

Essays on Market Microstructure

by
Jo Albertsen Saakvitne

A dissertation submitted to BI Norwegian Business School
for the degree of PhD

PhD specialisation: Finance

Series of Dissertations 2/2019

BI Norwegian Business School

Jo Albertsen Saakvitne
Essays on Market Microstructure:

© Jo Albertsen Saakvitne
2019

Series of Dissertations 2/2019
ISBN: 978-82-8247-147-3
ISSN: 1502-2099

BI Norwegian Business School
N-0442 Oslo
Phone +47 4641 0000
www.bi.no



Printed by Skipnes Kommunikasjon AS

Acknowledgements

Over the last four years, I have had the privilege to interact with an amazing community of researchers and colleagues, both at BI and in the wider academic world. I have benefited immensely from these interactions, formal and informal, and I am deeply grateful to all those who have graciously and generously given of their time. In particular I have found the Department of Finance at BI to be a place of open office doors, where faculty are always happy to discuss, advice and comment. To belong to such a place has been a significant asset for me as a PhD student.

There are some people in this community of academics whom I want to thank in particular:

Professor Dagfinn Rime, my advisor, for being a committed mentor throughout the entire process, providing guidance, challenge and encouragement in the right amounts and at the right times. Professor Giulia di Nunno, who have been an inspiration to me, first as a teacher and supervisor and later as a coauthor. My coauthors, Geir Bjønnes, Martin Evans, Neophytos Kathitziotis and Peter O'Neill, whose collaboration have been a continuous source of motivation and development, and from whom I have learned a lot. My small cohort of PhD students and office-mates, Vasyi Kotsovskyi, Irena Kustec, Jens Sørli Kværner, Maximilian Rohrer, and Espen Skretting. Our fellow intellectual struggles and collaborations have been a highly valued part of the last four years.

I have the great blessing of living a life where many fantastic opportunities are available to me, amongst those this endeavor of doing a PhD. None of this would have been possible were it not for the wisdom and support of all my parents. I am very grateful to all of you. Finally, I am thankful to Therese, for all your patience and love.

Contents

1	Fixing the Fix? Assessing the Effectiveness of the 4PM Fix Benchmark	9
1.1	Introduction	9
1.2	Background	12
1.2.1	Role of Benchmarks in Markets	12
1.2.2	The 4pm Fix and the FX market	14
1.2.3	Academic Literature	15
1.3	Data Description	17
1.3.1	Data Sources	17
1.3.2	Market Structure and Composition over Time	18
1.4	Benchmark Quality	30
1.4.1	Representativeness	31
1.4.2	Attainability	36
1.4.3	Robustness	43
1.5	Reference Market Liquidity	46
1.5.1	Methodology	47
1.5.2	Liquidity improves after dealer collusion revelations	49
1.5.3	Liquidity worsens after fix window lengthening	50
1.6	Implications for Benchmark Design	57
1.7	Conclusion	61
1.1	Methodology for Calculating the Fix	63
1.2	Participant Categorisation Details	64
1.3	Extended Attainability Simulation	67
1.4	Full Regression Tables	69
1.5	Miscellaneous Figures and Tables	89
2	Trading strategies and information flow around price benchmarks	117
2.1	Introduction	118
2.1.1	The 4pm currency fix and fill-at-fix orders	121

2.1.2	Related literature	123
2.2	Model	124
2.2.1	Exogenous price impact - a model of frontrunning	125
2.2.2	Endogenous price impact	126
2.3	Results	130
2.3.1	Baseline parameters	130
2.3.2	Sensitivity to variances	133
2.4	Discussion	136
2.A	Numerical solution method	137
2.A.1	Convergence results for the numerical solution algorithm	139
2.B	Additional model sensitivity checks	141
2.B.1	Correlation between benchmark orders (x) and the fair value (v)	141
2.B.2	Time-varying autocorrelated noise	142
3	A note on optimal trade times for financial quotes with a last look	144
3.1	Introduction	145
3.1.1	Background	148
3.1.2	Stylized facts of high-frequency prices	152
3.2	Model	155
3.2.1	The last look-quote	155
3.2.2	The price process	159
3.3	Solution to the optimal stopping-problem	163
3.3.1	Time-changed Skellam Process	163
3.3.2	Optimal stopping with infinite horizon and no time change	164
3.3.3	General case	165
3.4	Last look under alternative price processes	170
3.4.1	Brownian motion	170
3.4.2	Skellam process	171
3.5	Conclusions	172
3.A	Proofs	172
3.B	Algorithms	185
4	Order anticipation and large traders - evidence from FX markets	188
4.1	Introduction	189
4.1.1	Anticipatory trading	190
4.2	Model	194
4.2.1	Risk neutral dealer in a three-period model	197

4.2.2	The general model	199
4.3	Data	202
4.4	Methodology	203
4.4.1	Inventory analysis	203
4.4.2	Pre-trade price dynamics	205
4.5	Results	206
4.5.1	Pre-trade inventory effects	206
4.5.2	Pre-trade price effects	207
4.6	Robustness - are the timestamps reliable?	210
4.6.1	Time lag analysis - A method for validating time stamps in trade data	210
4.6.2	Results of time lag analysis	213
4.A	Descriptive statistics and additional robustness	214

Introduction

This dissertation consists of four papers: “Fixing the fix? Assessing the Effectiveness of the 4PM Fix Benchmark”, “Trading strategies and information flow around price benchmarks”, “A note on optimal trade times for financial quotes with a last look”, and “Order anticipation and large traders - evidence from FX markets”. In this section I provide a brief summary of each of them.

In the first paper, “Fixing the fix? Assessing the Effectiveness of the 4PM Fix Benchmark”, we examine the design and effectiveness of the 4pm Fix, the largest benchmark in FX markets. We study trading around the benchmark between 2012 and 2017 with a unique dataset that allows us to identify the actions of individual traders. These data provide new insights into how trading decisions affect the properties of the Fix benchmark, and how the presence of the Fix affects trading patterns. Two events are the particular focus of our analysis: the 2013 allegations that major banks had been colluding to rig the 4pm Fix, and the 2015 reform of the benchmark methodology.

In the second paper, “Trading strategies and information flow around price benchmarks”, I characterize equilibrium pricing and trading strategies in a competitive market where a subset of liquidity traders have a preference for executing their trades at a benchmark price. The model explains recent empirical evidence from foreign exchange markets, including those of my first chapter.

In the third paper, “A note on optimal trade times for financial quotes with a last look”, we model the option value embedded in “last look”-quotes, building on results from option pricing theory. We introduce the time-changed discrete Levy process as a model for price dynamics, in order to account for realities of high-frequency financial prices, and we show how the optimal stopping problem associated to the quote can be solved via Least Squares Monte Carlo. For various special cases we provide explicit formulae. We also solve the optimal stopping problem for the cases where the price process follows a Brownian motion, and a Skellam process.

In the fourth paper, “Order anticipation and large traders - evidence from FX markets”, we provide novel evidence on how a major FX dealer bank adjusts its inventory *before* particularly large customer orders are executed. This is to our knowledge the first time such evidence is presented. We also study pre-trade price dynamics, and show that the observed price patterns differ significantly before large customer trades than at other times. Our results indicates that pre-trade price impact is a significant source of indirect trading costs for these orders.

Chapter 1

Fixing the Fix? Assessing the Effectiveness of the 4PM Fix Benchmark

MARTIN D.D. EVANS, GEORGETOWN UNIVERSITY

PETER O'NEILL, FINANCIAL CONDUCT AUTHORITY

DAGFINN RIME, BI NORWEGIAN BUSINESS SCHOOL

JO SAAKVITNE, BI NORWEGIAN BUSINESS SCHOOL

Abstract

We examine the design and effectiveness of the 4pm Fix, the most important benchmark in FX markets, using a unique dataset of trader identified order-book data from an inter-dealer venue. We propose and examine new measures of benchmark quality and examine changes to market liquidity and trader behaviour. Benchmark quality, measured as price efficiency and robustness, improves after the lengthening of the fix window to 5 minutes, but comes at the cost of a significant increase in tracking error for users of the benchmark. We also find that quoted spreads and price impact increase following the window lengthening, with HFTs trading more aggressively during the fix.

1.1 Introduction

This study examines the most important benchmark in the foreign exchange (FX) market: the WM/R 4pm Closing Spot Rate, also known as ‘the 4pm fix’. We study trading around the benchmark between 2012 and 2017 with a unique dataset that allows us to identify the actions of individual traders. These data provide new insights into how trading decisions affect the properties of the fix benchmark, and how the presence of the fix affects trading patterns. Two events are the particular focus of our analysis: the 2013

allegations that major banks had been colluding to rig the 4pm fix, and the 2015 reform of the benchmark methodology.

Benchmarks have played a significant role in markets for centuries. They are particularly important in markets, such as FX, that are fragmented and characterised by a significant amount of bilateral trading. In these markets, a benchmark reduces information asymmetries between dealers and their clients, increasing price transparency and reducing search costs (Duffie et al., 2017). Benchmarks are also hugely important for reference purposes, the WM/R rate is used as an input in MSCI and FTSE indices that funds totalling \$6tn in net assets reference and track against (Mooney, 2016). Financial benchmarks are also widely used as reference rates to settle derivative contracts, and a broad range of participants rely on benchmarks as a fair and transparent price to execute, or for valuation purposes — to rebalance funds or portfolios.

The main contribution of this paper is to inform optimal benchmark design, through a characterisation of a benchmark’s effectiveness and the liquidity in the market it references around two significant events: the dealer collusion revelations in 2013, and the change to the benchmark calculation methodology in 2015. We utilise a unique data set that includes participant identities — a crucial requirement to examine fix-trading behaviours. Very little research has been done on this before, as the earlier academic research on FX benchmark rates has been focused on examining price patterns around the fix window and related manipulative practices. We also make a significant contribution to the FX microstructure literature as the first study, in recent years, to provide liquidity metrics for a major inter-dealer venue that can only be derived from full orderbook data.

Firstly, we classify and measure the usefulness of the fix rate along three dimensions: how closely it represents rates throughout the day (representativeness); the extent that market participants can replicate the fix rate through their own trading (attainability) and how resilient it is to manipulation (robustness). This paper is among the first to propose benchmark-effectiveness measures. Duffie and Dworczak (2018), in a theoretical model, examines robustness and estimation efficiency — which is an abstraction similar to our representativeness measure. We find that the representativeness of the benchmark has increased after the lengthening of the benchmark window in 2015. We also find that, after this lengthening, the robustness of the benchmark increased, but at the cost of a reduction in attainability.

A benchmark is **representative** if it accurately represents prices of the underlying asset

throughout the day. Representativeness is an important attribute of all financial benchmarks. Benchmark rates that often take on extreme values compared with rates at other times of the day are not very representative. Furthermore, price dynamics during and around the fix window should not exhibit clear signs of market inefficiencies such as short-term predictability and strong price reversals. We find that short-term price reversals in prices around the fix decrease steadily throughout our sample period, and disappear from 2015 onwards. This coincided with changes in trading behaviour of several types of market participants — dealer banks began doing relatively less trading *before* the fix and more *during* the fix, the total trading volume of dealers that were subsequently fined for rigging decreased by one fifth, and direct trading costs in the largest currencies in our sample decreased by 5 to 10% relative to other times of the day.

Attainability is a particular concern for trade-based benchmarks — benchmarks that are calculated by sampling trades on a reference market during a pre-defined window. Users of the benchmark may try to ‘attain’ the benchmark price by trading during this sampling window, but encounter tracking error when their trade prices vary from the benchmark price — due to factors such as the benchmark taking a median of a subset of trades. We find that the change to lengthen the reference window, which was recommended by the Financial Stability Board (FSB, 2014) and implemented by WM/R, reduced attainability (or tracking error) by a magnitude of between 2 and 5 times for the largest currencies in our sample. This significantly increases the tracking error of market participants, and thus trading costs, for those participants that use the benchmark for rebalancing purposes.

Robustness refers to the extent that a benchmark is susceptible to manipulation. We show that the changes implemented in 2015 to increase the fix window have increased robustness. We show that the introduction of outlier trades in a simulated price series, has half the impact with a 5-minute fix window in comparison to a 1-minute window. However, we also show that the impact in both settings is economically small, at less than 1 basis point. We suggest that this is because the existing benchmark design — its sampling method and use of medians — is highly robust to our method of simulating outlier (manipulative) trades.

Secondly, a well-functioning benchmark depends upon a **liquid reference market**. A useful and popular benchmark can also *cause* an agglomeration of liquidity (Duffie and Stein, 2015a). Liquidity is, therefore, both a determinant of a benchmark’s effectiveness and an outcome of it — for example, if a benchmark is more attainable, representative

and robust, then it encourages more participation, which begets liquidity and enhances its effectiveness further. We examine how liquidity has evolved during our sample period. After the revelations of rigging in 2013, we find that trading costs during the fix have decreased, in the form of lower quoted spreads. After the lengthening of the fix window in 2015, quoted spreads and price impact rose, while orderbook depth decreased. These aggregate effects coincided with changes in the trading patterns of participants, particularly an increase in ‘aggressive’ or ‘liquidity-taking’ trading behaviour of high frequency traders (HFTs) around the fix window.

Thirdly, we document that, despite much controversy following the dealer collusion revelations in 2013, the benchmark is still very important. Both trading volume and the composition of participant types are broadly unchanged over our sample period. However, we do observe significant adjustments in trading patterns after key events in our sample: collusion revelations in 2013 and after changes to the benchmark calculation methodology in 2015.

Lastly, the changes made to the 4pm benchmark that we examine in this paper highlight the general trade-off that exists between attainability and robustness. For example, the benchmark calculation method ensures uncertainty about which trades are selected in its sample, which makes the benchmark harder to manipulate but also harder to attain. We discuss several incremental changes to the benchmark methodology that might increase its attainability without significantly reducing its robustness.

Section 1.2 describes the role of benchmarks and details of the 4pm fix, and discusses academic literature. Section 1.3 describes our data and measures, and provides descriptive statistics. Section 1.4 assesses how the benchmark’s representativeness, attainability and robustness are affected by the 2013 media event, and the 2015 change in the window-calculation methodology. Section 1.5 assesses the change in liquidity of the underlying FX market around the fix. Section 1.6 relates our findings to the optimal design of benchmarks, and Section 1.7 concludes.

