, Y_{32400}\}$ series provides a benchmark of the estimator's output for a short-range dependent series with a structural break.

In all cases, the shape of the ECDFs for the cross-day series is similar to that for the $\{X_1, X_2, \ldots, X_{32400}\}$ series (see Figure 5.12), albeit with a larger step size because there are only 29 cross-day series in our data set. Therefore, the cumulative-sum change-point estimator performs very poorly at detecting the true locations of the daily boundaries in the cross-day series.

For all intra-day and cross-day limit order sign series and cancellation sign series, Berkes' change-point test strongly rejects the null hypothesis of a single structural break in favour of the alternative hypothesis of true long memory (see Figure 5.13). The test thereby provides strong evidence against the hypothesis that the apparent long-memory properties of these series are an artifact caused by structural breaks at daily boundaries. For EUR/USD and GBP/USD, the test also rejects the null hypothesis for the vast majority of intra-day and cross-day absolute-return series. The results are less clear for the EUR/GBP absolute-return series, but this is unsurprising given that the strength of long memory in absolute returns (and, therefore, the plausibility of the alternative hypothesis) is much weaker than for EUR/USD and GBP/USD. The test does not reject the null hypothesis for the intra-day or cross-day return series, but this is again unsurprising given that long memory in the return series is highly implausible.

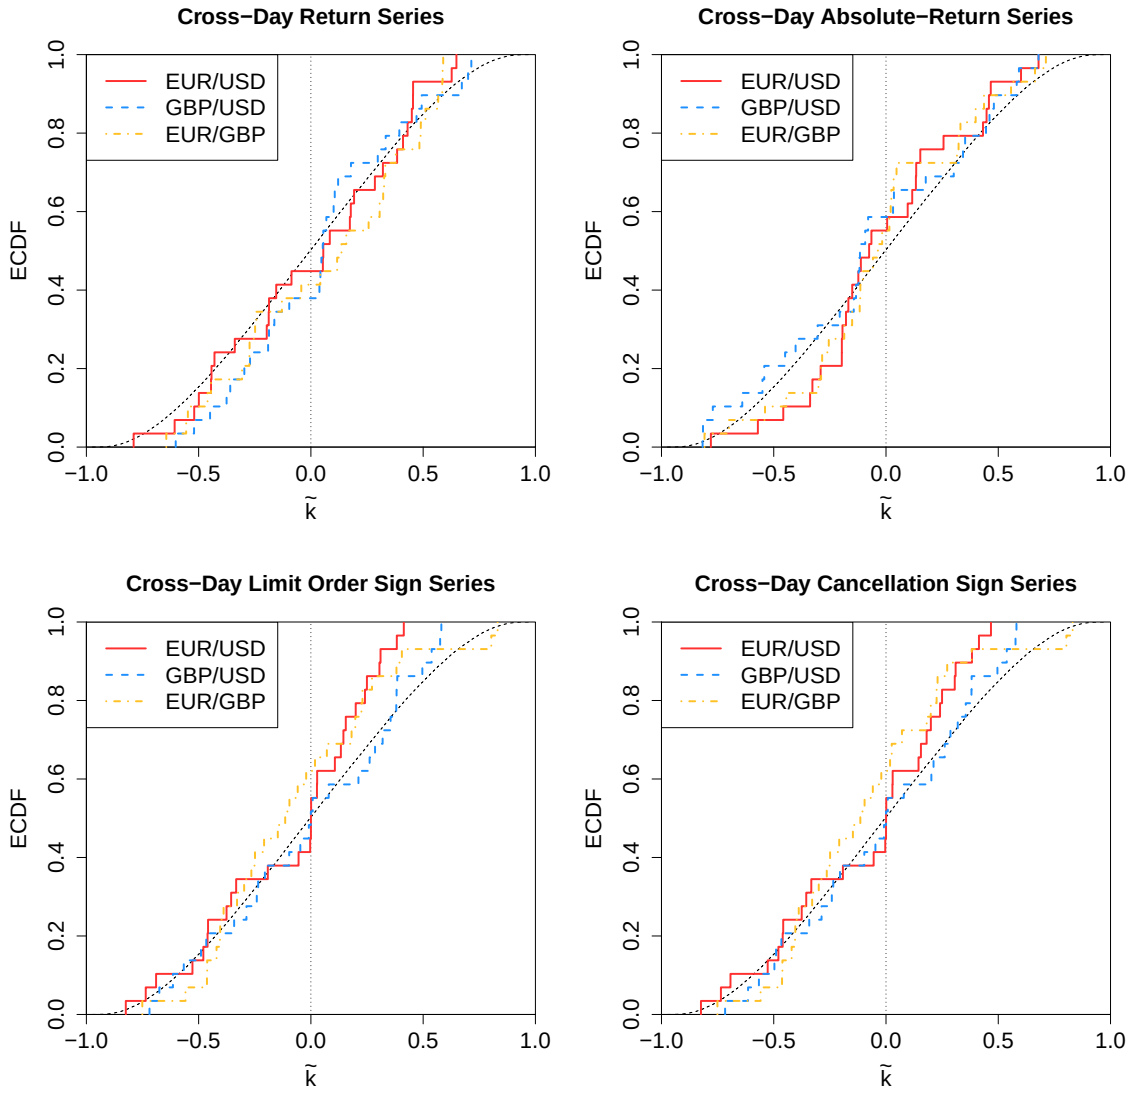

Figure 5.12: ECDFs of the normalized cumulative-sum change-point estimator (5.41) for the EUR/USD, GBP/USD, and EUR/GBP cross-day (top left) return series, (top right) absolute-return series, (bottom left) limit order sign series, and (bottom right) cancellation sign series. The dashed black curve illustrates the estimator's output for a second-order stationary series with no structural breaks and the dotted black line illustrates the estimator's output for a short-range dependent series with a structural break (see the description in the main text).

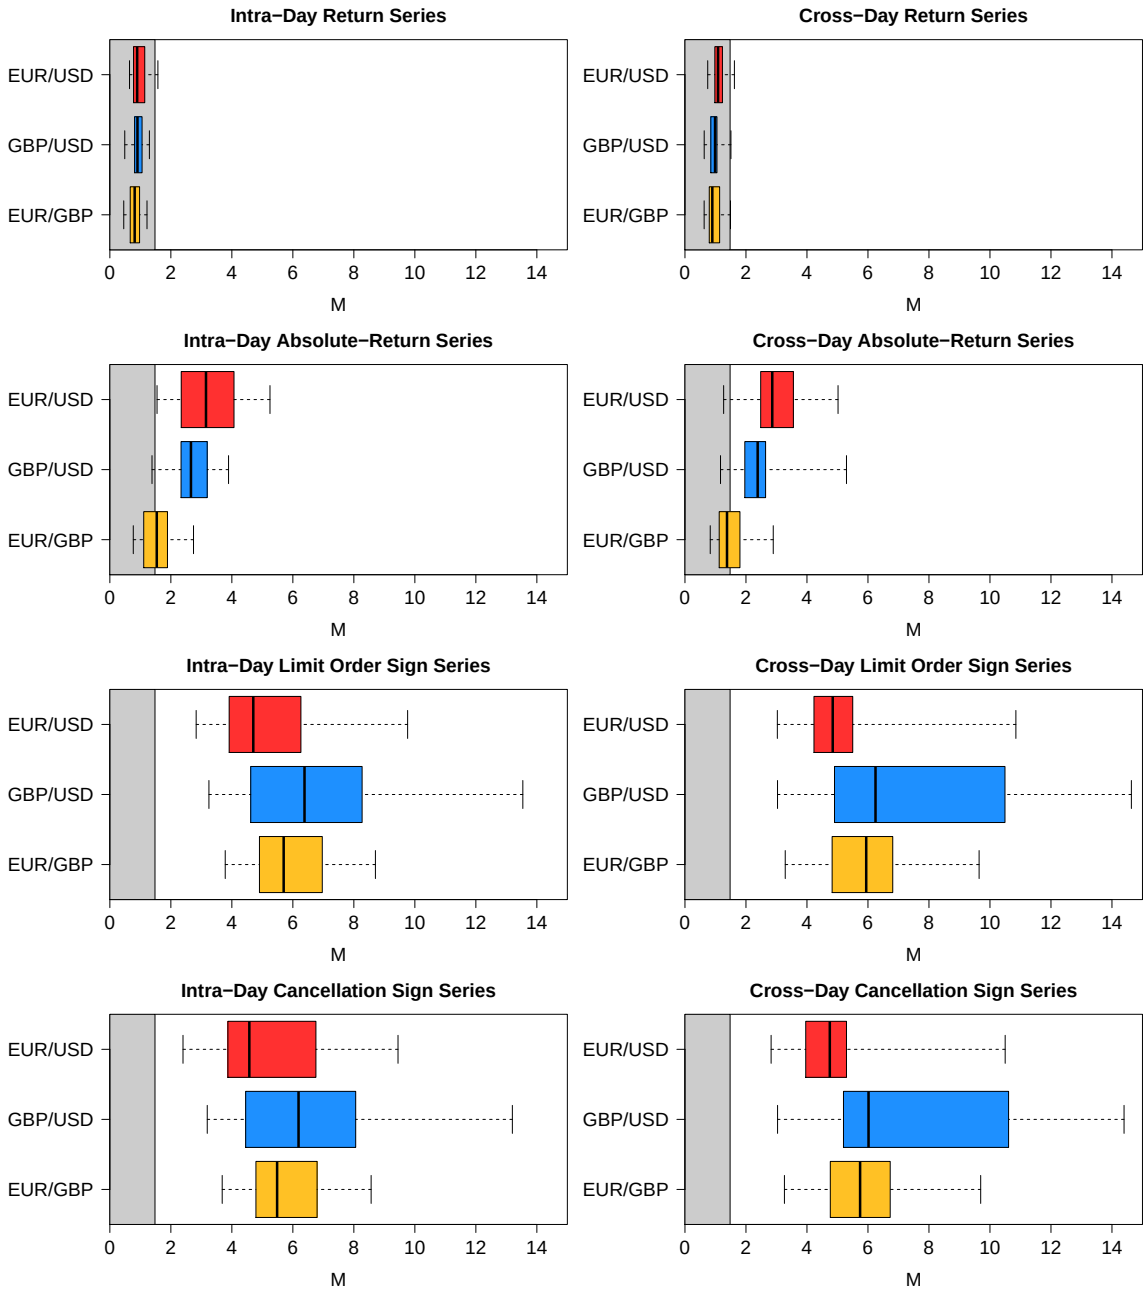

Figure 5.13: Box plots of Berkes' change-point test statistic M (see Section 5.3.5) for EUR/USD, GBP/USD, and EUR/GBP (top) return series, (second) absolute return series, (third) limit order sign series, and (bottom) cancellation sign series. The left plots show the results for the intra-day series and the right plots show the results for the cross-day series. The boxes indicate the lower quartile, median, and upper quartile of the values of M across all intra-day or cross-day series. The whiskers indicate the minimum and maximum of the values of M across all intra-day or cross-day series. For all series, we use Andrews' [5] data-driven plug-in estimator $\hat{q}$ from Equation (5.30) to calculate M. The grey shading indicates the critical region for Berkes' change-point test at the 5% significance level.

5.7 Discussion

5.7.1 Returns

For each of the three currency pairs that we study, we find no evidence for autocorrelations in the return series at any time lag $k \geq 1$ second (see Figure 5.3). Several other empirical studies (see Section 5.4.1) have reported low-order negative autocorrelations in returns in a wide range of different markets, but we do not find evidence of this effect on Hotspot FX. We propose two possible reasons why this could be the case.

First, the quasi-centralized structure of trading on Hotspot FX (see Chapter 2) may influence the statistical properties of the return series. Several authors [40, 55, 96] have argued that low-order negative autocorrelations in returns are at least in part due to "bid-ask bounce" (i.e., the tendency for consecutive trades to "bounce" between $b(t)$ and $a(t)$). The Hotspot FX trade-price series describes the prices of all trades that occur, but different traders on the platform have access to different trading opportunities at the same time. Superimposing the activity generated by different traders into a single trade-price series may eliminate the low-order negative autocorrelations that are present in data from centralized markets.

Second, the increasing popularity of electronic trading algorithms [12, 14, 125] may cause autocorrelations in returns to disappear more quickly in more recent data. Several authors [55, 56, 153] have argued that the lack of autocorrelations in returns can be explained by traders searching for statistical arbitrage opportunities.[9] Electronic trading algorithms are able to detect and capitalize on such opportunities in a tiny fraction of the time that it would require a human to perform the same tasks. Therefore, the increased use of such algorithms may reduce the length of time for which negative autocorrelations persist in the return series.

5.7.2 Absolute Returns

Our calculations suggest that the absolute-return series for EUR/USD and GBP/USD exhibit moderate long memory, with a Hurst exponent of $H \approx 0.6$ to $H \approx 0.7$ (see Table 5.2). For EUR/GBP, by contrast, long memory in the absolute-return series seems to be much weaker, with a Hurst exponent slightly above 0.5. We propose two possible reasons for why this may be the case. First, traders may employ different

[9] If traders detect predictability in returns, then they can implement simple strategies with positive expected earnings. By repeatedly exploiting these possibilities, such traders reduce the strength of the autocorrelations until they no longer persist.

strategies when trading EUR/GBP than when trading EUR/USD and GBP/USD. However, given that there are many strong similarities between the statistical properties of order flow for the different currency pairs (see Chapter 4), we do not find it plausible that this explanation would account for the relatively large differences in H that we observe. Second, the lower levels of trading activity for EUR/GBP may influence the estimates of H. The mean inter-trade time for EUR/GBP is several times larger than it is for both EUR/USD and GBP/USD (see Table 4.1), which causes the absolute-return series for EUR/GBP to contain a much larger fraction of 0 entries than do the corresponding series for EUR/USD and GBP/USD. This data sparsity may reduce the strength of long memory for EUR/GBP when compared to EUR/USD and GBP/USD.

To assess the plausibility of this explanation, we repeated our calculations for intra-day absolute-return series sampled once every 10 seconds (rather than once per second). This caused the variance among the estimates of H to increase considerably (which is unsurprising given that lower-frequency sampling produces series with fewer entries), but it also caused the mean of the estimates to increase. In particular, both DFA and log-periodogram regression produced mean estimates of $H \approx 0.7$ to $H \approx 0.75$ for EUR/USD and GBP/USD and $H \approx 0.6$ for EUR/GBP. Moreover, other empirical studies, which have all sampled absolute-return series on much lower frequencies than once per second, have all reported larger Hurst exponents than those that we estimate for Hotspot FX (see Section 5.4.2). Together, these results suggest that the long-memory properties of absolute returns are influenced by the number of trades that occur between successive samplings of the series.

5.7.3 Limit Orders and Cancellations

For each of the three currency pairs that we study, the limit order sign series and cancellation sign series exhibit strong long memory that persists over several thousands of events (see Figure 5.4). Our calculations suggest that the Hurst exponent for these series is $H \approx 0.7$ (see Table 5.2), which is similar to the estimates reported by several other empirical studies of order flow in different markets [34, 136, 167].

There are two main hypotheses for why order-flow series exhibit long memory. The first hypothesis is that traders monitor each other's actions and update their strategies by imitating those of their most successful competitors [135]. This hypothesis suggests that long memory in order flow is caused by traders who deliberately replicate previous market activity. The second hypothesis is that traders who wish to perform large trades split them into smaller chunks, which they submit over several days (or even

months) to minimize their market impact [33, 34, 137]. This hypothesis suggests that long memory in order flow is caused by individual traders submitting many different orders with the same sign.

Gerig [90] assessed the plausibility of these two hypotheses by studying order-flow data that identifies order ownership. He reported that the order-sign series for orders that originate from the same brokerage exhibit strong long memory but that order-sign correlations for orders that originate from different brokerages decay quickly to 0. He thereby argued that the second hypothesis provides a more plausible explanation of long memory in order flow than does the first. Importantly, the first hypothesis requires that traders are able to monitor the temporal evolution of the market, whereas the second hypothesis does not. Given that traders on Hotspot FX are only able to view the order flow generated by their trading partners and that the Hotspot FX platform does not disseminate any information regarding order ownership, we believe that our findings provide further evidence that long memory in order flow is due to order splitting rather than imitation.

We also find strong evidence of low-order, negative autocorrelations in each of the three currency pairs' limit order sign series and cancellation sign series (see Figure 5.3). To our knowledge, this behaviour has not been reported by other empirical studies of order flow in different markets. Understanding the origins of this behaviour could help to illuminate how order flow on Hotspot FX differs from that of other markets, and we therefore believe that further study of this phenomenon is an important avenue for future research (see Section 5.8.1).

5.7.4 Intra-Day versus Cross-Day Series

For each of the three currency pairs that we study, we find no evidence that the autocorrelation properties of the cross-day series differ from those of the intra-day series (see Section 5.6.2). Moreover, we find no evidence that the apparent long-memory properties of the data are an artifact caused by a structural break (see Figures 5.12 and 5.13).

In a recent study, Axioglou and Skouras [7] concluded that the apparent long memory of order flow on the LSE is an artifact caused by structural breaks between different trading days. This raises the interesting question of why our results for Hotspot FX should contrast with these findings. We conjecture that the answer to this question lies in the fact that trading on the LSE commences at 08:00:00 and ceases at 16:30:00 [199] whereas trading in the FX spot market occurs 24 hours a day. Although we restrict our attention to the peak trading hours of 08:00:00 – 17:00:00

GMT each day (see Section 3.3.2), the absence of a market-wide closing time has several important consequences for the way that traders act.

First, many financial institutions require that traders unwind their positions (i.e., rebalance their net daily holdings to 0) before the end of each trading day [148]. The LSE market closing at 16:30:00 constitutes a hard deadline by which any traders who seek to unwind their positions must fulfil this goal, even if doing so requires them to trade at unfavourable prices. This could cause the statistical properties of order flow late in the trading day to differ substantially from those early in the trading day, and could therefore result in a structural break at the daily boundaries when concatenating data from different days. In the FX spot market, by contrast, there is no market-wide closing time by which traders who seek to unwind their positions must complete this task.

Second, the absence of market opening and closing times in the FX spot market enables traders in different time zones to begin and end their trading days at different times. Hsieh and Kleidon [119] noted that many traders spend the early part of their trading day assessing the state of the market, the middle part of their trading day performing the majority of their trades, and the late part of their trading day resetting their net inventory to 0. In markets with specified opening and closing times (such as the LSE), all traders progress through this cycle in phase. In the FX spot market, by contrast, traders in different time zones are able to choose the length and timing of their trading days as they wish. The flux of traders from different time zones into and out of the FX spot market could cause the statistical properties of order flow to differ from those in markets where all traders' trading days are aligned by the market opening and closing times.

Third, Axioglou and Skouras [7] conjectured that most traders on the LSE reassess their trading strategy once per day, while markets are closed, then implement their chosen strategy throughout the next trading day. This is not a feasible description of the actions of traders in the FX spot market, which is always open.

5.8 Summary and Future Work

In this chapter, we have investigated the autocorrelation properties of returns, absolute returns, limit order signs, and cancellation signs on Hotspot FX. Due to the extremely high levels of activity on the platform, and in contrast to other empirical studies, we were able to study the long-memory properties of intra-day series without needing to aggregate data from different trading days. For each of the three currency

pairs that we studied, we found that autocorrelations in the return series disappear within less than 1 second. We found moderate long memory (with a Hurst exponent of $H \approx 0.6$ to $H \approx 0.7$) in the EUR/USD and GBP/USD absolute-return series, but we found only weak long memory (with a Hurst exponent slightly above 0.5) in the EUR/GBP absolute-return series. We found that all three currency pairs' limit order sign series and cancellation sign series possess relatively strong long memory, with a Hurst exponent ranging from $H \approx 0.7$ to $H \approx 0.8$. We also uncovered a low-order, negative autocorrelation in the limit order sign series and cancellation sign series, which (to our knowledge) has not been reported previously.

We found that all of our results for data that spans different trading days were similar to those for intra-day data, and we strongly rejected the hypothesis that the apparent long memory is an artifact caused by structural breaks at the daily boundaries. Therefore, we conclude that long memory is a robust statistical property of absolute returns, limit order signs, and cancellation signs that persists across daily boundaries.

5.8.1 Future Work

Our findings suggest several possible avenues for future research.

- **Investigations of other markets:** It is crucial that researchers perform further empirical investigations of recent, high-quality data from other markets. Such work will help to develop an up-to-date understanding of the short- and long-range autocorrelation properties of order flow and price formation. It will be interesting to see whether the low-order, negative autocorrelations that we detect in the limit order sign series and cancellation sign series on Hotspot FX are present in other markets.

- **Formulating and testing models:** There exist several models that attempt to explain long memory in financial markets [33, 34, 88, 109, 127, 135, 137, 146]. However, most such models rely on auxiliary assumptions or unobservable parameters that make it difficult to relate their output to market data. The development of new models that make falsifiable predictions and thereby enable quantitative comparisons of competing hypotheses will help to clarify the true origins of the various interacting forms of long memory in financial markets.

- **Understanding cross-day series:** We have conjectured that the differences between our findings and those reported by Axioglou and Skouras [7] for the LSE

are caused by differences in the structure of trading days in the two markets. Further empirical study of data from other markets will clarify whether our conjecture is correct and will help to illuminate the role of daily periodicities in financial markets.

- **The efficiency paradox:** It is an important open question to reconcile the so-called "efficiency paradox": how can return series remain unpredictable given that order flow exhibits long memory? To date, there exist two main hypotheses. Bouchaud *et al.* [34, 35] proposed that markets reside at a self-organized critical point in which liquidity takers cause long-range autocorrelations in order flow that exactly balance the long-range *negative* autocorrelations caused by liquidity providers. Alternatively, Lillo and Farmer [136] proposed that predictability in order flow is offset by a negative correlation with available liquidity. Although several authors have noted that these two hypotheses can be regarded as equivalent [33, 77, 90], Bouchaud *et al.*'s explanation portrays price impact as fixed and temporary, whereas Lillo and Farmer's explanation portrays price impact as variable and permanent [77]. To date, there exists no clear consensus as to which approach best describes the temporal evolution of real markets. Further empirical and theoretical study of the efficiency paradox is an important avenue for future research.

Chapter 6

Price Formation on Hotspot FX

In this chapter, we perform an empirical study to investigate how heterogeneity in traders' access to trading opportunities impacts price formation on Hotspot FX. The work in this chapter is also described in the jointly authored paper [101].

6.1 Introduction

During the past decade, price formation in LOBs has received considerable attention across a wide range of disciplines (see [98] for our recent review). To date, however, almost all work on this topic addresses centralized LOBs, in which all traders have access to the same trading opportunities at the same time. Although many electronic trading platforms (including the LSE Electronic Trading Service [199], NASDAQ [171], and the Euronext Universal Trading Platform [74]) utilize this mechanism, many others employ alternative LOB configurations with different trading rules.

One such alternative is a quasi-centralized LOB. Recall from Chapter 2 that a quasi-centralized LOB is an LOB in which each trader's orders can only match to those of a specified set of trading partners but in which the price of each trade is announced publicly. In the extreme case where a quasi-centralized LOB's trade-partnership network (see Section 2.4) is complete, all traders are able to trade with all others and price formation is equivalent to that in a centralized LOB. Otherwise, different traders in the same quasi-centralized LOB have access to different trading opportunities at the same time.

In this chapter, we investigate how heterogeneity in traders' access to trading opportunities impacts price formation for the EUR/USD, GBP/USD, and EUR/GBP currency pairs on Hotspot FX. Specifically, we calculate the difference between the price at which each trade occurred and the price at which it would have occurred if the owner of the relevant market order had had access to all trading opportunities on

the platform at its time of its submission. We call this difference the *bypassing price* of a trade. We calculate both the difference between the bypassing prices of successive trades and the difference between the best quotes at their times of execution, in order to calculate the relative importance (in a sense that we make precise in Section 6.4) of these factors for explaining changes in trade price.

Although we identify a small number of trades that have a large bypassing price, we find that the vast majority of market orders match at a price that is close to the best quotes at their time of execution. For each of the three currency pairs, we find that the relative importance of changes in the best quotes for explaining changes in the trade price is much higher than that of changes in the bypassing price. We conclude that most traders on Hotspot FX are sufficiently well-connected with other traders on the platform that the restriction of trading activity to their trading partners does not strongly affect the prices of their trades.

The remainder of this chapter is organized as follows. In Section 6.2, we discuss price formation in a quasi-centralized LOB and give several definitions that we use throughout the chapter. In Section 6.3, we describe our data and methodology. In Section 6.4, we discuss the concept of relative importance. We present our empirical results in Section 6.5 and discuss these results in Section 6.6. In Section 6.7, we summarize our findings and discuss several possible avenues for future research.

6.2 Price Formation in a Quasi-Centralized LOB

Since the widespread uptake of limit order trading, many empirical and theoretical studies have provided answers to important questions regarding price formation in centralized LOBs, such as the nature of large price movements [202], the price impact of market orders [34, 204], the lagged impact of trading [35], and the possible origins of widely observed stylized facts [55, 202] (see [33] and [98] for surveys). By contrast, the current literature on price formation in quasi-centralized LOBs consists of little more than technical descriptions of the matching mechanisms and trading interfaces on specific platforms [14, 85, 185, 189]. Although these publications describe how each trader's access to trading opportunities depends on his/her trading partnerships, they do not provide any empirical analysis of real markets. Given the extremely large trading volumes that they facilitate [13], understanding price formation in quasi-centralized LOBs is an important task.

6.2.1 The Bypassing Price of a Trade

Recall from Chapter 2 that whenever a trader in a centralized LOB submits a sell (respectively, buy) market order $x = (p_x, \omega_x, t_x)$, the order at least partially matches to the highest-priority active buy (respectively, sell) order, which resides at the price $\underline{b}(t_x)$ (respectively, $\underline{a}(t_x)$). In a quasi-centralized LOB, by contrast, an arriving sell (respectively, buy) market order x submitted by a trader $\theta_i \in \Theta$ matches to the highest-priority active buy (respectively, sell) order that is owned by a trader θ_j such that $\theta_i \leftrightarrow \theta_j$, but it *bypasses* any higher-priority active buy (respectively, sell) orders that are owned by traders θ_k such that $\theta_i \nleftrightarrow \theta_k$.

Throughout this chapter, we perform all of our calculations in *transaction time*, whereby we advance the clock by 1 unit whenever a trade occurs. Let

$$T^B = t_1^B, t_2^B, \ldots, t_{N^B}^B,$$
$$T^A = t_1^A, t_2^A, \ldots, t_{N^A}^A$$

denote respectively the times of N^B successive seller-initiated trades[1] and N^A successive buyer-initiated trades in a quasi-centralized LOB, and let $p_1^B, p_2^B, \ldots, p_{N^B}^B$ and $p_1^A, p_2^A, \ldots, p_{N^A}^A$ denote the prices of these trades.

Definition. *The* traded-bid-price series *is the time series*

$$\left\{p_t^B : t \in T^B\right\} = p_1^B, p_2^B, \ldots, p_{N^B}^B. \tag{6.1}$$

The traded-ask-price series *is the time series*

$$\left\{p_t^A : t \in T^A\right\} = p_1^A, p_2^A, \ldots, p_{N^A}^A. \tag{6.2}$$

Definition. *The* bid-price series

$$\left\{p_t^b : t \in T^B\right\} = p_1^b, p_2^b, \ldots, p_{N^B}^b, \tag{6.3}$$
$$p_i^b = \underline{b}(T_i^B), \qquad i = 1, 2, \ldots, N^B$$

is the time series of bid prices, sampled immediately before the times of the N^B seller-initiated trades. The corresponding definition for the ask-price series $\left\{p_t^a : t \in T^A\right\}$ *is similar, using the N^A buyer-initiated trades instead of the N^B seller-initiated trades.*

[1] A seller-initiated trade occurs when an incoming sell market order matches to an active buy order, which resides on the bid side of the LOB. Therefore, we use the superscript B to denote such trades. Throughout this chapter, we study buyer-initiated trades and seller-initiated trades separately, in order to avoid difficulties associated with bid-ask bounce [40, 55, 96].

Definition. *The* bid-bypassing-price series

$$\begin{aligned} \left\{q_t^B : t \in T^B\right\} &= q_1^B, q_2^B, \ldots, q_{N^B}^B, \\ q_i^B &= p_i^b - p_i^B, \qquad i = 1, 2, \ldots, N^B \end{aligned} \tag{6.4}$$

is the time series of differences between the traded bid price and the bid price. The ask-bypassing-price series

$$\begin{aligned} \left\{q_t^A : t \in T^A\right\} &= q_1^A, q_2^A, \ldots, q_{N^A}^A, \\ q_i^A &= p_i^A - p_i^a, \qquad i = 1, 2, \ldots, N^A \end{aligned} \tag{6.5}$$

is the time series of differences between the ask price and the traded ask price.

Observe the difference in signs between the bid-side and ask-side definitions: q_i^B measures how far behind the bid price an arriving sell market order matches (i.e., how much smaller p_i^B is than p_i^b) and q_i^A measures how far behind the ask price an arriving buy market order matches (i.e., how much larger p_i^A is than p_i^a). Therefore, all trades in a quasi-centralized LOB have a non-negative bypassing price. In the extreme case of a complete BTA network (where every pair of traders possesses a BTA), all trades have a bypassing price of zero, so price formation is equivalent to that in a centralized LOB.

6.2.2 Price Changes

Throughout this chapter, we calculate the (appropriately signed) changes between successive entries in the trade-price, bypassing-price, and quote-price series.

Definition. *Given the traded-bid-price series* $\left\{p_t^B : t \in T^B\right\}$*, the* change-in-traded-bid-price series *is the time series*

$$\begin{aligned} \left\{\Delta p_t^B : t \in T^B\right\} &= \Delta p_1^B, \Delta p_2^B, \ldots, \Delta p_{N^B-1}^B, \\ \Delta p_i^B &= p_i^B - p_{i+1}^B, \qquad i = 1, 2, \ldots, N^B - 1. \end{aligned} \tag{6.6}$$

Given the traded-ask-price series $\left\{p_t^A : t \in T^A\right\}$*, the* change-in-traded-ask-price series *is the time series*

$$\begin{aligned} \left\{\Delta p_t^A : t \in T^A\right\} &= \Delta p_1^A, \Delta p_2^A, \ldots, \Delta p_{N^A-1}^A, \\ \Delta p_i^A &= p_{i+1}^A - p_i^A, \qquad i = 1, 2, \ldots, N^A - 1. \end{aligned} \tag{6.7}$$

Definition. *Given the bid-price series* $\left\{p_t^b : t \in T^B\right\}$*, the* change-in-bid-price series *is the time series*

$$\begin{aligned} \left\{\Delta p_t^b : t \in T^B\right\} &= \Delta p_1^b, \Delta p_2^b, \ldots, \Delta p_{N^B-1}^b, \\ \Delta p_i^b &= p_i^b - p_{i+1}^b, \qquad i = 1, 2, \ldots, N^B - 1. \end{aligned} \tag{6.8}$$

Given the ask-price series $\left\{p_t^a : t \in T^A\right\}$*, the* change-in-ask-price series *is the time series*

$$\begin{aligned} \left\{\Delta p_t^a : t \in T^A\right\} &= \Delta p_1^a, \Delta p_2^a, \ldots, \Delta p_{N^A-1}^a, \\ \Delta p_i^a &= p_{i+1}^a - p_i^a, \qquad i = 1, 2, \ldots, N^A - 1. \end{aligned} \tag{6.9}$$

Observe the differences in signs between the bid-side and ask-side definitions: Δp_i^B measures how far behind the previous traded bid price an arriving sell market order matches (i.e., how much smaller p_{i+1}^B is than p_i^B) and Δp_i^A measures how far behind the previous traded ask price an arriving buy market order matches (i.e., how much larger p_{i+1}^A is than p_i^A). Similarly, Δp_i^b measures how far behind the previous bid price an arriving sell market order matches (i.e., how much smaller p_{i+1}^b is than p_i^b) and Δp_i^a measures how far behind the previous ask price an arriving buy market order matches (i.e., how much larger p_{i+1}^a is than p_i^a).