1.2 Background

1.2.1 Role of Benchmarks in Markets

Despite the importance of benchmarks to markets, only recently has academic research begun to examine them. [Duffie and Stein \(2015a\)](#) characterise the benefits that bench-

marks bring to markets, including lower search costs, higher market participation, better matching efficiency and lower moral hazard in delegated execution, and lower trading costs associated with higher liquidity at the benchmark. These benefits result in *agglomeration*, wherein participants choose to trade at the benchmark price, as the benefits of the benchmark outweigh their idiosyncratic reasons to trade without using it (to trade a time period away from it). Agglomeration then increases the benchmark's benefits, which then drives feedback effects. [Duffie et al. \(2017\)](#) propose a theory model in which the introduction of a benchmark in a bilateral OTC market improves liquidity by reducing market participant's search frictions. [Aquilina et al. \(2017\)](#) examine the reform of the ISDAFIX¹ interest rate swap benchmark in 2015, finding an improvement in liquidity, which they argue arises from increased transparency associated with a market-derived, rather than submission-based, benchmark.

There is comparatively more research on the manipulation of benchmarks, largely precipitated by the London Interbank Offered Rate (LIBOR) scandal, beginning with [Abrantes-Metz et al. \(2012\)](#), who examine the 1-month LIBOR rate. There have been some efforts to describe the characteristics, or optimal design, of effective benchmarks. [Duffie and Stein \(2015a\)](#) argue that benchmarks should be derived from actual transactions, and [Duffie and Dworzak \(2018\)](#) demonstrate that benchmarks are more susceptible to manipulation if their reference market is more thinly traded. In their model, they characterise the choice benchmark administrators must make when designing their benchmark: they must trade off its robustness to manipulation against its efficiency² of estimating an asset's value. The International Organization of Securities Commissions (IOSCO) proposed a set of 'Principles of Effective Benchmarks'³ in 2013, which include ensuring it is appropriate to the reference market's size, liquidity, and price dynamics; ensuring it is based on observable arm's length transactions; and that the methodology should be transparent.

¹International Swaps and Derivatives Association Fix.

²Efficiency, in this model, is defined as the extent to which the benchmark estimates the asset's value without error within the calculation window. We refer to a similar concept as *representativeness* in our paper, meaning the extent that the benchmark price is an accurate reflection of prices throughout the day.

³Most of these principles relate to governance procedures of benchmark administrators and submitters, rather than the design of benchmarks.

1.2.2 The 4pm Fix and the FX market

This study examines the largest benchmark price in the spot foreign exchange markets: the WM/R Closing Spot Rate (known as ‘the 4pm Fix’)⁴ and a market it sources prices from: Thomson Reuters Matching. The spot FX market is composed of inter-dealer and single-dealer venues. The dominant inter-dealer venues are Thomson Reuters Matching and EBS. In the determination of the 4pm Fix, rates are taken from these venues, as well as a third dealer-customer platform named Currenex for some currencies.

The 4pm Fix benchmark calculation methodology is published by Reuters (2017), but essentially consists of sourcing trades from the interdealer platforms during the fix window, as well as the quoted spread at the time of the trade. Median prices are then calculated separately for trades that execute at the bid (along with the opposing ask at the time), versus trades that execute at the ask/offer (along with the opposing bid at the time). The fix price is then taken as the mid-rate of these two medians. A bid and offer is also published, which is calculated as the higher of the median quoted spreads at the time of the trade, or a predefined minimum spread — this ensures the spread is always positive and economically significant. A single trade is captured each second from each of the reference platforms. Where there are insufficient trades, best bid and offer rates are instead captured.⁵ Prior to 15 February 15 2015 this was a 1-minute window: 3:59:30 to 4:00:30. The fix window is now a 5-minute period from 3:57:30 to 4:02:30 London local time. We present a more detailed explanation of the methodology in Section 1..1.

The FX market is the most heavily traded market in the world, with \$1.7tn executed in spot FX per day in April 2016, down from \$2tn in April 2013, according to the Bank of International Settlements (2016a). Around a third of total FX volume (\$5.1tn per day) is in spot, with the rest being swaps and other derivatives. The market is concentrated across certain currency pairs, in 2016 EURUSD accounted for 23% of all spot trading, USDJPY 17.7%, GBPUSD 9.2%, and AUDUSD 5.2%. The UK handles the majority of all FX market trading, with 37% of all volume in April 2016, down from 40.8% in 2013 (Bank of International Settlements, 2016a).

Trading is concentrated on these venues by currency pairs: in the major currencies

⁴For brevity we refer to this as the 4pm Fix throughout this paper.

⁵In practice this occurs with less liquid currency pairs — see Reuters (2017) for a detailed description of this methodology.

EBS has the majority of EURUSD,⁶ USDJPY and USDCHF trading, while Reuters has GBPUSD and AUDUSD and several smaller currencies. These concentrations are difficult to verify, as the venues do not publish statistics, but they are perhaps reflected in the WM/R Closing Price methodology, which sources rates only from Reuters for GBPUSD and AUDUSD.⁷

1.2.3 Academic Literature

The 4pm Fix

Research that has focused on the 4pm fix in FX markets specifically is largely concerned with its manipulation. [Osler et al. \(2016\)](#) propose a model of dealers colluding, and [Duffie and Dworczak \(2018\)](#) and [Saakvitne \(2016a\)](#) propose models where dealers do not collude. There have also been empirical examinations of the price dynamics around the fix by [Evans \(2017\)](#) and [Ito and Yamada \(2017\)](#), which find returns are consistent with collusive behaviour or individual manipulation or both.

Papers that examine the role and utility of the 4pm fix in markets begin with [Melvin and Prins \(2015a\)](#), which examine its important role in FX hedging⁸ by showing that equity market index movements predict end-of-month FX returns. [Ito and Yamada \(2017\)](#) find that trading volumes do not decrease after the extension to 5 minutes and that trading volume is more evenly distributed within the window. [Marsh et al. \(2017\)](#) examine the price discovery during the 4pm fix in the futures market versus the spot market in a recent sample. They find that inter-dealer trades have no price impact on average during the fix period. They explain this by demonstrating that order-flow is less directional in the fix than other intraday periods. Broker ITG examines the fix from an investor perspective by conducting transaction cost analyses of fix trades. They argue that the fix is one of the most volatile intraday periods to trade ([ITG, 2014](#)) with average returns of 10 to 25 basis points around the window, which they view as an economically significant implementation shortfall for asset managers. [Chochrane \(2015\)](#) argues that this is still a

⁶[Breedon and Vitale \(2010\)](#) estimate EBS' share of EURUSD as at least 88%.

⁷AUDUSD, USDCAD, USDCZK, USDDKK, GBPUSD, USDHKD, EURHUF, USDILS, USDMXN, USDNOK, NZDUSD, USDPLN, USDRON, USDSEK, USDSGD, USDTRY and USDZAR are sourced only from Thomson Reuters Matching. USDCNH and USDRUB are sourced from both Thomson Reuters Matching and EBS. EURUSD, USDCHF and USDJPY are sourced from EBS, Currenex and Thomson Reuters Matching ([Reuters, 2017](#)).

⁸The predecessor to this paper is an unpublished working paper from 2010 called: 'London 4pm fix: The most important FX institution you have never heard of', demonstrating the lack of historical focus.

concern after the extension to 5 minutes in 2015.

FX Market Microstructure

We also provide the first taxonomy of trading participants in this market. This extends the work of [Chaboud et al. \(2014\)](#), the first to document the rise of the high-frequency traders in the FX market and their improvements to the efficiency of prices. The nature and existence of private information in FX markets has been a significant research interest, in contrast to equities markets, where its existence is considered uncontroversial. [Peiers \(1997\)](#) finds that Deutsche Bank was an informed trader in the Deutschemark and [Ito et al. \(1998\)](#) and [Killeen et al. \(2006\)](#) also provide evidence for the existence of FX market informed trading.

Academic research on the FX market may also help understand the context of the 4pm fix scandal. [Menkhoff \(1998\)](#) portrays a widespread view among FX dealers that fundamental information is unimportant. This view, alongside the established importance of order flows in driving returns, may have contributed to the collusive behaviours that were uncovered — wherein dealers shared order-flow information ahead of the fix.

Research on liquidity in FX markets has focused on its unique two-tiered structure (interdealer market and dealer markets) and the role of dealers. [Melvin and Yin \(2000\)](#) show a positive relationship between inter-dealer quoted spreads and volume and volatility, and [Mende \(2006\)](#) shows spreads widened on the day of the September 11th attacks. [King et al. \(2013\)](#) summarises unique behaviours of interdealer spreads in comparison to other markets. Dealers do not adjust their quotes to reflect changes in inventory ([Bjonnes and Rime, 2005](#); [Osler et al., 2011](#)), and do not quote wider spreads to their informed customers ([Osler et al., 2011](#)) so that they can profit from their informed trades. [Mancini et al. \(2013\)](#) show FX liquidity has commonality across currencies with equity and bond markets.

1.3 Data Description

1.3.1 Data Sources

We use proprietary order-book data from Thomson Reuters Matching (TRM) in our analysis, which contains all order-book events from the venue’s matching engine (new orders, cancellations, executions — and subsets therein: hidden orders, non-resting orders, etc.).⁹ These events are ordered sequentially and timestamped to the millisecond. The trades contain volume information and directional identifiers. The participant responsible for each event is also included, which map to 838 different legal entities in our sample. The participant identifier is a four character Terminal Controller Identifier (Dealing) Code (TCID). This reconciles to the legal entity name of the trading firm as well as the location of its trading desk. These entities are classified as large broker dealers, commercial banks, asset managers, independent trading firms including HFTs, hedge funds and other participants. Participants can trade directly on TRM as clients of a prime broker (prime broker clients — PBCs), on their own account — as direct participants, or indirectly through their broker — engaging them to trade as their principal or agent.¹⁰ In our data, participants that trade through dealers as PBCs are separately identified. Trades that dealers perform on behalf of clients (whether principal or agency) are not separately identified from their own proprietary trades. The details of our classification methodology are presented in the Appendix, in Section 1.2.

Our sample period is approximately two and a half years from the 28 October 2010, to the 5 June 2015, and around 6 months from the 15 January 2017 to the 14 June 2017. This reflects the choice by the FCA for a sample period spanning the significant events for the fix, and a more recent period. This request excluded 2016 to reduce the collection burden on firms. The currency pairs in our sample are AUDUSD, EURHUF, EURSEK, EURUSD and GBPUSD. Reuters is one of the most important inter-dealer platforms for FX, and is the only reference market for the calculation of the WM/R fix in all of the pairs in our sample except EURUSD.¹¹ Trades on the inter-dealer venue are purely wholesale in nature as the minimum trade size is one million of the respective base currency: GBP, EUR or AUD. We remove trading holidays and weekends from our sample, as these periods have very low trading and liquidity. We source historical 4pm

⁹This data was obtained directly by the FCA for market monitoring and research purposes.

¹⁰These relationships are analogous to those found in equity markets: Direct Market Access (DMA) through member firms, member firms and clients of member firms.

¹¹EURUSD takes rates from the EBS and Currenex markets as well as TRM.

fix prices from Datastream.

We also incorporate data from Thomson Reuters Tick History for our control variables that measure changes in volatility, carry and the TWI of USD. Volatility is taken from the one-week implied volatility of OTC options contracts, carry¹² is taken from the Deutsche Bank ‘Balanced Currency Harvest USD’ and the TWI¹³ of USD is taken from the Deutsche Bank ‘Short USD Currency Portfolio Index - Total Return ETF’.¹⁴

We obtain macro news announcements from ‘FX Street’, which provides a complete history of all currency-related macro news, including central-bank announcements, speeches, economic news releases and confidence indices. Each release is assigned a ‘volatility rating’ of 1 to 3.¹⁵

1.3.2 Market Structure and Composition over Time

Liquidity Measures

The unique nature of our data allows us to compute measures of trading behaviour on the level of individual market participants (TCIDs). Using the classification scheme described in Section 1.2 we aggregate these measures into category-wide variables.

We implement a range of market quality measures in this paper, which are discussed in more detail in Section 1.5. We detail two of these measures here, that we calculate on a participant category basis. We also compute a range of other variables: number of messages, number of aggressive and passive trades, average life of quotes, flow, VWAP and trade imbalance for each individual participant TCIDs. Some of these merit closer

¹²Carry is the return obtained from holding an asset, which in an FX context refers to the a collection of assets that make up a ‘carry trade’. This trade involves borrowing a currency with a low interest rate and buying a currency with a high interest rate.

¹³Trade Weighted Index: An index that aims to measure the effective value of an exchange rate by compiling a weighted average of exchange rates of home versus foreign currencies, with the weight for each foreign country equal to its share in trade.

¹⁴The RIC codes for the OTC options contracts are: GBPSWO=, AUDSWO=, EURUSWO=, EURSEKSWO=, EURHUFSWO=. Short USD Currency Portfolio Index — Total Return ETF: DBUS-DXSI, Balanced Currency Harvest USD: DBHVBUSI.

¹⁵News rated 3 is the highest, and consists of official rates announcements, monetary policy meeting minutes, CPI releases, Bank Governor speeches, non-farm payrolls, etc.

explanation, which we provide below.

The *effective spread* is computed as the difference between the trade price and the midpoint multiplied by two. This is, effectively, the quoted spread prevailing in the market at the time of a trade. We use the quoted spread prevailing before the market order that triggered the trade arrived (otherwise the effective spread would typically be nil). The effective spread is computed as:

$$\text{EffectiveSpread} = 2q \left(\frac{p_\tau - m_\tau}{m_\tau} \right)$$

where p_τ is the transaction price, m_τ is the midpoint of the best bid and offer (BBO) at the time of the trade, and q indicates the direction of the trade (+1 for buyer-initiated trades and -1 for seller initiated trades) which is taken from the initiator identifier in the orderbook data.

Price impact is computed for each individual trade, as the midpoint prevailing m seconds after a trade i , minus the midpoint at the time of a trade. We compute price impacts for 1 millisecond, 1 second, 5 seconds, 1 minute and 5 minutes. When aggregating over periods we use volume-weighted means. We only compute price impact from the perspective of the aggressive side of the trade. Price impact is computed as:

$$\text{PriceImpact}_{i,t} = q_{i,t} (m_{i,t+m} - m_{i,t}) / m_{i,t}$$

Flow is the amount bought minus the amount sold by an individual TCID over a given time period. When aggregating flow over a given participant category, we sum the flow of the individual participants. Naturally, the flow summed across *all* TCIDs is always nil. We compute separate variables for aggressive and passive flow.