Definition. *Given the bid-bypassing-price series* $\left\{q_t^B : t \in T^B\right\}$*, the* change-in-bid-bypassing-price series *is the time series*

$$\begin{aligned} \left\{\Delta q_t^B : t \in T^B\right\} &= \Delta q_1^B, \Delta q_2^B, \ldots, \Delta q_{N^B-1}^B, \\ \Delta q_i^B &= q_{i+1}^B - q_i^B, \qquad i = 1, 2, \ldots, N^B - 1 \end{aligned} \tag{6.10}$$

The corresponding definition for the change-in-ask-bypassing-price series $\left\{\Delta q_t^A : t \in T^A\right\}$ *is similar, using the ask-bypassing-price series* $\left\{q_t^A : t \in T^A\right\}$ *instead of the bid-bypassing-price series* $\left\{q_t^B : t \in T^B\right\}$*.*

The signs of the bid-bypassing-price and ask-bypassing-price series are already aligned by (6.4) and (6.5), so no further sign modification is required.

6.3 Data

For each of the EUR/USD, GBP/USD, and EUR/GBP currency pairs, we use the Hotspot FX data to study market activity during the peak trading hours of 08:00:00 – 17:00:00 GMT on each of the 30 days in our sample.

6.3.1 Inferring matches

Recall from Chapter 3 that the Hotspot FX data consists of two separate files: the trade-date file, which describes all trades that occur, and the tick-data file, which lists all updates to $\mathcal{L}(t)$. Whenever a market order matches to a visible active order, this is reported in both the tick-data and trade-data files.

Both files record all events with millisecond time stamps, but (as is common with LOB data [114]) the two data files are recorded by different servers, so the time stamps for the two files differ by a small, variable amount. In order to calculate the quote-price and bypassing-price series that we study in this chapter, it is necessary to infer which updates in the trade-data file correspond to which updates to $\mathcal{L}(t)$, as reported in the tick-data file. We perform this inference using Algorithm 6.1.

For each trade listed in the trade-data file, Algorithm 6.1 searches for all possible updates to $\mathcal{L}(t)$ listed in the tick-data file that could correspond to the given trade and that occur within a given time interval of the trade time (as reported in the trade-data file). If the algorithm finds exactly one suitable update to $\mathcal{L}(t)$ within the given time window, then it records a *strong pairing* between the events listed in the two files. If the algorithm finds more than one suitable update, then it records a *weak pairing* for the event in the trade-data file. In order to minimize the error associated with incorrect pairings, in this chapter we study only the trades for which Algorithm 6.1 reports a strong pairing, and only the price changes for which Algorithm 6.1 reports a strong pairing for both the corresponding trades (i.e., the i^{th} and the $(i+1)^{\text{th}}$ trades).

6.3.1.1 Choosing τ

The output of Algorithm 6.1 depends on the parameter τ, which controls the width of the time window over which the algorithm searches for suitable pairings for each trade listed in the trade-data file. In order to identify a suitable choice of τ (i.e., a choice that results in a small number of trade-data file events for which Algorithm 6.1 identifies no suitable pairings but that also results in a high proportion of strong pairings among the trade-data file events that are paired) we perform the following two tests.

6.3.1.2 Test 1: Tracking Unique Trade Sizes

To perform the first test, we create two lists for each of the three currency pairs each day. The first list contains all trade sizes reported in the trade-data file. The second list contains all amounts by which the depth was modified due to cancel-order or

Fix $\tau \geq 0$.
for *each trade reported in the trade-data file* **do**
  Denote the trade time by t_i.
  Denote the trade price by p_i.
  Denote the trade size by ω_i.
  Let $j \leftarrow\!\shortmid 0$.
  for *each active order departure reported in the tick-data file* **do**
    if *the departure meets all of the following criteria:*

- *occurs within the time window* $[t_i - \tau, t_i + \tau]$*;*
- *has the opposite sign as* ω_i*;*
- *has size with magnitude less than*[a] *or equal to* $|\omega_i|$*;*
- *occurs at price* p_i*;*

    then
      Let $j \leftarrow\!\shortmid j + 1$.
    end
  end
  if $j = 0$ **then**
    Record that there exists no suitable pairing for the trade.
  end
  else if $j = 1$ **then**
    Record that the trade is *strongly paired*.
  end
  else
    Record that the trade is *weakly paired*.
  end
end

[a] We include depth modifications with size less than $|\omega_i|$ because part of the reported trade size may have matched to hidden liquidity, whereas the tick-data file only lists the portion of the modification that influences visible liquidity (see Section 3.2.2).

Algorithm 6.1: Algorithm to infer which entries in the tick-data file correspond to which entries in the trade-data file.

modify-order events (as reported in the tick-data file). We then identify any sizes that appear exactly once in each list and calculate the difference between the time reported for the corresponding events in the trade-data and tick-data files. Fewer than 5% of all trade-data file events can be paired with tick-data file activity in this way, but such pairings are likely to be correct because the small number of events with unique sizes in each of the two lists makes coincidence improbable. Figure 6.1 shows how the fraction of paired unique-sized market orders varies with τ.

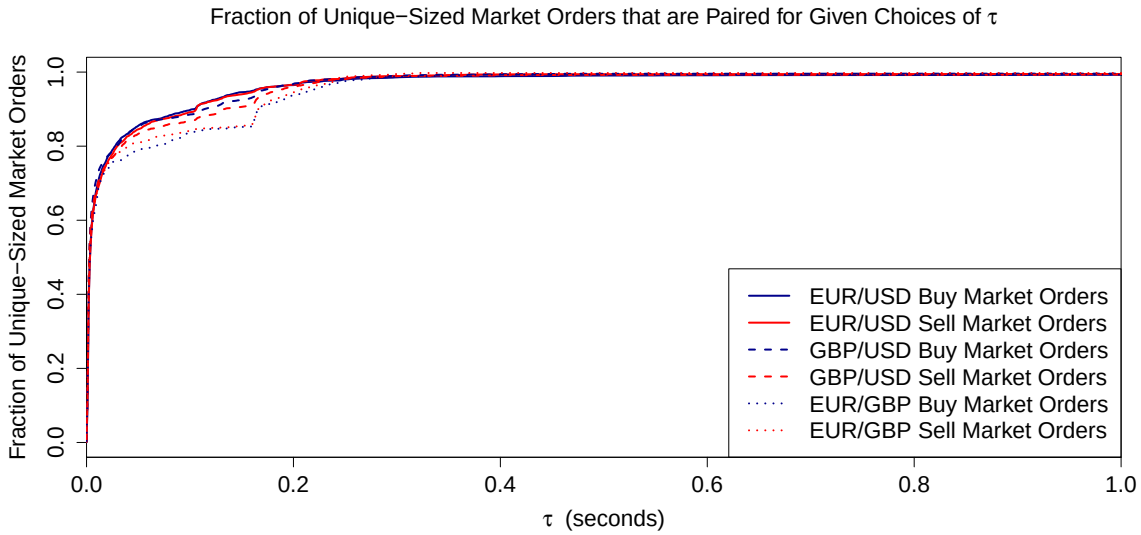

Figure 6.1: Fraction of unique-sized market orders that are paired for given choices of τ for EUR/USD, GBP/USD, and EUR/GBP.

6.3.1.3 Test 2: Fractions of Strong and Weak Pairings

To perform the second test, we compare the fraction of trade-data file events for which Algorithm 6.1 assigns strong and weak pairings for discrete choices of τ. Figure 6.2 shows these fractions for each $\tau \in \{0.1, 0.2, \ldots, 1.0\}$ seconds.

6.3.1.4 Analysis of the Tests

About 80%–90% of each currency pair's unique-size market orders identify their pairing within $\tau = 0.1$ seconds, and about 99% of each currency pair's unique-size market orders identify their pairing within $\tau = 0.3$ seconds (see Figure 6.1). The percentage of trade-data file events for which Algorithm 6.1 identifies no suitable pairings decreases monotonically as τ increases (see Figure 6.2), but choosing too large a value of τ increases the likelihood that pairings are recorded as weak rather than strong,

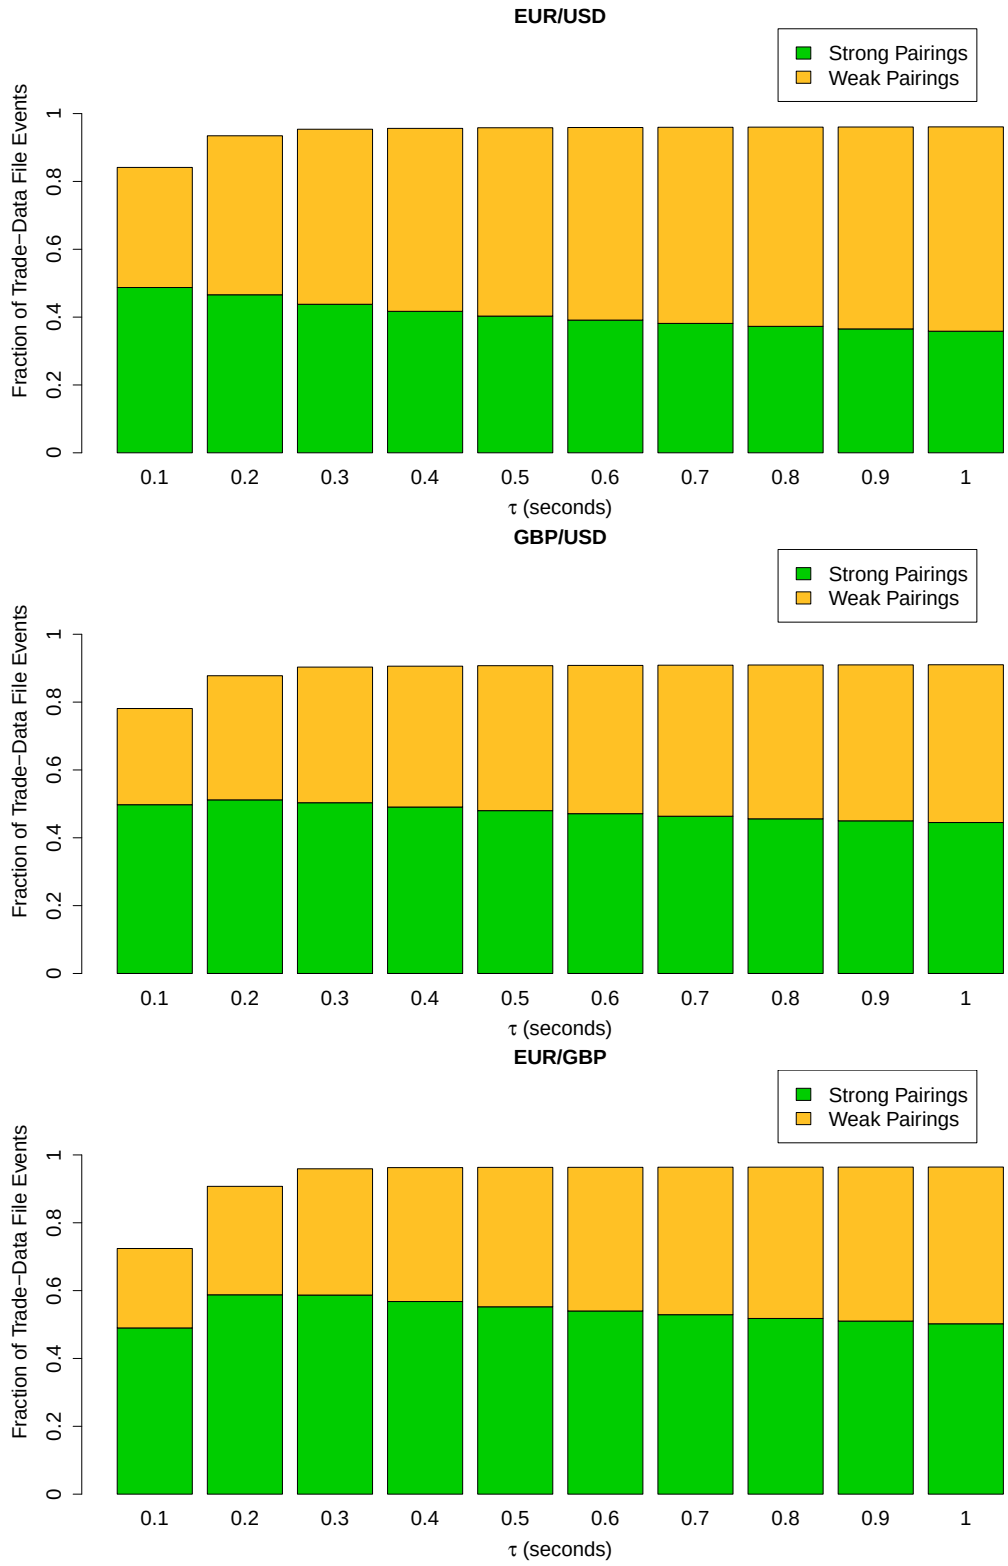

Figure 6.2: Fraction of events in the (top) EUR/USD, (middle) GBP/USD, and (bottom) EUR/GBP trade-data files that Algorithm 6.1 reports as strongly or weakly paired for each $\tau \in \{0.1, 0.2, \ldots, 1.0\}$ seconds.

because a wider range of tick-data file events are considered for pairing to each trade-data file event. Based on these results, we deem the optimal choice of τ to lie within the range $\tau \in [0.1, 0.5]$.

We perform all of our calculations throughout this chapter using three different choices of τ, namely $\tau = 0.1$, $\tau = 0.3$, and $\tau = 0.5$. We present the results for $\tau = 0.3$ in the main text and the results of our main statistical calculations for $\tau = 0.1$ and $\tau = 0.5$ in Appendix A.3. All of our results using $\tau = 0.1$ and $\tau = 0.5$ are qualitatively similar to those for $\tau = 0.3$.

6.3.2 Excluded Entries

Market orders that walk up the book (i.e., that match to multiple active orders across several different prices) pose two conflicting problems for our calculations. First, the partial matchings of such market orders are often separated by additional bypassed active orders. Therefore, it is not appropriate to assign a single bypassing price to the whole trade. Second, some market orders match to a very large number of active orders. The bypassing prices of each such partial matching are correlated, because each subsequent partial matching bypasses all the same active orders as its predecessors. Long sequences of correlated data points generated by a single market order obscure the statistical properties of other orders, so it is not appropriate to regard each partial matching as a separate market order.

In order to avoid these difficulties, we include in our calculations only the first partial matching of each market order. Cleaning the data in this way does not cause us to discard a large number of trade-data file events because relatively few market orders walk up the book (the percentage ranges from 8.54% for EUR/USD sell market orders to 3.93% for EUR/GBP buy market orders — see Table 4.1).

When calculating the price-change series, we also exclude any entries that reflect inter-day price changes (i.e., where the $(i+1)^{\text{th}}$ entry of the price series corresponds to a different trading day than the i^{th} entry).

6.4 Relative Importance

Throughout this chapter, we perform several multiple regressions to identify relationships between the price-change series that we study. The *relative importance* of the input variables in such a multiple regression describe their relative contributions to some specified aspect of the model's output [21, 65]. There are several different

measures of relative importance, and the type of relative importance measured in a given situation depends on the purpose of the analysis [1].

We seek to understand the proportion of the explained variance in the multiple regressions that is attributable to each of the input variables. This measure of relative importance is called the "dispersion importance" [1]. We provide a discussion of relative importance and a formal definition of dispersion importance in Appendix A.1.

Each relative importance calculation that we perform compares the dispersion importance of two input variables. For brevity, we normalize the dispersion importance of the two input variables so that they sum to 1, then report only the normalized relative importance $\tilde{I}_1$ of the first input variable. The normalized relative importance of the second variable is therefore equal to $\tilde{I}_2 = 1 - \tilde{I}_1$.

6.5 Results

6.5.1 Trade-Price Series

Figure 6.3 shows the change-in-traded-bid-price and change-in-traded-ask-price series for EUR/USD, GBP/USD, and EUR/GPB. The trade prices for each of the three currency pairs undergo many large changes, several of which exceed the sample standard deviation by a factor of more than 5 and some of which exceed the sample standard deviation by a factor of more than 10. It is not clear from studying these series in isolation whether the large price changes are due to market-wide volatility spikes that affect all traders' local LOBs, or whether they are due to some traders having poor access to liquidity due to their (lack of) trading partnerships. We return to this question several times throughout the chapter.

6.5.2 Bypassing Prices

Table 6.1 lists summary statistics about the bypassing prices of trades. More than half of the trades that we study have a bypassing price of 0 ticks, and the mean bypassing price for each of the three currency pairs is small. The bypassing price of some market orders is extremely large (particularly for EUR/USD), but the bypassing-price ECDFs and survivor functions[2] (see Figure 6.4) reveal that most market orders match at a price that is close to the best quotes at their time of execution.

[2]Given a set of observations $x_1, x_2, \ldots, x_N$ with ECDF $F(x)$, the survivor function is $1 - F(x) = \frac{1}{N}\sum_{i=1}^{N} \mathbb{I}_{x_i > x}$, where $\mathbb{I}$ is the indicator function.

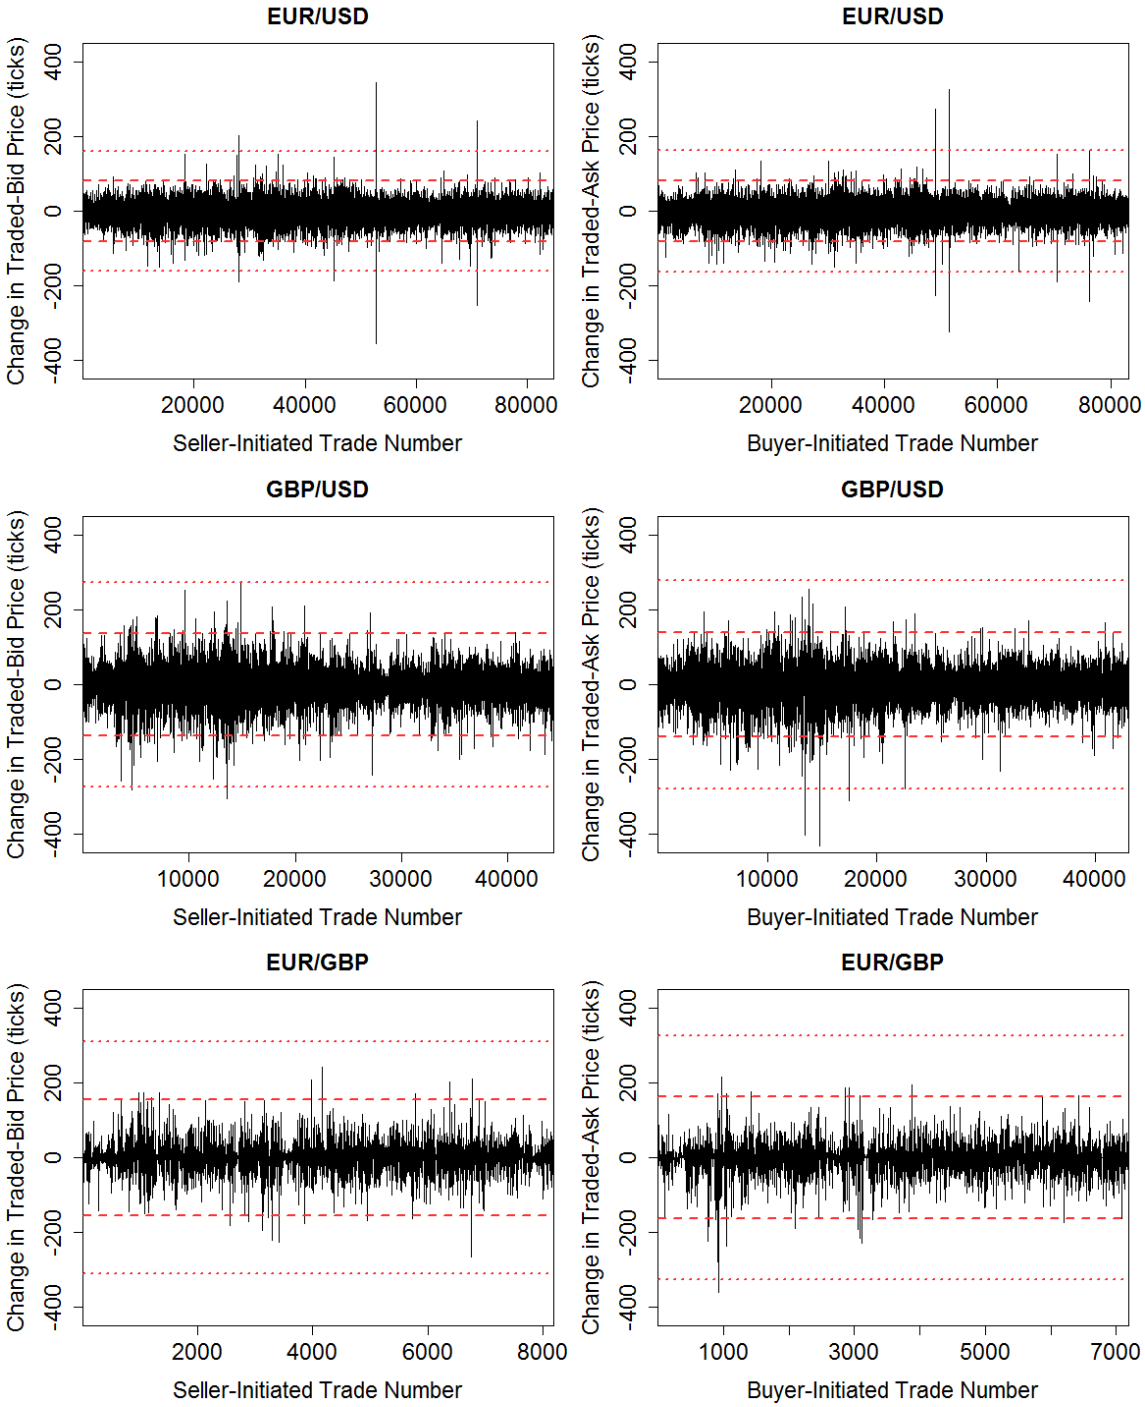

Figure 6.3: Change-in-traded-bid-price and change-in-traded-ask-price series for (top) EUR/USD, (middle) GBP/USD, and (bottom) EUR/GBP. The dashed (respectively, dotted) red lines denote 5 (respectively, 10) standard deviations of each series.

	EUR/USD		GBP/USD		EUR/GBP	
	q_i^B	q_i^A	q_i^B	q_i^A	q_i^B	q_i^A
Mean (ticks)	2.58	2.48	3.19	3.17	2.00	1.80
Standard Deviation (ticks)	5.46	5.92	6.41	6.89	4.30	3.89
Maximum (ticks)	363	335	119	142	73	43
Percentage Non-Zero	44.5%	43.9%	40.2%	37.4%	34.6%	32.6%

Table 6.1: Summary statistics for bid-side bypassing prices q_i^B and ask-side bypassing prices q_i^A for EUR/USD, GBP/USD, and EUR/GBP.

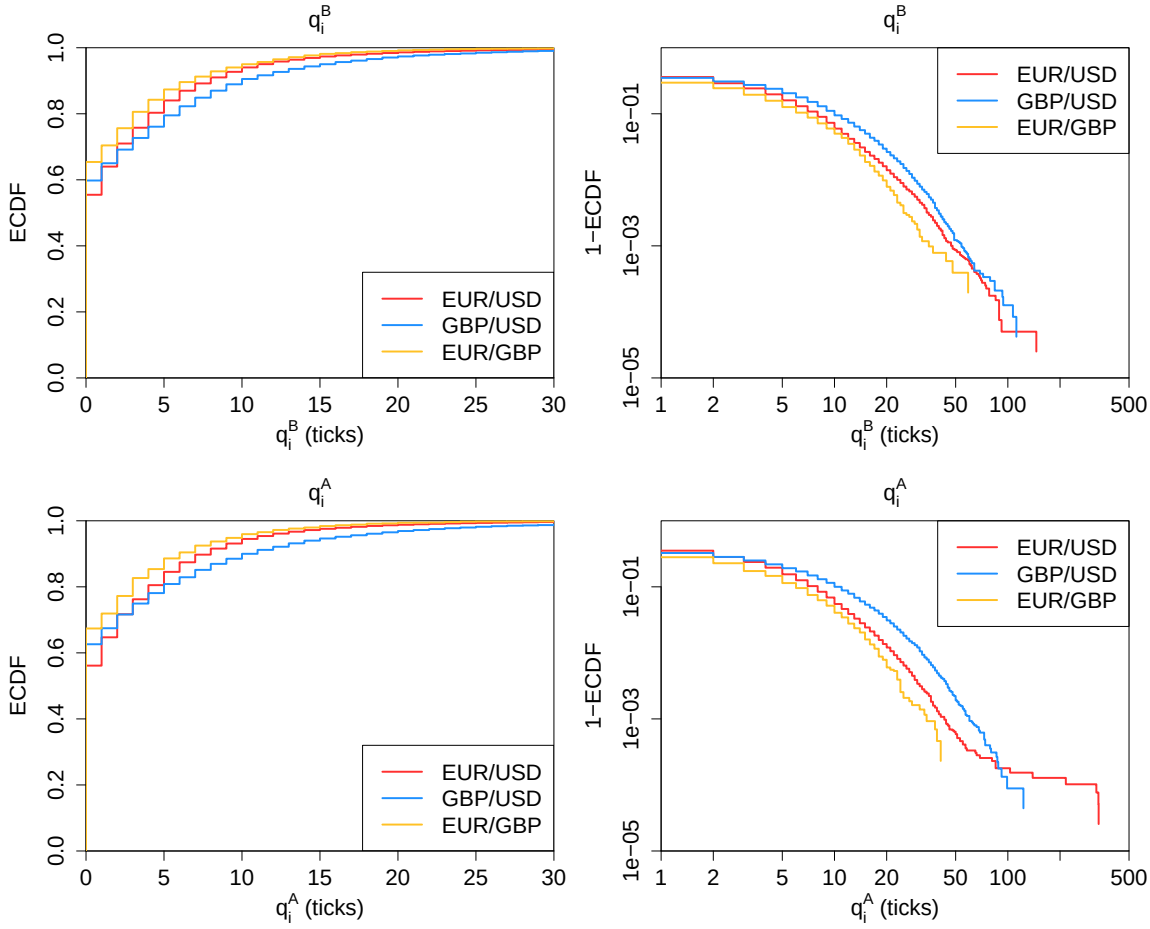

Figure 6.4: (Left panels) ECDFs for (top) bid-bypassing prices q_i^B and (bottom) ask-bypassing prices q_i^A. (Right panels) survivor functions $1 - \text{ECDF}$ for the same series, plotted in doubly logarithmic coordinates.

For each of the three currency pairs, the ECDF of bid-bypassing prices is qualitatively similar to that of ask-bypassing prices, except that EUR/USD has a small number of buyer-initiated trades with extremely large ask-bypassing prices for which there are no corresponding seller-initiated trades with such large bid-bypassing prices. Bypassing prices tend to be smallest for EUR/GBP, slightly larger for EUR/USD, and largest for GBP/USD.

6.5.3 Price-Change Distributions

Table 6.2 shows the percentage of trades that cause non-zero change in trade price, bypassing price, and quote price for each of the three currency pairs. Non-zero price changes are more common in the trade-price and quote-price series than they are in the bypassing-price series. Figure 6.5 shows the ECDFs for changes in trade price, bypassing price, and quote price for each of the three currency pairs. Each bid-side ECDF is qualitatively similar to the corresponding ask-side ECDF, but there are several interesting differences between the different price-change series and the different currency pairs.

	EUR/USD		GBP/USD		EUR/GBP	
	Bid Side	Ask Side	Bid Side	Ask Side	Bid Side	Ask Side
Trade Price	78.7%	79.4%	79.5%	81.1%	76.3%	81.4%
Bypassing Price	55.0%	54.7%	51.6%	48.9%	44.0%	45.2%
Quote Price	77.5%	77.8%	80.1%	80.8%	75.4%	81.0%

Table 6.2: Percentage of non-zero entries in the change-in-trade-price, change-in-bypassing-price, and change-in-quote-price series for EUR/USD, GBP/USD, and EUR/GBP.

To help identify similarities between changes in trade price, bypassing price, and quote price, we construct quantile-quantile (Q-Q) plots to compare the three price-change distributions (see Figure 6.6 for seller-initiated trades; results for buyer-initiated trades are similar). For each of the three currency pairs, the distribution of changes in trade price is similar to that of changes in quote price. By contrast, the magnitude of changes in bypassing price tends to be smaller than the magnitude of changes in the other two series.

We also normalize each price-change series by subtracting the sample mean and dividing by the sample standard deviation. Figure 6.7 shows the ECDFs for normalized changes in traded bid price, bid-bypassing price, and bid price (results for buyer-initiated trades are similar). In all cases, the normalized distributions have

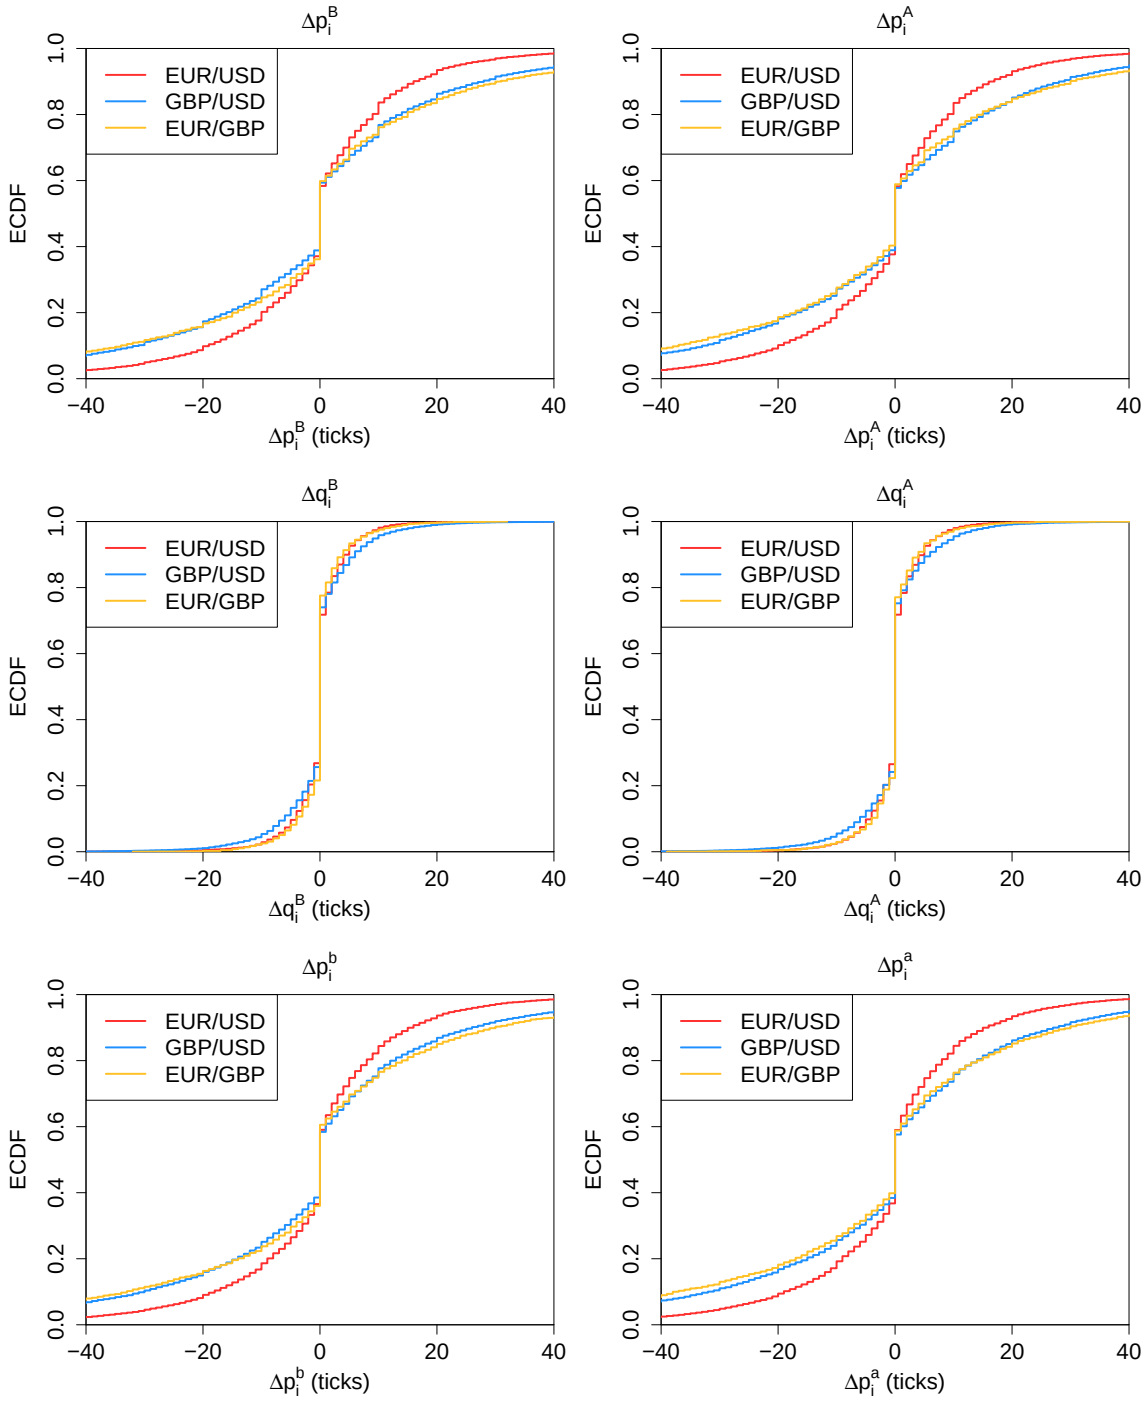

Figure 6.5: ECDFs for (top left) changes in traded bid price Δp_i^B, (top right) changes in traded ask price Δp_i^A, (middle left) changes in bid-bypassing price Δq_i^B, (middle right) changes in ask-bypassing price Δq_i^A, (bottom left) changes in bid price Δp_i^b, and (bottom right) changes in ask price Δp_i^a.