VWAP is the volume-weighted average transaction price attained by all TCIDs in a given category over a given time period.

Trade imbalance is the ratio of flow to volume, computed for each individual participant TCID. It is a measure of the one-sidedness of a participant's trading activity: if all trades are in the same direction, the trade imbalance is 1. If the participants buys and sells in equal amounts over a given time interval, the trade imbalance is 0. When aggregating trade imbalance over a participant category, we volume-weight the individual

imbalances of the constituent TCIDs.

Market Structure

The WM/R 4pm fix is a benchmark price, which has two broad categories of users: firstly, those that use the fix as a valuation price for constructing indexes (for example, [MSCI \(2018\)](#)) that comprise bonds, equities or instruments in different currencies. This means that passive investment managers and ETFs will incur fund tracking errors unless they trade at this fix price. [Melvin and Prins \(2015a\)](#) cite several surveys that show asset managers hedge most of their exchange-rate exposures.

Second, the benchmark is popular with investors and corporates who may not have FX trading capabilities, or a desire to manage intraday positions, such that a single transparent benchmark price is preferable. Such firms may issue a ‘standing instruction’ to the custodian of their investments to automatically execute FX positions at the benchmark ([DuCharme, 2013](#)) or to their brokers as ‘trade at fix orders’.

Despite the 4pm fix’s importance, there is no information available on which participants use it, how they access or trade with it, and what prices they receive. In this section we provide this information, for the first time, by currency pair, over time and by participant type.

Fix volumes: despite much controversy in recent years, and while volumes traded over our sample spanning 2012 to 2017 in the broader FX market have trended downward, fix volumes appear constant, as detailed in [Figure 1.1](#). We also find that the composition of traders in the fix remains predominantly unchanged ([Figure 1.3](#)), though there does appear to be a reduction in share of trading by the major dealers (‘Dealer-R’). [Figure 1.2](#) and [Table 1.3](#) shows the composition of participants in the fix window compared with the control window. The most striking difference is that HFTs have a much lower market share in the fix window than at other times of the day, at 14 and 30% respectively. Dealers, agency brokers and custodians, on the other hand, have a higher share of total volume in the fix window than in our control window.

Composition of fix traders: The most prominent trend in the market share of the different participant groups is the steady decline in the trading volume of the largest dealers ([Figure 1.3](#)). In particular, it is interesting to note a sharp decline in the trading vol-

Figure 1.1: Total Volumes - Fix and Control Periods - 2012 to 2017 - GBPUSD and AUDUSD
 This chart presents the total volume of trades each month, for GBPUSD and AUDUSD in the 12 to 2pm control period and the fix window period.





Figure 1.2: Trading Volume by Participant Categories - Fix and Non-Fix

This chart presents the proportion of total volume for each participant class, in the 12 to 2pm control period and the fix window period, calculated by the pooling GBPUSD and AUDSUSD in the entire 2012-2017 sample period.



ume of the dealers that were later fined for illegal trading practices, and a corresponding increase in the volume of other dealers, from the second quarter to the fourth quarter of 2013, around the time when the first news stories about rigging of the 4pm fix was published. It is not possible to determine if this decline is prompted from the dealers themselves reducing their fix-related trading or their clients switching dealers.

Figure 1.3: Trading Volume % by Participant Categories - by Month

This chart presents the proportion of total volume for each participant class in the fix window period, calculated by the pooling GBPUSD and AUDSUSD each month in the 2012 to 2017 sample period.

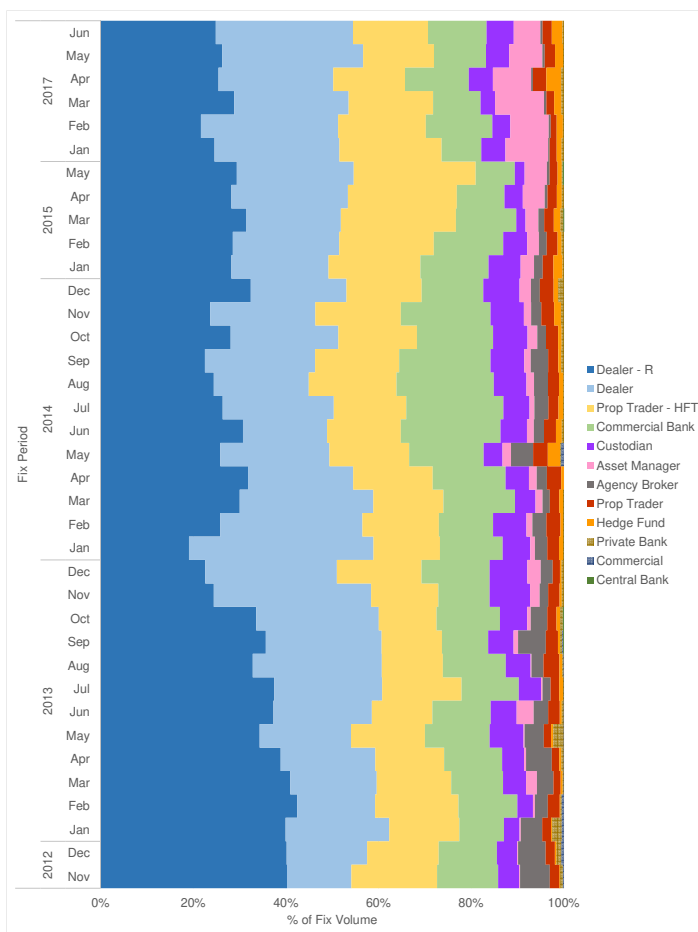


Table 1.1: Summary Statistics During the Fix by Currency-Year

Volume is total volume during the fix window. Depth is computed as the average of depth at bid and offer sides of the book (at the best bid and offer and the top 10 levels or all levels). Mean number of messages ('#msg'), quote life ('q.lif'), unique TCIDs ('#TCIDs'), number of trades ('#trades') and number of aggressor trades ('#agr.trades') are calculated across all currency-dates. #agr.trades is smaller than #trades because it doesn't include the component orders that make up a trade - of which there are least 2. Quoted spread ('Qtd.Sprd') is time-weighted, effective spreads ('Eff.Sprd.') and price impact (PI) is volume-weighted in basis points.

Pair	Year	# TCIDs	Volume	Depth best	Depth top	Depth total	#msg	q-life	# trades	#agr. trades	Qtd. sprd	Eff. sprd	PI lms	PI ls	PI 5s
	2012	54.2	557.7	17.2	167	336.2	1623.4	93.43	251.7	97.8	1.1	1	0.58	0.62	0.65
	2013	51.4	604.5	14.6	122.9	288.8	1454.8	85.52	261.8	104.8	1.25	4	0.71	0.76	0.89
audusd	2014	45.6	463	13.5	117.5	298.9	1375.6	104.56	217.3	88.6	1.37	1.2	0.69	0.78	0.95
	2015	48.4	442.3	7	85.5	277.8	3454.6	69.27	314	136.7	1.93	1.5	0.96	1.2	1.38
	2017	43.2	298.1	7.4	92.3	319.5	3873.5	93.31	221.4	95.4	1.77	1.4	0.96	1.19	1.23
	2012	9.2	12.8	1.8	4.7	33.3	53.7	494.75	9.8	6.1	4.43	2.9	1.79	1.82	2.22
	2013	9.8	13.8	1.9	4.9	33.1	78.1	445.24	11	6.3	4.36	2.7	2.04	2.35	2.57
eurhuf	2014	10.5	13.5	2	4.9	31	93.7	282.61	11.1	6.2	4	2.4	1.91	1.96	1.89
	2015	15.9	27.3	2	4.6	28.5	315	193.59	24	12.3	6.09	4	2.19	2.39	2.37
	2017	16.1	26.7	2.2	8.4	39.1	789	139.75	24.1	12.9	3.76	2.6	1.16	1.23	1.4
	2012	28.1	153.2	3.2	6.4	52.3	471.1	43.92	93.3	42.5	1.76	1.4	0.99	0.92	0.97
	2013	27.3	139.9	3.1	6	52.9	573.8	23.56	85.2	39.1	1.94	1.5	1.11	1.12	1.34
eursek	2014	26.2	135.2	3.2	5.9	42.7	577.8	21.32	83.6	38.3	2.24	1.5	1.31	1.3	1.48
	2015	30.1	171.8	2.7	4.9	50	2330.6	12.91	128	60.4	3.22	2.2	1.31	1.47	1.76
	2017	31.2	137.7	3.2	15.6	104	1837.8	34.4	115.4	55.2	1.97	1.3	0.78	1.1	1.22
	2012	16.8	11.8	2.8	30.7	49	731.5	44.9	7.9	4.6	1.21	0.9	0.43	0.61	0.95
	2013	12.3	9	2.5	36.8	60	495	70.31	6.3	4.3	1.7	1.2	0.96	0.93	0.91
eurusd	2014	9.5	7.8	2.3	38.9	60	326.9	149.65	5.7	3.5	1.62	1.1	0.87	0.82	0.82
	2015	15.4	19.8	2	37.5	55	1659.8	23.85	14.3	7.6	2.26	1.4	1.14	1.3	1.21
	2017	16.5	29.9	1.4	15.8	54.2	2509.2	21.94	26.8	13	1.19	0.9	0.75	0.81	1
	2012	48.7	413.2	9.2	99.8	173.1	1602.7	45.54	213.1	87	0.85	0.8	0.58	0.58	0.63
	2013	49.3	449.5	8.7	85.9	188.1	1598.4	48.45	230.4	97.9	0.96	0.8	0.58	0.62	0.68
gbpusd	2014	44.1	345.1	8.7	76.9	174.6	1425.1	69.52	191.8	82.5	0.91	0.7	0.48	0.53	0.57
	2015	46.5	394.8	4.8	49.7	161.8	3578.6	37.71	299.9	135.8	1.33	0.9	0.64	0.8	0.92
	2017	48.7	463.5	4.6	58.3	299.3	5000.6	47.72	363.6	159.4	1.28	0.9	0.7	0.93	1

Table 1.2: Mean daily fix volume as share of control window volume. The mean trading volume in the fix is calculated for the currency-year and divided by the mean trading volume in the control window of 12 to 2pm.

year	audusd	eurhuf	eursek	eurusd	gbpusd
2012	0.20	0.06	0.20	0.02	0.20
2013	0.18	0.05	0.21	0.02	0.16
2014	0.19	0.05	0.26	0.02	0.16
2015	0.16	0.11	0.23	0.05	0.17
2017	0.23	0.18	0.31	0.12	0.29

Fix volume shares across time: most aspects of our data set feature large variation between currency pairs and across time. Table 1.1 shows that GBPUSD and AUDUSD are by far the most active currency pairs in our sample, with average daily volume at the fix of 480 and 400m units of base currency, respectively. EURSEK is at a third place with 140m in daily fix volume, while the trading volume of both EURHUF and EURUSD is much less, at 17 and 14m. EURUSD is of course the most active currency pair in the world in general, but trading is concentrated to other platforms. Of the currencies on our sample, EURSEK is the one that sees the largest relative increase in volume during the fix, with trading volume at the fix being 28% of volume during the control window on average. For AUDUSD and GBPUSD the share is 25 and 23% respectively, while it is only 10 and 6% for EURHUF and EURUSD. Table 1.2 shows the breakdown by year. Trading volume has been steadily declining for AUDUSD over time, while it has been standing still or growing for the other currency pairs.

Fix utilisation: Table 1.3 show the average trading imbalance, which is a measure of directionality of trading or what proportion of trades are in the same direction. This measure proxies for the extent a participant category utilises the fix as a benchmark, with high directionality implying greater utilisation. This measure is calculated for individual participants and averaged by category, and is higher during the fix window, with aggregate averages of 0.85 at the fix versus 0.58 during the control window. HFTs, prop traders and asset managers have lower directionality than participants from other categories, but their directionality is still higher during the fix. HFTs have a particularly low directionality, at 0.23 during the control and 0.63 during the fix. This demonstrates that the fix is (still) very much a mechanism for conducting large rebalancing flows, as described in e.g. [Melvin and Prins \(2015a\)](#); [Evans \(2017\)](#). It also demonstrates that there are participants active in the fix that are not utilising it for benchmark purposes:

HFTs, proprietary traders and asset managers. The trading pattern of HFTs is consistent with different trading strategies, such as market making, going after short-term profit opportunities or high-frequency arbitrage.

Table 1.3: Mean effective spreads, price impacts, volume shares and average imbalances by participant category for the fix and control window. All currencies pooled. Volume share is computed as the sum of traded quantity across all TCIDs in a participant category, divided by all trades in the control window (12 to 2pm) or in the fix window. Average trading imbalance ('Imbal.') is first calculated individually for all TCIDs in each category, and then reported as a mean for each category across all currency dates. Effective spread ('Eff.sprd') is: $2q \left(\frac{p_\tau - m_\tau}{m_\tau} \right)$ where p_τ is trade price, m_τ is the midpoint and q indicates the direction of the trade, expressed in basis points and weighted at the day-currency volume level and then meaned across all currency dates for the participant group. Price impact ('PI') is computed as the change in midpoint after x seconds, divided by the midpoint at the time of the trade in basis points. Price impact is volume-weighted and aggregated in the same manner as effective spread.

Category	Eff. Sprd.	PI 1ms	PI 1s	PI 5s	Volm. (ctrl)	Volm. (fix)	Imbal. (ctrl)	Imbal. (fix)
Agency Broker	1.9	0.51	0.61	0.72	0.0111	0.0361	0.78	0.93
Asset Manager	0.9	0.81	0.94	0.97	0.0366	0.0442	0.40	0.69
Central Bank					0.0040	0.0000	0.88	1.00
Commercial	0.8	0.32	0.54	0.47	0.0016	0.0040	0.79	0.93
Commcl. Bank	1.7	0.90	0.95	1.10	0.1448	0.1606	0.69	0.97
Custodian	1.5	0.58	0.67	0.70	0.0286	0.0683	0.68	0.92
Dealer - R	1.7	0.99	1.06	1.20	0.2379	0.2811	0.57	0.92
Dealer	1.7	0.87	0.97	1.05	0.1885	0.2169	0.58	0.91
Hedge Fund	1.1	0.62	0.76	0.80	0.0167	0.0120	0.70	0.93
Private Bank	1.1	0.33	0.42	0.33	0.0072	0.0080	0.77	0.97
Prop - HFT	1.3	0.88	1.08	1.15	0.2959	0.1446	0.23	0.63
Prop Trader	1.3	0.50	0.78	0.94	0.0270	0.0241	0.56	0.82

Informed order flow: dealers, HFTs and commercial banks have the highest 1- and 5-second price impact of any participants during the fix. Their price impact ranges between 1 to 1.2 basis points (Table 1.3). Hedge funds, commercials, agency brokers and custodians all have a lower price impact, ranging from 0.5 to 0.8 basis points at 1- and 5-seconds.