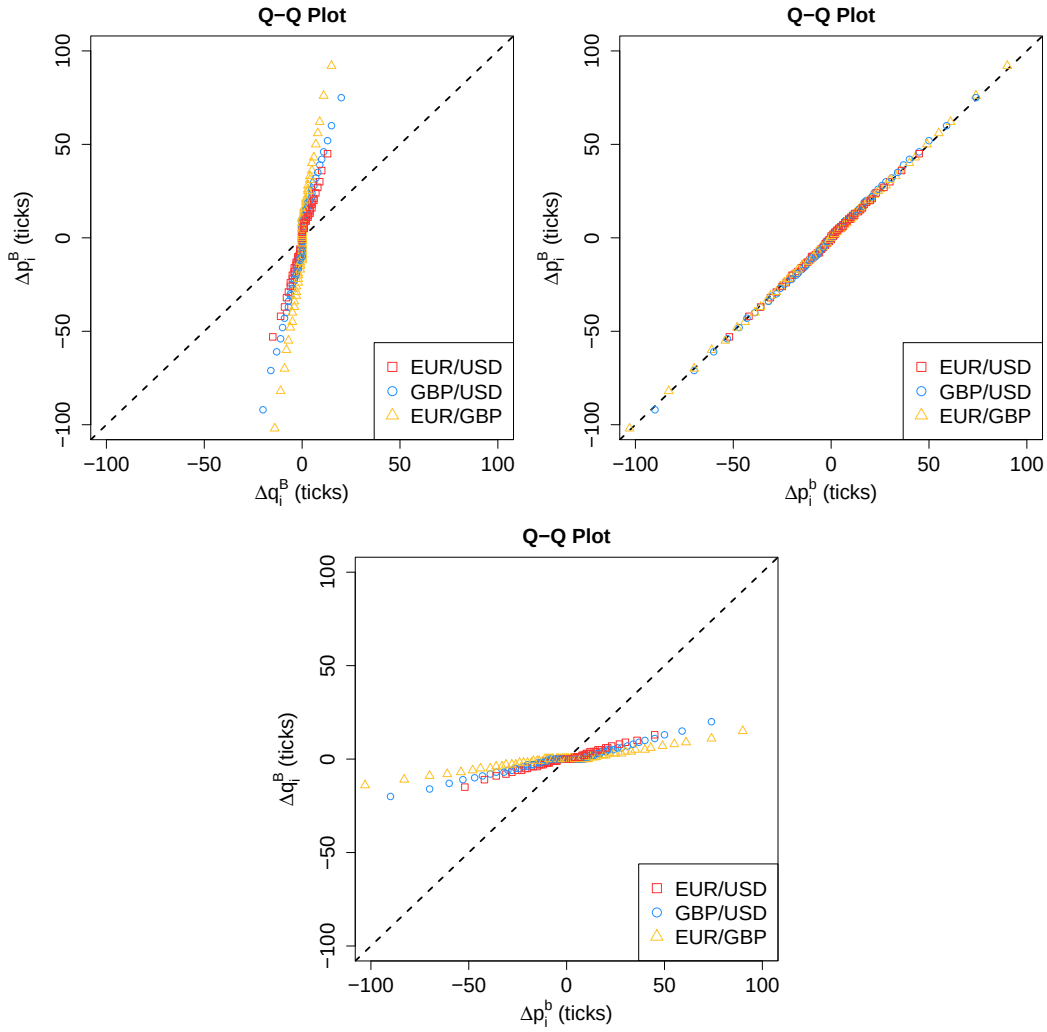

Figure 6.6: Quantile-quantile (Q-Q) plots for (top left) changes in traded bid price Δp_i^B against changes in bid-bypassing price Δq_i^B, (top right) changes in traded bid price Δp_i^B against changes in bid price Δp_i^b, and (bottom) changes in bid-bypassing price Δq_i^B against changes in bid price Δp_i^b. The points denote the $0.01, 0.02, \ldots, 0.99$ quantiles of the empirical distributions. The dashed black lines denote the diagonal. The corresponding plots for buyer-initiated trades are qualitatively similar.

greater mass close to 0 and heavier tails than a standard normal distribution. The distribution of normalized changes in bypassing price has heavier tails than the other two normalized distributions.

6.5.4 Relative Importance of Price Changes

In this section, we calculate the dispersion importance (see Section 6.4) of the input variables Δq_i^B, Δp_i^b, Δq_i^A, and Δp_i^a in the multiple regressions

$$\hat{y}_i^B = \beta_0 + \beta_1 \Delta q_i^B + \beta_2 \Delta p_i^b, \qquad i = 1, 2, \ldots, N^B - 1, \tag{6.11}$$

$$\hat{y}_i^A = \alpha_0 + \alpha_1 \Delta q_i^A + \alpha_2 \Delta p_i^a, \qquad i = 1, 2, \ldots, N^A - 1, \tag{6.12}$$

where $\hat{y}_i^B$ and $\hat{y}_i^A$ are the predictions of the output variables Δp_i^B and Δp_i^A, respectively, and $\beta_0, \beta_1, \beta_2, \alpha_0, \alpha_1, \alpha_2 \in \mathbb{R}$ are the least-squares parameter estimates.[3]

For each seller-initiated trade $i = 1, 2, \ldots, N^B - 1$, observe that

$$\begin{aligned} \Delta q_i^B + \Delta p_i^b &= p_{i+1}^b - p_{i+1}^B - p_i^b + p_i^B + p_i^b - p_{i+1}^b, \\ &= p_i^B - p_{i+1}^B, \\ &= \Delta p_i^B. \end{aligned}$$

Similarly, for each buyer-initiated trade $i = 1, 2, \ldots, N^A - 1$,

$$\begin{aligned} \Delta q_i^A + \Delta p_i^a &= p_{i+1}^A - p_{i+1}^a - p_i^A + p_i^a + p_{i+1}^a - p_i^a, \\ &= p_{i+1}^A - p_i^A, \\ &= \Delta p_i^A. \end{aligned}$$

Therefore, the regression equation (6.11) has the exact solution

$$\beta_0 = 0, \quad \beta_1 = \beta_2 = 1 \tag{6.13}$$

for all $i = 1, 2, \ldots, N^B$, and the regression equation (6.12) has the exact solution

$$\alpha_0 = 0, \quad \alpha_1 = \alpha_2 = 1 \tag{6.14}$$

for all $i = 1, 2, \ldots, N^A$. This implies that both multiple regressions (6.11) and (6.12) have a coefficient of determination $R^2 = 1$.

In order to calculate the dispersion importance, we first perform separate ordinary least-squares (OLS) regressions between each pair of price-change series (see Figure 6.8). Table 6.3 lists the coefficient of determination R^2 for each of these types of

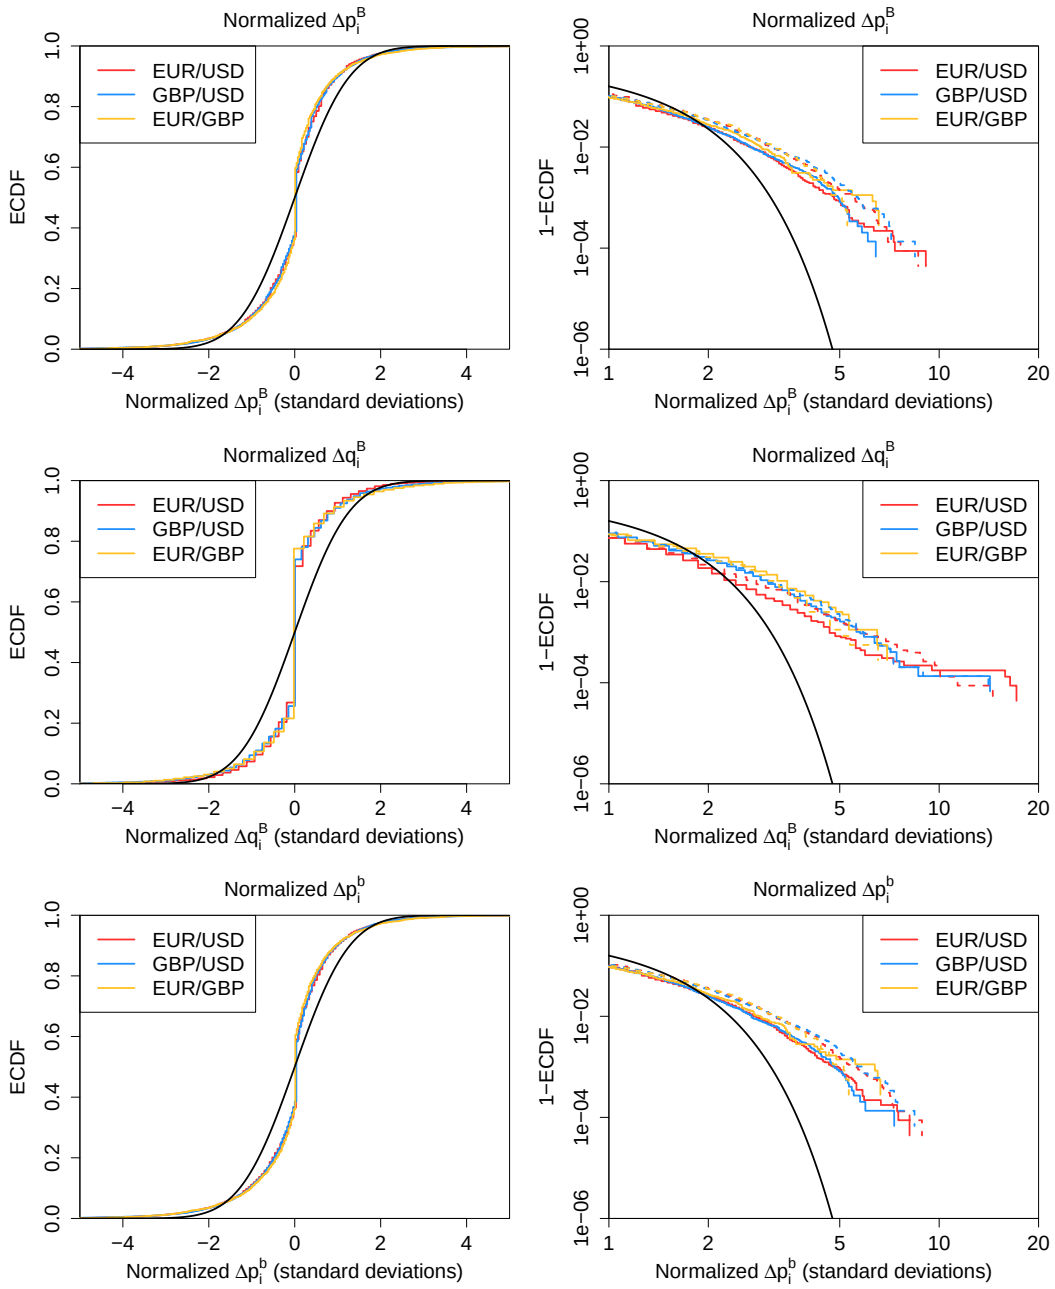

Figure 6.7: (Left panels) ECDFs for (top) changes in traded bid price Δp_i^B, (middle) changes in bid-bypassing price Δq_i^B, and (bottom) changes in bid price Δp_i^b, after first normalizing the data by subtracting its mean and dividing by its standard deviation. (Right panels) survivor functions $1 - \text{ECDF}$ for the same series, plotted in doubly logarithmic coordinates. In the right panels, the solid lines denote the normalized data (i.e., they show the upper tails of the normalized distribution) and the dashed lines denote the normalized data multiplied by -1 (i.e., they show the lower tails of the normalized distribution). The solid black lines denote the standard normal distribution. The corresponding plots for buyer-initiated trades are qualitatively similar in each case.

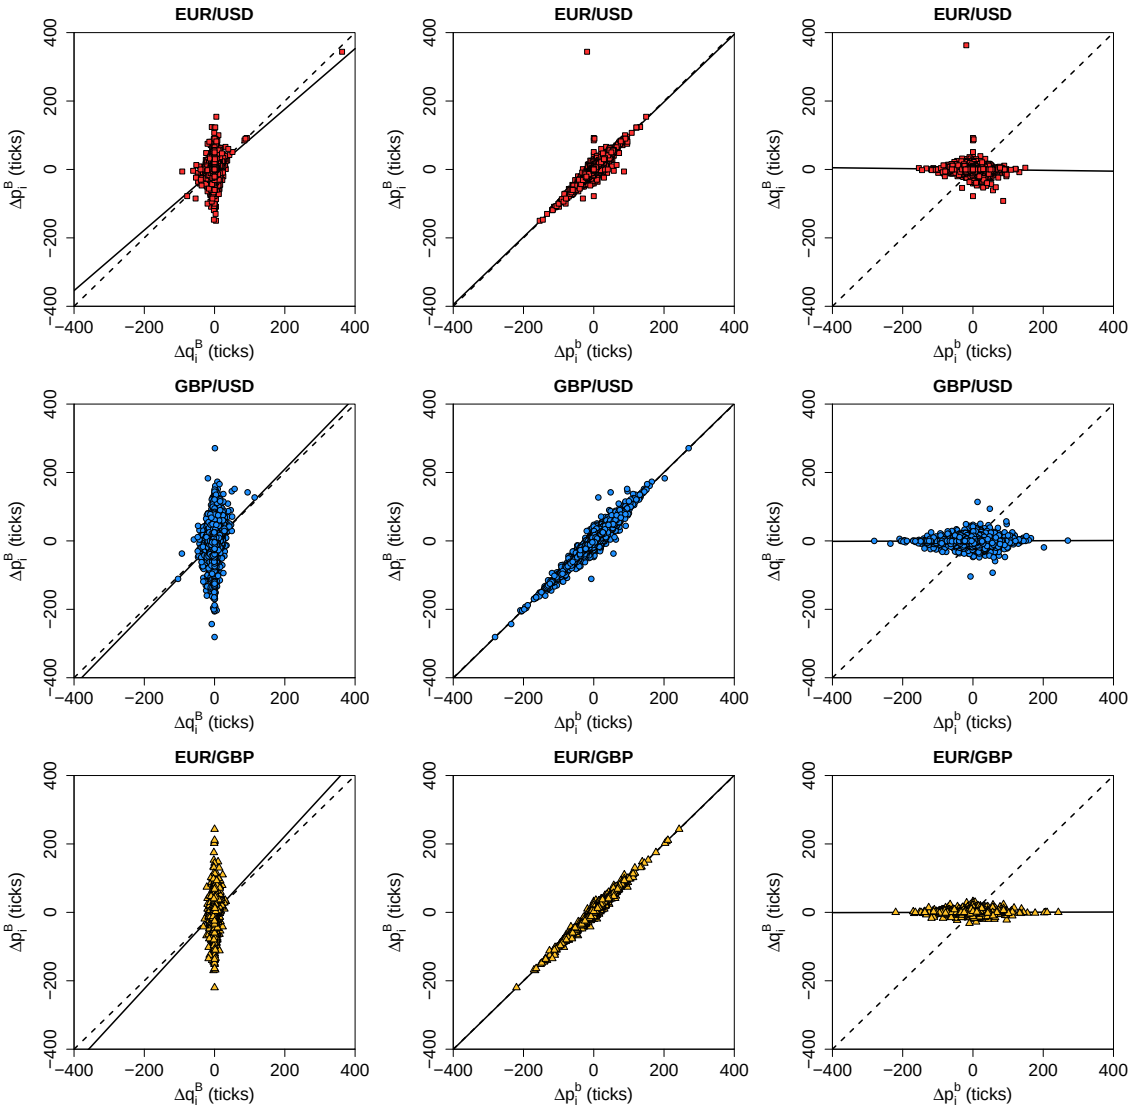

Figure 6.8: Scatter plots of (left panels) changes in traded bid price Δp_i^B against changes in bid-bypassing price Δq_i^B, (middle panels) changes in traded bid price Δp_i^B against changes in bid price Δp_i^b, and (right panels) changes in bid-bypassing price Δq_i^B against changes in bid price Δp_i^b for (top) EUR/USD, (middle) GBP/USD, and (bottom) EUR/GBP. In each plot, the dashed black line denotes the diagonal and the solid black line denotes the ordinary least-squares regression line for the data. The corresponding plots for the ask side are qualitatively similar.

	EUR/USD		GBP/USD		EUR/GBP	
	Bid side	Ask side	Bid side	Ask side	Bid side	Ask side
Multiple Regression	1	1	1	1	1	1
First OLS Regression	0.078	0.087	0.059	0.048	0.022	0.017
Second OLS Regression	0.900	0.885	0.947	0.949	0.982	0.978
Third OLS Regression	0.001	0.002	0.000	0.000	0.000	0.000

Table 6.3: Coefficients of determination R^2 for regressions for EUR/USD, GBP/USD, and EUR/GBP price changes. The multiple regressions are given by Equations (6.11) and (6.12). The first OLS regressions are for change in trade price onto change in bypassing price (see left panels of Figure 6.8). The second OLS regressions are for change in trade price onto change in quote price (see middle panels of Figure 6.8). The third OLS regressions are for change in bypassing price onto change in quote price (see right panels of Figure 6.8).

regressions. The R^2s for the regressions of changes in quote price onto changes in trade price are much larger than those for the other two types of regression.

Table 6.4 lists the normalized dispersion importance $\tilde{I}_1$ of the change in bypassing price in the multiple regressions (6.11) and (6.12). We use two different techniques to estimate the standard error of $\tilde{I}_1$. First, we use a standard, non-parametric bootstrap [118]. Second, we use a heuristic called the variance-plot method [18, 78], which is similar to the bootstrap but which takes into account possible serial correlations in the data. We provide details of both of these techniques in Appendix A.2.

	EUR/USD		GBP/USD		EUR/GBP	
	Bid side	Ask side	Bid side	Ask side	Bid side	Ask side
$\tilde{I}_1$	0.089	0.101	0.056	0.050	0.020	0.019
Bootstrap s.e.	0.013	0.012	0.002	0.002	0.002	0.003
Variance-Plot s.e.	0.022	0.027	0.008	0.007	0.004	0.005

Table 6.4: Normalized dispersion importance $\tilde{I}_1$ of change in bypassing price for change in trade price, for EUR/USD, GBP/USD, and EUR/GBP. The normalized dispersion importance of change in quote price is given by $1 - \tilde{I}_1$ in each case. The second row of the table denotes the standard deviation of 1000 bootstrap samples of $\tilde{I}_1$. The third row of the table denotes the variance-plot estimate of the standard error of $\tilde{I}_1$ (see Appendix A.2.2).

For both buyer-initiated and seller-initiated trades and for each of the three currency pairs, the estimates of $\tilde{I}_1$ are small. Therefore, changes in bypassing price

[3]For a detailed discussion of multiple regressions, see [24].

explain a much smaller fraction of the variance of changes in trade price than do changes in quote price. In all cases, the variance-plot estimates of the standard errors are larger than the corresponding bootstrap estimates, which suggests that the data contains serial correlations that are not reflected in the the bootstrap estimates. Nevertheless, even the variance-plot estimates of the standard errors are small, so we are able to draw our conclusion about relative importance with a high degree of confidence.

6.6 Discussion

Although more than half of all trades that we study occur at the best quotes (see Table 6.1), the bypassing prices of the remaining trades vary over several orders of magnitude (see Figure 6.4). This suggests that some traders on Hotspot FX have access to relatively few trading opportunities on the platform.

There are two possible reasons that a trade could have a large bypassing price. First, the owner of the market order involved in the trade could have few trading partners, causing their order to bypass active orders across many prices and match deep into the LOB. Second, a trader with few trading partners could place an active buy order far above the previous bid price (respectively, an active sell order far below the previous ask price) without receiving a matching, then a subsequent sell (respectively, buy) market order from another trader could bypass this active order and match at a price near the previous bid (respectively, ask) price. Figure 6.9 shows a schematic of these two possibilities.

It is difficult to disentangle these effects for the large number of trades that cause small changes in trade price, bypassing price, or quote price, because it is common for these series to fluctuate slightly between subsequent market order arrivals due to the small changes in $\mathcal{L}(t)$ that occur frequently throughout the trading day. However, we are able to study more clearly the trades that cause large price changes. If the second scenario arose frequently on Hotspot FX, then some of the trades that cause a large change in bypassing price would also cause a large change in quote price. We find no evidence of any such trades for any of the three currency pairs (see right panels of Figure 6.8). Therefore, we conclude that trades with large bypassing price are typically caused by the arrival of an extreme-priced market order whose owner has few trading partners.

The vast majority of market orders that we study match at or near the best quotes at their time of execution (see Section 6.5.2). Moreover, changes in quote

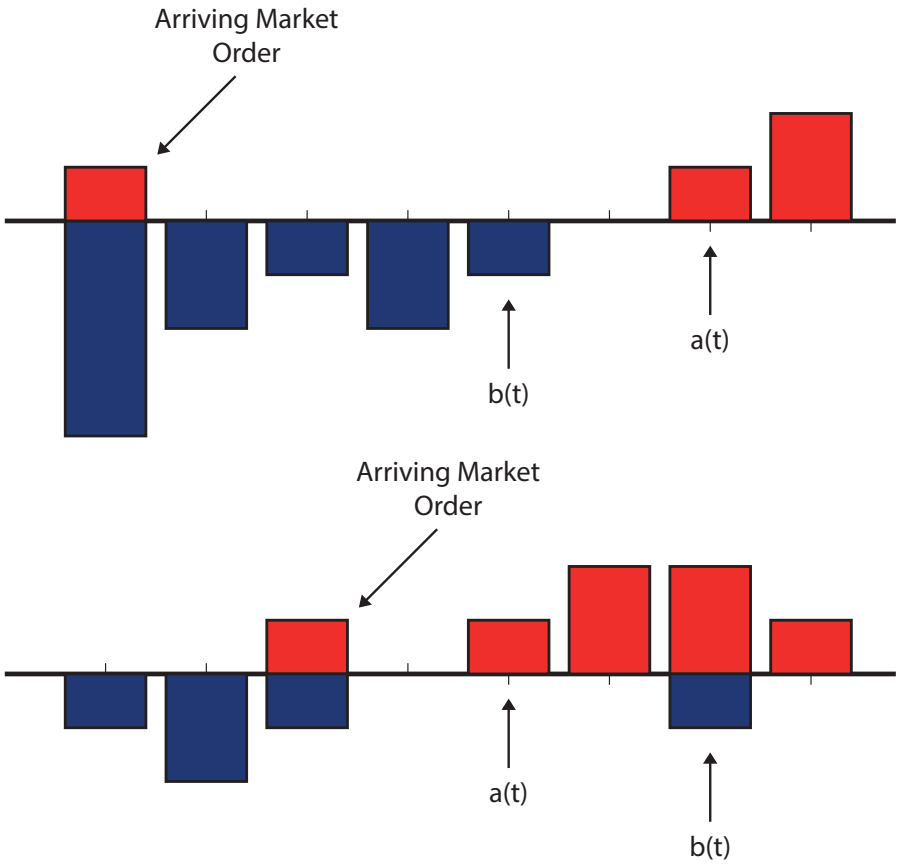

Figure 6.9: Schematic to illustrate the possible causes of large bypassing prices. (Top) An incoming sell market order owned by a trader with few trading partners bypasses active buy orders across many different prices and matches deep into the LOB. (Bottom) A trader with few trading partners places an active buy order far above the previous bid price, then a subsequent sell market order from another trader matches at a price near the previous bid.

price account for a much higher fraction of the variance of changes in trade price than do changes in bypassing price (see Table 6.4). These findings imply that the heterogeneity in traders' access to trading opportunities on Hotspot FX has a small aggregate impact on price formation. However, we also identify a handful of market orders that have very large bypassing prices (see Figure 6.4).

Although traders are not able calculate the exact bypassing prices of their trades, they are able to infer whether they are trading at worse prices than those available to their better-connected peers by monitoring the trade-price series. This prompts the question of why some traders submit market orders with large bypassing prices when they know this to be the case. We believe that the answer lies in fact that Hotspot FX provides direct market access to a wide variety of different financial institutions, many of whom do not have access to the traditional inter-dealer market. For example, small financial institutions such as hedge funds are not able to trade on the Reuters Dealing 3000 or EBS platforms (see Chapter 1), nor do they have access to inter-dealer voice-broking or direct-trading services [43]. Therefore, they may regard the relatively large bypassing prices that they experience on Hotspot FX to be an acceptable cost of trading, particularly if the bypassing prices are smaller than the costs that they would incur if they traded via a different mechanism, such as a retail aggregator platform [14].

An alternative possibility is that market orders with large bypassing prices originate from traders who implement complicated strategies that involve the submission of several different orders. A popular example of one such strategy is a volume-weighted average price (VWAP) strategy [151], in which a trader submits several different orders to achieve a target average price. Traders implementing a VWAP strategy may not scrutinize all constituent orders that they submit, which could cause a small fraction of their trading volume to match at an unfavourable price.

6.7 Summary and Future Work

In this chapter, we have investigated how the heterogeneity in traders' access to trading opportunities on Hotspot FX influences the price of trades on the platform. We decomposed each change in trade price into two constituent parts: the change in bypassing price and the change in quote price. We found that although some trades bypass active orders across many prices and match deep into the LOB, the vast majority of arriving market orders match at or close to the best quotes at their time of execution. Therefore, the vast majority of trades are conducted by traders

who have sufficiently many trading partners on the platform to experience little or no impact on the price, compared to the best quotes.

We found that changes in bypassing price had much lower relative importance for explaining the variance of changes in trade price than did changes in quote price. However, we also identified a small number of trades that had a large bypassing price, which we argued could be regarded as a substantial cost to the owners of the relevant market orders. We conjectured that such market orders were likely to have originated from financial institutions who do not have direct access to the traditional inter-dealer market, and who therefore regard the high bypassing price of their trades as an acceptable cost of trading.

6.7.1 Future Work

Our findings suggest several possible avenues for future research.

- **Understanding large price movements:** Several empirical studies from a wide range of different markets have reported that the tails of the distributions of price changes decay like a power law with exponent $\alpha \approx 3$ [64, 94, 95, 123, 144, 145]. To our knowledge, however, all empirical publications on this topic have studied data from centralized markets. Assessing whether this is also the case on Hotspot FX would be an interesting task, because many of the largest price movements in the Hotspot FX data are caused by trades that have a particularly large bypassing price, whereas all trades in centralized markets have a bypassing price of 0. Understanding the functional form of the tails of price-change distributions is important to both investors and regulators, because risk calculations and hedging strategies depend on the probability of large price changes.

- **Investigating other currency pairs:** Our data describes only the EUR/USD, GBP/USD, and EUR/GBP LOBs, all of which have large numbers of trades each day. It is plausible that traders of less heavily traded currency pairs could have far fewer trading partners than traders of these three currency pairs, which could cause bypassing prices to play a more prominent role in the prices of trades.

- **Understanding autocorrelations:** Due to the structure of the Hotspot FX data, we use Algorithm 6.1 to infer which updates to $\mathcal{L}(t)$ correspond to which trades. In order to minimize the error associated with incorrect pairings, we

study only the entries of the quote-price and bypassing-price series (and the corresponding changes in these series) for which Algorithm 6.1 reports a strong pairing. Therefore, it is not possible for us to study the autocorrelation properties of these series due to their many missing entries. Comparing the temporal properties of bypassing-price and quote-price series to those of trade-price series could help further illuminate their interactions.

Chapter 7

An Agent-Based Model of Trade on a Quasi-Centralized Trading Platform

In this chapter, we introduce an agent-based model of trade on a quasi-centralized trading platform[1] and investigate the volatility of the trade-price series that the model generates. The work in this chapter is also described in the jointly authored paper [97].

7.1 Introduction

Loosely speaking, "volatility" is a measure of the variability of the price of a traded asset [15]. All else held equal, a high-volatility asset is expected to undergo larger price changes than a low-volatility asset during a given time interval . Volatility is an important consideration for market practitioners, as it often plays a central role in optimal portfolio construction [163], forecasting future price movements [26], and the pricing of derivatives contracts [180].

The measurement and modelling of volatility in financial markets have remained active areas of research for several decades (see [15, 180, 184] for recent surveys). The literature contains a wide selection of different volatility estimators, and the exact form of volatility studied in a given situation often depends on both the data available and the purpose of performing the estimation [190]. The possible sources of volatility in financial markets have been discussed extensively; common topics in

[1]Recall from Chapter 2 that a quasi-centralized trading platform is an electronic trading platform on which each trader can only trade with a specified set of trading partners but in which the price of each trade is announced publicly.

this area include macroeconomic or political activity [63], news arrivals [124], and the market impact of trades [34, 204].

On a quasi-centralized trading platform, each trader can only access the trading opportunities that are offered by his/her trading partners. As a result, the volatility that traders on such a platform experience depends on the network of trading partnerships between them. Understanding how a trade-partnership network's structure influences volatility is an important task, because it could provide insight into several important properties of price formation and illuminate the origins and nature of large price changes. However, very little is known about trade-partnership networks on real trading platforms, so empirical study of historical market data does not provide clear insight into this problem.

In this chapter, we introduce an agent-based model (ABM) of trade on a quasi-centralized trading platform to investigate this question in an artificial market. Our model's temporal evolution is governed by the agents' trade-partnership network and a trade-price-reversion parameter, which controls how agents regard the previous trade price when deciding how to act. We design our ABM to be sufficiently simple to enable us to derive several analytical results about its behaviour. We simulate our ABM using different trade-partnership networks and different values of the trade-price-reversion parameter, and we calculate three different measures of volatility from the resulting trade-price series.

We find that volatility in trade prices is lowest when the trade-partnership network is complete (which corresponds to a centralized trading platform, where all traders are able to trade with all others). In most cases, we find that reducing the total number of edges in the trade-partnership network leads to an increase in all three volatility measures that we study, but we also uncover some network configurations for which this is not the case. This implies that the topology of the network, and not just the density of edges, influences price formation in the model. We find that particularly small choices of the trade-price-reversion parameter tend to produce large trade-price volatility. We note that the case in which the trade-price-reversion parameter equals zero corresponds to a situation in which traders are not able to view the trade-price series in real time, and we thereby argue that providing this information to traders constitutes a strong stabilizing force.

The remainder of this chapter is organized as follows. In Section 7.2, we discuss volatility on quasi-centralized trading platforms and introduce the three volatility measures that we study throughout the chapter. In Section 7.3, we discuss several other ABMs of price formation in financial markets. In Section 7.4, we introduce

our ABM and discuss several of its features. In Section 7.5, we derive several results about the temporal evolution of our model. We present our simulation results in Section 7.6 and discuss these results in Section 7.7. In Section 7.8, we summarize our findings and discuss several possible avenues for future research.

7.2 Volatility on A Quasi-Centralized Trading Platform

In contrast to centralized trading platforms, in which all traders face the same trading opportunities, the quasi-centralized[2] structure of multi-institution trading platforms in the FX spot market causes different traders to face different trading opportunities at the same time. This heterogeneity in access to trading opportunities makes it difficult to estimate (or even define) volatility on such a platform, because the volatility that each trader experiences depends on his/her trading partnerships.

We propose two possible approaches to this issue: to calculate an individual volatility estimate for each trader or to calculate a single volatility estimate for a price series that is common to all traders. By comparing the volatility that different traders experience, the first approach may help to quantify how a trade-partnership network's structure influences local price variability. However, this approach involves studying a high-dimensional, computationally intensive problem for all but small numbers of traders, and empirical investigations using this approach would require knowledge of both the full trade-partnership network and the identities of the traders associated with specific market activity. The second approach is much less computationally intensive, but the market-wide volatility estimate may not reflect the volatility that individual traders experience.

These two approaches lead to different statistical and modelling frameworks, both of which have advantages and drawbacks. Given the large number of traders in the FX spot market (see Section 7.6.3), detailed study of local volatility measures would require extremely intensive computations that involve large numbers of interacting components. Therefore, we adopt the second approach. Specifically, we design and implement an ABM of trade on a quasi-centralized trading platform, calculate several statistics related to the variability of the trade-price series that it generates, and

[2]Recall from Chapter 2 that a quasi-centralized trading platform is an electronic trading platform in which each trader can only access trading opportunities that are offered by his/her trading partners but in which the price of each trade is announced publicly.

perform direct comparisons between the values of these statistics for different trade-partnership networks.

In contrast to our calculations in the previous chapters, our model does not distinguish between buyer-initiated and seller-initiated trades. In this chapter, we therefore study all trades together, irrespective of their direction. We write $\phi(t)$ to denote the price of the most recent trade at or before time t. All traders on a quasi-centralized trading platform are able to calculate $\phi(t)$ in real time, so their perceptions of the price variability of a given currency pair may be influenced by their observations of this series.

There are several other time series (such as the global-mid-price series $m(t)$, the global-bid-price series $b(t)$, and the global-ask-price series $a(t)$) that we could use to study volatility. However, our results in Chapter 4 suggest that some traders submit and rapidly cancel large numbers of orders near the best quotes, and thereby cause the values of $m(t)$, $b(t)$, and $a(t)$ to fluctuate rapidly without any trades occurring. These fluctuations could result in misleading volatility estimates. Moreover, traders on a quasi-centralized trading platform do not know the values of $m(t)$, $b(t)$, or $a(t)$ in real time, so their perceptions of price variability cannot be influenced by observations of these series.

7.2.1 Our Volatility Estimates

In this section, we describe the statistics that we calculate to asses the volatility of the trade-price series. Let

$$\{P_t\} = \{P_1, P_2, \ldots\} \tag{7.1}$$

denote the times series of trade prices, and let

$$\begin{aligned} \{D_t\} &= \{D_1, D_2, \ldots\}, \\ D_i &= |P_{i+1} - P_i|, \ i = 1, 2, \ldots, \end{aligned} \tag{7.2}$$

denote the time series of absolute price differences between successive trades. Given a finite-length realization $\{d_1, d_2, \ldots, d_N\}$ of $\{D_t\}$, we calculate the following three statistics:

- The *mean absolute price change*

$$v_1 = \frac{1}{N} \sum_{i=1}^{N} d_i, \tag{7.3}$$

- The *standard deviation of absolute price changes*

$$v_2 = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(d_i - v_1)^2}, \tag{7.4}$$

- The *percentage of non-zero price changes*

$$v_3 = \frac{100}{N}\sum_{i=1}^{N}\mathbb{I}_{d_i \neq 0}, \tag{7.5}$$

  where $\mathbb{I}_{d_i \neq 0}$ is the indicator function

$$\mathbb{I}_{d_i \neq 0} = \begin{cases} 0, \text{ if } d_i = 0, \\ 1, \text{ otherwise,} \end{cases} \qquad i = 1, 2, \ldots, N. \tag{7.6}$$

Our empirical results in Chapter 6 suggest that ECDFs of absolute price changes on real quasi-centralized trading platforms exhibit a roughly exponential shape (see Figure 6.5). The statistics v_1, v_2, and v_3 describe, respectively, the first and second moments and the vertical intercept of such distributions, and thereby facilitate simple comparisons between important statistical properties of the trade-price volatility generated by our model using different trade-partnership networks.

7.3 Agent-Based Models of Price Formation

An *agent-based model (ABM)* [28, 46, 49] is a model in which a large number of possibly heterogeneous *agents* interact in a specified way. For a given choice of input parameters, a single run of an ABM generates a (possibly stochastic) realization of the model's output, according to the rules governing the behaviour and interactions of the individual agents.

Several important features of ABMs make them attractive for modelling price formation in financial markets. First, ABMs do not require assumptions regarding market stability or efficiency. Second, they do not require assumptions that price series are ergodic or converge to some form of equilibrium as $t \to \infty$ [76]. Third, ABMs enable explicit specification of heterogeneity in agents' behaviour, and can thereby provide insight into the consequences of complex, heterogeneous interactions between agents [38]. Fourth, ABMs facilitate study of both the behaviour of individual agents and the emergent, system-level properties. Therefore, ABMs of price formation facilitate quantitative comparisons between different market scenarios [76]. Fifth, agents'

strategies can mimic complex, real-world behaviours that are difficult to study using other types of model [49].

During the past 50 years, the study of ABMs has helped to illuminate many aspects of price formation in financial markets. The literature on this topic is very large, and several recent surveys have discussed specific models in detail. Chakraborti *et al.* [46] presented a chronological record of major developments in the field, from the earliest ABMs of financial markets in the 1960s to the market-microstructure models that have become popular in recent years. Chen *et al.* [49] presented a taxonomy of ABMs, reviewed a wide selection of ABMs of financial markets, and discussed which existing models are able to reproduce several widely observed stylized facts [55]. In our recent survey [98], we discussed how ABMs provide insight into several important elements of trading via LOBs. In this section, we review a small selection of ABMs most relevant to the present work.

7.3.1 Agent Heterogeneity

Several publications have addressed how heterogeneity between agents' actions affects price formation. Some authors (see, e.g., [10, 146]) have studied ABMs with a small number of different agent types and investigated how the statistical properties of price formation vary according to the fraction of each agent type. An ABM with k different types of agents is often called a *k-type ABM* [49]. For example, Bak *et al.* [10] studied how the fraction of "rational traders" (who seek to optimize personal utility functions) and "noise traders" (who randomly update their private valuations of the asset) affect the statistical properties of the trade-price series in a 2-type ABM. Other authors (see, e.g., [69, 134]) have studied ABMs in which all agents are different. For example, LeBaron [134] studied an ABM in which each agent decides how to act based on m previous observations of the price series, where m varies across agents.

Cont and Bouchaud [58] studied an ABM in which agent interactions are heterogeneous. Specifically, they assumed that agents form groups (which the authors modelled using a random graph [27]) and that each agent aligns his/her decisions to buy or sell according to the actions of his/her group. The authors provided a quantitative link between the strength of agents' herd behaviour and the price-change distribution that the model generates.

7.3.2 Price Formation in the FX Spot Market

Perraudin and Vitale [179] studied how knowledge of other agents' actions affects price formation in a highly stylized ABM of the FX spot market. In their model, each agent is either a customer or a dealer (see Section 1.1). A specified fraction of customers are informed about the fundamental value of the currency pair; the remaining customers and all dealers are not. The authors compared the temporal evolution of the model in two different scenarios: one in which information about all dealers' trades is available to all other dealers, and one in which it is not. They reported that their model generates wider bid-ask spreads and higher volatility when dealers do not have access to the information.

7.3.3 Perfect Rationality Versus Zero Intelligence

Different ABMs adopt different modelling assumptions with regard to the actions of individual agents. Some authors assume that agents are *perfectly rational* decision-makers who seek to achieve trading goals or maximize utility [80, 131, 177, 187]. These models help to identify possible links between agents' strategies and the temporal evolution of the market. However, perfect-rationality assumptions have come under scrutiny because utility maximization is often inconsistent with empirical observations of real traders' behaviour [93, 147].

At the other extreme lies a *zero-intelligence* approach, in which agents are assumed to make their decisions according to specified stochastic processes. In this way, agents' actions can be regarded as a consequence of them following a set of rules without strategic considerations. Much like those regarding perfect rationality, zero-intelligence assumptions are extreme simplifications that are inconsistent with empirical observations. However, a zero-intelligence framework has the appeal of leading to models that can yield quantifiable and falsifiable predictions without the need for auxiliary assumptions. It is, therefore, a useful starting point for building models.

Between the two extremes of perfect rationality and zero intelligence lies a broad range of other approaches that make weaker assumptions about agent behaviour, often at the cost of resulting in models that are more difficult to study. Although such models still motivate interesting observations, it is often difficult to trace how their input parameters influence specific elements of their output [46].

7.3.4 Difficulties of Studying ABMs

Although ABMs of price formation provide many important benefits, they also suffer from several drawbacks. First, it is difficult to encode a quantitative set of rules that describe the complex and interacting strategies used by real traders [46]. Second, finding a set of rules that produces a specific behaviour provides no guarantee that such a set of rules is the only one to do so [182]. Third, the large number of interacting components in some ABMs makes it difficult to track how a specified input parameter affects the output [46]. Fourth, many ABMs are intractable, so it is often impossible to derive results about their temporal evolution or to prove that their output will satisfy certain conditions [84]. For example, unexpected interactions between agents could prevent trades from occurring. It is important to consider these issues when designing an ABM, in order to understand their potential impact on the model's output. We return to some of these issues in the context of our ABM in Section 7.4.5.

7.4 Our ABM

7.4.1 An Introduction to our Model

In this section, we present a qualitative description of the main features of our model. Consider trade for a single asset on a quasi-centralized trading platform. Let $\Theta = \{\theta_1, \theta_2, \ldots, \theta_M\}$ denote the set of M agents who trade on the platform and assume that each agent has at least 1 trading partner in Θ. For each $t = 0, 1, 2, \ldots$, let $\phi(t) \in \mathbb{Z}$ denote the price of the most recent trade at or before time t, and for each $i = 1, 2, \ldots, M$, let $p_{-i}(t), p_i(t) \in \mathbb{Z}$ denote, respectively, agent θ_i's *buy price* and *sell price* (i.e., the prices at which agent i is willing to buy and sell the traded asset at time t).[3]

In our model, each agent θ_i updates his/her buy price $p_{-i}(t)$ and sell price $p_i(t)$ according to a discrete-time random walk with drift on a discrete pricing grid. For each price, the strength of the drift is determined by the *trade-price-reversion parameter* κ and the (signed) difference between the given price and $\phi(t)$ (see Section 7.4.3). At any time, if there exists an agent θ_i whose buy price is equal to the sell price of another agent θ_j such that $\theta_i \leftrightarrow \theta_j$, then θ_i and θ_j perform a trade. Any pair

[3]The temporal evolution of our model and the statistics that we calculate all correspond to price differences, not to prices themselves. Therefore, our model's output is price-translation invariant. We simulate our model using an integer price grid to avoid difficulties associated with prices reaching a boundary. An alternative approach is to simulate the model using a natural-number price grid, but to choose the initial conditions in Equation (7.7) such that there is negligible probability that any trader's buy or sell price reaches 0.

of agents $\theta_i \nleftrightarrow \theta_j$ (i.e., any pair of agents that are not trading partners) are unable to trade with each other.

Flowchart 7.1 describes the main features of our ABM's temporal evolution. We list several definitions and present a formal description of the model in the next section.

7.4.2 Our ABM

At any time $t = 0, 1, 2, \ldots$, the state of our ABM is determined by the M agents' buy and sell prices and the most recent trade price.

Definition. *For each time $t = 0, 1, 2, \ldots$, the* model state at time t *is the vector* $\mathcal{M}(t) = (p_{-M}(t), \ldots, p_{-2}(t), p_{-1}(t), p_1(t), p_2(t), \ldots, p_M(t), \phi(t))$.

For constants b_0, a_0 such that $b_0 < a_0$, the initial model state $\mathcal{M}(0)$ is given by

$$\begin{aligned} p_{-i}(0) &= b_0, \\ p_i(0) &= a_0, \end{aligned} \tag{7.7}$$

for each $\theta_i \in \Theta$ and

$$\phi(0) = (b_0 + a_0)/2. \tag{7.8}$$

At each time step $t = 1, 2, \ldots$, choose a number $u_1(t)$ uniformly at random from the set $\{1, 2, \ldots, M\}$. Call the agent $\theta_{u_1(t)}$ the *active agent at time t* and call all other agents the *inactive agents at time t*. Independently, choose a number $u_2(t)$ uniformly at random from the set $\{-1, +1\}$ and calculate the *active index* $u(t) = u_1(t)u_2(t)$. Call $p_{u(t)}$ the *active price at time t* and call all other prices the *inactive prices at time t*.

Calculate the *revised price at time t*

$$R(t) = p_{u(t)}(t-1) + G(t), \tag{7.9}$$

where

$$G(t) = \begin{cases} -1 & \text{with probability } P(w(t)), \\ +1 & \text{with probability } 1 - P(w(t)) \end{cases} \tag{7.10}$$

is the *price revision at time t*, the function $P : \mathbb{Z} \to [0, 1]$ is the *price-update function* (see Section 7.4.3), and

$$w(t) = p_i(t-1) - \phi(t-1) \tag{7.11}$$

is the *price discrepancy at time t*.

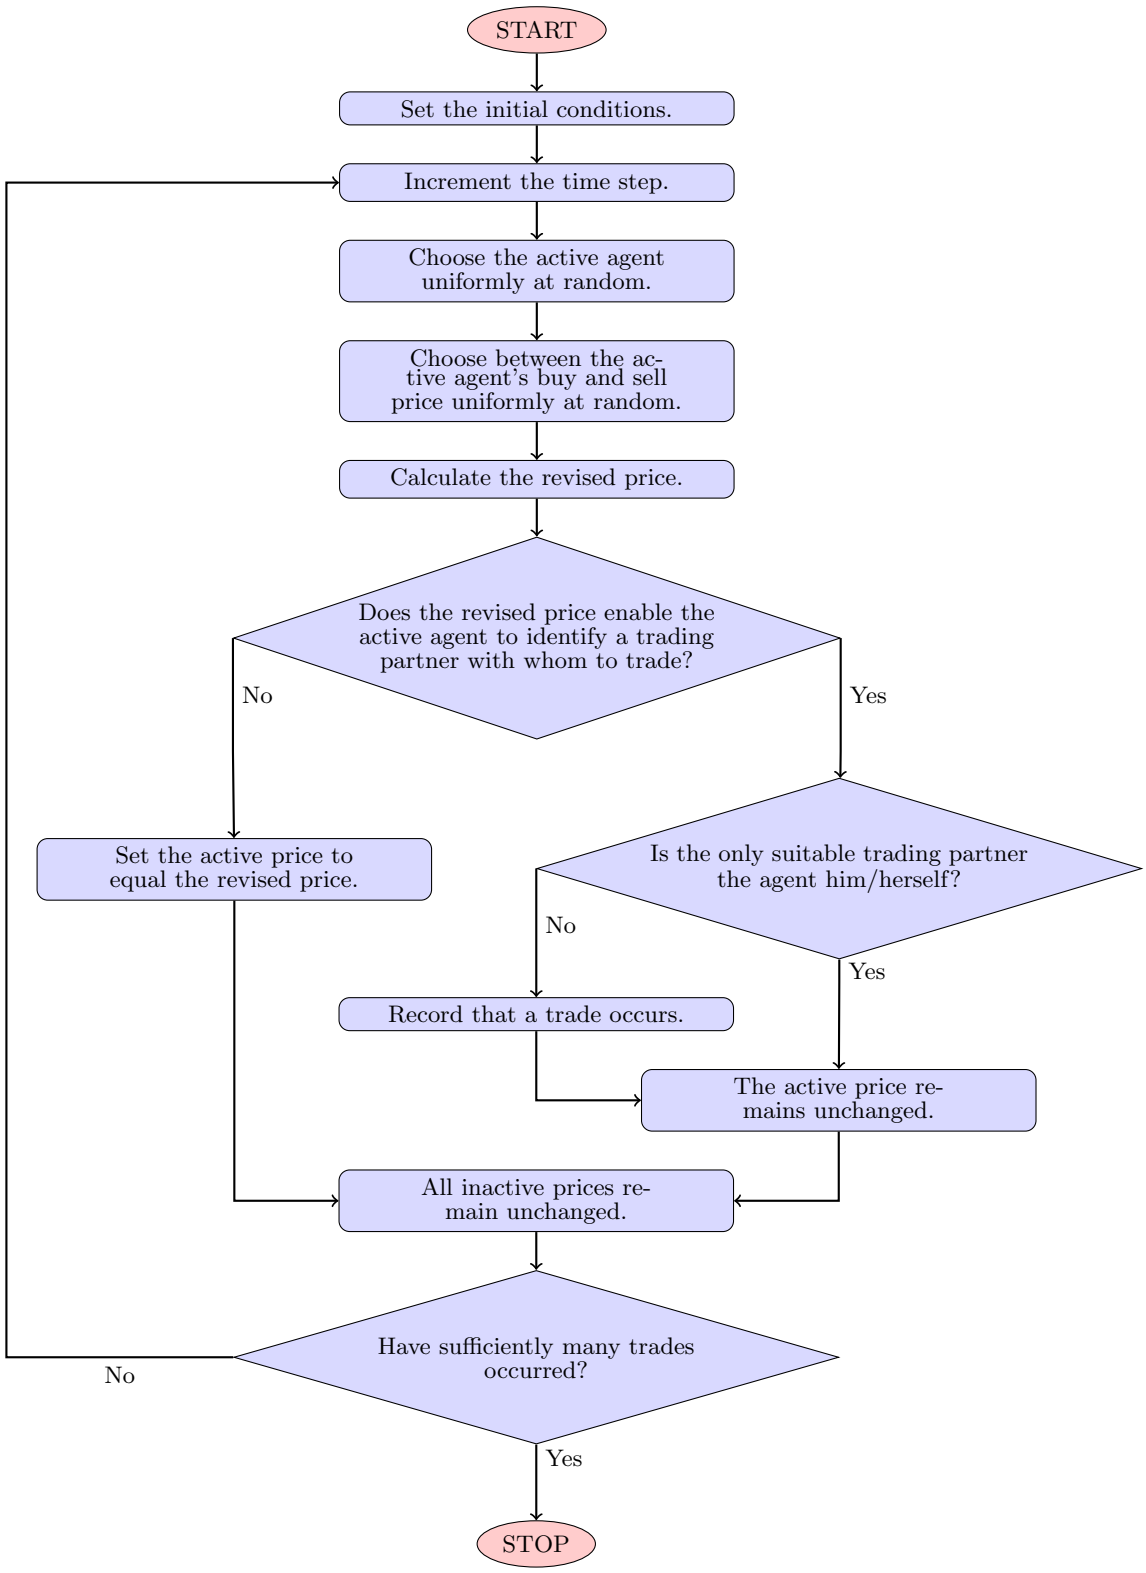

Flowchart 7.1: Qualitative description of the main features of the temporal evolution of our ABM.

Calculate the set

$$C(t) = \begin{cases} \left\{ i \mid R(t) = p_i(t-1), \theta_{u_1(t)} \leftrightarrow \theta_i, \theta_i \in \Theta \right\}, & \text{if } u_2(t) = -1, \\ \left\{ i \mid R(t) = p_{-i}(t-1), \theta_{u_1(t)} \leftrightarrow \theta_i, \theta_i \in \Theta \right\}, & \text{if } u_2(t) = 1. \end{cases} \tag{7.12}$$

Consider the following three cases.

- If $C(t)$ is empty, then the price revision does not enable the active agent to identify any suitable trading partners, so no trade occurs. The active agent adopts the revised price as his/her new buy or sell price.

- If $C(t)$ contains only the element $u_1(t)$, then the price revision does not enable the active agent to identify any suitable trading partners, so no trade occurs. However, adopting the revised price would cause the active agent's buy and sell prices to become equal, so the active agent reverts to his/her previous price $p_{u(t)}(t) = p_{u(t)}(t-1)$.

- If $C(t)$ contains at least one element not equal to $u_1(t)$, then the price revision enables the active agent to trade. After trading, the active agent reverts to his/her previous price $p_{u(t)}(t) = p_{u(t)}(t-1)$.

At the end of each time step, all inactive prices adopt the same value as they took in the previous time step. Repeat this process until the model has generated the required number of trades (see Algorithm 7.1).

7.4.3 The Price-Update Function

In our ABM, we do not explicitly model the decision-making process that the active agent undertakes when choosing the value of the price revision. Instead, we model the outcome of this process using the binary random variable $G(t)$, for which the probability mass function (7.10) depends on the price-update function P. Throughout this chapter, we use the logistic sigmoid function

$$P(w(t)) = \frac{1}{1 + e^{-\kappa w(t)}} \tag{7.13}$$

as the price-update function for our model, where $\kappa > 0$ is the trade-price-reversion parameter. Figure 7.1 illustrates the shape of P for several different choices of κ.

We choose the logistic sigmoid function as our price-update function for several reasons. First, the function's range is $[0, 1]$, which ensures that its output is always a valid probability. Second, the function has a simple analytical form that enables the

Fix $\kappa > 0$.
Set the initial model state $\mathcal{M}(0)$.
for *each* $t = 1, 2, \ldots$ **do**
  Choose $u_1(t)$ uniformly at random from the set $\{1, 2, \ldots, M\}$.
  Choose $u_2(t)$ uniformly at random from the set $\{-1, 1\}$.
  Let $u(t) = u_1(t)u_2(t)$.
  Let $w(t) = p_{u(t)}(t-1) - \phi(t-1)$.
  Calculate $G(t)$ according to Equation (7.10).
  Let $R(t) = p_{u(t)}(t-1) + G(t)$.
  Calculate $C(t)$ according to Equation (7.12).
  if $C(t) = \emptyset$ **then**
    Let $\phi(t) = \phi(t-1)$.
    Let $p_{u(t)}(t) = R(t)$.
  end
  else if $C(t) = \{u_1(t)\}$ **then**
    Let $\phi(t) = \phi(t-1)$.
    Let $p_{u(t)}(t) = p_{u(t)}(t-1)$.
  end
  else
    Let $\phi(t) = R(t)$.
    Let $p_{u(t)}(t) = p_{u(t)}(t-1)$.
  end
  Let $p_{-i}(t) = p_{-i}(t-1)$ and $p_i(t) = p_i(t-1)$ for all inactive prices at time t.
end

Algorithm 7.1: Our agent-based model of trade on a quasi-centralized trading platform.

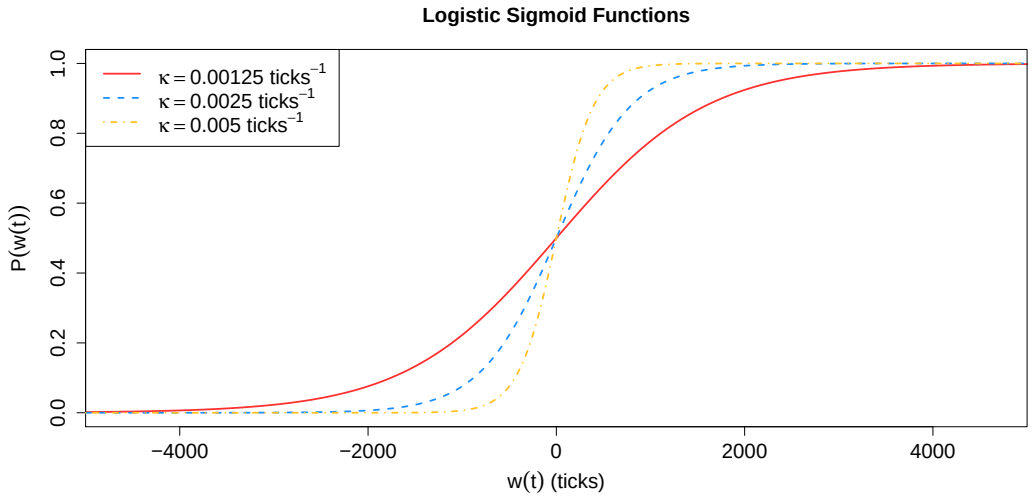

Figure 7.1: Logistic sigmoid price-update functions for several different choices of κ.

derivation of several results about the temporal evolution of our model. Third, for any fixed $\kappa > 0$, it holds that

$$w(t) > 0 \Rightarrow P(w(t)) > 1/2,$$
$$w(t) = 0 \Rightarrow P(w(t)) = 1/2,$$
$$w(t) < 0 \Rightarrow P(w(t)) < 1/2.$$

Therefore, if the price discrepancy is positive (respectively, negative), then the price revision is more likely to be negative than positive (respectively, positive than negative). Fourth, for any fixed $\kappa > 0$, the function P is smooth, monotonic, and satisfies

$$\lim_{w(t)\to\infty} P(w(t)) = 1,$$
$$\lim_{w(t)\to-\infty} P(w(t)) = 0.$$

Therefore, prices with larger price discrepancies experience stronger drift towards $\phi(t)$. Fifth, P satisfies

$$\lim_{\kappa\to\infty} P(w(t)) = 1 \text{ for } w(t) > 0,$$
$$\lim_{\kappa\to\infty} P(w(t)) = 1/2 \text{ for } w(t) = 0,$$
$$\lim_{\kappa\to\infty} P(w(t)) = 0 \text{ for } w(t) < 0,$$
$$\lim_{\kappa\to 0} P(w(t)) = 1/2 \text{ for all } w(t).$$

Therefore, for any non-zero price discrepancy, larger values of κ increase the probability of a price revision towards $\phi(t)$.

Using the logistic sigmoid price-update function (7.13) is equivalent to assuming that the log odds of downward and upward revisions of the active price satisfy

$$\log\left(\frac{P(w(t))}{1-P(w(t))}\right) = -\kappa w(t).$$

Although we regard our ABM as a zero-intelligence model, the logistic sigmoid function also arises as the solution of a decision problem in which agents are assumed to be perfectly rational [165]. We provide details of this formulation in the context of our ABM in Appendix A.4.

7.4.4 An Alternative Price-Update Function

Our ABM shares many similarities with a special case of the model studied by Bak *et al.* [10], in which agents revise their buy and sell prices independently and at random according to the price-update function

$$Q(w(t)) = \begin{cases} (1+\tilde{\kappa})/2, & \text{if } w(t) > 0, \\ 1/2, & \text{if } w(t) = 0, \\ (1-\tilde{\kappa})/2, & \text{if } w(t) < 0, \end{cases} \tag{7.14}$$

for a drift constant $\tilde{\kappa} \in (0, 1]$. We also simulated the temporal evolution of our ABM using Bak *et al.*'s price-update function Q, and we found that our results were qualitatively similar to those for the logistic sigmoid price-update function (see Section 7.6).

7.4.5 Discussion of the Model

Figure 7.2 shows an example simulation of 20000 time steps of our ABM with $M = 12$ agents, $\kappa = 0.0025$, and the trade-partnership network shown in Figure 7.3. The model's temporal evolution is governed by the trade-price-reversion parameter κ, the trade-partnership network, and the agents' price revisions (which arise according to the price-update function P). Although all agents act according to the same rules, each agent's access to trading opportunities is governed by his/her trading partnerships, which vary across agents in all but the extreme case of a complete trade-partnership network (which corresponds to a centralized trading platform).

The purpose of our investigation is to examine how the structure of the trade-partnership network and the choice of κ influence volatility in the trade-price series generated by the model. We seek to keep our model as simple as possible while retaining the key features that characterize a quasi-centralized trading platform: each

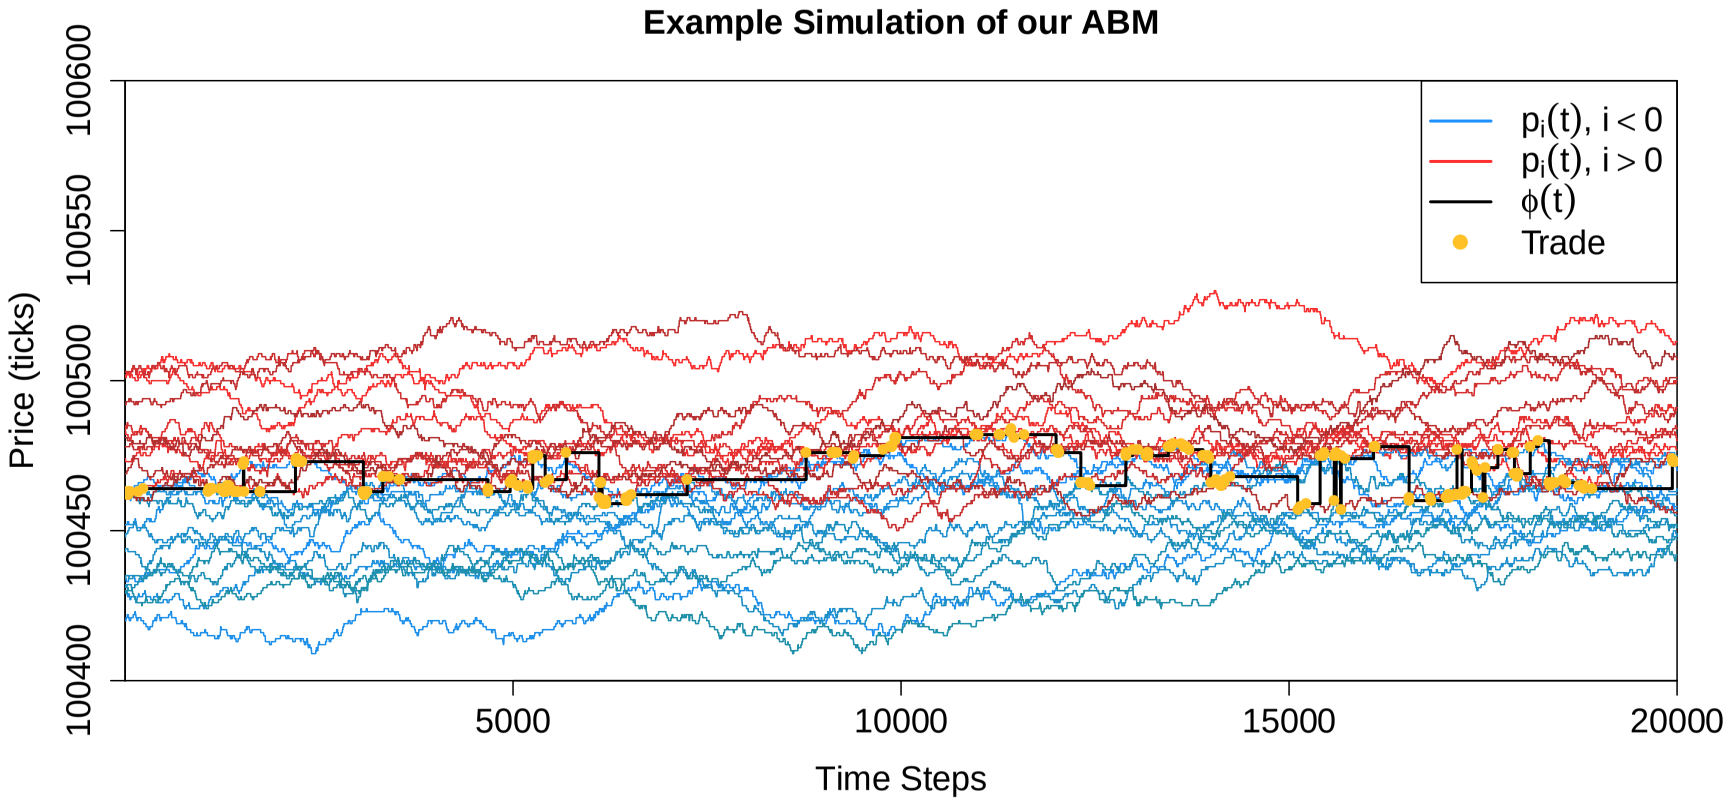

Figure 7.2: Simulation of 20000 time steps of our ABM with $M = 12$, $\kappa = 0.0025$, and the trade-partnership network shown in Figure 7.3. The black curve denotes the trade-price series and the yellow circles denote trades. The blue and red curves denote, respectively, the agents' buy and sell prices (the different shades denote the different agents). We use the initial conditions $b_0 = 100000$ and $a_0 = 101000$, and we discard the first 10000 trades as a burn-in period to remove the transient behaviour of the model (see Section 7.6.1).

agent can only trade with his/her trading partners and all agents know $\phi(t)$ at all times t.