Table 1.4: Correlation of flows (net position change) during the fix, for GBPUSD and AUDUSD. Net position change is computed as the sum of signed trade volume across all TCIDs in each category, using trades in the fix window only.

	Broker	Ass.mngr	Cm.bank	Cstd	Dealer	Dealr-R	Hedge	Prop
Agency Broker								
Asset Manager	0.04							
Commercial Bank	-0.08	-0.24						
Custodian	0.03	-0.04	0.02					
Dealer	-0.04		-0.31	-0.27				
Dealer - R	-0.31	-0.23	-0.18	-0.12	-0.51			
Hedge Fund	0.08	0.37	-0.34	-0.11	0.06	-0.18		
Prop Trader	0.26	0.31	-0.24	-0.11	0.02	-0.42	0.41	
Prop Trader - HFT	0.26	0.35	-0.23	-0.07	-0.08	-0.46	0.41	0.69

Correlated order flow: Table 1.4 shows how the flows (net position changes) of the participant groups are correlated. The flows of dealers and commercial banks are negatively correlated with the other participants, again consistent with these participants performing traditional market-making and liquidity provision during the fix. HFTs, hedge funds and prop traders have highly correlated flows, with correlation coefficients ranging between 0.4 to 0.7.

Tracking error (fix attainability): in Table 1.5 we compare the volume-weighted average price (VWAP) attained by participants in each category with the daily WM/R 4pm fixing rate. The comparison is done by computing the root mean square difference (RMSE) between the VWAP and the fix rate. We compute VWAPs for all trades done during the day (daily VWAP), and for trades done during the fix window only (fix VWAP). The table shows only GBPUSD and AUDUSD, as these are the largest and most liquid currencies in our sample. Ranking participant groups by their RMSE against the fix rate indicates to what extent the participants are ‘matching’ the WM/R fix rate in their trading. Asset managers, agency brokers and hedge funds are all trading at relatively low RMSE’s of 1.2 to 2.0 basis points. Prop traders, HFTs and dealers have a much higher RMSE of 3.3 to 3.4 basis points. Custodians have the highest fix-window RMSE of all participants.

Liquidity provision: we also observe significant differences in *how* participants trade, as shown in Table 1.1. Asset managers conduct 90% of their trading using marketable orders (labelled ‘aggressive trades’), followed by proprietary traders and HFTs at 76 and

Table 1.5: Root mean square error (RMSE) between volume-weighted average price (VWAP) and the WM/R benchmark rate, by participant category. The daily RMSE is a comparison with the daily VWAP, the fix RMSE is a comparison with the fix-window VWAP. Currencies used are GBPUSD and AUDUSD. Units: basis points.

Participant	Daily RMSE	Fix RMSE
Asset Manager	22.78	1.24
Agency Broker	23.56	1.59
Hedge Fund	26.26	1.95
Commercial Bank	24.92	2.29
Dealer	20.65	3.16
Prop Trader	23.04	3.28
Prop Trader - HFT	21.90	3.31
Dealer - R	22.25	3.35
all	22.00	3.35
Custodian	20.07	3.67

70%. This liquidity consumption by HFTs is high in comparison to equities markets, where they are considered important market-makers,¹⁶ though this might be the case merely on the inter-dealer market in our sample. In comparison, dealers and commercial banks provide a large amount of liquidity, with 35 to 40% of their trading volume conducted using marketable orders. Custodians and agency brokers also conduct a large share of their trading using passive limit orders, at 65 to 60% of their total trading. Proprietary traders and HFTs have a significantly higher number of messages going to the trading platform relative to the number of trades they do, compared with most other participants.

Liquidity measures across time: quoted spreads are lowest for GBPUSD, AUDUSD and EURUSD at 1.0, 1.4 and 1.6 basis points, respectively (Table 1.1). Both quoted and effective spreads have increased from 2012 to 2015 for all currency pairs. Also 1 and 5-second price impact have on average increased from 2012 to 2015. These changes could be specific to the TRM trading platform, or they could be part of a wider trend.

¹⁶Menkveld (2013) finds that around 80% of all HFT trading is passive and Baron et al. (2017) finds that 50% is, in a more recent sample.

1.4 Benchmark Quality

This study aims to track the evolution of the effectiveness of the fix over the last five years. The reform of the fix was a protracted and gradual process, with several events that we detail below. In this paper, we focus on two discrete events that have the most significant impact in our sample. Firstly, the initial revelations by Bloomberg on 12 June 2013 of dealer collusion and, secondly, the lengthening of the fix window on 15 February 2015. We refer to these events as the ‘media’ and ‘window’ events.

On 12 June 2013, [Liam Vaughan and Choudhury \(2013\)](#) published the first story that detailed a practice of collusion between major dealers to share client order information ahead of the fix. The shared information was used to infer the direction of buying and selling imbalances during the fix, allowing the colluding dealers to trade ahead of their clients. These revelations were unexpected, and prompted subsequent investigations by multiple securities regulators. Therefore, we expect the event to precipitate a change in participant behaviour in our data and refer to this as ‘**the media event**’. On 12 November 2014 the FCA fined five banks a total of £1.1 billion for ‘failing to control business practices in their G10 spot foreign exchange (FX) trading operations’.¹⁷

In response to concerns about the benchmark, the Financial Stability Board (FSB) formed a working group that published a set of recommendations, on 30 September 2014 to improve the integrity of the benchmark, including widening the fix window from 1 minute to 5 minutes ([FSB, 2014](#)). These changes were implemented by WM/R on 15 February 2015.¹⁸ The Fair and Efficient Markets Review, authored by the Bank of England, the FCA and HM Treasury ([FEMR, 2015](#)) said that the lengthening of the window would: ‘Reduce the opportunity for manipulation’ and ‘increas[e] the range of FX trades captured during the fixing window, giving a more representative and resilient fix.’ We examine this event as ‘**the window event**’.

On 1 April 2015 the FCA brought the WM/R 4pm fix into its regulatory regime¹⁹ along with six other benchmarks, and following the regulation of the LIBOR in April 2013. In addition, the Market Abuse Regulation (MAR), introduced in July 2016, designated the

¹⁷See: [FCA fines five banks £1.1 billion for FX failings and announces industry-wide remediation programme](#) The Commodity Futures Trading Commission (CFTC) also issued a \$1.bn fine to the same banks. Barclays was later fined £284m by the FCA on the 20th of May, 2015.

¹⁸For the less liquid ‘non-traded’ currencies the change was from 2 minutes to 5 minutes.

¹⁹See: [FCA PS 15/6: Bringing additional benchmarks into the regulatory and supervisory regime](#)

manipulation of regulated benchmarks as a civil offence for the first time. We do not examine this event, as we view it as merely establishing into law the behavioural changes enacted through supervisory and enforcement actions. We will analyse the benchmark's effectiveness across three dimensions: analysing representativeness for each event, and its attainability and robustness for the window event.

1.4.1 Representativeness

The 4pm fix is perceived to be a daily 'closing price' for a market that does not actually close. It arises from the importance of daily closing prices in equities markets and the institutional infrastructure that surrounds it — funds calculate net asset values (NAV) using the closing price, and then calculate FX exposures using the 4pm fix. If the closing price is not representative of, or far away from, intraday prices, it does not represent an effective benchmark. Of course, differences will arise between the closing price and intraday prices as the value of assets change over time — the 4pm fix is the value *as at* 4pm. Users of the benchmark recognise that it is a *snapshot in time*, but they would like this snapshot not to be systemically at odds with intraday prices.

To be representative, the benchmark must accurately represent prices throughout the day, and the price dynamics around it should not have clear signs of market inefficiencies such as short-term predictability and price reversals. To operationalise this definition, we first take a daily volume-weighted average transaction price (daily VWAP) value, and investigate the deviation between this price and the 4pm benchmark rates. We then test how representativeness has changed around: the first revelations of rigging, and the lengthening of the reference window period to 5 minutes.

Any change in measured representativeness can, in principle, be divided into two components — a 'mechanical' effect arising purely from a change in the benchmarking methodology, and an 'endogenous' effect arising from changes to how market participants adapt to the new regime. We disentangle these two effects using two methods. We also investigate the price dynamics around the benchmark time for evidence of market inefficiencies.

Mechanical effect of increasing benchmark window length:

We attempt to isolate the possible mechanical effect that increasing the fix window

from 1 to 5 minutes has, from any endogenous effects stemming from changes in the behaviour of participants. A mechanical effect may arise because the median of prices sampled over 5 minutes may be different to those sampled over 1 minute. To isolate such an effect, we study the statistic:

$$M_t = \sqrt{\frac{\sum_{n=1}^N (\tilde{b}_{t,n,5} - v_t)^2}{\sum_{n=1}^N (\tilde{b}_{t,n,1} - v_t)^2}}$$

Where we have employed the notation:

M_t : A measure of the mechanical effect of changing fix window from 1 to 5 minutes for day t

$\tilde{b}_{t,n,5}$: Synthetic 5-minute fix rate calculated for random time window n

$\tilde{b}_{t,n,1}$: Synthetic 1-minute fix rate calculated for random time window n

v_t : Volume-weighted average transaction price for day t

N : Number of random time windows per date

For a given t , the synthetic windows used for computing $\tilde{b}_{n,1}$ and $\tilde{b}_{n,5}$ have the same starting point — the first window extends 1 minute forward in time, while the second window extends 5 minutes forward. Moreover, no data from the actual fix window is used to compute M . The reason is that price dynamics in the actual window are affected by the endogenous effects of participants adapting their behaviour to the new regime. We use the same control period of 12 to 2pm, which excludes the fix. We draw random 150 days from 2013 – 2015, and compute $N = 1000$ random time windows for each day. The measure M_t is thus not a measure that depends on a before-after separation of the data.

After computing the measure M_t , we find that the mean of M_t is very close to one (0.99 ± 0.03). This means that the change in the benchmarking procedure, when examined by itself, would not have a material effect on the representativeness of the benchmark rates.²⁰

Endogenous effect of increasing benchmark window length:

We now isolate any changes to representativeness driven purely by changes in the behaviour of market participants by controlling for time variation in volatility. We aim

²⁰It is possible that our result may be biased due to the presence of any macroeconomics news, since the likelihood of the 5-minute period overlapping with macroeconomic news is higher than the 1-minute period. However, this would bias in favour of finding a difference in the measures, which we do not find.

to measure the benchmark's representativeness as the variation between the fix rate and underlying market prices throughout the day.

We study the statistic $D_{t,p}$, defined as,

$$D_{t,p} = \sqrt{\frac{(b_{t,p} - v_{t,p})^2}{N^{-1} \sum_{n=1}^N (\tilde{b}_{t,p,n} - v_{t,p})^2}}$$

Where we have used the notation:

$D_{t,p}$: a measure of the behaviour-driven effect of changing the fix window

b_t : the actual benchmark rate for day t and currency-pair p

$\tilde{b}_{t,p,n}$: a synthetic benchmark rate calculated for random time window n , currency-pair p , on day t

$v_{t,p}$: Volume-weighted average transaction price for day t and currency-pair p

N : Number of random time windows in each day

The nominator measures the error of the benchmark rate as a proxy for the daily VWAP rate. The denominator is an adjustment for two things: i) the mechanical increase in the efficiency of the estimator from the window increasing from 1 to 5 minutes, and ii) time variation in price volatility unrelated to the fix methodology. The synthetic benchmarks are computed using data from between 12pm and 2pm, and in accordance with WM/R methodology.

We compute $D_{t,p}$ for all days in 2013 to 2015, using $N = 1000$ random windows for each day. The result is a time series spanning days both before and after the lengthening of the calculation window. When calculating the synthetic benchmark, we extend the length of the calculation window n after the 15 February 2015 in accordance with the actual change in methodology.

After computing representativeness on each date, to assess any statistically significant differences we estimate a regression model, with $D_{t,p}$ as our dependent variable. The control variables are: *afterWindow* takes the value of one after the window event date, *volatility* is the implied FX options volatility for the currency pair at the time of the fix, *monthend* takes the value of one if the pair-date is the last trading weekday, and *macro* takes the value of one if there is a major macro news announcement from 2pm until the end of the fix. We also use currency and weekday fixed effects.

We estimate the three month period before and after the window event and find no change in representativeness after the event. We also find no change after the media event. We also estimate a similar model across all dates in the sample period from 2012 to 2015, with a timetrend variable *datecount* that increments one for each date in our sample, with month-fixed effects. This points to no gradual increase in representativeness in the sample period. We do find the benchmark becomes significantly less representative on month-end dates, with the ratio increasing 137%, higher even than with macro news announcements at 46% (see Table 1.16). The results are reported in Section 1.4.

Price dynamics (market efficiency):

It has been documented in the existing literature that price dynamics around the 4pm fix have been different from other times of the day, and, in particular, that prices have exhibited short-term spikes and subsequent reversals (Evans, 2017).

We examine short-term reversals through a correlation analysis. Specifically, let v_1, v_2, v_3 denote the market-wide VWAPs in the 15 minutes before the fix, during the fix, and the 15 minutes after the fix, respectively. We compute the correlations in 'currency returns', meaning:

$$r = \text{cor} \left(\frac{v_2 - v_1}{v_1}, \frac{v_3 - v_2}{v_2} \right)$$

We pool together all currencies in the sample and compute r by quarter.


We find a negative and statistical significant correlation coefficient r for most quarters in the period 2012 to 2014. There is a visible change around the time the fix window was lengthened (the first quarter of 2015), and from 2015 onwards the correlations are generally insignificant. The correlation coefficients and corresponding p-values are shown in Figure 1.4.