Importantly, we do not seek to reproduce the statistical properties of volatility from empirical data. Therefore, we do not incorporate other sources of volatility — such as macroeconomic activity or news announcements — into our ABM. Instead, we model all price revisions (irrespective of their cause) using the stochastic process $G(t)$ (see Section 7.4.2). Similarly, we do not consider the trading mechanism that agents use to conduct trades, nor do we consider the size of trades that occur. We design our model in this way to highlight how the trade-partnership network and trade-price-reversion parameter κ influence volatility in $\phi(t)$, without incorporating more complicated aspects of trading that could make our results harder to interpret.

Designing our model in this way helps to address two common criticisms of ABMs (see Section 7.3.4). First, it provides a quantitative measure of how changes in the trade-partnership network and κ impact volatility, and thereby helps us to identify links between the model's inputs and outputs without the influence of other factors. Second, it enables us to derive analytical results that guarantee the existence of several important statistical quantities and highlight interesting features of the model's temporal evolution (see Section 7.5). For example, for any trade-partnership network and any choice of κ, we are able to prove that the expected length of time between a successive pair of trades is finite, and that all moments of the price-change distribution exist.

7.5 Mathematical Results

In this section, we list several analytical results about the temporal evolution of our ABM.

Definition. *A model state $\mathcal{M}(t)$ is* feasible *if for all traders $\theta_i, \theta_j \in \Theta$ such that $\theta_i \leftrightarrow \theta_j$, trader θ_i's buy price $p_{-i}(t)$ and trader θ_j's sell price $p_j(t)$ satisfy $p_{-i}(t) < p_j(t)$.*

Proposition 7.1. *For any trade-partnership network and any $\kappa > 0$, the model state $\mathcal{M}(t)$ is feasible for all times $t = 0, 1, 2, \ldots$.*

Proof. By the initial conditions, $\mathcal{M}(0)$ is feasible.

Assume that $\mathcal{M}(t)$ is feasible for some $t \geq 0$. It follows that $p_{-i}(t) < p_j(t)$ for all $\theta_i, \theta_j \in \Theta$ such that $\theta_i \leftrightarrow \theta_j$. Construct the set $C(t+1)$ according to Algorithm 7.1. Consider the following two cases:

1. If $C(t+1) \neq \emptyset$, then $p_{u(t+1)}(t+1) = p_{u(t+1)}(t)$. For all inactive prices, $p_{-i}(t+1) = p_{-i}(t)$ and $p_i(t+1) = p_i(t)$. Therefore, all prices are the same at time $t+1$ as they were at time t. By assumption, $\mathcal{M}(t)$ is feasible. Therefore, $\mathcal{M}(t+1)$ is feasible.

2. If $C(t+1) = \emptyset$, then consider the following two cases.

   (a) If $u_2(t+1) = -1$, then $R(t+1) < p_j(t)$ for all $\theta_j \in \Theta$ such that $\theta_{u_1(t+1)} \leftrightarrow \theta_j$.

   (b) If $u_2(t+1) = 1$, then $R(t+1) > p_{-i}(t)$ for all $\theta_i \in \Theta$ such that $\theta_{u_1(t+1)} \leftrightarrow \theta_i$.

   In either case, $p_{u(t+1)}(t+1) = R(t+1)$. For all inactive prices, $p_{-i}(t+1) = p_{-i}(t)$ and $p_i(t+1) = p_i(t)$, so $p_{-i}(t+1) < p_j(t+1)$ for all $\theta_i, \theta_j \in \Theta$ such that $\theta_i \leftrightarrow \theta_j$. Therefore, $\mathcal{M}(t+1)$ is feasible.

In both cases, $\mathcal{M}(t+1)$ is feasible. The result follows by induction. □

7.5.1 Temporal Evolution of the Model

Lemma 7.1. *For any trade-partnership network and any $\kappa > 0$, the unconditional expectation of $\phi(t)$ at each time $t = 0, 1, 2, \ldots$ exists and is given by*

$$\mathbb{E}\left(\phi(t)\right) = \phi(0).$$

Proof. Let t_1 denote the time of the first trade. For all times $t < t_1$, $\phi(t) = \phi(0) < \infty$. At each time step t, the active agent revises his/her active price according to $G(t)$, which takes the values $+1$ or -1. Therefore,

$$|p_{-i}(t) - b_0| \leq t,$$
$$|p_j(t) - a_0| \leq t$$

for all $i, j = 1, 2, \ldots, M$. A trade can only occur at a price at which either an agent's buy price or an agent's sell price resides, so for all times $t \geq t_1$, it follows that

$$|\phi(t) - b_0| \leq t,$$
$$|\phi(t) - a_0| \leq t.$$

Hence, for all times $t = 0, 1, 2, \ldots$, it holds that

$$b_0 - t \leq \phi(t) \leq a_0 + t,$$

so $\phi(t)$ is bounded. Therefore, $\mathbb{E}\left(\phi(t)\right)$ exists for all $t = 0, 1, 2, \ldots$.

Fix a time t. For a given simulation of t steps of the model, let

$$E = [u(1), G(1)], [u(2), G(2)], \ldots, [u(t), G(t)]$$

denote the sequence of active indices and price revisions at times $1, 2, \ldots, t$. Let Υ denote the set of all possible sequences of this type that correspond to feasible model states at all times up to and including t. For every such sequence $E \in \Upsilon$, there exists a sequence $\tilde{E} \in \Upsilon$ such that

$$\tilde{E} = [-u(1), -G(1)], [-u(2), -G(2)], \ldots, [-u(t), -G(t)].$$

For each element in the sequence E that describes an agent revising his/her buy price upwards (respectively, downwards), the corresponding element in $\tilde{E}$ describes an agent revising his/her sell price downwards (respectively, upwards). Therefore, if the sequence of events in E produces the price change $\phi(t) = \phi(0) + k$, then the sequence of events in $\tilde{E}$ produces the price change $\phi(t) = \phi(0) - k$. By symmetry, the probability of observing the sequence E during a given simulation of t steps of the model is equal to that of observing the sequence $\tilde{E}$.

All possible sequences of feasible model states $\mathcal{M}(0), \mathcal{M}(1), \ldots, \mathcal{M}(t)$ can be described by a sequence of the form $E \in \Upsilon$. For all possible paths such that $\phi(t) = \phi(0) + k$ (of which there are countably many because there are countably many agents), there exists a corresponding path $\tilde{E} \in \Upsilon$ that occurs with the same probability as E and such that $\phi(t) = \phi(0) - k$. Therefore,

$$\mathbb{P}(\phi(t) = \phi(0) + k) = \mathbb{P}(\phi(t) = \phi(0) - k) \tag{7.15}$$

for all $k \in \mathbb{Z}$.

For any $t = 0, 1, 2, \ldots,$

$$\begin{aligned}
\mathbb{E}(\phi(t)) &= \sum_{k \in \mathbb{Z}} (\phi(0) + k)\, \mathbb{P}(\phi(t) = \phi(0) + k) \\
&= \phi(0)\mathbb{P}(\phi(t) = \phi(0)) + \sum_{k>0} (\phi(0) + k)\, \mathbb{P}(\phi(t) = \phi(0) + k) \\
&\quad + \sum_{k<0} (\phi(0) + k)\, \mathbb{P}(\phi(t) = \phi(0) + k) \\
&= \phi(0)\mathbb{P}(\phi(t) = \phi(0)) + \sum_{k>0} (\phi(0) + k)\, \mathbb{P}(\phi(t) = \phi(0) + k) \\
&\quad + \sum_{k'>0} (\phi(0) - k')\, \mathbb{P}(\phi(t) = \phi(0) - k').
\end{aligned}$$

By Equation (7.15), $\mathbb{P}(\phi(t) = \phi(0) + k) = \mathbb{P}(\phi(t) = \phi(0) - k)$. Therefore,

$$\begin{aligned}
\mathbb{E}(\phi(t)) &= \phi(0)\mathbb{P}(\phi(t) = \phi(0)) + \sum_{k>0} (\phi(0) + k + \phi(0) - k)\,\mathbb{P}(\phi(t) = \phi(0) + k) \\
&= \phi(0)\mathbb{P}(\phi(t) = \phi(0)) + \sum_{k>0} 2\phi(0)\mathbb{P}(\phi(t) = \phi(0) + k) \\
&= \phi(0) \sum_{k\in\mathbb{Z}} \mathbb{P}(\phi(t) = \phi(0) + k) \\
&= \phi(0), \text{ as required.}
\end{aligned}$$

□

Definition. *The* trade times T_ϕ *is the set of times at which trades occur.*

Proposition 7.2. *For any trade-partnership network and any $\kappa > 0$, given any time t such that $t + 1 \notin T_\phi$, then*

$$\mathbb{E}\Big[\, |p_{-i}(t+1) - \phi(t+1)| \,\Big|\, p_{-i}(t), \phi(t)\Big] < |p_{-i}(t) - \phi(t)|$$

for all $\theta_i \in \Theta$ such that $p_{-i}(t) \neq \phi(t)$ and

$$\mathbb{E}\Big[\, |p_j(t+1) - \phi(t+1)| \,\Big|\, p_j(t), \phi(t)\Big] < |p_j(t) - \phi(t)|$$

for all $\theta_j \in \Theta$ such that $p_j(t) \neq \phi(t)$.

Proof. We prove the result for buy prices. The result for sell prices follows by symmetry. For buy prices, all expectations are conditional expectations given $p_{-i}(t)$ and $\phi(t)$. To keep the notation uncluttered, we do not display this conditioning.

For any $\theta_i \in \Theta$,

$$\begin{aligned}
\mathbb{E}[p_{-i}(t+1) - \phi(t+1)] &= \mathbb{E}[p_{-i}(t+1) - \phi(t) + \phi(t) - \phi(t+1)] \\
&= \mathbb{E}[p_{-i}(t+1) - \phi(t)] + \mathbb{E}[\phi(t) - \phi(t+1)].
\end{aligned}$$

By the statement of the proposition, a trade does not occur at time $t+1$. Therefore, $\phi(t+1) = \phi(t)$, so $\mathbb{E}[\phi(t) - \phi(t+1)] = 0$ and

$$\mathbb{E}[p_{-i}(t+1) - \phi(t+1)] = \mathbb{E}[p_{-i}(t+1) - \phi(t)]. \tag{7.16}$$

Consider the following two cases.

- If $p_{-i}(t)$ is an inactive price, which occurs with probability $(2M-1)/2M$, then $p_{-i}(t+1) = p_{-i}(t)$.

- If $p_{-i}(t)$ is the active price, which occurs with probability $1/2M$, then

$$p_{-i}(t+1) = \begin{cases} p_{-i}(t) - 1, & \text{with probability } P(w(t)), \\ p_{-i}(t) + 1, & \text{with probability } 1 - P(w(t)). \end{cases}$$

Therefore,

$$\begin{aligned} \mathbb{E}\left[p_{-i}(t+1) - \phi(t)\right] &= \frac{2M-1}{2M}\left(p_{-i}(t) - \phi(t)\right) \\ &\quad + \frac{1}{2M}\left(p_{-i}(t) - \phi(t) - 1\right)\frac{1}{1+e^{-\kappa(p_{-i}(t)-\phi(t))}} \\ &\quad + \frac{1}{2M}\left(p_{-i}(t) - \phi(t) + 1\right)\left(1 - \frac{1}{1+e^{-\kappa(p_{-i}(t)-\phi(t))}}\right) \\ &= p_{-i}(t) - \phi(t) + \frac{1}{2M}\left(1 - \frac{2}{1+e^{-\kappa(p_{-i}(t)-\phi(t))}}\right). \end{aligned} \tag{7.17}$$

When $p_{-i}(t) - \phi(t) > 0$, we seek to show that $\mathbb{E}\left[p_{-i}(t+1) - \phi(t+1)\right] < p_{-i}(t) - \phi(t)$. By (7.16) and (7.17), this is equivalent to

$$\frac{1}{2M}\left(1 - \frac{2}{1+e^{-\kappa(p_{-i}(t)-\phi(t))}}\right) < 0,$$

which is satisfied for any $\kappa > 0$ and any $p_{-i}(t) - \phi(t) > 0$.

When $p_{-i}(t) - \phi(t) < 0$, we seek to show that $\mathbb{E}\left[p_{-i}(t+1) - \phi(t+1)\right] > p_{-i}(t) - \phi(t)$. By (7.16) and (7.17), this is equivalent to

$$\frac{1}{2M}\left(1 - \frac{2}{1+e^{-\kappa(p_{-i}(t)-\phi(t))}}\right) \geq 0,$$

which is satisfied for any $\kappa > 0$ and any $p_{-i}(t) - \phi(t) < 0$. □

7.5.2 Movement and Trade Times

Definition. *For each $\theta_i \in \Theta$, the* buy movement times

$$T_{-i} = \{t_h | u(t_h) = -i\} \tag{7.18}$$

is the set of times at which trader θ_i's buy price is the active price, and the sell movement times

$$T_j = \{t_h | u(t_h) = j\} \tag{7.19}$$

is the set of times at which trader θ_j's sell price is the active price.

Lemma 7.2. *For any trade-partnership network, any $\kappa > 0$, any $\theta_i, \theta_j \in \Theta$, and any time $(t_h + 1) \in T_{-i} \cup T_j$ such that $p_{-i}(t_h) \neq p_j(t_h)$ and $t_h + 1 \notin T_\phi$, it holds that*

$$\mathbb{P}\left(|p_{-i}(t_h+1) - p_j(t_h+1)| < |p_{-i}(t_h) - p_j(t_h)|\right) > 1/2. \tag{7.20}$$

Proof. By the statement of the lemma, $p_{-i}(t_h) \neq p_j(t_h)$. Without loss of generality, assume that $p_{-i}(t_h) < p_j(t_h)$. Given that $(t_h + 1) \in T_{-i} \cup T_j$, it follows that either $p_{-i}(t_h + 1) \neq p_{-i}(t_h)$ or $p_j(t_h + 1) \neq p_j(t_h)$. Therefore, there are four possible combinations of values of $p_{-i}(t_h + 1)$ and $p_j(t_h + 1)$:

$$\begin{aligned} p_{-i}(t_h + 1) &= p_{-i}(t_h) + 1 \text{ and } p_j(t_h + 1) = p_j(t_h), \\ p_{-i}(t_h + 1) &= p_{-i}(t_h) - 1 \text{ and } p_j(t_h + 1) = p_j(t_h), \\ p_{-i}(t_h + 1) &= p_{-i}(t_h) \text{ and } p_j(t_h + 1) = p_j(t_h) + 1, \\ p_{-i}(t_h + 1) &= p_{-i}(t_h) \text{ and } p_j(t_h + 1) = p_j(t_h) - 1. \end{aligned} \tag{7.21}$$

Let the probabilities of these four events be denoted $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3$, and $\mathcal{P}_4$, respectively. Because the active agent is chosen uniformly at random and given that $(t_h + 1) \in T_{-i} \cup T_j$, it follows that

$$\mathcal{P}_1 + \mathcal{P}_2 + \mathcal{P}_3 + \mathcal{P}_4 = 1 \tag{7.22}$$

and

$$\mathcal{P}_1 + \mathcal{P}_2 = \mathcal{P}_3 + \mathcal{P}_4 = 1/2. \tag{7.23}$$

Observe that $\mathbb{P}\left(|p_{-i}(t_h + 1) - p_j(t_h + 1)| < |p_{-i}(t_h) - p_j(t_h)|\right) = \mathcal{P}_1 + \mathcal{P}_4$. Therefore, we seek to show that

$$\mathcal{P}_1 + \mathcal{P}_4 > \frac{1}{2}. \tag{7.24}$$

By Equations (7.22) and (7.23), this inequality holds if and only if

$$\mathcal{P}_1 + \mathcal{P}_4 > \mathcal{P}_2 + \mathcal{P}_3. \tag{7.25}$$

Consider the following three cases.

1. If $p_{-i}(t_h) < \phi(t_h) < p_j(t_h)$, then $p_j(t_h) - \phi(t_h) > p_{-i}(t_h) - \phi(t_h)$. For any $c_1, c_2 \in \mathbb{Z}$ and any $\kappa > 0$, observe that

   $$c_1 > c_2 \Leftrightarrow P(c_1) < P(c_2). \tag{7.26}$$

   Therefore, $\mathcal{P}_1 > \mathcal{P}_3$ and $\mathcal{P}_4 > \mathcal{P}_2$, so $\mathcal{P}_1 + \mathcal{P}_4 > \mathcal{P}_2 + \mathcal{P}_3$, as required.

2. If $\phi(t_h) \leq p_{-i}(t_h) < p_j(t_h)$, then by Equation (7.13):

   $$\begin{aligned} \mathcal{P}_1 + \mathcal{P}_4 &= \frac{1}{2}\left(1 - \frac{1}{1 + e^{-\kappa(p_{-i}(t_h) - \phi(t_h))}}\right) + \frac{1}{2}\left(\frac{1}{1 + e^{-\kappa(p_j(t_h) - \phi(t_h))}}\right) \\ &= \frac{e^{-\kappa(p_{-i}(t_h) - \phi(t_h) + p_j(t_h) - \phi(t_h))} + 2e^{-\kappa(p_{-i}(t_h) - \phi(t_h))} + 1}{2\left(1 + e^{-\kappa(p_{-i}(t_h) - \phi(t_h))}\right)\left(1 + e^{-\kappa(p_j(t_h) - \phi(t_h))}\right)}. \end{aligned} \tag{7.27}$$

Similarly,

$$\mathcal{P}_2 + \mathcal{P}_3 = \frac{1}{2}\left(\frac{1}{1+e^{-\kappa(p_{-i}(t_h)-\phi(t_h))}}\right) + \frac{1}{2}\left(1 - \frac{1}{1+e^{-\kappa(p_j(t_h)-\phi(t_h))}}\right)$$

$$= \frac{e^{-\kappa(p_{-i}(t_h)-\phi(t_h)+p_j(t_h)-\phi(t_h))} + 2e^{-\kappa(p_j(t_h)-\phi(t_h))} + 1}{2\left(1+e^{-\kappa(p_{-i}(t_h)-\phi(t_h))}\right)\left(1+e^{-\kappa(p_j(t_h)-\phi(t_h))}\right)}. \tag{7.28}$$

The right hand sides of Equations (7.27) and (7.28) differ only in the middle term in the numerator. Because $p_j(t_h) > p_{-i}(t_h)$, it follows that $\mathcal{P}_1 + \mathcal{P}_4 > \mathcal{P}_2 + \mathcal{P}_3$, as required.

3. If $p_{-i}(t_h) < p_j(t_h) \leq \phi(t_h)$, then the same result as in case 2 applies by symmetry.

□

Proposition 7.3. *For any trade-partnership network, any $\kappa > 0$, and any feasible model state $\mathcal{M}(t)$, the expected length of time until the next trade is finite.*

Proof. Let $\tau \geq 0$ denote the time of the first trade at or after time t. We seek to show that $\mathbb{E}(\tau - t) < \infty$.

If $t \in T_\phi$, then $\tau = t$, so $\mathbb{E}(\tau - t) = 0 < \infty$.

If $t \notin T_\phi$, then consider any buy price $p_{-i}(t)$ and any sell price $p_j(t)$ such that $\theta_i \leftrightarrow \theta_j$. Define

$$T_{(-i,j,t)} = \{t_h | t_h \in (T_{-i} \cup T_j), t_h \geq t\}. \tag{7.29}$$

For each time $t_h \in T_{(-i,j,t)}$, let

$$\psi(t_h) = p_j(t_h) - p_{-i}(t_h).$$

If $\psi(t_h) = 1$, then θ_i and θ_j trade at time t_h if either $t_h \in T_i$ and $G(t_h) = +1$ or $t_h \in T_j$ and $G(t_h) = -1$.

Observe that τ is a stopping time [174]. Note also that until the next trade occurs, the evolution of $\psi(t_h)$ is a discrete-time Markov process on the state space $\mathbb{N}$ with transitions

$$\psi(t_{h+1}) = \begin{cases} \psi(t_h) - 1 & \text{with probability } \mathcal{P}_1 + \mathcal{P}_4, \\ \psi(t_h) + 1 & \text{with probability } \mathcal{P}_2 + \mathcal{P}_3, \end{cases}$$

where $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3$, and $\mathcal{P}_4$ are defined by Equation (7.21).

Let $\tilde{\psi}(t)$ be a discrete-time Markov process on the state space $\mathbb{N}$ with transitions

$$\tilde{\psi}(t+1) = \begin{cases} \tilde{\psi}(t) + 1 & \text{with probability } \tilde{\mathcal{P}}, \\ \tilde{\psi}(t) - 1 & \text{with probability } 1 - \tilde{\mathcal{P}}. \end{cases}$$

Let $\tilde{\psi}(0) = n$ for some $n \in \mathbb{N}$, and let τ_f denote the first time that $\tilde{\psi}$ visits the state f for each $f = 0, 1, 2, \ldots, n$. Observe that $\tau_n = 0$ and that each of $\tau_0, \tau_1, \ldots, \tau_{n-1}$ are stopping times. By standard results from Markov chains (see, e.g., [174]), if $\tilde{\mathcal{P}} < 1/2$, then

$$\mathbb{E}(\tau_{n-1} - \tau_n) = \frac{1}{2(1 - \tilde{\mathcal{P}}) - 1}.$$

Therefore, applying the strong Markov property at each of $\tau_0, \tau_1, \ldots, \tau_{n-1}$ and by linearity of the expectation, it follows that

$$\mathbb{E}(\tau_0) = \frac{n}{2(1 - \tilde{\mathcal{P}}) - 1}.$$

Therefore, $\mathbb{E}(\tau_0) < \infty$ for all $\tilde{\mathcal{P}} < 1/2$.

Unlike $\tilde{\psi}$, the transition probabilities of ψ are not constant because they depend on the values of $p_{-i}(t_h)$ and $p_j(t_h)$. However, by Lemma 7.2, it follows that $\mathcal{P}_1 + \mathcal{P}_4 > 1/2$ and $\mathcal{P}_2 + \mathcal{P}_3 < 1/2$ for all times $t_h \in T_{(-i,j,t)}$. Therefore, if no trade between any other pair of agents occurs before the next trade between agent θ_i's buy price and agent θ_j's sell price, then τ denotes the time of this trade and $\mathbb{E}(\tau - t)$ is finite. Otherwise, a trade occurs between some other pair of buy and sell prices before a trade occurs between agent θ_i's buy price and agent θ_j's sell price, so $\mathbb{E}(\tau - t)$ is finite. Therefore, $\mathbb{E}(\tau - t)$ is finite in both cases. □

7.5.3 Price-Change Moments

Proposition 7.4. *For any trade-partnership network, any $\kappa > 0$ and any successive trade times $t_h, t_{h+1} \in T_\phi$, it holds that*

$$\mathbb{E}\left[(\phi(t_{h+1}) - \phi(t_h))^n\right] < \infty \tag{7.30}$$

for all $n = 1, 2, \ldots$.

Proof. By Proposition 7.3, the expected length of time between any successive pair of trades is finite. By Proposition 7.2, the distance that any agent moves away from the previous trade price before the next trade occurs is bounded. Therefore, $(\phi(t_{h+1}) - \phi(t_h))$ is always finite, so

$$\mathbb{E}\left[(\phi(t_{h+1}) - \phi(t_h))^n\right] < \infty$$

for all $n = 1, 2, \ldots$. □

7.6 Simulation Results

In this section, we study simulated output from our ABM to investigate how volatility in the trade-price series depends on both the structure of the trade-partnership network and the trade-price-reversion parameter κ.

7.6.1 Methodology

Before we begin our simulations, we fix 100 different seeds (i.e., initial conditions) $\Omega_1, \Omega_2, \ldots, \Omega_{100}$ for the pseudo-random number generator. For each trade-partnership network and each value of κ that we study, we simulate 100 independent repetitions of our ABM, where the i^{th} repetition uses seed Ω_i, for $i = 1, 2, \ldots, 100$. We use the same seeds $\Omega_1, \Omega_2, \ldots, \Omega_{100}$ for all trade-partnership networks and all values of κ to ensure that differences between the simulation outputs reflect only differences in the model's inputs rather than differences in the (pseudo-)random number draws.

For the results that we present in this section, we use the initial conditions $b_0 = 100000$ and $a_0 = 101000$. For each repetition, we simulate 20000 trades, then we discard the first 10000 trades as a burn-in period to remove the transient behaviour of the model.[4] For each repetition $i = 1, 2, \ldots, 100$, we use the remaining 10000 trades to calculate the mean absolute price change, the standard deviation of absolute price changes, and the percentage of non-zero price changes (see Section 7.2), which we label v_1^i, v_2^i, and v_3^i. In this section, we investigate the means and standard deviations of the v_j^i,

$$\overline{v}_j = \frac{1}{100} \sum_{i=1}^{100} v_j^i, \tag{7.31}$$

$$\sigma_j = \sqrt{\frac{1}{99} \sum_{i=1}^{100} \left(v_j^i - \overline{v}_j\right)^2}, \tag{7.32}$$

for $j = 1, 2, 3$.

We also repeated all of our simulations using different numbers of trades in the range 500 to 100000 and different numbers of independent repetitions of the model in the range 10 to 200. We found that our results were qualitatively similar for all simulations with more than about 1000 trades and more than about 20 independent

[4] We calculated the values of v_1, v_2, and v_3 using a rolling window, and we found that the transient behaviour (characterized by very small or very large values of v_1, v_2, and v_3) was no longer detectable after approximately 1000 trades. We use a burn-in period of 10000 trades to be prudent.

repetitions of the model. Therefore, we deem our simulations for 10000 trades and 100 independent repetitions of the model to produce statistically stable results.

7.6.2 The Networks

For a given number of agents M, there are

$$\binom{M}{2} = \frac{M!}{2!\,(M-2)!}$$

possible edges in the corresponding trade-partnership network. Therefore, there are $2^{\binom{M}{2}}$ different trade-partnership networks (up to isomorphism) for a given choice of M. Even for small values of M, it is computationally infeasible to perform an exhaustive search across all such networks.

Instead, we study the simulated output of our ABM using three different classes of trade-partnership networks, which we call "gangs networks", "tiers networks", and "core-periphery networks". For each of these three classes, a scalar-valued network parameter determines the entire trade-partnership network for a given M. We describe these network classes in Sections 7.6.2.1, 7.6.2.2, and 7.6.2.3, respectively.

We chose these network classes because they are extremely simple, deterministic, 1-parameter models that help to illustrate how volatility of the trade-price series varies according to the topological properties of the trade-partnership network. Although they represent extreme simplifications of the structure of real trade-partnership networks, we provide some simple motivations for our choices of these model classes in the subsequent sections.

7.6.2.1 Gangs Networks

Definition. *For a given number of agents M and choice of network parameter $\eta_G \in \{1, 2, \ldots, M/2\}$, a* gangs network *is a trade-partnership network that consists of η_G equal-sized* gangs,[5] *such that each node is connected to all other nodes in its gang but is not connected to any nodes not in its gang.*

Figure 7.3 shows a schematic of a gangs network. We use this network class to model a quasi-centralized trading platform on which each trader only trades within a fully-connected group. For example, if a quasi-centralized trading platform is populated by equal numbers of large banks and small banks, and if

[5] If η_G does not divide M exactly, then the gangs are as close to equal-sized as possible (i.e., no pair of gangs differ in size by more than 1).

- all large banks trade with all other large banks,
- all small banks trade with all other small banks,
- no large banks trade with any small banks,

then the resulting trade-partnership network is a gangs network with 2 gangs.

When $\eta_G = 1$, there exists only one gang, so the trade-partnership network is complete and all agents are able to trade with all others. When $\eta_G = M/2$, a gangs network consists of $M/2$ pairs of agents, each of whom is only able to trade with his/her single trading partner.

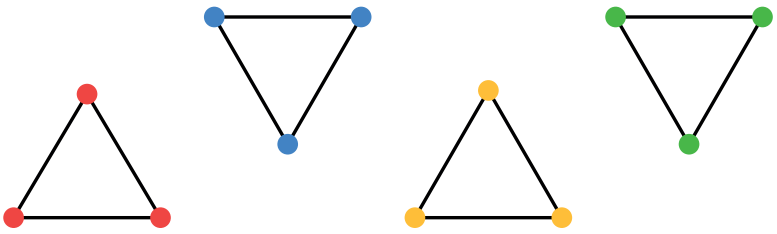

Figure 7.3: Schematic of a gangs network with $M = 12$ and $\eta_G = 4$. The nodes that belong to the first, second, third, and fourth gang are coloured red, blue, yellow, and green, respectively.

7.6.2.2 Tiers Networks

Definition. *For a given number of agents M and choice of network parameter $\eta_T \in \{1, 2, \ldots, M\}$, a* tiers network *is a trade-partnership network that consists of a sequence of η_T equal-sized* tiers,[6] *such that each node is connected to all other nodes in its tier and in the neighbouring tiers, but is not connected to any other nodes.*

Figure 7.4 shows a schematic of a tiers network. We use tiers networks to model a quasi-centralized trading platform on which traders reside in one of several tiers (which could, for example, reflect a trader's size or perceived creditworthiness) and in which each trader only trades with traders in their own tier or a neighbouring tier. For example, if a quasi-centralized trading platform is populated by equal numbers of large banks, small banks, and small financial institutions, and if:

- all large banks trade with all other large banks and all small banks,

[6] If η_T does not divide M exactly, then the tiers are as close to equal-sized as possible (i.e., no pair tiers differ in size by more than 1).

- all small banks trade with all large banks, all other small banks, and all small financial institutions,
- all small financial institutions trade with all small banks and all other small financial institutions,
- no large banks trade with any small financial insitutions,

then the resulting trade-partnership network is a tiers network with 3 tiers.

When $\eta_T = 1$, there exists only one tier, so the trade-partnership network is complete and all agents are able to trade with all others. Similarly, when $\eta_T = 2$, all agents are in either the first tier or the second tier, but these tiers are adjacent so the trade-partnership network is complete. When $\eta_T = M$, a tiers network corresponds to a straight line in which each agent is only able to trade with his/her immediate neighbours.

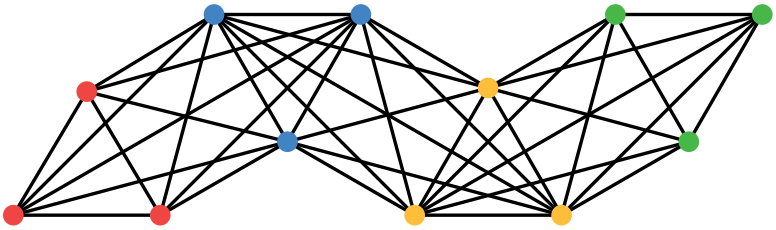

Figure 7.4: Schematic of a tiers network with $M = 12$ and $\eta_T = 4$. The nodes that belong to the first, second, third, and fourth tier are coloured red, blue, yellow, and green, respectively.