Significant serial dependence in price changes is not something one would expect to observe in an efficient market.²¹ Therefore, market efficiency around the fix has improved significantly in our sample period. It is beyond the scope of this paper to provide any causal inferences for this improvement. While the disappearance of collusive behaviour is one potential cause, another is the lengthening of the fix providing for a longer period of time for liquidity shocks to dissipate.

²¹A deeper analysis of these price patterns can be found in Evans (2017), including Sharpe ratios of possible trading strategies.

Figure 1.4: Price reversals around the fix — correlation and p-values in currency returns between the fix window and the 15 minutes after the fix period.

This Figure shows the correlation coefficient in currency returns, based on market-wide VWAPs. All currency returns are pooled together, giving 5 observations per date, and then grouped by quarter. The numbers show the p-value associated with the correlation coefficient, computed from a Fisher Z-transformation.



1.4.2 Attainability

Attainability refers to the extent to which the benchmark price can be replicated by a market participant who implements a trading pattern that matches the benchmarking procedure. This dimension is only relevant for benchmarks that input prices from a reference market, and for participants that wish to ‘attain’ or replicate the benchmark by trading on this reference market.²²

The design of the WM/R, and indeed most benchmarks, is such that there is a degree of unpredictability in the selection of prices during the window. This is by intention and, as explored in [Duffie and Dworczak \(2018\)](#), this unpredictability is designed to make the benchmark less susceptible to manipulation (more robust), but it also makes it less attainable.²³ We examine the impact that key factors in the design of benchmark, such as the length of the window, have on attainability, and the impact of exogenous factors like volatility.

We empirically measure the attainability of 4pm fix rates by comparing the volume-weighted average price of all trades within the fix window (‘within-window VWAP’) with the fix rate. The difference between the within-window VWAP and the fix rate reflects the tracking error of the average market participant. We expect the average tracking error to be zero, but the variability of the tracking error is of interest, particularly how this variability changes when the calculation window is lengthened in 2015. We measure the attainability A of the fix rate as the root mean square error of the tracking error against the within-window VWAP:

$$A = \sqrt{K^{-1} \sum_{t=0}^K (f_t - w_t)^2}$$

Where K is a given number of trading days, and f_t and w_t are the fix rate and within-window VWAP at day t . We also calculate this by participant category to examine any heterogeneous effects on different participants.

The lengthening of the calculation window has a near-mechanical effect on the variability of the tracking error on a trading strategy that aims to replicate the fix rate. We

²²A precondition for attainability is that the benchmarking procedure is sufficiently transparent for participants to know how to replicate it. For example, participants must know the time period that benchmark prices are drawn from so as to then trade in that period. There are other factors that determine attainability, which we explore in this paper.

²³And practically impossible to attain in practice exactly.

illustrate this effect through a simulation exercise. The simulation works as follows: we generate a large number of simulated price paths, modelled as realisation of a Brownian motion process with volatility σ over the interval $[0, T]$, where T is the length of the fix window in seconds. We set $B(0) = 0$ for simplicity. For each price path we sample the simulated spot rate at the end of each second, and set the simulated fix rate f to be the median of this sample,

$$f = \text{median}\{B_1, B_2, \dots, B_T\}$$

To simulate a trading strategy of a hypothetical trader trying to replicate the fix rate, we also sample the spot rate at N equally spaced points in the interval $[0, T]$. These N points represent the trades of our hypothetical trader. We are interested in the average transaction price p of this trader,

$$p = N^{-1} \sum_{n \in \mathcal{N}} B_n$$

Where the set of time points \mathcal{N} are equally spaced,

$$\mathcal{N} = \left\{ \frac{T}{N+1}, \frac{2T}{N+2}, \dots, \frac{NT}{N+1} \right\}$$

For each run u of the simulation we compute the ‘tracking error’ e_u ,

$$e_u = f_u - p_u$$

We compute tracking errors for a large number U of simulations, e_1, \dots, e_U , and examine their distribution. The theoretical error e has mean zero, and so we concentrate on the RMSE (standard deviation) of e .

This analysis assumes that each second within the fix window has a trade observation. In practice, during the 1-minute regime this was 44% of trades for GBPUSD in 2014, for example, and 27% in the 5-minute (see Table 1.35). We conduct an additional simulation which account for this, described in the Annex.

The results of the simulation exercise, presented in Figures 1.5 and 1.6, show that the lengthening of the window to 5 minutes increases tracking error by a factor of about 2.2 times for one replicating trade, and 2.17 for $N = 20$ replicating trades. A hypothetical participant that splits their order into multiple trades will greatly reduce their tracking error. Moving from 1 to 2 trades reduces tracking error by 36.2% in the 5-minute regime, and moving from 2 to 5 reduces it again by 37.3%. This effect diminishes with further

splitting, with 5 to 10 reducing it by 6.6% and 10 to 20 by 1.2%. Most of the reduction occurs between 1 to 5 trades — a 60% reduction. This relationship is substantially the same for the 5-minute regime.

In our simulation, we hold constant parameters that are also important determinants of attainability: volatility and the spread. But when we examine actual tracking error around the window event, the window-length effect is large enough to dominate. The predicted relationship between tracking error and window length in our simulations is borne out in actual trading outcomes in our data. Figure 1.7 and Tables 1.6 and 1.7 show how the root mean square deviation between the within-window VWAP and daily WM/R benchmark rates increased after the fix window was lengthened in 2015. The variability of participants' tracking errors went up, and thus attainability decreased. The observed increase is larger than suggested by our simulation results — when pooling all currency pairs and all participant types we find that variability of the tracking error increased more than fivefold. One reason why empirical attainability decreased by much more than in our simulations is that spreads also increased in most currencies after the window change. All else being equal, when participants are trading at higher spreads, their tracking error against a midpoint fix rate also increases. This latter effect is endogenous, or behaviour-driven. The total change in observed attainability is thus a combination of an endogenous and a direct effect. Our empirical analysis of attainability is conducted on the ex-post decisions participants have made, in relation to trade timing and order splitting, which we examine earlier in our simulations.

The construction of the benchmarking procedure introduces a lower bound on the variability of the tracking error, as measured by the RMSE. The WM/R benchmarking procedure defines the fix rate as a median price, while the VWAP is an average price. The expected difference between an average and a median is zero for most reasonable models of spot exchange rate, but the expected square deviation is positive. We return to this design choice when we discuss robustness and possible improvements of the benchmarking methodology in Section 1.6.

In this section we have demonstrated that participants can improve their tracking error by splitting their trades across more seconds in the window, but is this feasible in reality? We find that the mean trade size during the fix is 2.82 for AUDUSD and 2.92 for GBPUSD in 2015 (see Table 1.1). However, when we examine individual participant classes (see Table 1.8), smaller participants have average trade sizes of close to 1: (commercial bank: 1.22, private bank: 1.09, agency broker: 1.28) and even the smaller dealer

Figure 1.5: Attainability Simulation — Varying n for 1-min and 5-min Fix Windows

This Figure reports the results of the simulation exercise which computes tracking error, reported on the y-axis in pips, for the 1 minute and 5 minute window lengths across a varying number of replicating trades (n), reported on the x-axis. Per-second volatility is calculated assuming a yearly volatility of 0.2.




Figure 1.6: Attainability Simulation — Varying σ for 1-min and 5-min Fix Windows

This Figure reports the results of the simulation exercise which computes tracking error, reported on the y-axis in pips, for the 1 minute and 5 minute window lengths across a varying volatility parameter (σ), reported on the x-axis. σ refers to the yearly volatility, which is calculated per-second for the purposes of the simulation.




Figure 1.7: Root mean square tracking error — variability of deviation between within-window VWAP and daily WM/R rate.

This Figure shows the root mean square deviation between within-window VWAP and daily WM/R rates, for all currency pairs pooled. Only participant categories averaging 5 or more trades in the fix window are shown, as well as the 'all' category which pools together all trades in the window. The 'before' sample ranges back to 3 months before the window change, the 'after' sample to 3 months after the change.




Table 1.6: RMSE of normalised tracking error (deviation between fix-window VWAP and benchmark rate, divided by the benchmark rate), by participant type. The category-specific VWAP is computed as an volume-weighted average price of all trades done by a TCID within that category. Only participant types with 5 or more trades in fix window are included. Category ‘all’ is all types pooled, including types with less than 5 trades. Unit: basis points.

Participant type	Before	After
Asset Manager	1.35	2.08
Commercial Bank	1.37	6.40
Custodian	1.39	10.00
Dealer	1.43	7.58
Dealer - R	1.57	7.27
Prop Trader	1.50	8.32
Prop Trader - HFT	1.53	7.45
all	1.22	7.07

Table 1.7: RMSE of normalised tracking error (deviation between fix-window VWAP and benchmark rate, divided by the benchmark rate), by currency pair. Unit: basis points.

Pair	Before	After
audusd	1.06	11.57
eurhuf	2.24	2.71
eursek	2.47	1.65
eurusd	1.38	3.22
gbpusd	0.81	6.71

category has a mean of 1.97. This implies that the smaller participant classes are unable to split their orders,²⁴ with even the largest participants in this market facing splitting constraints, and no categories able to split their orders up to the optimal levels of 5 or more. This constraint only exists because of the large minimum trade size of 1m USD on the inter-dealer platform, which appears to be too high to allow for optimal order splitting. It is possible that participants have total order sizes that exceed the constraint, but decide to execute a portion of this on other execution venues during the window, such that we overestimate their tracking error. It is also possible that they execute a portion before the window, but it is unclear what impact that would have on their tracking error.

²⁴ Available liquidity is not a determinant here as best bid or offer depth is typically much higher than average trade sizes (see Table 1.1).

Table 1.8: Mean number of messages, quote life, unique TCIDs, number of trades and number of aggressor trades in the fix period, by participant category. Calculated as a mean across the total values for each measure (except q.life which is a mean) for each currency-date in our sample.

Category	#msg	q.life	#TCIDs	#trades	#agr.trades
Asset Manager	61.0	96.50	1.0	8.3	7.5
Commercial Bank	54.1	1556.38	6.7	23.1	8.2
Custodian	25.1	182.88	1.8	10.0	4.2
Dealer	172.4	150.41	6.0	32.7	11.8
Dealer - R	136.3	311.43	5.5	38.2	14.7
Hedge Fund	40.7	20.93	1.3	3.4	2.0
Private Bank	3.7	3974.92	1.1	2.4	1.2
Prop Trader	211.9	12.91	2.7	7.6	5.8
Prop Trader - HFT	749.1	60.86	7.0	34.4	24.2
Agency Broker	11.8	190.33	1.8	6.6	2.3
Central Bank	2.1	6954.92	1.0	1.0	
Commercial	76.0	155.50	1.0	1.9	1.4

1.4.3 Robustness

A benchmark is robust if it resistant to manipulation. We adopt a simulation-based methodology to assess the extent that the benchmark is resistant to a few ‘outlier’ trades. These outliers can be thought of as trades engineered with the purpose of affecting the fix rate. The method measures how much the benchmark deviates in the presence of such outliers, compared with when calculated on a dataset without outliers.

It is important to note that our method is limited in scope, and does not measure robustness against all possible forms of manipulation. Examples of other manipulation techniques include illegal sharing of customer information among liquidity providers, trading strategies based on exploiting short-term price impact and the spreading of false news. As such, our quantitative results on robustness are partial in nature.

Our simulation method is based on generating two price series: one ‘clean’ and one ‘dirty’. The dirty series differ from the clean in that a certain number of outlier obser-

variations are inserted. We compare the fix rate computed on the dirty series, with the one computed on the clean series. The benchmarking procedure is considered robust when the deviation between the two simulated benchmarks is small.

We implement this methodology as follows. Let B_t be the clean price series, which we model as a random walk:

$$B_t = \sum_{n=1}^t z_n$$

$$z_n \sim N(0, \sigma_z) \text{ i.i.d.}$$

We assume that trades indexed $\mathcal{M} = \{t_1, t_2, \dots, t_m\}$ have been manipulated, and model the dirty price series \tilde{B}_t as,

$$\tilde{B}_t = \begin{cases} B_t & \text{if } t \notin \mathcal{M} \\ B_t + y_t & \text{if } t \in \mathcal{M} \end{cases}$$

where the ‘manipulation term’ y_t has ten times the variance of the clean trades,

$$y_n \sim N(0, 10\sigma_z) \text{ i.i.d.}$$

For a given calculation window of T seconds, we compute benchmark rates on the clean and dirty data sets, f and \tilde{f} , as the median prices in the interval $[0, T]$:

$$f = \text{median}(B_1, \dots, B_T)$$

$$\tilde{f} = \text{median}(\tilde{B}_1, \dots, \tilde{B}_T)$$

We compute these simulated benchmark rates $L = 1000$ times, and measure robustness R as the mean square error,


$$R = \sqrt{L^{-1} \sum_{l=1}^L (f - \tilde{f})^2}$$

We use $m = 5$ outlier price observations and recompute benchmark rates for $L = 1000$ simulation runs, with a length $T = 60$, or 1 minute.

First, Figure 1.8 shows that the benchmark computed on ‘dirty’ data, meaning a data set with outliers, deviates very little from the one computed on ‘clean’ data. The root mean square deviation between the two benchmarks, denoted by R , ranges from 0.05 pips to

Figure 1.8: Robustness and price volatility.

This Figure shows robustness measure R plotted against yearly volatilities of 5 to 40 %. Unit: pips.



0.35 pips for yearly volatilities ranging from 5 to 40%. ‘Pips’ refer to the minimum allowable price increment in our (and most FX) markets, which is the price quoted to 4 decimal places. For a currency pair that is traded at a price of \$1, 1 pip represents a 0.01% (or 1 basis point) change in prices.

The small effect that we find likely relates to the general property of medians — it takes a large number of observations to significantly affect the median. In statistical robustness theory, this is referred to as the *breakdown point* — the proportion of ‘incorrect’ observations an estimator (such as the median) can handle before affecting the result. The median has the highest possible breakdown point of any location estimator, while the mean has the lowest possible breakdown point. In this respect, the median and mean represent two extreme choices in benchmark design.

Second, Figure 1.8 shows that the improvement in robustness from lengthening the fix window (moving from the red to the blue line) is highly dependent on the volatility of the price-generating process, but also that the overall improvement in robustness is small.

Our analysis makes several assumptions, the most important of which is that of non-permanent price impact. While we find the price impact to be lower in the fix period than the control period, we still find it to be non-zero (See Table 1.3). The introduction of price impact would decrease the benchmark’s robustness under both the 1-minute and the 5-minute window, as ‘dirty trades’ would affect subsequent trades. But the effect that incorporating price impact would have on the move to 5-minute window is unclear. If we assume the level of price impact is exogenous to the change to 5 minutes, introducing price impact would decrease robustness more for the 1-minute period than the 5-minute period, as the longer window allows more time for price impact to dissipate.

1.5 Reference Market Liquidity

In Section 1.4 we examined the quality — or the effectiveness — of the benchmark itself. In this chapter, we examine the liquidity of the underlying FX market during, and around, the fix calculation period. Because the underlying FX market is an input to the fix calculation, its liquidity is endogenously related to the benchmark’s effectiveness —

it is both a determinant of and an outcome of it. For example, a decrease in the representativeness of the benchmark may prompt market participants to stop using it, reducing trading volumes and liquidity. The reduction in trading volumes will then further reduce the benchmark’s representativeness. A reduction in the attainability of the benchmark may also prompt participants to avoid using it, or to decide to trade outside of the reference window. Therefore, we expect liquidity to change in the reference market due to endogenous feedback effects from the changes in the benchmark we characterise in Section 1.4, i.e. the media event and the window event.

1.5.1 Methodology

To assess the liquidity of the reference market, we compute a range of market liquidity measures using tick-by-tick orderbook data and then compute volume and time-weighted means for each currency-date. We compute these over two time periods: the fix window, and a control window.

We then estimate these market liquidity measures using a regression model. To control for changes in liquidity exogenous to the fix itself, we express fix liquidity as a log ratio of the control observation.

$$\ln y_t = \alpha + \beta(\text{after})_t + \Gamma^T C_t + w_t$$

y_t is the fix-to-control ratio of a given liquidity measure, meaning the fix observation divided by the control observation. C_t is a vector of control variables and *after* is an indicator variable taking the value one after the relevant event. We use subscript t to index time, and assume (w_t) to be a white-noise sequence. We estimate this model separately for each currency pair.

The reason for expressing the dependent variable as a ratio of the fix observation to the control window observation is the autocorrelation and seasonality effects present in the untransformed levels.²⁵

²⁵In general we find very little to no evidence for autocorrelation in these ratio-measures. In the levels there is significant serial dependence. We have also modelled the level of each liquidity measure using ARMA time-series models with seasonal effects and exogenous controls (SARMAX models). The SARMAX models give the same conclusions as the regression models reported in this paper.

Market Liquidity Measures

Volume is recorded separately for both the passive and active sides of each trade. This means that when one unit is traded, the daily trade volume will increase by two. Total volume is computed by summing all trade volume for TCIDs in a given category and a given time period. Volume forms a dependent variable in our regressions as log ratio of the total volume in the fix to total volume in the control period (12 to 2pm) for a given currency date.

The *quoted spread* is calculated with respect to each event in our orderbook data: such as limit order placements, cancellations, trades and order amendments. We then calculate a time-weighted average quoted spread, weighted according to the time interval a spread is active. The spreads reported in this paper are relative, meaning that we compute the difference between best buy and sell price,²⁶ and divide by the midpoint. Quoted spread forms a dependent variable in our regressions as log ratio of the time-weighted average quoted spread in the fix to the time-weighted average quoted spread in the control period for a given currency date.

Depth of the orderbook is also measured for each order book message. This is computed as a time-weighted mean of the sum of the buy and sell sides of the orderbook. We compute three depth measures: depth at the best bid and offer level, depth at the top ten levels, and depth at all levels. Depth forms a dependent variable in our regressions as the log ratio of the time-weighted average quoted depth in the fix to time-weighted average quoted depth in the control period for a given currency date.

The *effective spread* is computed as the quoted spread prevailing in the market at the time of a trade, and *price impact* is computed for each individual trade, as described in more detail in Section 1.3.2. Effective spread and price impact form dependent variables in our regressions as log ratios of the same measures in their respective control periods for a given currency date.

Control variables

volatility measures changes in volatility and is calculated as a log ratio (or the log return) of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. We also include a *vollevel*

²⁶The ‘best bid’ is the highest buy price on the orderbook at a given time, and the ‘best ask’ is the lowest sell price.

control, which is the time-weighted average of the options midpoint price.

We include *carry* and *shortUSD* calculated as the log ratio (or the log return) of the time weighted average of the index values in the control period, versus the fix window. The index prices are updated every 10 seconds. The Carry index is a proxy for changes in global carry strategies which may impact underlying liquidity in the currency pairs. The Short USD is a TWI index designed to proxy for the value of the USD against all other currency pairs (as opposed to as a cross rate), which may also impact underlying liquidity in individual crosses.

To control for outsized trading that occurs on month-end dates for valuation purposes (Melvin and Prins, 2015a), we calculate *monthend*, which takes the value of one on the last trading weekday of the month, factoring in trading holidays.

To control for macroeconomic news announcements, we construct *macro*, which is an indicator variable that takes a value of one for macro news events with the highest volatility rating of three that occur in the period 9am to 4pm UK local time.

We scale all of our measures by a ‘control window’, which is a time period selected to be reflective of broader market liquidity, but is also exogenous to any changes in the fix. We select 12pm-2pm to avoid any pre-fix trading, as well as the US open and major macroeconomic news announcements that occur around it. We also use a control window of 9am to 11am to ensure our results are robust to control period selection.

1.5.2 Liquidity improves after dealer collusion revelations

Regression results are reported in Table 1.9 in a consolidated format that reports only the estimates for the *AfterDummy* variable for each dependent liquidity variable. Control variable estimates are omitted, with the full regression estimates reported in the Annex in Tables 1.17 to 1.32. Table 1.9 first reports the results of regressions of market liquidity over the three months before and after the 12 June 2013, which is the first time dealer collusion behaviour was published in the media (the ‘media event’).

For the two major currencies in our sample, AUDUSD and GBPUSD, we find a statistically and economically significant decrease in quoted and effective spreads relative to the control window of 10 and 11%, respectively. This could be explained from a

decrease in collusion related adverse selection costs for liquidity providers, assuming the collusive behaviour ceased following its disclosure. The lack of findings on currencies other than AUDUSD and GBPUSD could be explained by EURSEK traders placing more importance on the ECB fix in comparison to WM/R. EURHUF could be explained by the comparatively small amount of funds and indexes holding HUF and thus the relative unimportance of the fix for HUF.

We examine the trading behaviours of participants in and around the fix to explain our liquidity findings. We measure no significant changes to trading patterns *within* the fix (Table 1.34), but the allocation of trading volume between the pre-fix, fix and post-fix periods did change significantly for some of the participant categories (Table 1.10). Dealers did less of their trading *before* the fix and more of their trading *during* the fix after the media event, with the smaller dealers reducing their pre-fix aggressor trading by a third. The total volume of the largest dealers decreased by 19%, which we can speculate reflected a reduction in customer orders. Also, HFTs increased relative amount of trading done during the fix window after the media event, but not by as much as the dealers.

1.5.3 Liquidity worsens after fix window lengthening

The fix window was lengthened from 1 to 5 minutes on 15 February 2015. We find evidence of a significant worsening of market liquidity conditions in the 3 months after this date in comparison to the 3 months before,²⁷ in the form of wider quoted spreads and lower depth. This coincided with sharp changes in the trading behaviour of HFTs and dealers.

Market liquidity results

Table 1.9 shows that total traded volume increased for GBPUSD, EURUSD and EURHUF after the change. It might be tempting to explain this increase in volume with the extension of the fix volume from one to five minutes, but fix volumes should be driven by benchmark-related execution requirements that should be exogenous to the window length — the need to rebalance FX exposures to avoid tracking error in a passive index fund, for example. An increase in volume could instead be explained by trader decisions to reallocate fix volumes from before or around the fix, to during the fix after

²⁷Our findings are robust to the choice of a 1.5 month window.

Table 1.9: Regressions of Liquidity Variables - Media and Window Events

This table reports coefficient estimates for the window and media event studies using the specification in Formula 1.5.1 for the regressions in each currency pair. For each dependent variable in the first column, only the estimates for the *AfterDummy* variable are reported. The control variable estimates are omitted, with the full regression estimates reported in the Annex. *AfterDummy* is a dummy variable, which takes the value of one for the time period after the 12 June 2013 for ‘the media event’ and after the 15 February 2015 for the ‘the window event’ regressions. *Volume* is the log ratio of the total volume traded in the fix, versus the control period, for a given a currency-date. *Quotedspread* is the log ratio of the time-weighted quoted spread in the fix, versus the control period, for a given a currency-date. *Depthatbest* is the log ratio of the time-weighted depth at the best bid and offer in the fix, versus the control period, for a given a currency-date. *DepthatTop10* is the log ratio of the time-weighted depth at the best 10 bid and offer price levels in the fix, versus the control period, for a given a currency-date. *Effectivespread* is the log ratio of the volume-weighted spread at the time a trade occurs, versus the control period, for a given a currency-date. The dependent variable, *Priceimpact(1sec)* is the log ratio of the volume-weighted price impact of a trade over a 1-second period, versus the control period, for a given a currency-date. Robust t-statistics are reported in parentheses.

Media Event	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
Volume	0.08	0.35*	0.1	0.18	0.16
	0.65	1.81	0.55	0.83	1.2
Quoted Spread	-0.10***	-0.01	0.1	0.09**	-0.11***
	-3.85	-0.16	1.39	1.97	-4.48
Depth at Best	-0.1	-0.003	0.06	-0.02	0.13
	-1.03	-0.04	0.48	-0.39	1.43
Depth at Top 10	-0.01	-0.21**	-0.03	0.01	0.11*
	-0.23	-2.39	-0.31	0.67	1.89
Effective Spread	-0.12***	-0.05	0.18**	0.01	-0.10***
	-3.72	-0.41	2.48	0.09	-4.07
Price Impact (1 Second)	-0.09	0.23	0.1	-0.39*	-0.05
	-1.08	1	0.83	-1.92	-0.6
Window Event	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
Volume	0.08	0.34*	0.19	0.61***	0.26**
	0.6	1.76	1.18	3.1	2.01
Quoted Spreads	0.07***	0.33***	-0.02	0.11***	0.06*
	2.77	4.49	-0.42	5.03	1.91
Depth at Best	-0.25***	-0.04	-0.22**	-0.0002	-0.26***
	-4.02	-0.51	-2.46	-0.004	-3.09
Depth at Top 10	-0.20***	0.56	-0.15*	0.06***	-0.07
	-4.08	1.63	-1.79	2.6	-1.31
Effective Spread	0.02	0.36***	-0.20***	-0.12*	-0.10***
	0.83	3.94	-3	-1.8	-3.28
Price Impact (1 Second)	0.22***	-0.07	0.02	0.29	0.22***
	3.63	-0.44	0.23	0.73	3.25

*p<0.1; **p<0.05; ***p<0.01

Table 1.10: Media event: Mean volume of aggressor and passive trades, by pre-fix-post windows, before and after (*) the event, for GBPUSD and AUDUSD. Volume is first summed across all TCIDs in each participant group, for each currency-date-fix quarter combination, and then averaged across currency pairs and dates. Absolute value ('Tot') in million of base currency, others as share of total. P-values of two-sample t-test for difference in means for the ratios pre/total and fix/total.

Participant	Before				After				p pre	p fix
	Tot	Pre	Fix	Post	Tot*	Pre*	Fix*	Post*		
Aggr. Trades:										
Agency Broker	22.8	0.21	0.50	0.29	18.3	0.15	0.56	0.30	0.10	0.29
Asset Manager	8.8	0.27	0.51	0.23	21.2	0.19	0.74	0.07		
Commercial Bank	67.1	0.25	0.54	0.21	51.7	0.22	0.56	0.22	0.81	0.55
Custodian	30.3	0.23	0.53	0.24	28.8	0.22	0.56	0.22	0.48	0.53
Dealer	101.8	0.27	0.53	0.19	101.3	0.18	0.59	0.22	0.00	0.00
Dealer - R	196.7	0.24	0.55	0.21	158.9	0.21	0.60	0.19	0.10	0.01
Hedge Fund	9.6	0.21	0.42	0.38	10.5	0.26	0.44	0.29	0.67	0.99
Private Bank	17.0	0.26	0.54	0.20	9.9	0.20	0.20	0.60		
Prop Trader	17.8	0.30	0.47	0.22	24.6	0.26	0.50	0.24	0.07	0.43
Prop Trader - HFT	163.5	0.35	0.33	0.31	158.0	0.33	0.36	0.31	0.02	0.02
Passive Trades:										
Agency Broker	42.7	0.14	0.66	0.20	28.0	0.19	0.56	0.25	0.37	0.08
Asset Manager	12.6	0.32	0.52	0.16	8.8	0.54	0.25	0.21	0.95	
Commercial Bank	85.0	0.31	0.44	0.25	85.4	0.27	0.49	0.24	0.52	0.14
Custodian	38.5	0.20	0.57	0.23	29.7	0.21	0.50	0.29	0.80	0.04
Dealer	124.3	0.29	0.47	0.24	140.1	0.22	0.53	0.24	0.00	0.01
Dealer - R	187.7	0.23	0.56	0.21	180.2	0.22	0.59	0.19	0.57	0.19
Hedge Fund	8.6	0.35	0.28	0.37	8.8	0.24	0.40	0.35	0.74	0.37
Private Bank	14.7	0.29	0.45	0.26	7.0	0.33	0.28	0.38	0.65	0.06
Prop Trader	10.5	0.29	0.40	0.31	9.8	0.21	0.47	0.32	0.86	0.00
Prop Trader - HFT	106.6	0.39	0.31	0.30	65.8	0.34	0.34	0.32	0.01	0.22

the window change. Volume is significantly higher at month-ends for all currencies, in line with findings of the previous literature (Melvin and Prins, 2015a; Evans, 2017; Marsh et al., 2017). The macro dummy is negative where significant, meaning that fix volume relative to the control window tends to be lower on days with important macro announcements. Table 1.26 shows that quoted spreads at the fix relative to the control window are significantly higher in all currencies except EURSEK after the change. The depth of the orderbook tends to be lower after the change, both at the best level and the top ten levels. This applies to all currencies except EURHUF, for which depth is not affected, and EURUSD, where depth at top ten levels actually increase slightly at the fix relative to the control.