7.6.2.3 Core-Periphery Networks

Definition. *For a given number of agents M and choice of network parameter $\eta_C \in [0, 1)$, a* core-periphery network *is a trade-partnership network that consists of $M\eta_C$* periphery nodes *and $M(1 - \eta_C)$* core *nodes,*[7] *such that all core nodes are connected to all other core nodes and all periphery nodes are connected to exactly one core node, and such that the degree of no two core nodes differs by more than 1.*

Figure 7.5 shows a schematic of a core-periphery network. We use core-periphery networks [62] to model a quasi-centralized trading platform on which all agents are either core agents or periphery agents. All core agents are trading partners with all

[7] If $M\eta_C$ and $M(1 - \eta_C)$ are not integers, then the periphery consists of $\lfloor M\eta_C \rfloor$ nodes and the core consists of $\lceil M(1 - \eta_C) \rceil$, where $\lfloor M\eta_C \rfloor$ and $\lceil M(1 - \eta_C) \rceil$ denote, respectively, the largest integer smaller than $M\eta_C$ and the smallest integer larger than $M(1 - \eta_C)$.

others, but each periphery agent is trading partners with exactly one core agent. For example, if a quasi-centralized trading platform is populated by large banks and small banks, and if:

- each small bank only trades with one large bank,
- all large banks trade with all other large banks and a subset of small banks,

then the resulting trade-partnership network is a core-periphery network.

When $\eta_C = 0$, all nodes reside in the core, so the trade-partnership network is complete and all agents are able to trade with all others. For all $\eta_C \geq (M-1)/M$, the core consists of a single node, so the trade-partnership network is a star network.

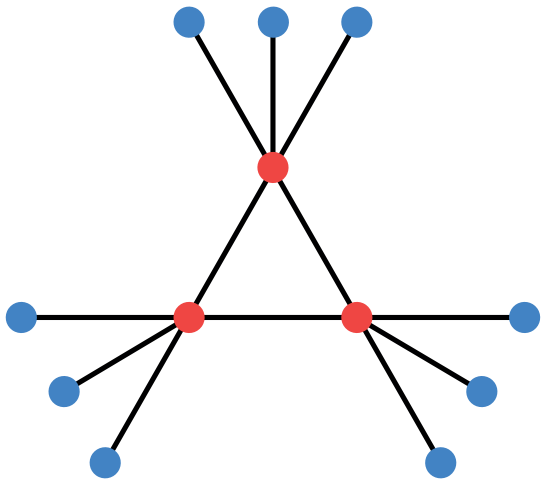

Figure 7.5: Schematic of a core-periphery network with $M = 12$ and $\eta_C = 0.75$. The nodes in the core are coloured red and the nodes in the periphery are coloured blue.

7.6.3 Number of Agents

The aim of this chapter is to help understand the relationship between the structure of the trade-partnership network, the trade-price-reversion parameter κ, and the volatility of the trade-price series. Therefore, we fix the same value of M for all of our simulations to ensure that our results do not reflect changes in the number of agents.

To identify a suitable choice of M, we first performed all of our simulations with different numbers of agents in the range $M \in [10, 1000]$. The shapes of the model-output surfaces that we present in Section 7.6.4 are qualitatively similar for all choices of M greater than about $M = 100$. For $M < 100$, small-sample effects play a substantial role in the temporal evolution of the model. However, given the large size

of the FX market[8] and the widespread use of Hotspot FX [176], we do not believe that $M < 100$ is a sensible choice. For the simulation results that we present in Section 7.6.4, we use $M = 200$ agents.

With this number of agents, the possible choices of network parameters are

$$\begin{aligned}\eta_G &\in \{1, 2, \ldots, 100\}, \\ \eta_T &\in \{1, 2, \ldots, 200\}, \\ \eta_C &\in [0, 1).\end{aligned}$$

To help facilitate comparisons between the different network classes, we parameterize η_G, η_T, and η_C using the scalar parameter $z \in \{1, 2, \ldots, 100\}$, according to

$$\begin{aligned}\eta_G &= z, \\ \eta_T &= z, \\ \eta_C &= (z - 1)/100.\end{aligned}$$

For example, the value $z = 5$ corresponds to a gangs network with 5 gangs, a tiers network with 5 tiers, or a core-periphery network in which 4% of agents are core agents. Although a tiers network with $M = 200$ agents can have up to 200 tiers, we restrict our attention to $\eta_T \leq 100$ to enable comparisons with the corresponding gangs networks.

7.6.4 Results

In this section, we compare the simulated output of our ABM for the gangs, tiers, and core-periphery networks, with different choices of z and κ. Figures 7.6, 7.7, and 7.8 show how the values of $\overline{v}_1$, $\overline{v}_2$, and $\overline{v}_3$, respectively, vary according to z and κ. Figure 7.9 shows how the values of $\overline{v}_1$, $\overline{v}_2$, $\overline{v}_3$, σ_1, σ_2, and σ_3 vary along the hyperplanes $\kappa = 0.0025$ and $z = 50$.

For each of the three network classes, $z = 1$ corresponds to a complete trade-partnership network. For a given choice of $\kappa > 0$, the gangs, tiers, and core-periphery networks all produce identical values of $\overline{v}_1$, $\overline{v}_2$, $\overline{v}_3$, σ_1, σ_2, and σ_3 in this case, because

[8] It is difficult to estimate the number of active FX spot traders worldwide. To protect the privacy of their clients, electronic trading platforms do not release information regarding their numbers of users, and large-scale market surveys such as the BIS Triennial Central Bank Survey [13] restrict their attention to large banks. A very rough proxy for the number of traders that use electronic FX spot trading platforms is the number of subscribers to Continuous Linked Settlement (CLS) [52], which is a popular market infrastructure that facilitates settlements of FX transaction payments. At present, the number of institutions that subscribe to CLS is in the hundreds.

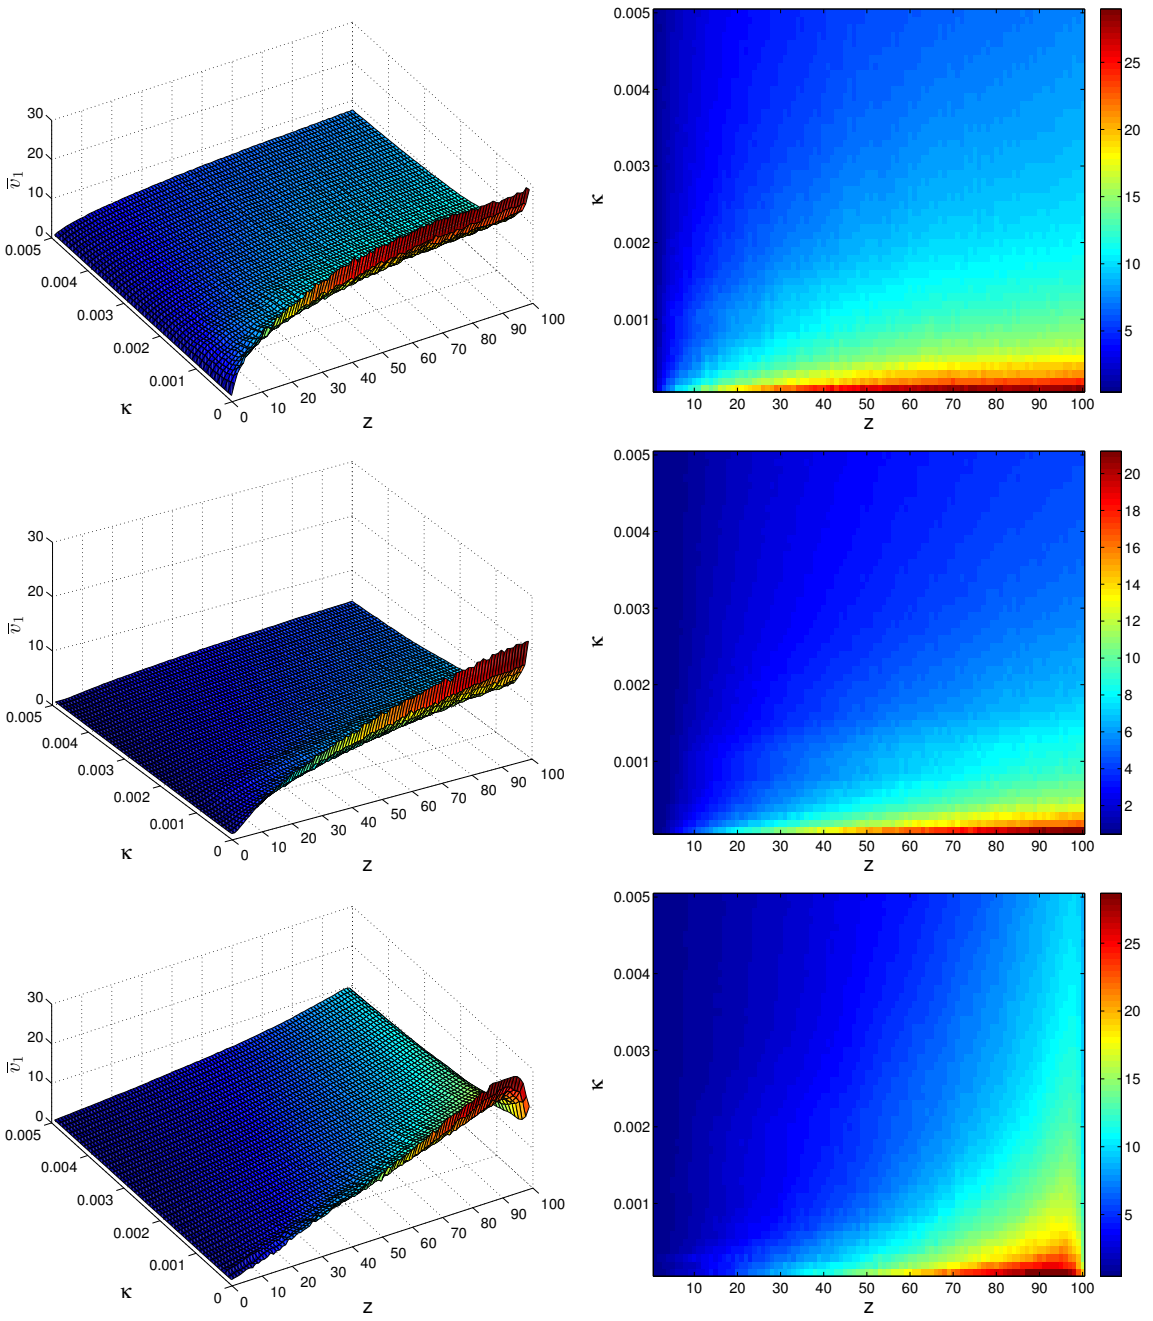

Figure 7.6: (Left column) Surface plots and (right column) heat maps of the values of $\overline{v}_1$ for the (top row) gangs, (middle row) tiers, and (bottom row) core-periphery networks with given values of z and κ.

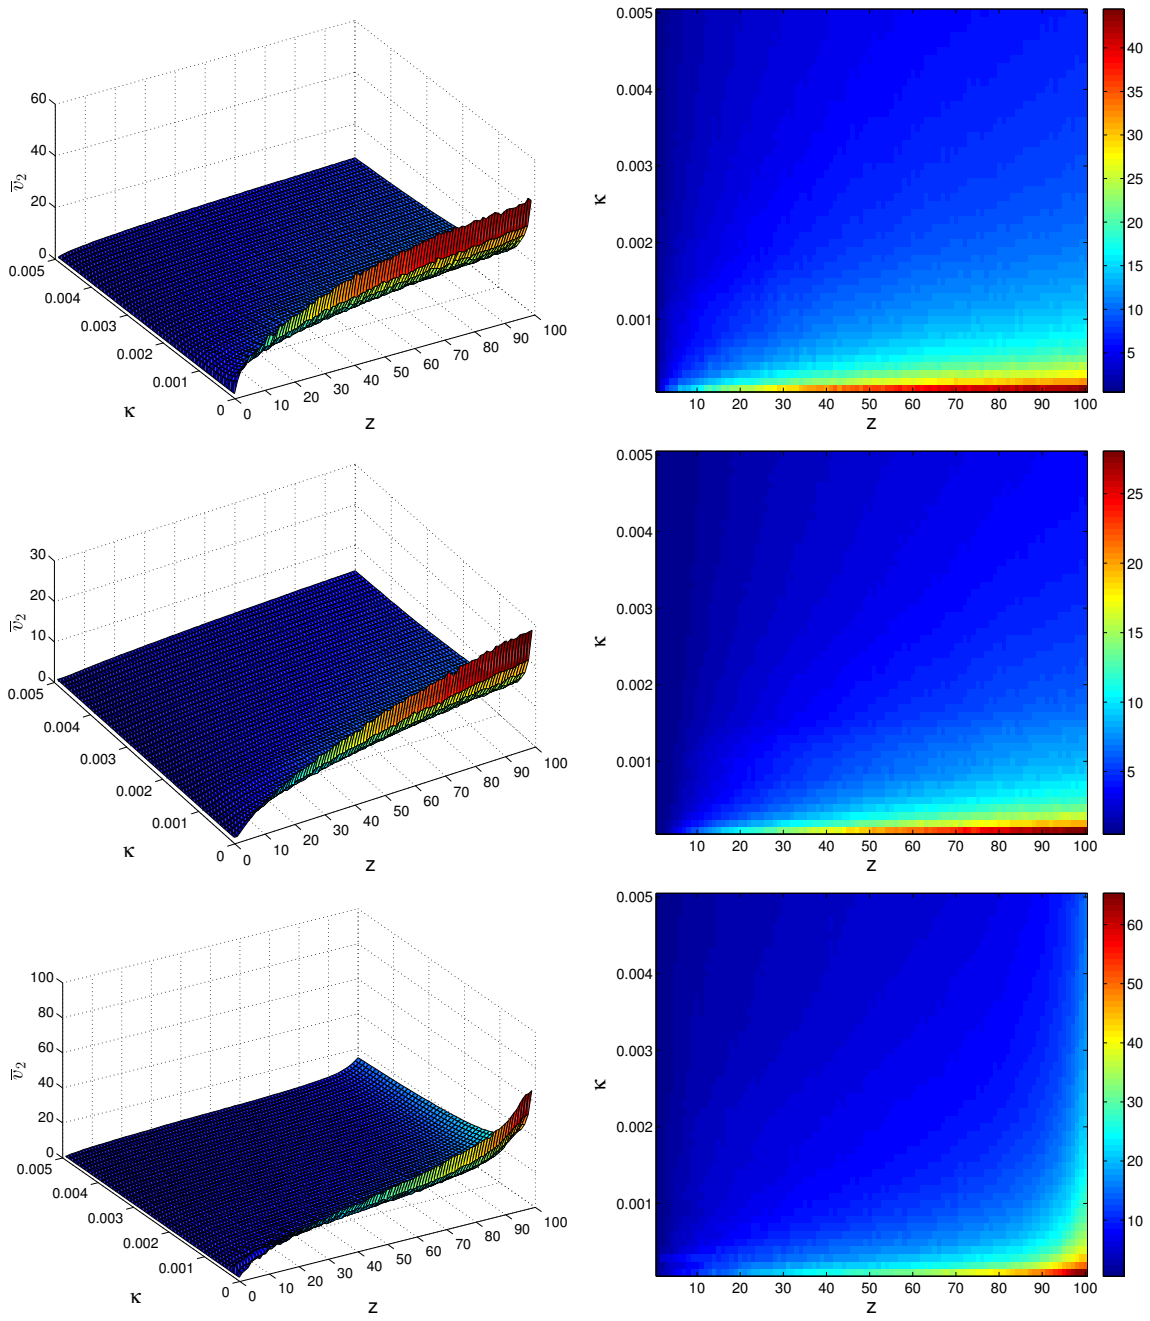

Figure 7.7: (Left column) Surface plots and (right column) heat maps of the values of $\overline{v}_2$ for the (top row) gangs, (middle row) tiers, and (bottom row) core-periphery networks with given values of z and κ.

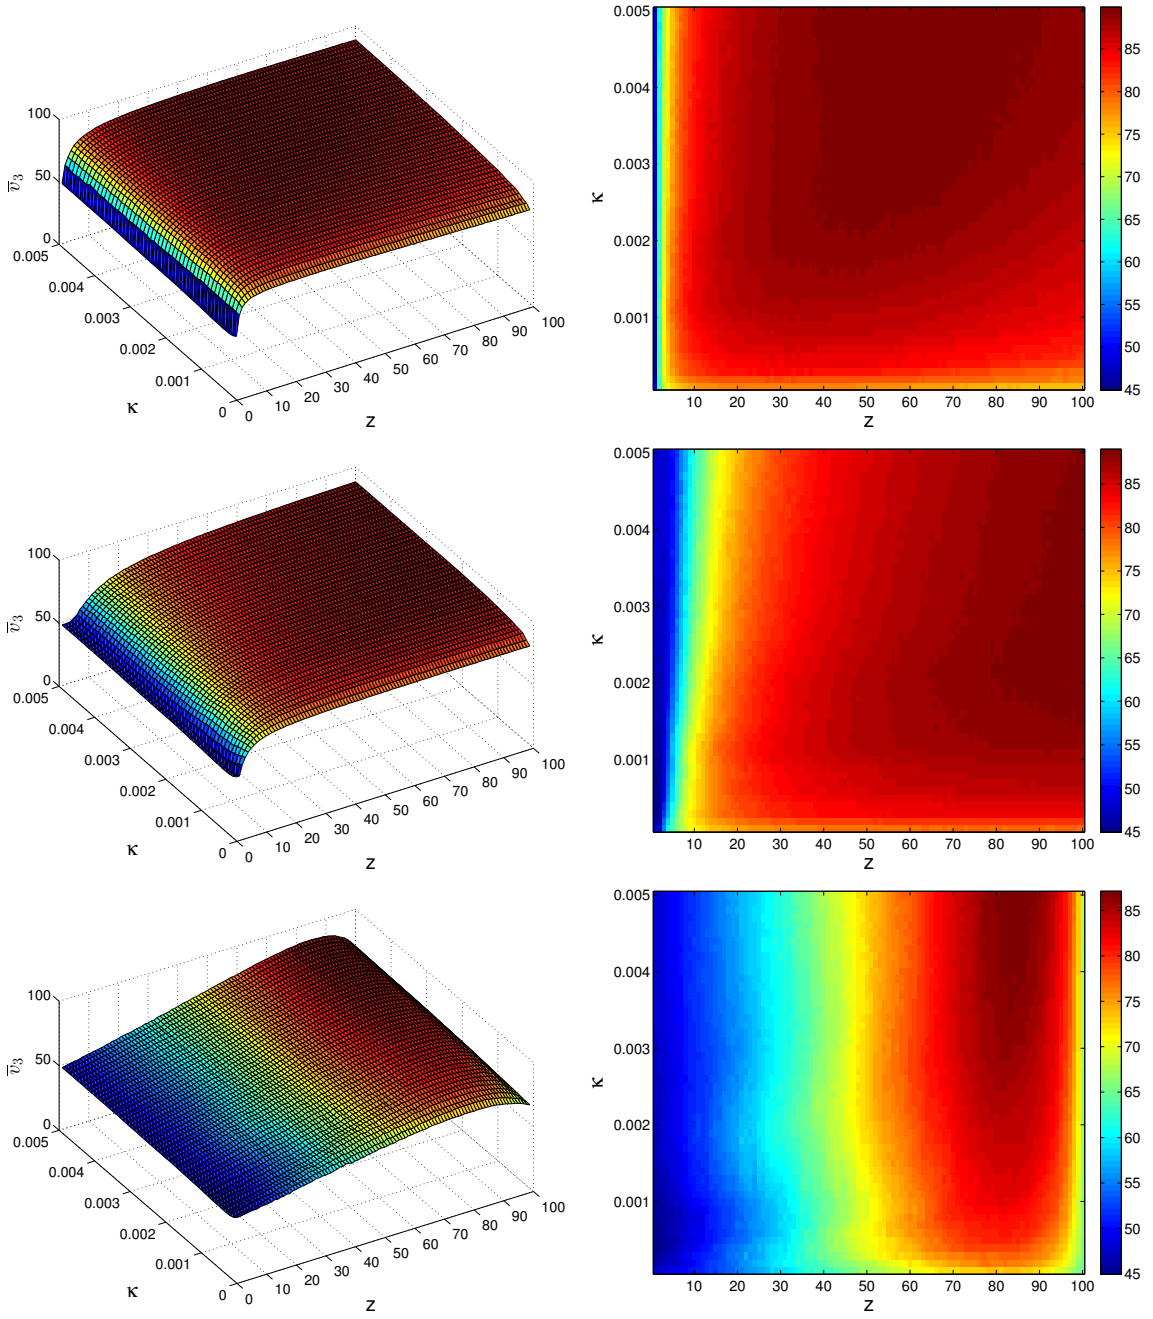

Figure 7.8: (Left column) Surface plots and (right column) heat maps of the values of $\overline{v}_3$ for the (top row) gangs, (middle row) tiers, and (bottom row) core-periphery networks with given values of z and κ.

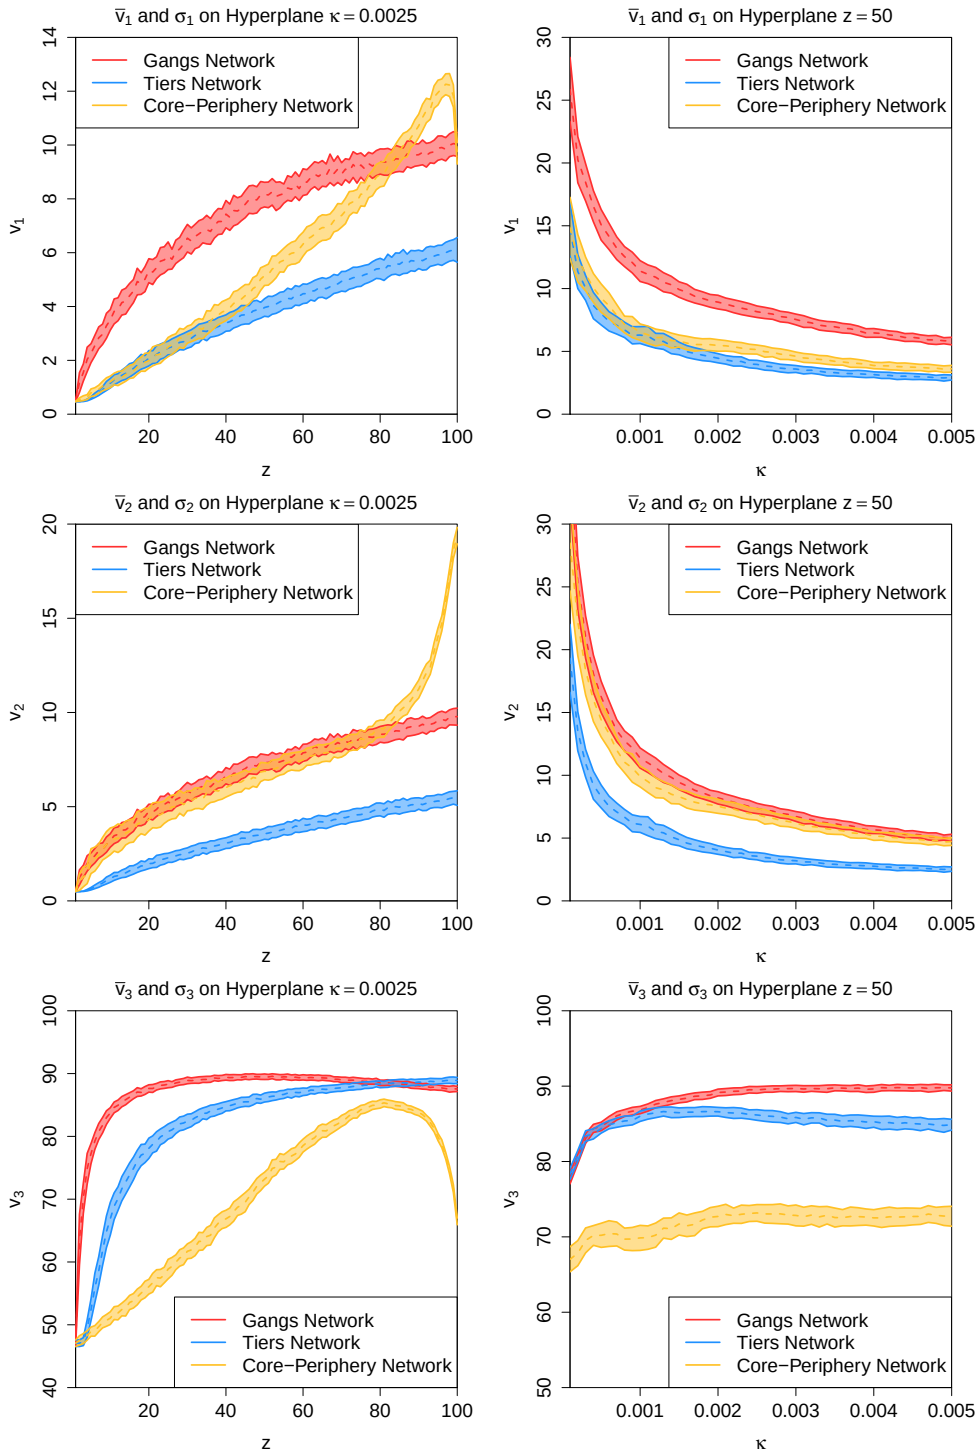

Figure 7.9: Values of (top row) $\overline{v}_1$ and σ_1, (middle row) $\overline{v}_2$ and σ_2, and (bottom row) $\overline{v}_3$ and σ_3 for the gangs, tiers, and core-periphery networks. The left column shows the volatility statistics along the hyperplane $\kappa = 0.0025$ and the right column shows the volatility statistics along the hyperplane $z = 50$. In each plot, the dashed lines denotes the sample mean and the shaded regions denote the sample mean plus or minus one sample standard deviation.

we use the same 100 seeds for our simulations (see Section 7.6.1). When the trade-partnership network is complete, $\overline{v}_1 \approx \overline{v}_2 \approx 0.5$ ticks for all $\kappa > 0$.

For all other choices of z and for all $\kappa > 0$, the values of $\overline{v}_1$ and $\overline{v}_2$ for the gangs networks are greater than or equal to the corresponding values for the tiers networks. This suggests that the inter-tier edges in the tiers networks constitute a stabilizing feedback mechanism for the trade-price series. The buy (respectively, sell) prices of agents in a given tier are bounded by the sell (respectively, buy) prices of agents in the same tier and in the adjacent tiers. This prevents the different tiers from drifting apart and thereby from generating large price changes when trades occur. In a gangs network, by contrast, the buy (respectively, sell) prices of agents in a given tier are only bounded by the sell (respectively, buy) prices of agents in the same tier. The only feedback mechanism between different gangs is the common trade price $\phi(t)$, and the behaviour of agents in different gangs is conditionally independent, given $\phi(t)$.

Despite this difference, both the $\overline{v}_1$ and $\overline{v}_2$ surfaces have qualitatively similar shapes for the gangs and tiers networks. Increasing z (which corresponds to the agents dividing into a larger number of gangs/tiers) tends to produce larger values of $\overline{v}_1$ and $\overline{v}_2$. Decreasing κ (which corresponds to the agents undergoing weaker reversion towards $\phi(t)$) tends to produce larger values of $\overline{v}_1$ and $\overline{v}_2$. The values of $\overline{v}_1$ and $\overline{v}_2$ are particularly large when κ is close to 0.

The behaviour of $\overline{v}_1$ and $\overline{v}_2$ for the core-periphery networks is less straightforward. Decreasing κ tends to produce larger values of $\overline{v}_1$ and $\overline{v}_2$, but the shape of the $\overline{v}_1$ and $\overline{v}_2$ surfaces differ because of their different behaviour for large z. For any fixed $\kappa > 0$, increasing z tends to produce larger values of $\overline{v}_1$ until $z \approx 95$. For values of z larger than about 95, increasing z tends to produce *smaller* values of $\overline{v}_1$. By contrast, increasing z tends to produce larger values of $\overline{v}_2$ for all values of z.

When the trade-partnership network is complete, $\overline{v}_3 \approx 50$, which implies that approximately half of all trades occur at the same price as the previous trade.[9] For the gangs networks, increasing z tends to produce larger values of $\overline{v}_3$ until $z \approx 20$, beyond which $\overline{v}_3$ reaches a plateau then decreases slightly. This suggests that when agents split into large numbers of gangs, further increasing z tends to increase the percentage of trades that occur at the same price as the previous trade. For the tiers networks, increasing z tends to produce larger values of $\overline{v}_3$ for all choices of z. For the core-periphery networks, increasing z tends to produce larger values of $\overline{v}_3$ until

[9]For all $\kappa > 0$, the exact percentage of non-zero price changes in this case is slightly below 50% because active buy prices $p_{-i}(t) \neq \phi(t)$ and active sell prices $p_j(t) \neq \phi(t)$ undergo strictly positive drift towards $\phi(t)$, whereas any active buy or sell prices that equal $\phi(t)$ have 0 drift.

$z \approx 80$. For values of z larger than about 80, increasing z tends to produce smaller values of $\overline{v}_3$.

7.7 Discussion

In this section, we discuss several interesting features of our simulation results.

7.7.1 Complete Trade-Partnership Networks

All three volatility measures that we study are smallest when the trade-partnership network is complete. In this case, approximately half of all trades occur at the same price as the previous trade and $\overline{v}_1 \approx \overline{v}_2 \approx 0.5$ ticks. By Proposition 7.1, $\mathcal{M}(t)$ is feasible at each time $t = 0, 1, 2, \ldots$ for any trade-partnership network. Given the discreteness of the pricing grid, for $\mathcal{M}(t)$ to be feasible when the trade-partnership network is complete requires $p_j(t) - p_{-i}(t) \geq 1$ for all $\theta_i, \theta_j \in \Theta$, so $s(t) = 1$ is a lower bound on the bid-ask spread in this case. Therefore, our results suggest that when all agents are able to trade with all others, the vast majority of the variability in $\phi(t)$ is due to bid-ask bounce [96] (i.e., the tendency for consecutive trades to "bounce" between $b(t)$ and $a(t)$). This case provides a useful benchmark against which to compare our results for other trade-partnership networks.

7.7.2 The Role of κ

In all cases that we study, decreasing κ tends to produce higher volatility in the trade-price series. Particularly small choices of κ produce high trade-price volatility. In the perfect-rationality interpretation of our model that we discuss in Appendix A.4, smaller choices of κ correspond to agents assigning less utility to the value of $w(t)$. Although we restrict our attention to choices of κ that are strictly positive,[10] we note that the case $\kappa = 0$ can also be interpreted as the agents not knowing the value of $\phi(t)$ (and thus being unable to use this information). Therefore, our results suggest that providing traders with information of $\phi(t)$ creates a strong stabilizing force on the trade-price series on a quasi-centralized trading platform. Although our model is quite different from theirs, Perraudin and Vitale [179] arrived at a similar conclusion from their ABM of trade in the FX spot market (see Section 7.3.2).

[10]For any configuration of the trade-partnership network, if $\kappa = 0$ then the expected length of time between a successive pair of trades is infinite (see the result for a simple symmetric random walk in [174]). Therefore, we do not study this case.

7.7.3 Large-z Behaviour of Core-Periphery Networks

In the core-periphery networks, when z becomes close to 100, further increasing its size leads to an increase in $\overline{v}_2$ but a decrease in both $\overline{v}_1$ and $\overline{v}_3$. Given that all absolute price changes are, by definition, non-negative, this result suggests that core-periphery networks with very small cores produce more trades with very small (or zero) and very large price changes than do those with slightly larger cores.

Although we are unable to provide a rigorous proof due to the complicated, non-linear temporal evolution of the model, we conjecture that this behaviour occurs due to the antagonistic interaction of two opposite tendencies. On the one hand, increasing z reduces the total number of edges in a core-periphery network. This tends to increase the values of $\overline{v}_1$, $\overline{v}_2$, and $\overline{v}_3$. On the other hand, all trades in a core-periphery network must involve the buy or sell price of at least one core agent. We call such a price a *core price*. Increasing z reduces the number of agents in the core, and therefore reduces the total number of core prices.