The rise in quoted spreads and decline in depth does not, however, result in higher *explicit* trading costs, measured at the time of trade (effective spreads), for the *average* market participant. We see an increase in the proportion of liquidity-taking trades by HFTs following the event, which could explain the increase in quoted spreads.

Interestingly, effective spreads move in the opposite direction of the quoted spread in EURSEK, EURUSD and GBPUSD, with effective spreads in AUDUSD being unaffected. The effective spread measure is a function of trade timing decisions by participants, if participants are able to execute comparatively more of their trades when the quoted spread is lower, effective spreads may decrease. Price impact is higher after the change in both the major currencies of our sample (AUDUSD and GBPUSD), as shown in Table 1.9. These tables also show that price impact tends to be higher on month-end dates.

Attainability, as discussed in Section 1.4.2, is affected by liquidity. The WM/R benchmark is reported and utilised as a median price. However, participants that want to replicate it are not able to obtain a median price, but must execute at the best bid and ask prices. Therefore, the tracking error for a fix with no price changes is at least the half-spread. The effect of the spread is not included in our simulation exercise, where we do not model the bid-ask spread for simplicity. Increases in quoted spreads will increase on the tracking error faced by participants, unless they are able to alter their trading strategies to obtain liquidity when the spread is comparatively lower. These timing effects will be reflected in the effective spread (the spread on actual trades). For the window event, we find an increase in quoted spreads across all currencies, but this does not impact effective spreads. Effective spreads for GBPUSD and EURSEK actually decline by 10 and 20% respectively. This may explain why empirical attainability

decreases for GBPUSD less than AUDUSD (8 times versus 11 times) and EURSEK attainability increases (See Table 1.7). Though this reconciliation is incomplete without consideration to changes in volatility.

Changes in trading of individual participants

There was significant discussion in the FSB's 2015 post-implementation report about changes to the behaviour of participants in response to the window lengthening. In this section, we examine these changes to corroborate these discussions and explain our market liquidity results.

Table 1.11 shows volume for each of the four quarters of the fix, broken down by participant category. Each quarter is 15 seconds in duration in the before period, and 75 seconds in duration in the after period. Before the change, there was a tendency for volume to be concentrated in the first part of the fix, as evident in Table 1.11. Large dealers typically did 50% more aggressor volume in the first quarter of the fix than the last, with an even larger difference for passive trades (Table 1.11). Volume is more evenly distributed over the fix window in the new regime, although it does still tend to tail off somewhat during the final quarter, and this effect is statistically significant across most categories.²⁸

Changes in price impact of individual participants is of interest to test whether the lengthening of the fix window results in more obvious trading signals — measured as increased price impact. The first-quarter aggressor trading of major dealers had the largest price impact during the fix in the old regime, with an average of 1.3 basis points. The price impact of the large dealers falls during the remainder of the fix. Prop traders have the opposite pattern in price impact — it started low, at 0.3 basis points on average, and increased to 1.2 basis points by the end of the fix. In the new regime, price impact for the major dealers (Dealer - R) still falls slightly, from 1.0 to 0.8 basis points, while for other dealers the price impact is constant at 1.0. HFTs have the largest price impact under the new regime, of 1.3 to 1.4 basis points.

The p-value in Table 1.12 refers to a two-sample t-test of whether the mean price im-

²⁸The p-values in Table 1.11 refer to a two-sample t-test of difference in means of the ratio (trading volume in first half of the fix)/(trading volume in second half of the fix). Given the data, it can be read as a statistical test of whether trading volume is more evenly spread out in the after period.

impact during the fix is different before and after the fix window was lengthened. The categories with changes in price impact that are significant at the 5% level are: dealers, prop traders and HFTs. The average price impact of dealers *across the fix* has risen from 0.79 to 1.01 basis points, possibly due to the shift in the distribution of volume within the fix. The average price impact for prop traders has risen from 0.76 to 1.12 basis points, while for HFTs it has risen from 1.05 to 1.33 basis points.

We next examine the proportion of trading that happens before and after the fix. This is of interest as it is possibly a measure of the extent to which participants choose to avoid trading at the fix, by trading before it. Participants may trade ahead of the fix to avoid fix-related price volatility. The FSB's 2015 post-implementation report says that fix trades by dealers have become significantly more automated through the use of agency execution algorithms — shifting fix trading flows from voice trading desks to automated trading desks. This appears to be in response to increased scrutiny of fix trades by front-office management and compliance staff (FSB, 2015). Our results do indicate that trading volume is more evenly distributed across the fix window now than before, and we find evidence of a change in the trade execution strategies of these participants. These results are detailed in Table 1.11, with other research by Ito and Yamada (2017) finding a similar distributional change. The switch to algorithms may also explain our observed reduction in effective spreads. An increase in timing ability, stemming from an algorithm being more capable to trade when the spread is comparatively narrower, would result in a reduction in effective spreads.

However, this change in behaviour by dealers has prompted some to suggest that other participants, namely HFTs, are more able to detect and trade on the order flow imbalance signal of these dealers, such as Ito and Yamada (2017) and Pragma (2015). We find some evidence to support this claim, with increased price impact during the fix window and a larger proportion of HFT volumes. Table 1.13 shows volume for the different participant categories, for the period 5 minutes before the fix, during the fix, and 5 minutes after the fix. These tables show a striking change in how much some important participants are trading right before relative to *at* the fix. Dealers used to do around 25% of their total volume in this time window during the pre-fix period. After the change, major and other dealers do 13% of their aggressor volume during the pre-period, and 17% of their passive volume. HFTs display a similar change — they have gone from doing on average 37 to 62% of their total aggressor volume during this period at the fix. Moreover, whilst the major dealers (Dealer -R) have reduced the absolute amount of aggressor trading done during time period, from 81.6 to 75.3m, HFTs have increased

from 123.5 to 142.5m. If we sum up all trading volume, aggressor and passive, HFTs, smaller dealers and asset managers have increased their trading volume in this time window by 8, 10 and 24% respectively. Larger dealers have reduced their volume by 5%, and agency brokers, commercial banks, custodians and hedge funds have reduced their volume even more (from 25 to 40% reductions).

Table 1.11: Window event: Mean volume of aggressor and passive trades, by fix quarter, before and after (*) the event, for GBPUSD and AUDUSD. Volume is first summed across all TCIDs in each participant group, for each currency-date-fix quarter combination, and then averaged across currency pairs and dates. Absolute value ('Tot') in million of base currency, quarterly volume ('Q') as share of total. P-value of two-sample t-test for difference in mean of the ratio (first half)/(second half).

Participant	Before					After					p-value
	Tot	Q1	Q2	Q3	Q4	Tot*	Q1*	Q2*	Q3*	Q4*	
Aggr. Trades:											
Agency Broker	13.7	0.21	0.17	0.30	0.32	10.1	0.19	0.21	0.35	0.25	0.60
Asset Manager	14.7	0.27	0.27	0.25	0.22	23.5	0.23	0.25	0.28	0.23	0.02
Commercial Bank	36.2	0.37	0.23	0.22	0.18	24.6	0.28	0.23	0.30	0.18	0.01
Custodian	27.4	0.31	0.34	0.24	0.11	21.5	0.23	0.25	0.24	0.28	0.00
Dealer	42.4	0.26	0.31	0.23	0.20	46.6	0.20	0.25	0.28	0.26	0.00
Dealer - R	48.3	0.31	0.25	0.23	0.21	55.4	0.28	0.25	0.26	0.21	0.02
Hedge Fund	10.1	0.20	0.40	0.10	0.30	6.0	0.24	0.24	0.29	0.23	0.94
Private Bank	3.5		0.43	0.29	0.29	6.8	0.32	0.18	0.15	0.35	
Prop Trader	12.2	0.31	0.27	0.23	0.19	12.1	0.25	0.24	0.24	0.26	0.06
Prop Trader - HFT	48.0	0.28	0.28	0.26	0.18	89.1	0.25	0.26	0.26	0.22	0.00
Passive Trades:											
Agency Broker	12.3	0.26	0.26	0.28	0.21	12.0	0.25	0.27	0.23	0.25	0.02
Asset Manager	6.0		0.67	0.17	0.17	15.0	0.07	0.33	0.40	0.20	
Commercial Bank	40.3	0.31	0.28	0.22	0.19	39.9	0.22	0.23	0.30	0.25	0.05
Custodian	26.6	0.28	0.26	0.26	0.20	19.1	0.21	0.25	0.24	0.30	0.20
Dealer	41.3	0.28	0.27	0.24	0.20	69.0	0.26	0.26	0.27	0.22	0.00
Dealer - R	65.2	0.30	0.30	0.23	0.16	85.4	0.25	0.28	0.26	0.21	0.00
Hedge Fund	10.9	0.25	0.25	0.23	0.26	12.0	0.23	0.25	0.31	0.22	0.25
Private Bank	10.5	0.10	0.10	0.48	0.33	8.1	0.24	0.29	0.22	0.25	
Prop Trader	8.9	0.29	0.29	0.20	0.22	7.9	0.29	0.21	0.21	0.29	0.28
Prop Trader - HFT	27.7	0.36	0.20	0.23	0.21	29.4	0.29	0.22	0.25	0.24	0.01

Table 1.12: Window event: Mean price impact (5sec), by fix quarter, before and after (*), for GBPUSD and AUDUSD. Volume is first summed across all TCIDs in each participant group, for each currency-date-fix quarter combination, and then averaged across currency pairs and dates. Basis points. P-value for two-sample t-test of difference in mean price impact across the entire fix.

Participant	Before				After				p-value
	Q1	Q2	Q3	Q4	Q1*	Q2*	Q3*	Q4*	
Agency Broker	-0.6	0.1	0.2	-0.0	1.5	0.8	0.7	0.7	0.72
Asset Manager	0.4	1.0	1.2	1.0	1.5	1.1	1.1	1.2	0.70
Commercial Bank	0.8	0.8	0.8	0.9	0.5	0.7	0.6	0.4	0.84
Custodian	0.4	0.8	0.5	0.6	0.8	0.7	0.3	0.9	0.23
Dealer	0.7	1.0	0.8	1.0	0.9	1.0	0.7	1.0	0.02
Dealer - R	1.3	1.0	0.8	0.9	0.9	0.9	0.8	0.8	0.54
Hedge Fund	-0.4	0.2	-0.3	0.8	-0.4	0.2	0.8	1.1	0.24
Private Bank		1.0	-1.3	-0.2	1.3	-0.2	1.6	1.2	
Prop Trader	0.3	1.3	0.8	1.1	0.8	0.9	1.2	1.2	0.02
Prop Trader - HFT	1.0	1.3	1.1	1.3	1.1	1.4	1.1	1.3	0.00

1.6 Implications for Benchmark Design

The findings in this study have several implications for the design of the 4PM Spot Closing Rate, and for benchmarks more generally. These are in respect to appropriate window length, minimum trade sizes, and sampling and weighting decisions.

We find that an increase in the size of the window of inputs used to calculate a benchmark results in increased tracking error (or reduced attainability) for participants trying to replicate a benchmark price. Therefore, benchmark administrators and regulators should be mindful that efforts to increase robustness must be weighed against attainability costs. We find that participants can significantly reduce their tracking error by splitting their fix orders over the reference window, but they may be unable to do so due to the large minimum trade size requirement of the reference market. As discussed in Section 1.4.2, the average trade size in the fix is between 1m and 2m, for participants that utilise it. This means that participants are already splitting orders as much as the minimum trade size of 1m USD allows them to. The large minimum trade size also means that smaller trading participants experience larger tracking error than larger participants.

Table 1.13: Window event: Mean volume of aggressor and passive trades, by pre-fix-post windows, before and after (*) the event, for GBPUSD and AUDUSD. Volume is first summed across all TCIDs in each participant group, for each currency-date-fix quarter combination, and then averaged across currency pairs and dates. Absolute value ('Tot') in million of base currency, others as share of total. P-values of two-sample t-test for difference in means for the ratios pre/total and fix/total.

Participant	Before				After				p pre	p fix
	Tot	Pre	Fix	Post	Tot*	Pre*	Fix*	Post*		
Aggr. Trades:										
Agency Broker	12.8	0.20	0.45	0.34	9.7	0.22	0.41	0.37	0.75	0.02
Asset Manager	23.0	0.29	0.41	0.31	32.5	0.20	0.57	0.22	0.00	0.00
Commercial Bank	46.3	0.24	0.56	0.20	28.4	0.18	0.62	0.20	0.39	0.01
Custodian	32.4	0.33	0.52	0.15	18.5	0.21	0.52	0.27	0.15	0.00
Dealer	73.2	0.24	0.52	0.24	68.0	0.17	0.65	0.18	0.00	0.00
Dealer - R	81.6	0.23	0.54	0.23	75.3	0.13	0.69	0.17	0.00	0.00
Hedge Fund	9.9	0.28	0.35	0.37	6.4	0.35	0.37	0.29	0.66	0.19
Private Bank	4.9	0.33	0.24	0.43	6.2	0.28	0.30	0.42		
Prop Trader	15.5	0.26	0.46	0.28	15.1	0.22	0.54	0.24	0.01	0.00
Prop Trader - HFT	123.5	0.32	0.37	0.31	142.5	0.18	0.62	0.19	0.00	0.00
Passive Trades:										
Agency Broker	14.9	0.25	0.38	0.37	10.2	0.28	0.50	0.22	0.35	0.00
Asset Manager	10.2	0.27	0.59	0.14	8.8	0.11	0.55	0.34		
Commercial Bank	67.5	0.26	0.52	0.23	57.1	0.20	0.62	0.17	0.06	0.00
Custodian	31.7	0.26	0.54	0.20	20.8	0.19	0.56	0.25	0.84	0.00
Dealer	87.0	0.28	0.45	0.27	108.0	0.17	0.63	0.20	0.00	0.00
Dealer - R	120.6	0.25	0.52	0.23	117.8	0.14	0.72	0.15	0.00	0.00
Hedge Fund	16.0	0.26	0.36	0.38	12.1	0.23	0.51	0.25	0.74	0.00
Private Bank	7.0	0.23	0.39	0.38	8.2	0.36	0.44	0.20		
Prop Trader	10.7	0.27	0.41	0.32	8.0	0.26	0.47	0.27	0.28	0.00
Prop Trader - HFT	55.2	0.33	0.39	0.28	50.5	0.24	0.52	0.25	0.00	0.00


Our simulation results for attainability and robustness highlight the tension that exists between these properties. We have shown that a lengthening of the calculation window decreases attainability substantially, but only improves our robustness measure by a small amount. Several of our results stem from the use of a median in the WM/R benchmarking methodology, and it is therefore natural to ask whether the trade off between attainability and robustness can be improved upon by using another location estimator in the benchmarking procedure. To quantify the choice we can use the simulation methodology developed in Section 1.4.3, but instead of studying the deviation between the ‘clean’ and ‘dirty’ benchmarks f and \tilde{f} using medians as the benchmarking procedure, we consider different benchmarking procedures and study the deviation between the dirty benchmark \tilde{f} and the mean clean price in the fix window, $\bar{B} = T^{-1} \sum_{t=1}^T B_t$. In statistical terms, we examine the (*statistical*) *efficiency* of various location estimators.