We conjecture that in a given run of the model, the number of pairs of subsequent trades that involve the same core price is likely to be larger for larger values of z, and that the absolute price change caused by such pairs of subsequent trades is likely to be smaller (and more likely to be 0) than for other pairs of subsequent trades. For example, assume that a pair of subsequent trades both involve core agent θ_j selling the asset, and that the first such trade occurs at time t'. Immediately after the first trade, the absolute price difference between $p_j(t')$ and $\phi(t')$ is at most 1. This is not necessarily the case for the other core prices. Therefore, given that the next trade also involves p_j, the next trade is likely to occur at or close to $\phi(t')$.

With $M = 200$ agents, the parameterization $\eta_C = (z-1)/100$ (see Section 7.6.3) implies that the total number of edges in a core-periphery network is given by

$$\binom{200 - 2(z-1)}{2} + 2(z-1), \tag{7.33}$$

where the first term counts the edges between the core agents and the second term counts the edges that connect the periphery agents to the core. For small values of z, increasing z causes a large decrease in the total number of edges, but for large values of z, increasing z causes a much smaller decrease in the total number of edges. Similarly, when z is small, it seems likely that increasing z would only slightly increase the number of pairs of subsequent trades that involve the same core price, because there are still many other core prices. When z is large, by contrast, it seems likely that increasing z would lead to a much greater increase in the number of pairs of

subsequent trades that involve the same core price. Therefore, the influence of the first effect is strong when z is small and small when z is large, whereas the opposite seems to be true for the second effect. We believe that the reduced influence of the first effect and the increased influence of the second effect explains the behaviour that we observe for $\overline{v}_1$, and $\overline{v}_3$ for large values of z.

7.8 Summary and Future Work

In this chapter, we introduced and studied an ABM of trading on a quasi-centralized trading platform. In our model, agents revise their buy and sell prices according to a zero-intelligence update rule based on their observations of the trade-price series $\phi(t)$ and a trade-price-reversion parameter κ. Importantly, the only sources of volatility in our model are updates to agents' private valuations and the agents' heterogeneous access to trading opportunities. We derived several analytical results about the state and temporal evolution of our model, and we simulated our model using several different classes of trade-partnership networks and with several different choices of the trade-price-reversion parameter κ.

We identified several links between both the trade-partnership network and trade-price-reversion parameter κ and volatility in the trade price series. In some cases, we found that decreasing the total number of edges in the trade-price network produced lower volatility, but in other cases we found that the opposite was true. Moreover, we identified clear differences between the volatility produced by the three different classes of trade-partnership networks that we studied. Therefore, our results illustrate that the topology of the trade-price network, and not just the total number of edges, plays an important role in the volatility of the trade-price series generated by the model.

Our results suggest that even when agents do not seek to optimize "utility" or achieve trading goals, providing them with access to the trade-price series can help to reduce volatility. This finding is particularly interesting given that our results from Chapter 4 suggest that traders on Hotspot FX regard the trade-price series to contain important information.

7.8.1 Future Work

Our results suggest several possible avenues for future research.

- **Different trade-partnership networks:** For our simulations, we restricted our attention to three simple classes of trade-partnership networks. It would be interesting to investigate other classes of trade-partnership networks to help understand how other configurations influence volatility. Studying random graphs or networks with less regular structure than those that we have studied would be a particularly interesting area for future research.

- **Real trade-partnership networks:** If it were possible to obtain information regarding the structure of a real-world trade-partnership network on an electronic trading platform, then it would be interesting to examine how the volatility of the trade-price series generated by our model compares to that of real data. For confidentiality reasons, exchanges do not disclose any information regarding trade partnerships between traders. However, a small survey of financial institutions could help to illuminate some local properties of the trade-partnership network, and could thereby provide a useful starting point.

- **Local volatility:** Although our calculations highlight several interesting features of the volatility of the trade-price series, they do not provide insight into the local volatility that individual traders experience. It would be interesting to investigate how local volatility varies across traders and to understand how volatility depends on the structure of a trade-partnership network.

- **Other sources of volatility:** Several factors influence volatility in financial markets (see Section 7.1). Extending our model to better reflect these factors could help it to reflect more closely the volatility observed in real markets.

- **Agent heterogeneity:** All agents in our model update their buy and sell prices according to the same set of rules and using the same trade-price-reversion parameter κ. Incorporating heterogeneity between different agents could help to generate more realistic price dynamics.

- **Trading mechanisms:** Our ABM does not consider the trading mechanism by which agents interact. It would be interesting to extend our model to consider a specific trading mechanism, such as an LOB, to examine whether this affects volatility in the trade-price series.

Chapter 8

Conclusions

The introduction of electronic trading has been one of the most important changes in the global FX market since the collapse of the Bretton Woods Agreement in the early 1970s [53, 85]. The rapid rise in popularity of ECNs has blurred the lines between the traditional customer and dealer roles and thereby catalysed profound structural change: as barriers to entry have fallen, a wide new class of financial institutions have begun to participate in a market that is more accessible, more transparent, and more liquid than ever before [125].

The aim of this thesis was to address several questions regarding the use of multi-institution trading platforms in the FX spot market. In Chapter 2, we discussed the microstructure of such platforms. We highlighted several important features of quasi-centralized LOBs that distinguish them from the centralized LOBs that are commonly used in other financial markets, and we introduced an undirected network representation of the trading partnerships on a given platform. In Chapter 3, we listed several important features of the Hotspot FX platform and provided a detailed discussion of the data that we studied in the subsequent chapters. We also described our methods for pre-processing and cleaning the Hotspot FX data and for reconstructing the global LOB at a given moment in time.

In Chapter 4, we performed an empirical study of order flow and LOB state on Hotspot FX. We uncovered several statistical regularities in both quote-relative and trade-relative coordinates, and we highlighted how the quasi-centralized structure of trading on Hotspot FX contrasts our findings to those reported by empirical studies of centralized LOBs in other markets. The extremely high levels of activity on Hotspot FX enabled us to calculate several interesting statistics and distributions on each trading day separately and to identify differences between the results for each day and the results for the whole period. We argued that these findings bring into question whether models that are designed to reproduce long-run statistics generate

useful output on shorter time horizons. We found substantial differences between the distributions of limit order arrivals and cancellations on different days, but we also found that rescaling the data to account for differences in location and scale caused the distributions from each day to collapse onto a single curve. We investigated the strength of this collapse by comparing the out-of-sample performance of two statistical models. We concluded that the first two moments of the distributions provide significant explanatory power for daily order flow.

In Chapter 5, we investigated the autocorrelation properties of returns, absolute returns, and order flow on Hotspot FX. We found that returns do not exhibit any significant autocorrelation at even a 1-second lag, but we found strong, statistically significant evidence for long memory in absolute returns and order flow. We tested and rejected the hypothesis that this apparent long memory was an artifact caused by structural breaks at daily boundaries. This result contrasts to the findings of a recent study of order flow on the LSE [7]. We conjectured that this contrast could be caused by differences between the structure of trading days in the two markets. We argued that the quasi-centralized structure of trading on Hotspot FX undermines an argument that long memory in order flow is a consequence of different traders mimicking each other [135], and instead favours the hypothesis that long memory is a consequence of strategic order splitting [33, 34, 137].

In Chapter 6, we performed an empirical study of price formation on Hotspot FX. We found that the bypassing price of the vast majority of trades on the platform was small or zero, and we argued that the handful of trades with large bypassing prices were caused by market order submissions from traders with very few trading partners. We showed that changes in the trade-price series can be decomposed exactly into corresponding changes in the quote-price and bypassing-price series, and we found that the dispersion importance of changes in the quote-price series for explaining changes in the trade-price series was much higher than that of changes in the bypassing-price series. We thereby concluded that most traders on Hotspot FX are sufficiently well-connected with other traders on the platform that the restriction of trading only with their trading partners does not strongly affect the prices of their trades.

In Chapter 7, we introduced and studied an ABM of trade on a quasi-centralized trading platform. In our model, the temporal evolution of the agents' buy and sell prices is governed by the network of trading partnerships and a scalar trade-price-reversion parameter. Importantly, the only sources of volatility in our model are the updates to agents' private valuations and the heterogeneity in agents' access to trading opportunities. Therefore, our ABM enabled us to investigate how changes in

the trade-partnership network influence the volatility of the trade-price series. We designed our model to be sufficiently simple to enable us to derive several results regarding its temporal evolution. We simulated our ABM using several different trade-partnership networks and several different choices of the trade-price-reversion parameter. We found that the topology of the network, and not just its number of edges, influences volatility in the model. We noted that small choices of the trade-price-reversion parameter tend to produce particularly high volatility, and we thereby concluded that providing traders on quasi-centralized trading platforms with real-time access to the trade-price series constitutes a strong stabilizing force.

Throughout this thesis, we have employed a wide range of techniques, estimators, and models to help illuminate several important aspects of trading via multi-institution trading platforms in the FX spot market. Our results highlight many important differences between these platforms and the electronic trading platforms used in other markets. Although our ABM suggests that the quasi-centralized structure of trading could exacerbate trade-price volatility in extreme situations, our empirical calculations indicate that the vast majority of traders on Hotspot FX have sufficiently many trading partners on the platform to experience little or no impact on price formation when compared to that in a centralized LOB. We thereby conclude that multi-institution trading platforms in the FX spot market retain many desirable features of centralized markets while providing traders with explicit control over their personal trading partnerships.

8.1 Future Research

As we have noted in Chapters 4–7, our results suggest many possible avenues for future research. To conclude, we discuss a selection of open problems related to our findings.

Developing market safeguards: Although there are many clear benefits to electronic trading, there are also some important drawbacks. For example, some electronic trading algorithms have caused so-called "flash crashes" [126], during which markets have undergone extreme price changes over short time intervals. Although some authors have studied possible regulatory interventions (such as "limit up/limit down" rules [66, 110] and circuit breakers [133]) that could help to safeguard against possible instabilities, much more work in this area is needed, given the potential disruption that these market malfunctions can cause.

The technology arms race: Our results in Chapter 5 suggest that electronic trading algorithms are extremely quick to detect and capitalize on statistical arbitrage opportunities. Competition between traders who seek to identify such opportunities has led to many financial institutions investing enormous sums of money in order to shave mere milliseconds from their trading latency [50]. Given that many traders cannot afford the large investments required to compete in this technology arms race, some authors have claimed that continuous-time LOB trading is a flawed market design [6, 39]. By contrast, others have concluded that such electronic trading algorithms provide a positive impact on markets by accelerating the price-discovery process and increasing liquidity [44, 115]. Detailed study of the impact of such trading algorithms could provide useful input into the ongoing regulatory debate regarding this issue.

Understanding prime brokerage: In recent years, a new market service has arisen to provide traders with few trading partnerships the ability to access better prices on quasi-centralized trading platforms. A *prime broker* is a large trading institution that, in exchange for a fee, allows smaller institutions to trade using its credit lines, thus widening the smaller institutions' access to trading opportunities [43, 54]. According to the 2013 BIS Triennial Central Bank Survey [13], about 16% of all FX transactions occur via a prime-brokerage relationship. Given the extremely large turnover of the market, this represents an enormous volume of trade. Understanding how prime brokers impact the market is therefore an important question for future research [54].

Temporal evolution of the trade-partnership network: Throughout this thesis, all of our investigations have assumed that the structure of the trade-partnership network remains fixed. However, both the number of traders who access a given multi-institution trading platform and the trading partnerships between traders are likely to change over time. At present, temporal networks are an active area of study in several different disciplines [117]. Understanding the temporal evolution of trade-partnership networks presents a rich set of modelling challenges regarding how and why traders forge and sever trading partnerships with their peers.

Emergent instabilities: It remains an important open question to understand how market-wide instabilities emerge from the microscopic-scale interactions between individual traders. This question is particularly pertinent for electronic trading platforms in the FX spot market because the widespread uptake of such platforms has encouraged market participation from a broad new class of financial institutions,

many of whom use highly leveraged strategies and are therefore particularly vulnerable to market turbulence. Models that explicitly incorporate heterogeneity in different traders' motivations and trading styles are required to help understand the mechanisms by which changes in local interactions amplify into system-wide destabilizing forces that impact market stability.

Understanding market structure: The notion that price formation can depend on market structure is a central theme of the market microstructure literature [189]. Many studies assume that markets are either fully centralized, with all traders able to trade with all others, or fully decentralized, with disjoint pockets of local liquidity. In reality, however, trade for most assets occurs via a complicated hybrid of different styles that reside somewhere between these two extremes. For example, trade for most equities occurs primarily on a central exchange such as the LSE or NYSE, but many equities are also traded off-exchange, in over-the-counter markets [167]. Moreover, many assets are traded on several different trading platforms simultaneously. Even if each such trading platform is fully centralized, the overall market for the asset is not. Investigating the wide spectrum of structures in modern financial markets presents many opportunities for future research.

Other currency pairs and other platforms: All of our empirical calculations have been based on the activity of the EUR/USD, GBP/USD, and EUR/GBP currency pairs on Hotspot FX. Further empirical study could help to illuminate similarities and differences with the behaviour of other currency pairs and other multi-institution trading platforms. It would be particularly interesting to compare our estimates of dispersion importance in Chapter 6 to similar estimates for less liquid currency pairs or for activity on the Reuters Dealing 3000 and EBS platforms (which are more heavily geared towards the traditional inter-dealer market).

This selection of open questions represents just a small subset of those related to the ongoing transformations in global financial markets. In the coming years, it will be interesting to see how new technologies shape the markets of the future, and, in turn, what new challenges these developments will present for future research.

Appendix A

Appendices

A.1 Relative Importance

A common goal of fitting a multiple-regression model is to quantify the relative contributions of the input variables to some specified aspect of the model's output [24, 65]. The relative contribution of each input variable is called its *relative importance* [21]. There are several different measures of relative importance (see [21] for a recent review), and the type of relative importance measured in a given situation depends on the purpose of the analysis. In this appendix, we provide details of how to calculate the dispersion importance, which is the type of relative importance that we study in Chapter 6. The dispersion importance measures the fraction of the variance of the output variable that each input variable explains. For a discussion of several other types of relative importance, see [1].

Consider a multiple regression between k input variables $x_1, x_2, \ldots, x_k$ and one output variable y. Given a sample of $i = 1, 2, \ldots, N$ observations $(x_{i,1}, x_{i,2}, \ldots, x_{i,k}, y_i)$, let

$$\hat{y}_i = \alpha_0 + \alpha_1 x_{i,1} + \alpha_2 x_{i,2} + \ldots + \alpha_k x_{i,k}, \tag{A.1}$$

where $\alpha_0, \alpha_1, \ldots, \alpha_n \in \mathbb{R}$ are the least-squares parameter estimates that minimize the mean squared error

$$\frac{1}{N} \sum_{i=1}^{N} (\hat{y}_i - y_i)^2 .$$

Let

$$\bar{y} = \frac{1}{N} \sum_{i=1}^{N} y_i \tag{A.2}$$

denote the sample mean of the y_i and let

$$SS_{mod} = \sum_{i=1}^{N} (\hat{y}_i - \bar{y})^2 , \tag{A.3}$$

$$SS_{tot} = \sum_{i=1}^{N} (y_i - \bar{y})^2 \tag{A.4}$$

denote respectively the multiple regression's *model sum of squares* and *total sum of squares*. The *coefficient of determination*

$$R^2 = \frac{SS_{mod}}{SS_{tot}} \tag{A.5}$$

measures the fraction of the sample variance of the output variable that the multiple regression equation (A.1) explains.

We use a technique called the Lindeman–Merenda–Gold (LMG) measure [138] to compute the dispersion importance of the input variables in the multiple regressions that we study in Chapter 6. Calculating the LMG measure is computationally infeasible for multiple regressions with large numbers of independent variables, but each of our multiple regressions contains only two independent variables, so we are able to perform the necessary calculations with relative ease.

Let $S = \{x_1, x_2, \ldots, x_k\}$ denote the set of k input variables in the full multiple regression. Let $s \subseteq S$ denote a subset of S, and let R_s^2 denote the coefficient of determination from the multiple regression using only the input variables in s. For example, if $S = \{x_1, x_2, x_3, x_4\}$ and $s = \{x_1, x_2\}$, then R_s^2 is the coefficient of the determination from the multiple regression

$$\hat{y}_i = \beta_0 + \beta_1 x_{i,1} + \beta_2 x_{i,2}. \tag{A.6}$$

For a set $s \subset S$ and an independent variable $x_j \in S$ but $x_j \notin s$, the *sequential sum of squares for x_j given s* is

$$\hat{R}_s^2 (x_j) = R_{\{s \cup x_j\}}^2 - R_s^2. \tag{A.7}$$

That is, $\hat{R}_s^2 (x_j)$ measures the additional variance of the output variable that a multiple regression including the input variables in $s \cup x_j$ explains, compared to that of a multiple regression including only the input variables in s.

Let r denote a permutation of the integers $1, 2, \ldots, k$, and let $s_j(r)$ denote the set of input variables whose indices appear before the index j in r. For example, if $k = 5$ and $r = (2, 5, 3, 1, 4)$, then $s_1(r) = \{x_2, x_5, x_3\}$.

For each input variable x_j, the *dispersion importance* [21, 138]

$$I_j = \frac{1}{k!} \sum_r \hat{R}^2_{s_j(r)}(x_j) \tag{A.8}$$

is the mean sequential sum of squares, averaged over all possible permutations of the input variables in the full multiple regression.

A.2 Estimating Standard Error

In this section, we provide details of the two methods that we use to estimate the standard error of the dispersion importance $\tilde{I}_1$ in Chapter 6.

A.2.1 The Non-Parametric Bootstrap

The first technique that we use to estimate the standard error is the non-parametric bootstrap [118]. The technique is performed as follows. For each $d = 1, 2, \ldots, d_{\max}$, draw (with replacement) a *bootstrap sample* (i.e., a random sample of size N from the N observations $(x_{i,1}, x_{i,2}, \ldots, x_{i,k}, y_i)$). Let $\tilde{I}_j^{\,d}$ denote the normalized dispersion importance for the input variable x_j given the d^{th} bootstrap sample. For each $j = 1, 2, \ldots, k$, the non-parametric bootstrap estimate of the standard error of $\tilde{I}_j$ is the standard deviation of the $\tilde{I}_j^{\,d}$ across $d = 1, 2, \ldots, d_{\max}$.

For our calculations in Chapter 6, we use $d_{\max} = 1000$ bootstrap samples. To verify that this choice produces a statistically stable estimate of the standard error, we repeated all of our calculations using only $d_{\max} = 500$ bootstrap samples. In all cases, we found that our estimates of the standard error using $d_{\max} = 500$ and $d_{\max} = 1000$ were equal to 3 decimal places.

A.2.2 The Variance-Plot Method

A possible problem with using the non-parametric bootstrap to estimate the standard error of an estimator is that the bootstrap samples do not posses the same autocorrelation structure as the original observations. Therefore, the non-parametric bootstrap can produce misleading estimates of an estimator's standard error if the original observations are serially correlated.

If the original observations contain only short-range autocorrelations (see Section 5.2.1), then so-called "block bootstrapping" techniques can help to retain the autocorrelation structure in the bootstrap samples (see [132] for a review). If, however, the original observations contain long-range autocorrelations (which our results in

Chapter 5 suggest could be the case), then block-bootstrap methods also produce misleading estimates [111].

To help overcome this problem, Beran [18] introduced a heuristic called the "variance-plot method".[1] The method is performed as follows. For each $n = 3, 4, \ldots$, divide the original N observations into N/n non-overlapping blocks.[2] For each input variable x_j, estimate $\tilde{I}_j$ on each block separately and calculate the sample standard deviation of the estimates across the N/n blocks. Plot this standard deviation against n/N on doubly logarithmic axes, and fit an OLS regression line to the plot. The variance-plot estimate for the standard error is given by the value of the OLS line at the point $N/n = 1$. Figure A.1 illustrates this process for the bid-side calculations in Table 6.4.

A.3 Results for Alternative Window Lengths

The results that we present in the main text of Chapter 6 correspond to the window length $\tau = 0.3$ seconds in Algorithm 6.1. Tables A.1 and A.2 list the corresponding results of our main statistical calculations from the chapter for $\tau = 0.1$ and $\tau = 0.5$, respectively.

A.4 The Logistic Sigmoid Function

Throughout Chapter 7, we present our ABM as a zero-intelligence model in which agents update their buy and sell prices according to the stochastic process $G(t)$. However, the model formulation that we adopt also arises as the solution of a decision problem in which agents are assumed to be perfectly rational [165]. In this appendix, we provide details of this result for the specific case of our ABM, in which agents choose between binary outcomes (i.e., a price revision of ± 1), the only property of the model state that influences agents' utility of a given alternative is $w(t)$, and all agents use the same utility function. McFadden [165] derives the solution for the general case, which results in the popular *multinomial logit model.*

Recall that $P(w(t))$ denotes the probability of a downward revision to the active price. Let $U(w(t))$ denote the utility that the active agent experiences for revising

[1] Farmer *et al.* [78] also provide a detailed discussion of this technique.

[2] If n does not divide N exactly, ensure that the blocks are as close to equal-sized as possible (i.e., that no two blocks differ in size by more than 1).

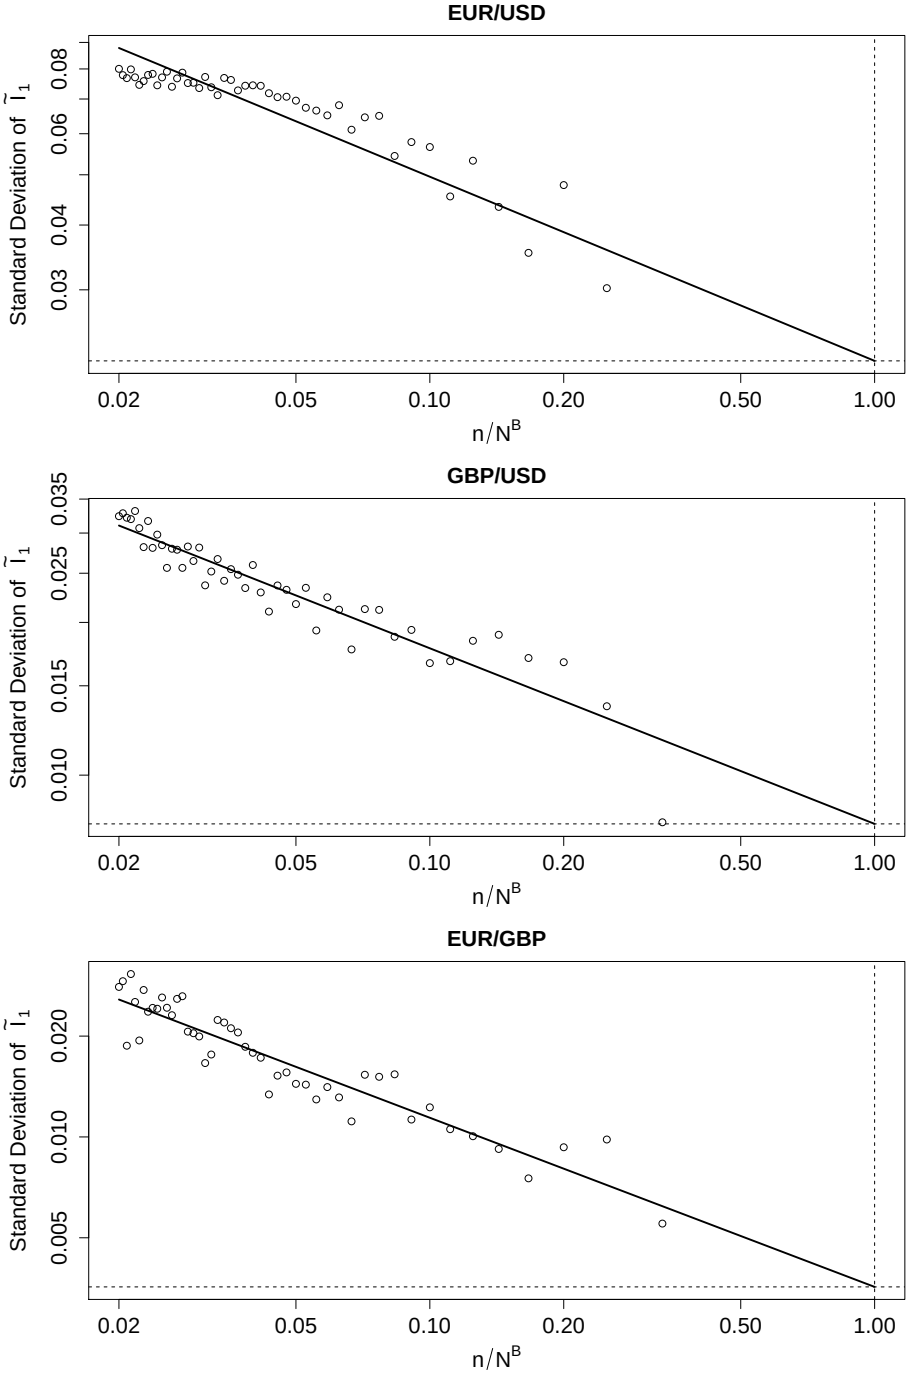

Figure A.1: Variance plots for the standard error of the normalized dispersion importance $\tilde{I}_1$ of the change in bypassing price for the change in trade price for (top) EUR/USD, (middle) GBP/USD, and (bottom) EUR/GBP seller-initiated trades. The points denote the standard deviation of the estimates of $\tilde{I}_1$ among blocks of the given length. The solid black line denotes the OLS regression line. The corresponding plots for the buyer-initiated trades are qualitatively similar.

		EUR/USD		GBP/USD		EUR/GBP	
		Bid side	Ask side	Bid side	Ask side	Bid side	Ask side
Bypassing Price	Mean (ticks)	2.73	2.64	3.30	3.36	2.47	2.27
	St. Dev. (ticks)	5.59	5.97	6.62	7.33	4.86	4.54
	Maximum (ticks)	363	335	107	123	73	43
	Percentage Non-Zero	45.8%	45.1%	39.3%	36.3%	37.3%	35.7%
Percentage Non-Zero Changes	Trade Price	76.7%	77.3%	74.9%	75.8%	61.4%	66.2%
	Bypassing Price	55.0%	54.6%	48.9%	45.0%	41.3%	41.0%
	Quote Price	76.0%	76.2%	76.8%	76.1%	61.4%	67.7%
OLS R^2	First OLS Regression	0.083	0.088	0.055	0.053	0.030	0.023
	Second OLS Regression	0.895	0.884	0.942	0.944	0.974	0.975
	Third OLS Regression	0.001	0.002	0.000	0.000	0.000	0.000
Dispersion Importance	$\tilde{I}_1$	0.094	0.102	0.056	0.054	0.028	0.024
	Bootstrap s.e.	0.012	0.011	0.002	0.002	0.004	0.004
	Variance-Plot s.e.	0.018	0.025	0.011	0.009	0.003	0.003

Table A.1: Results of our main statistical calculations from Chapter 6 when using a window length of $\tau = 0.1$ in Algorithm 6.1. The results regarding bypassing price, percentage of non-zero changes, OLS R^2, and dispersion importance correspond to Tables 6.1, 6.2, 6.3, and 6.4, respectively (for which we use $\tau = 0.3$).

		EUR/USD		GBP/USD		EUR/GBP	
		Bid side	Ask side	Bid side	Ask side	Bid side	Ask side
Bypassing Price	Mean (ticks)	2.53	2.46	3.15	3.12	1.92	1.77
	St. Dev. (ticks)	5.40	6.01	6.37	6.87	4.20	3.88
	Maximum (ticks)	363	335	119	142	73	43
	Percentage Non-Zero	44.0	43.5	40.0	37.0	34.0	32.0
Percentage Non-Zero Changes	Trade Price	78.0	79.1	79.6	81.4	76.2	80.9
	Bypassing Price	54.4	53.9	51.5	48.5	43.2	44.2
	Quote Price	77.1	77.7	80.4	81.2	75.6	80.7
OLS R^2	First OLS Regression	0.081	0.085	0.056	0.048	0.024	0.016
	Second OLS Regression	0.899	0.883	0.947	0.950	0.983	0.979
	Third OLS Regression	0.001	0.003	0.000	0.000	0.000	0.000
Dispersion Importance	$\widetilde{I}_1$	0.091	0.101	0.054	0.049	0.020	0.019
	Bootstrap s.e.	0.015	0.014	0.002	0.002	0.003	0.003
	Variance-Plot s.e.	0.030	0.031	0.009	0.006	0.006	0.005

Table A.2: Results of our main statistical calculations from Chapter 6 when using a window length of $\tau = 0.5$ in Algorithm 6.1. The results regarding bypassing price, percentage of non-zero changes, OLS R^2, and dispersion importance correspond to Tables 6.1, 6.2, 6.3, and 6.4, respectively (for which we use $\tau = 0.3$).

the active price downwards. If there exists a constant κ such that at each time t, the utility function $U(w(t))$ can be written in the form

$$U(w(t)) = \kappa w(t) + \varepsilon(t), \tag{A.9}$$

where the $\varepsilon(t)$ are independent Logistic$(0, 1)$ random variables and such that the agent chooses a downward price revision if and only if $U(w(t)) > 0$, then

$$\begin{aligned} P(w(t)) &= \mathbb{P}\left(U(w(t)) > 0\right), \\ &= \mathbb{P}\left(\kappa w(t) + \varepsilon(t) > 0\right), \\ &= \mathbb{P}\left(\varepsilon(t) > -\kappa w(t)\right), \\ &= 1 - \mathbb{P}\left(\varepsilon(t) \leq -\kappa w(t)\right), \\ &= 1 - \frac{1}{1 + e^{\kappa w(t)}}, \\ &= \frac{1}{1 + e^{-\kappa w(t)}}. \end{aligned}$$

Therefore, if all agents in our model are assumed to be perfectly rational and have a common utility function of the form shown in Equation (A.9), then the price-update function P is given by a logistic sigmoid function with parameter $\kappa w(t)$, where κ is determined by the relationship in (A.9).

References

[1] C. H. Achen. *Interpreting and Using Regression*. Sage, Newbury Park, CA, USA, 1982.

[2] Y. Aït-Sahalia, P. A. Mykland, and L. Zhang. Ultra high frequency volatility estimation with dependent microstructure noise. *Journal of Econometrics*, 160(1):160–175, 2011.

[3] A. Alfonsi, A. Fruth, and A. Schied. Optimal execution strategies in limit order books with general shape functions. *Quantitative Finance*, 10(2):143–157, 2010.

[4] R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3(2):5–39, 2001.

[5] D. W. K. Andrews. Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica*, 59:817–858, 1991.

[6] S. Arnuk and J. Saluzzi. Latency arbitrage: The real power behind predatory high frequency trading. Technical report, Themis Trading, available at `http://www.themistrading.com/article_files/0000/0519/THEMIS_TRADING_White_Paper_--_Latency_Arbitrage_--_December_4__2009.pdf`, 2009.

[7] C. Axioglou and S. Skouras. Markets change every day: evidence from the memory of trade direction. *Journal of Empirical Finance*, 18(3):423–446, 2011.

[8] R. T. Baillie. Long memory processes and fractional integration in econometrics. *Journal of Econometrics*, 73(1):5–59, 1996.

[9] R. T. Baillie and P. C. McMahon. *The Foreign Exchange Market: Theory and Econometric Evidence*. Cambridge University Press, Cambridge, UK, 1990.

[10] P. Bak, M. Paczuski, and M. Shubik. Price variations in a stock market with many agents. *Physica A*, 246(3):430–453, 1997.

[11] Bank for International Settlements. Triennial central bank survey: Foreign exchange and derivatives market activity in 2001. Technical report, Bank for International Settlements, available at `http://www.bis.org/publ/rpfx02t.pdf`, 2001.

[12] Bank for International Settlements. Triennial central bank survey: report on global foreign exchange market activity in 2010. Technical report, Bank for International Settlements, available at `http://www.bis.org/publ/rpfxf10t.pdf`, 2010.

[13] Bank for International Settlements. Foreign exchange turnover in April 2013: preliminary global results. Technical report, Bank for International Settlements, available at `http://www.bis.org/publ/rpfx13fx.pdf`, 2013.

[14] W. Barker. The global foreign exchange market: growth and transformation. *Bank of Canada Review*, 4:3–12, 2007.

[15] O. E. Barndorff-Nielsen and N. Shephard. Volatility. In R. Cont, editor, *Encyclopedia of Quantitative Finance*, pages 1898–1901. Wiley, Chichester, UK, 2010.

[16] M. Baron, J. Brogaard, and A. Kirilenko. The trading profits of high frequency traders. Technical report, The U.S. Commodity Futures Trading Commission, available at `http://faculty.chicagobooth.edu/john.cochrane/teaching/35150_advanced_investments/Baron_Brogaard_Kirilenko.pdf`, 2012.

[17] J. Beran. Statistical methods for data with long-range dependence. *Statistical Science*, 7:404–416, 1992.

[18] J. Beran. *Statistics for Long-Memory Processes*, volume 61. Chapman and Hall, New York, NY, USA, 1994.

[19] I. Berkes, L. Horváth, P. Kokoszka, and Q. M. Shao. On discriminating between long-range dependence and changes in mean. *The Annals of Statistics*, 34:1140–1165, 2006.

[20] R. N. Bhattacharya, V. K. Gupta, and E. Waymire. The Hurst effect under trends. *Journal of Applied Probability*, 20:649–662, 1983.

[21] J. Bi. A review of statistical methods for determination of relative importance of correlated predictors and identification of drivers of consumer liking. *Journal of Sensory Studies*, 27(2):87–101, 2012.

[22] B. Biais, T. Foucault, and S. Moinas. Equilibrium high-frequency trading. Technical report, Toulouse School of Economics, available at `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1834344`, 2011.

[23] B. Biais, P. Hillion, and C. Spatt. An empirical analysis of the limit order book and the order flow in the Paris Bourse. *The Journal of Finance*, 50(5):1655–1689, 1995.

[24] N. H. Bingham and J. M. Fry. *Regression: Linear Models in Statistics*. Springer, London, UK, 2010.

[25] G. H. Bjønnes and D. Rime. Dealer behavior and trading systems in foreign exchange markets. *Journal of Financial Economics*, 75(3):571–605, 2005.

[26] T. Bollerslev, R. Y. Chou, and K. F. Kroner. ARCH modeling in finance: a review of the theory and empirical evidence. *Journal of Econometrics*, 52(1):5–59, 1992.

[27] B. Bollobás. *Random Graphs*. Cambridge University Press, Cambridge, UK, 2001.

[28] E. Bonabeau. Agent-based modeling: methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, 99(Suppl. 3):7280–7287, 2002.

[29] G. G. Booth and F. R. Kaen. Gold and silver spot prices and market information efficiency. *Financial Review*, 14(1):21–26, 1979.

[30] G. G. Booth, F. R. Kaen, and P.E. Koveos. R/S analysis of foreign exchange rates under two international monetary regimes. *Journal of Monetary Economics*, 10(3):407–415, 1982.

[31] M. D. Bordo. The Bretton Woods international monetary system: a historical overview. In M. D. Bordo and B. Eichengreen, editors, *A Retrospective on the Bretton Woods System: Lessons for International Monetary Reform*, pages 3–108. The University of Chicago Press, Chicago, IL, USA, 1993.

[32] J. P. Bouchaud. Economics needs a scientific revolution. *Nature*, 455:1181, 2008.

[33] J. P. Bouchaud, J. D. Farmer, and F. Lillo. How markets slowly digest changes in supply and demand. In T. Hens and K. R. Schenk-Hoppé, editors, *Handbook of Financial Markets: Dynamics and Evolution*, pages 57–160. North–Holland, Amsterdam, The Netherlands, 2009.

[34] J. P. Bouchaud, Y. Gefen, M. Potters, and M. Wyart. Fluctuations and response in financial markets: the subtle nature of 'random' price changes. *Quantitative Finance*, 4(2):176–190, 2004.

[35] J. P. Bouchaud, J. Kockelkoren, and M. Potters. Random walks, liquidity molasses and critical response in financial markets. *Quantitative Finance*, 6(02):115–123, 2006.

[36] J. P. Bouchaud, M. Mézard, and M. Potters. Statistical properties of stock order books: empirical results and models. *Quantitative Finance*, 2(4):251–256, 2002.

[37] J. P. Bouchaud and M. Potters. *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management*. Cambridge University Press, Cambridge, UK, 2003.

[38] M. Buchanan. A little more conversation. *Nature Physics*, 4(9):667, 2008.

[39] E. Budish, P. Cramton, and J. Shim. The high-frequency trading arms race: frequent batch auctions as a market design response. Technical report, University of Chicago Booth School of Business,, available at `http://faculty.chicagobooth.edu/eric.budish/research/HFT-FrequentBatchAuctions.pdf`, 2013.

[40] J. Y. Campbell, A. W. Lo, A. C. MacKinlay, and R. F. Whitelaw. *The Econometrics of Financial Markets*. Princeton University Press, Princeton, NJ, USA, 1997.

[41] A. Carbone, G. Castelli, and H. E. Stanley. Time-dependent Hurst exponent in financial time series. *Physica A*, 344(1):267–271, 2004.

[42] M. Catanzaro and M. Buchanan. Network opportunity. *Nature Physics*, 9(3):121–123, 2013.

[43] Celent. Electronic platforms in foreign exchange trading. Technical report, Available at http://www.e-forex.net/Files/surveyreportsPDFs/Celent%20FX%20report.pdf, 2007.

[44] A. Chaboud, B. Chiquoine, E. Hjalmarsson, and C. Vega. Rise of the machines: Algorithmic trading in the foreign exchange market. Technical report, Federal Reserve, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1501135, 2011.

[45] A. Chakraborti, I. M. Toke, M. Patriarca, and F. Abergel. Econophysics review I: empirical facts. *Quantitative Finance*, 11(7):991–1012, 2011.

[46] A. Chakraborti, I. M. Toke, M. Patriarca, and F. Abergel. Econophysics review II: agent-based models. *Quantitative Finance*, 11(7):1013–1041, 2011.

[47] D. Challet and R. Stinchcombe. Analyzing and modeling 1+1d markets. *Physica A*, 300(1–2):285–299, 2001.

[48] C. Chatfield. *Time-Series Forecasting*. Chapman and Hall, Boca Raton, FL, USA, 2000.

[49] S. H. Chen, C. L Chang, and Y. R. Du. Agent-based economic models and econometrics. *The Knowledge Engineering Review*, 27(2):187–219, 2012.

[50] G. Christodoulou. The foreign exchange and over-the-counter derivatives markets in the United Kingdom. *The Bank of England Quarterly Bulletin*, 4:548–563, 2007.

[51] A. Clauset, C. R. Shalizi, and M. E. J. Newman. Power-law distributions in empirical data. *SIAM Review*, 51(4):661–703, 2009.

[52] CLS Group. Retrieved 23 November 2013 from http://www.cls-group.com.

[53] Foreign Exchange Committee. A survey assessing the impact of electronic broking on the foreign exchange market. Technical report, Foreign Exchange Committee, available at http://www.newyorkfed.org/fxc/elecbrok.pdf, 1997.

[54] Foreign Exchange Committee. Foreign exchange prime brokerage: product overview and best practice recommendations. Technical report, Foreign Exchange Committee, available at http://www.ny.frb.org/fxc/2005/fxc051219a.pdf, 2005.

[55] R. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001.

[56] R. Cont. Long-range dependence in financial markets. In J. Lévy-Véhel and E. Lutton, editors, *Fractals in Engineering*, pages 159–179. Springer, London, UK, 2005.

[57] R. Cont. Statistical modeling of high-frequency financial data. *Signal Processing Magazine, IEEE*, 28(5):16–25, 2011.

[58] R. Cont and J. P. Bouchaud. Herd behavior and aggregate fluctuations in financial markets. *Macroeconomic Dynamics*, 4(2):170–196, 2000.

[59] R. Cont, M. Potters, and J. P. Bouchaud. Scaling in stock market data: stable laws and beyond. In B. Dubrulle, F. Graner, and D. Sornette, editors, *Les Houches Workshop*, pages 75–85, New York, NY, USA, 1997. Springer.

[60] R. Cont, S. Stoikov, and R. Talreja. A stochastic model for order book dynamics. *Operations Research*, 58(3):549–563, 2010.

[61] H. Cramér. On the composition of elementary errors. *Scandinavian Actuarial Journal*, 1928(1):13–74, 1928.

[62] P. Csermely, A. London, L. Y. Wu, and B. Uzzi. Structure and dynamics of core/periphery networks. *Journal of Complex Networks*, 1(2):93–123, 2013.

[63] D. M. Cutler, J. M. Poterba, and L. H. Summers. What moves stock prices? *The Journal of Portfolio Management*, 15(3):4–12, 1989.

[64] M. M. Dacorogna, U. A. Müller, O. V. Pictet, and C. G de Vries. The distribution of extremal foreign exchange rate returns in extremely large data sets. Technical report, Olsen and Associates, available at `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=6377`, 1995.

[65] R. B. Darlington. Multiple regression in psychological research and practice. *Psychological Bulletin*, 69(3):161–182, 1968.

[66] F. DeLuca. High frequency trading. *Review of Banking and Finance*, 32:62–193, 2012.

[67] M. M. Deza and E. Deza. *Dictionary of Distances*. Elsevier, Amsterdam, The Netherlands, 2006.

[68] F. X. Diebold and G. D. Rudebusch. On the power of Dickey–Fuller tests against fractional alternatives. *Economics Letters*, 35(2):155–160, 1991.

[69] C. Diks and R. van der Weide. Herding, asynchronous updating and heterogeneity in memory in a CBS. *Journal of Economic Dynamics and Control*, 29(4):741–763, 2005.

[70] L. Ding and J. Hiltrop. The electronic trading systems and bid-ask spreads in the foreign exchange market. *Journal of International Financial Markets, Institutions and Money*, 20(4):323–345, 2010.

[71] Z. Ding, C. W. J. Granger, and R. F. Engle. A long memory property of stock market returns and a new model. *Journal of Empirical Finance*, 1(1):83–106, 1993.

[72] EBS. `https://emea.ebsspot.com/SHARED/HELP/userguide.pdf`, 2011.

[73] B. Eichengreen. Three perspectives on the Bretton Woods system. In M. D. Bordo and B. Eichengreen, editors, *A Retrospective on the Bretton Woods System: Lessons for International Monetary Reform*, pages 621–658. The University of Chicago Press, Chicago, IL, USA, 1993.

[74] Euronext. Retrieved 23 November 2013 from `https://europeanequities.nyx.com/connecting/universal-trading-platform`, 2013.

[75] M. D. D. Evans. *Exchange-Rate Dynamics*. Princeton University Press, Princeton, NJ, USA, 2011.

[76] J. D. Farmer and D. Foley. The economy needs agent-based modelling. *Nature*, 460(7256):685–686, 2009.

[77] J. D. Farmer, A. N. Gerig, F. Lillo, and S. Mike. Market efficiency and the long-memory of supply and demand: is price impact variable and permanent or fixed and temporary? *Quantitative Finance*, 6(02):107–112, 2006.

[78] J. D. Farmer, P. Patelli, and I. I. Zovko. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences*, 102(6):2254–2259, 2005.

[79] R. P. Flood and M. P. Taylor. Exchange rate economics: what's wrong with the conventional macro approach? In J.A. Frankel, G. Galli, and A. Giovannini, editors, *The Microstructure of Foreign Exchange Markets*, pages 261–302. The University of Chicago Press, Chicago, IL, USA, 1996.

[80] T. Foucault. Order flow composition and trading costs in a dynamic limit order market. *Journal of Financial Markets*, 2(2):99–134, 1999.

[81] J. A. Frankel, G. Galli, and A. Giovannini. *The Microstructure of Foreign Exchange Markets*. The University of Chicago Press, Chicago, IL, USA, 1996.

[82] J. A. Frankel and A. K. Rose. Empirical research on nominal exchange rates. In G. M. Grossman and K. Rogoff, editors, *Handbook of International Economics*, volume 3, pages 1689–1729. Elsevier, Amsterdam, The Netherlands, 1995.

[83] International Monetary Fund. World economic outlook. Technical report, International Monetary Fund, available at `http://www.imf.org/external/pubs/ft/weo/2013/01/pdf/text.pdf`, 2013.

[84] J. M. Galán, L. R. Izquierdo, S. S. Izquierdo, J. I. Santos, R. Del Olmo, A. López-Paredes, and B. Edmonds. Errors and artefacts in agent-based modelling. *Journal of Artificial Societies and Social Simulation*, 12(1):1, 2009.

[85] P. Gallardo and A. Heath. Execution methods in foreign exchange markets. *Bank for International Settlements Quarterly Review*, 1:83–91, 2009.

[86] P. M. Garber. The collapse of the Bretton Woods fixed exchange rate system. In M. D. Bordo and B. Eichengreen, editors, *A Retrospective on the Bretton Woods System: Lessons for International Monetary Reform*, pages 461–494. The University of Chicago Press, Chicago, IL, USA, 1993.

[87] M. B Garman. Market microstructure. *Journal of Financial Economics*, 3(3):257–275, 1976.

[88] A. Gaunersdorfer, C. H. Hommes, and F. O. O. Wagener. Bifurcation routes to volatility clustering under evolutionary learning. *Journal of Economic Behavior and Organization*, 67(1):27–47, 2008.

[89] Á. Gereben and N. Kiss. A brief overview of the characteristics of interbank Forint/Euro trading. *Magyar Nemzeti Bank Bulletin*, 1(2):21–26, 2010.

[90] A. N. Gerig. *A Theory for Market Impact: How Order Flow Affects Stock Price*. PhD thesis, University of Illinois at Urbana-Champaign, Champaign, IL, USA, 2007.

[91] J. Geweke and S. Porter-Hudak. The estimation and application of long-memory time-series models. *Journal of Time Series Analysis*, 4(4):221–238, 1983.

[92] L. Giraitis, P. S. Kokoszka, and R. Leipus. Testing for long memory in the presence of a general trend. *Journal of Applied Probability*, 38(4):1033–1054, 2001.

[93] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.

[94] P. Gopikrishnan, M. Meyer, L. A. N. Amaral, and H. E. Stanley. Inverse cubic law for the distribution of stock price variations. *The European Physical Journal B*, 3(2):139–140, 1998.

[95] P. Gopikrishnan, V. Plerou, L. A. N. Amaral, M. Meyer, and H. E. Stanley. Scaling of the distribution of fluctuations of financial market indices. *Physical Review E*, 60(5):5305–5316, 1999.

[96] T. F. Gosnell, A. J. Keown, and J. M. Pinkerton. The intraday speed of stock price adjustment to major dividend changes: bid- ask bounce and order flow imbalances. *Journal of Banking and Finance*, 20(2):247–266, 1996.

[97] M. D. Gould, N. Hautsch, S. Williams, M. McDonald, S. D. Howison, and M. A. Porter. An agent-based model of trade on a quasi-centralized trading platform. *In preparation*, 2013.

[98] M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, and S. D. Howison. Limit order books. *Quantitative Finance*, 13(11):1709–1742, 2013.

[99] M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, and S. D. Howison. Long memory in the foreign exchange spot market. *In preparation*, 2013.

[100] M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, and S. D. Howison. Statistical properties of limit order books in the foreign exchange spot market. *In preparation*, 2014.

[101] M. D. Gould, S. Williams, M. McDonald, S. D. Howison, and M. A. Porter. Price formation on a multi-institution trading platform in the foreign exchange spot market. *In preparation*, 2013.

[102] C. Gouriéroux, J. Jasiak, and G. Le Fol. Intra-day market activity. *Journal of Financial Markets*, 2(3):193–226, 1999.

[103] C. W. J. Granger. Long memory relationships and the aggregation of dynamic models. *Journal of Econometrics*, 14(2):227–238, 1980.

[104] C. W. J. Granger and Z. Ding. Stylized facts on the temporal and distributional properties of daily data from speculative markets. Technical report, UCSD Department of Economics, available at `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=5766`, 1999.

[105] C. W. J. Granger and N. Hyung. Occasional structural breaks and long memory with an application to the S&P 500 absolute stock returns. *Journal of Empirical Finance*, 11(3):399–421, 2004.

[106] M. T. Greene and B. D. Fielitz. Long-term dependence in common stock returns. *Journal of Financial Economics*, 4(3):339–349, 1977.

[107] G. F. Gu, W. Chen, and W. X. Zhou. Empirical regularities of order placement in the Chinese stock market. *Physica A*, 387(13):3173–3182, 2008.

[108] G. F. Gu, W. Chen, and W. X. Zhou. Empirical shape function of limit-order books in the Chinese stock market. *Physica A*, 387(21):5182–5188, 2008.

[109] D. M. Guillaume, M. M. Dacorogna, R. R. Davé, U. A. Müller, R. B. Olsen, and O. V. Pictet. 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, 1997.

[110] A. D. Hall and P. Kofman. Limits to linear price behavior: Futures prices regulated by limits. *Journal of Futures Markets*, 21(5):463–488, 2001.

[111] P. Hall, B. Y. Jing, and S. N. Lahiri. On the sampling window method for long-range dependent data. *Statistica Sinica*, 8(4):1189–1204, 1998.

[112] J. Hasbrouck and G. Saar. Limit orders and volatility in a hybrid market: The Island ECN. Technical report, NYU, available at `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1294561`, 2002.

[113] U. Hassler and J. Wolters. On the power of unit root tests against fractional alternatives. *Economics Letters*, 45(1):1–5, 1994.

[114] N. Hautsch. *Econometrics of Financial High-Frequency Data*. Springer, Heidelberg, Germany, 2012.

[115] T. Hendershott, C. M. Jones, and A. J. Menkveld. Does algorithmic trading improve liquidity? *The Journal of Finance*, 66(1):1–33, 2011.

[116] B. Hollifield, R. A. Miller, and P. Sandås. Empirical analysis of limit order markets. *The Review of Economic Studies*, 71(4):1027–1063, 2004.

[117] P. Holme and J. Saramäki. Temporal networks. *Physics Reports*, 519(3):97–125, 2012.

[118] J. L. Horowitz. The bootstrap. In J. J. Heckman and E. E. Leamer, editors, *Handbook of Econometrics*, pages 3159–3228. Elsevier, Amsterdam, The Netherlands, 2001.

[119] D. A. Hsieh and A. W. Kleidon. Bid-ask spreads in foreign exchange markets: implications for models of asymmetric information. In J.A. Frankel, G. Galli, and A. Giovannini, editors, *The Microstructure of Foreign Exchange Markets*, pages 41–72. The University of Chicago Press, Chicago, IL, USA, 1996.

[120] H. Hurst. Long term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, 116:770–799, 1951.

[121] P. Isard. *Exchange Rate Economics*. Cambridge University Press, Cambridge, UK, 1995.

[122] T. Ito and Y. Hashimoto. 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, 2006.

[123] D. W. Jansen and C. G. de Vries. On the frequency of large stock returns: putting booms and busts into perspective. *The Review of Economics and Statistics*, 73:18–24, 1991.

[124] C. M. Jones, G. Kaul, and M. L. Lipson. Information, trading, and volatility. *Journal of Financial Economics*, 36(1):127–154, 1994.

[125] M. King and D. Rime. The $4 trillion question: what explains FX growth since the 2007 survey? *Bank for International Settlements Quarterly Review*, 4:27–42, 2010.

[126] A. Kirilenko, A. S. Kyle, M. Samadi, and T. Tuzun. The flash crash: the impact of high-frequency trading on an electronic market. *Working Paper, SSRN eLibrary ID 1686004*, 2011.

[127] A. Kirman and G. Teyssiere. Microeconomic models for long memory in the volatility of financial time series. *Studies in Nonlinear Dynamics and Econometrics*, 5(4):281–302, 2002.

[128] Knight Capital Group. `http://www.hotspotfx.com/overview/index.jsp`, 2014.

[129] Knight Capital Group. `http://www.hotspotfx.com/products/hotspot_volumes.jsp`, 2014.

[130] Knight Capital Group. `http://www.hotspotfx.com/download/userguide/HSFX/HSFX_UserGuide_wrapper.html`, 2014.

[131] A. S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.

[132] S. N. Lahiri. Theoretical comparisons of block bootstrap methods. *The Annals of Statistics*, 27(1):386–404, 1999.

[133] B. Lauterbach and U. Ben-Zion. Stock market crashes and the performance of circuit breakers: Empirical evidence. *The Journal of Finance*, 48(5):1909–1925, 1993.

[134] B. LeBaron. Agent-based financial markets: Matching stylized facts with style. In D. Colander, editor, *Post Walrasian Macroeconomics: Beyond the DSGE Model*, pages 221–235. Cambridge University Press, Cambridge, UK, 2006.

[135] B. LeBaron and R. Yamamoto. Long-memory in an order-driven market. *Physica A*, 383(1):85–89, 2007.

[136] F. Lillo and J. D. Farmer. The long memory of the efficient market. *Studies in Nonlinear Dynamics and Econometrics*, 8(3):1–33, 2004.

[137] F. Lillo, S. Mike, and J. D. Farmer. Theory for long memory in supply and demand. *Physical Review E*, 71(6):066122, 2005.

[138] R. H. Lindeman, P. F. Merenda, and R. Z. Gold. *Introduction to bivariate and multivariate analysis*. Scott Foresman, Glenview, IL, USA, 1980.

[139] M. Liu. Modeling long memory in stock market volatility. *Journal of Econometrics*, 99(1):139–171, 2000.

[140] Y. Liu, P. Cizeau, M. Meyer, C. K. Peng, and H. E. Stanley. Correlations in economic time series. *Physica A*, 245(3):437–440, 1997.

[141] A. W. Lo. Long-term memory in stock market prices. *Econometrica*, 59(5):1279–1313, 1991.

[142] A. W. Lo and A. C. MacKinlay. *A Non-Random Walk Down Wall Street*. Princeton University Press, Princeton, NJ, USA, 2001.

[143] I. Lo and S. G. Sapp. Order aggressiveness and quantity: how are they determined in a limit order market? *Journal of International Financial Markets, Institutions and Money*, 20(3):213–237, 2010.

[144] F. M. Longin. The asymptotic distribution of extreme stock market returns. *Journal of Business*, 63:383–408, 1996.

[145] T. Lux. The limiting extremal behaviour of speculative returns: an analysis of intra-daily data from the Frankfurt Stock Exchange. *Applied Financial Economics*, 11(3):299–315, 2001.

[146] T. Lux and M. Marchesi. Volatility clustering in financial markets: a microsimulation of interacting agents. *International Journal of Theoretical and Applied Finance*, 3(04):675–702, 2000.

[147] T. Lux and F. Westerhoff. Economics crisis. *Nature Physics*, 5(1):2–3, 2009.

[148] R. K. Lyons. Tests of microstructural hypotheses in the foreign exchange market. *Journal of Financial Economics*, 39(2):321–351, 1995.

[149] R. K. Lyons. *The Microstructure Approach to Exchange Rates*. MIT Press, Cambridge, MA, USA, 2001.

[150] A. Madhavan. Market microstructure: a survey. *Journal of Financial Markets*, 3(3):205–258, 2000.

[151] A. N. Madhavan. VWAP strategies. *Trading*, 2002(1):32–39, 2002.

[152] B. B. Mandelbrot. The variation of certain speculative prices. *The Journal of Business*, 36(4):394–419, 1963.

[153] B. B. Mandelbrot. When can price be arbitraged efficiently? A limit to the validity of the random walk and martingale models. *The Review of Economics and Statistics*, 53(3):225–236, 1971.

[154] B. B. Mandelbrot. Limit theorems on the self-normalized range for weakly and strongly dependent processes. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 31(4):271–285, 1975.

[155] B. B. Mandelbrot and M. S. Taqqu. Robust R/S analysis of long-run serial correlation. *Bulletin of the International Statistical Institute*, 48(2):59–104, 1979.

[156] B. B. Mandelbrot and J. W. Van Ness. Fractional Brownian motions, fractional noises and applications. *SIAM review*, 10(4):422–437, 1968.

[157] B. B. Mandelbrot and J. R. Wallis. Noah, Joseph, and operational hydrology. *Water Resources Research*, 4(5):909–918, 1968.

[158] B. B. Mandelbrot and J. R. Wallis. Computer experiments with fractional Gaussian noises: part 1, averages and variances. *Water Resources Research*, 5(1):228–241, 1969.

[159] B. B. Mandelbrot and J. R. Wallis. Computer experiments with fractional Gaussian noises: part 2, rescaled ranges and spectra. *Water Resources Research*, 5(1):242–259, 1969.

[160] B. B. Mandelbrot and J. R. Wallis. Robustness of the rescaled range R/S in the measurement of noncyclic long run statistical dependence. *Water Resources Research*, 5(5):967–988, 1969.

[161] B. B. Mandelbrot and J. R. Wallis. Some long-run properties of geophysical records. *Water Resources Research*, 5(2):321–340, 1969.

[162] R. N. Mantegna and H. E. Stanley. *An Introduction to Econophysics: Correlations and Complexity in Finance*. Cambridge University Press, Cambridge, UK, 1999.

[163] H. Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.

[164] J. Maskawa. Correlation of coming limit price with order book in stock markets. *Physica A*, 383(1):90–95, 2007.

[165] D. McFadden. Conditional logit analysis of qualitative choice behavior. In P. Zarembka, editor, *Frontiers in Econometrics*, pages 105–142. Academic Press, New York, NY, USA, 1974.

[166] R. A. Meese and K. Rogoff. Empirical exchange rate models of the seventies: do they fit out of sample? *Journal of International Economics*, 14(1):3–24, 1983.

[167] S. Mike and J. D. Farmer. An empirical behavioral model of liquidity and volatility. *Journal of Economic Dynamics and Control*, 32(1):200–234, 2008.

[168] T. Mikosch and C. Stărică. Long-range dependence effects and ARCH modeling. In P. Doukhan, G. Oppenheim, and M.S. Taqqu, editors, *Theory and Applications of Long-Range Dependence*, pages 439–459. Birkhäuser Boston, New York, NY, USA, 2002.

[169] M. Mitchell. *Complexity: A Guided Tour*. Oxford University Press, Oxford, UK, 2009.

[170] G. H. Mu, W. Chen, J. Kertész, and W. X. Zhou. Preferred numbers and the distributions of trade sizes and trading volumes in the Chinese stock market. *The European Physical Journal B*, 68(1):145–152, 2009.

[171] NADSAQ. `http://business.nasdaq.com/trade/markets/index.html`, 2014.

[172] New York Stock Exchange. `http://www.nyxdata.com/nysedata/asp/factbook/main.asp`, 2014.

[173] W. K. Newey and K. D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55:703–708, 1987.

[174] J. R. Norris. *Markov Chains*. Cambridge University Press, Cambridge, UK, 1998.

[175] R. Nurkse. *International Currency Experience: Lessons of the Interwar Period*. League of Nations, Geneva, Switzerland, 1944.

[176] E. Pan. FX trading and technology in 2012. Technical report, StreamBase Systems, available at `http://www.tradersmagazine.com/news/fxsurveysb.php`, 2012.

[177] C. A. Parlour. Price dynamics in limit order markets. *Review of Financial Studies*, 11(4):789, 1998.

[178] C. K. Peng, S. V. Buldyrev, S. Havlin, M. Simons, H. E. Stanley, and A. L. Goldberger. Mosaic organization of DNA nucleotides. *Physical Review E*, 49(2):1685–1689, 1994.

[179] W. Perraudin and P. Vitale. Interdealer trade and information flows in a decentralized foreign exchange market. In J.A. Frankel, G. Galli, and A. Giovannini, editors, *The Microstructure of Foreign Exchange Markets*, pages 73–106. The University of Chicago Press, Chicago, IL, USA, 1996.

[180] S. H. Poon and C. W. J. Granger. Forecasting volatility in financial markets: A review. *Journal of Economic Literature*, 41(2):478–539, 2003.

[181] M. Potters and J. P. Bouchaud. More statistical properties of order books and price impact. *Physica A*, 324:133–140, 2003.

[182] T. Preis, S. Golke, W. Paul, and J. J. Schneider. Statistical analysis of financial returns for a multiagent order book model of asset trading. *Physical Review E*, 76(1):016108, 2007.

[183] W. Rea, L. Oxley, M. Reale, and J. Brown. Estimators for long-range dependence: an empirical study. *arXiv:0901.0762*, 2009.

[184] R. Rebonato. *Volatility and Correlation: The Perfect Hedger and The Fox*. Wiley, West Sussex, UK, 2004.

[185] D. Rime. New electronic trading systems in foreign exchange markets. In D. C. Jones, editor, *New Economy Handbook*, pages 469–504. Academic Press, San Diego, CA, USA, 2003.

[186] P. M. Robinson. Rates of convergence and optimal spectral bandwidth for long range dependence. *Probability Theory and Related Fields*, 99(3):443–473, 1994.

[187] I. Roşu. A dynamic model of the limit order book. *Review of Financial Studies*, 22(11):4601–4641, 2009.

[188] M. J. Sager and M. P. Taylor. Under the microscope: the structure of the foreign exchange market. *International Journal of Finance and Economics*, 11(1):81–95, 2006.

[189] L. Sarno and M. P. Taylor. The microstructure of the foreign-exchange market: a selective survey of the literature. *Princeton Studies in International Economics*, 89, 2001.

[190] N. Shephard, editor. *Stochastic Volatility*. Oxford University Press, Oxford, UK, 2005.

[191] N. V. Smirnov. On the estimation of the discrepancy between empirical curves of distribution for two independent samples. *Moscow University Mathematics Bulletin*, 2(2):3–26, 1939.

[192] E. Smith, J. D. Farmer, L. Gillemot, and S. Krishnamurthy. Statistical theory of the continuous double auction. *Quantitative Finance*, 3(6):481–514, 2003.

[193] J. A. Stephan and R. E. Whaley. Intraday price change and trading volume relations in the stock and stock option markets. *The Journal of Finance*, 45(1):191–220, 1990.

[194] M. S. Taqqu, V. Teverovsky, and W. Willinger. Estimators for long-range dependence: an empirical study. *Fractals*, 3(4):785–798, 1995.

[195] M. P. Taylor. The economics of exchange rates. *Journal of Economic Literature*, 33(1):13–47, 1995.

[196] S. J. Taylor. *Modelling Financial Time Series*. World Scientific Publishing, Singapore, 2008.

[197] T. Teräsvirta. An introduction to univariate GARCH models. In T. G. Andersen, R. A. Davis, J. P. Kreiss, and T. Mikosch, editors, *Handbook of Financial Time Series*, pages 17–42. Springer, Berlin, Germany, 2009.

[198] V. Teverovsky, M. S. Taqqu, and W. Willinger. A critical look at Lo's modified R/S statistic. *Journal of Statistical Planning and Inference*, 80(1):211–227, 1999.

[199] The London Stock Exchange. `www.londonstockexchange.com/products-and-services/trading-services/sets/sets.htm`, 2014.

[200] Thomson–Reuters. `https://dxtrapub.markets.reuters.com/docs/Matching_Rule_Book.pdf`, 2011.

[201] I. M. Toke. "Market making" in an order book model and its impact on the spread. In F. Abergel, Chakraborti A., B. K. Chakrabarti, and Mitra M., editors, *Econophys-Kolkata V*, pages 49–64, Milan, Italy, 2011. Springer.

[202] B. Tóth, J. Kertész, and J. D. Farmer. Studies of the limit order book around large price changes. *The European Physical Journal B*, 71(4):499–510, 2009.

[203] B. Tóth, Y. Lempérière, C. Deremble, J. De Lataillade, J. Kockelkoren, and J. P. Bouchaud. Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, 1(2):021006, 2011.

[204] M. Wyart, J. P. Bouchaud, J. Kockelkoren, M. Potters, and M. Vettorazzo. Relation between bid-ask spread, impact and volatility in order-driven markets. *Quantitative Finance*, 8(1):41–57, 2008.

[205] L. Xu, P. C. Ivanov, K. Hu, Z. Chen, A. Carbone, and H. E. Stanley. Quantifying signals with power-law correlations: a comparative study of detrended fluctuation analysis and detrended moving average techniques. *Physical Review E*, 71(5):051101, 2005.

[206] L. Zhao. *A Model of Limit Order Book Dynamics and a Consistent Estimation Procedure*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, USA, 2010.

[207] I. Zovko and J. D. Farmer. The power of patience: a behavioral regularity in limit order placement. *Quantitative Finance*, 2(5):387–392, 2002.