Figure 1.9 shows the deviation between mean clean price and various benchmarking procedures, for calculation windows of 60 and 300 seconds. It is clear that the median does not perform very well in comparison to the other estimators. In statistical terms, the median suffers from low efficiency. The large increase in the RMSE of the median when lengthening the window from 60 to 300 seconds further underlines the poor efficiency properties of the median. In comparison, both the trimmed and the winsorised mean perform much better, also compared with the more complicated Hodges-Lehman and Tukey biweight estimators. The winsorised mean performs especially well under the longer fix window, taking on the lowest variability of the benchmarking procedures under consideration. These results highlight that the median is, in a sense, an extreme choice of benchmarking methodology — it has good robustness properties but very poor efficiency, and that alternatives exist with almost equally good robustness but much better efficiency.

The choice of sampling only one trade per second for the benchmark is also a choice that improves the robustness of the benchmark, as a would-be manipulator’s trades cannot guarantee their trades are selected, but this choice also diminishes attainability. However, we think that increasing the number of trades sampled within a second would not improve attainability significantly, as the intra-second volatility is small relative to the inter-second volatility, and trades that consume multiple levels of liquidity are rare. The same is true of the choice to not use volume weighting in the benchmark, though we argue that the large number of single share executions means that this does not impact attainability significantly.

Figure 1.9: Efficiency of different benchmarking procedures.

This Figure shows the root mean square deviation between the mean ‘clean’ price over the fix window, B_T , and various statistical location estimators (‘benchmarking procedures’). The estimators used are the median, the mean with 5% of the data trimmed away on either side, the mean with 5% of the data winsorised on either side, the Hodges-Lehman estimator and Tukey’s biweighted robust mean estimator. Unit: pips.



Duffie and Dworczak (2018) explore the trade-off between efficient estimation and the robustness to manipulative strategies in a theoretical setting. In their model they argue that the optimal transaction-based methodology is a capped volume weighted mean price, which weights transactions linearly up to an optimal size, after which weightings are constant.

Reductions in attainability are important, as they may lead to participants deciding not to trade at the benchmark, which then result in negative liquidity agglomeration effects that then diminish the benchmark's representativeness. There is already evidence of this in the case of the popularity of NEX market's 'eFIX' pre-fix netting product²⁹ — an 'independent netting and execution facility' (FSB, 2015) that arranges matches between counterparties ahead of and at the yet-to-be-determined fix price. This means that flow that would have been executed within the fix window is instead executed outside of it. The netting facility relies on, but does not contribute to, fix price discovery. This is a similar case to dark pool venues in equity markets that reference the lit market price to match orders.

1.7 Conclusion

In this paper we have examined the effectiveness of the 4pm fix, the largest benchmark in FX markets. We proposed three dimensions along which the effectiveness of the benchmark can be evaluated: how closely the benchmark rate represents rates throughout the day (representativeness), the extent that market participants can replicate the fix rate through their own trading (attainability) and how resilient it is to manipulation (robustness). We also examine the liquidity conditions in the reference market, as liquidity is both a determinant and an outcome of benchmark effectiveness. Our unique dataset, consisting of event-by-event orderbook data identifying individual participants, allows us to connect these aggregate effects with changes in the trading patterns of different market participants around the benchmark.

Our first finding on representativeness is that price reversals after the fix mostly or wholly disappeared when the fix window was lengthened in 2015. It is not clear what would cause such a price pattern in the first place. Strong short-term predictability in prices is not something one would expect to see in a well-functioning financial market, although we are not the first to document it in the context of the 4pm fix, see Evans

²⁹www.nexmarkets.com/~media/Files/E/EBSBrokerTec/info-sheets/NEXMARKETS_eFix_Matching_v1.pdf

(2017) and Ito and Yamada (2017). We do conclude that the disappearance of price reversals indicate an improved functioning of the reference market during the fix window. This constitutes an improvement in the representativeness of the fix. This evidence is consistent with a curtailing of certain disruptive trading practices (see e.g. Osler et al. (2016) and Saakvitne (2016a) for examples of such trading practices), which would have the effect of lowering the prevalence of extreme price movements around the fix, which in turn improves the benchmark's representativeness.

Our second finding on representativeness relates to the impact of the window lengthening in 2015, and the media reports of rigging allegations in 2013. The deviation between fix rates and daily volume-weighted average rates does not decrease for these events, after controlling for variables such as macroeconomic news and month-end dates. We do not find an effect on representativeness, despite the changes we observe in participant behaviour: a 20% reduction in the fix trading volume of the dealers that were involved in the scandal, as well as a change in the relative share of trading volume that dealers executed before, during and after the fix.

Our findings on attainability and robustness highlight the general trade-off that exists in benchmark design; the efforts to improve robustness, by lengthening the fix window to 5 minutes, came at the cost of as much as a fivefold increase in tracking error for users of the benchmark. We believe the increase in tracking error had two causes: a near-mechanical effect, which we illustrate and quantify through simulations; and a behaviour-driven effect, in the form of higher explicit trading costs (spreads). We do not take a stance on whether the overall effect amounts to an improvement or not, but it seems likely that there are reasonable benchmark designs that perform better than the current fix methodology on at least one of these two dimensions, without performing notably worse on the other.

The question of optimal benchmark design is one that the literature has only recently begun to examine — particularly how trade-based benchmarks should be designed. The history of the 4pm fix shows that robustness is an important consideration for trade-based benchmarks. It also brings to the forefront the issue of attainability, which is a unique concern for trade-based benchmarks. Overall, the 4pm fix's methodology of sampling transaction prices over a very short and predetermined time interval does facilitate attainability more than alternative designs, such as closing auctions in equity markets with randomised clearing times and other mechanisms known only to participants ex-post.

Our findings also underline the importance of market liquidity as both a determinant of and an outcome of the effectiveness of a benchmark. For example, explicit trading costs (spreads) directly impact the attainability of a benchmark through the tracking error it impose on participants trying to replicate the benchmark rate, and changes to the benchmark design that feed back into market liquidity will therefore indirectly affect the benchmark's attainability, as in the case of lengthening the fix window. Similarly, both the robustness and representativeness of the benchmarking procedure relies on having a liquid reference market, and so these aspects of benchmark effectiveness can also be inadvertently affected by changes to benchmark design through endogenous feedback effects. Therefore, a proposed change to a benchmark should not be examined in isolation, without taking into account the likely adaptations by market participants.

1.1 Methodology for Calculating the Fix

If the number of trades in the 1- or 5-minute window exceeds currency's threshold³⁰ proceed with **trade methodology** below, otherwise proceed with the **order methodology** in the next section:

Trade methodology:

1. For each second in the window, sample³¹ a single trade and record:
 - (a) the trade price
 - (b) whether it is a bid/offer trade (buy/sell)
 - (c) the opposing side best bid or offer price³²
2. At the end of the window, pool together all:
 - (a) bid trade prices with opposing bid prices
 - (b) offer trade prices with opposing offer prices

³⁰This threshold is predefined by WM/R and is not published.

³¹Trades are selected using a confidential sampling process. If the trade sampled is not a 'valid' trade (if it is outside of the BBO at the time of the trade) then the second is discarded.

³²For example, if the trade sampled is an offer trade (a buy trade) obtain the best bid at the time of the trade. If the bid and offer are crossed at this point in time, the second is discarded.

3. Calculate a median of each of the bid and offer pools
4. Calculate the midpoint of these two medians
5. This is the WM/R 4pm fix Rate

If the number of trades is less than the currency's threshold, proceed with order methodology:

Order methodology:

1. For each second in the window, sample³³ the best bid and offer orders.
2. At the end of the window, pool together all:
 - (a) bid orders
 - (b) offer orders
3. Calculate a median of each of the bid and offer pools
4. Calculate the midpoint of these two medians
5. This is the WM/R 4pm fix Rate

1.2 Participant Categorisation Details

All order and trade events are identified by a four character Terminal Controller Identifier (Dealing) Code (TCID). This reconciles to the legal entity name of the trading firm as well as the location of its trading desk. A major dealer in our sample will have several TCIDs, but these are not separated by desk (e.g. treasury vs forwards), nor by flow (agency vs proprietary). There are 838 unique active TCIDs in our sample period and we assigned them to 11 different categories of participants. We did this by first sourcing additional information about the participants from Orbis' Bureau van Dijk database of private and public companies. We then examined the websites of the companies to understand the nature of their business. Where information was scarce, as was the case for e.g. proprietary trading firms, we examined alternative sources, like the LinkedIn pages of their executives. This enabled us to discern if the business was principally focused in

³³Trades are selected using a confidential sampling process. If the order sampled are not 'valid' orders (if they are crossed) then the second is discarded.

e.g. asset management, agency broking or proprietary trading.

'Dealers' consist of the investment banking firms that are dominant in the FX market as dealers. They number 14 in total, and must have an average rank in 2012 to 2017 Euromoney surveys in the top 18³⁴ in three or more of the flow categories: Non-Financials, Real Money, Leveraged Funds and (non-dealing) Banks. We also identify a separate category of dealers, being the 7 that were fined for abuses within the WM/R 4pm fix, and this consists of almost all the top 5 dealers in practice. The FX market is highly concentrated, with Euromoney's 2017 survey estimating that the top 5 banks account for 41.05% and 44.79% of trading in 2017 and 2016 respectively.³⁵ In applying our categories, we apply the same category to all global entities: if a bank is a major dealer with 20 TCIDs in different global offices, we apply the same category to all. In practice over 95% of the trading of the major dealers is done in London, New York, Tokyo and Singapore. For the 4pm benchmark, this will consist of London and New York.

Custodians are firms that list the provision of custodian and fund administration as their core functions. While they might also provide fund management and other services, the vast majority of their funds under management are as custodian or administrator. While firms like J.P. Morgan and Citi have large custodian businesses, which, unfortunately, are not separately identified from their dealer businesses, we aim to classify firms according to their most dominant economic functions. Commercial banks are banks that are not Dealers, and typically self-described as 'Commercial Banks'; these form the vast majority of TCIDs in our sample by number.

We have exercised judgement in the application of the categories so that they are as informative as possible. 'Private Bank' is differentiated from 'Commercial Bank' where the firm describes itself as such, or lists 'Wealth Management' as its primary function. 'Hedge Fund' is differentiated from Asset Management where the firm describes itself as such, or references 'global macro strategies', 'FICC³⁶ trading' or 'quantitative analysis' as a primary focus. In contrast, asset managers make reference to managing pension or mutual funds without a FICC focus. We have vastly more hedge funds in our sample than asset managers.

'Agency Broker' is differentiated from 'Commercial Bank' where the firm is primarily

³⁴Most satisfy this requirement if the threshold is also set at 10.

³⁵See: [Euromoney Survey Release 2017](#).

³⁶Fixed Income Currency & Commodities

focused on agency execution services and describes itself foremost as a broker and does not provide commercial banking services. This category includes firms which provide FX spot and derivatives execution services to retail clients.

‘Prop Trading’ firms are those which manage trade using their own capital. ‘Prop Trader — HFT’ are a subset of this category, which employ strategies that are high frequency in nature. Many of these firms are self-described HFTs, but we also apply an objective criteria that the 1% left tail of the distribution of all order resting times in the sample for a TCID is below 200 milliseconds. We apply this as a secondary identification procedure for firms that are not self-described as HFTs.

Lastly, the ‘Commercial’ category consists of firms of an entirely non-financial nature — there are just a handful of these TCIDs in our sample. ‘Central Bank’ consists of central bank trading.

Of course, while we have categorised the nature of these businesses, we are unable to disaggregate their own flows from their clients, and to disaggregate their client flows. Most of the firms in our sample are dealers or banks. The asset managers that we see trading on their own account are a small subset of the total number of asset managers that are engaging their dealers to trade on their behalf.

[Chaboud et al. \(2014\)](#) utilises a dataset from EBS that contains identifiers for human and non-human automated trades, tracking the rise in computerised trading in FX markets from 2003 to 2007. They find that the percentage of trades with at least one automated counterparty rises almost linearly from 0% to between 60 and 80%. We do not have an identifier for algorithmic trades in our data, but we suspect the use of algorithms to submit orders (whether as programmed by a human or by an automated algorithm)³⁷ to be highly pervasive in our sample from 2012 to 17, with [Schaumburg \(2014\)](#) and [Arnold and Schaefer \(2014\)](#) finding evidence this trend persists. [Chaboud et al. \(2014\)](#) are not able to distinguish HFT trades, as we are able to in our data.

³⁷For example, much trading is now handled by humans programming smart order routers that provide order splitting functionality across multiple trading venues.

1.3 Extended Attainability Simulation

The attainability simulation conducted in the paper assumes that every second in the benchmark window has a valid trade observation. The WM/R methodology excludes seconds that do not have trades, and in practice it is common for a significant number of seconds to have zero trades.

In this simulation we extend the first by accounting for the fact that the WM/R benchmark rate typically is not computed upon 60 trades, but a much smaller number.

We model the total number of trades in the fix window m as the state of the Poisson process $M(T)$, and the trade times t_1, \dots, t_m as the associated arrival times. This is a natural generalization of the previous model, as it implies that the trade times t_1, \dots, t_m are uniformly scattered across the interval $[0, T]$.

The intensity of the Poisson process $M(t)$ can be estimated from the data, as the mean number of trades in the fix window. Clearly it is appropriate to use a different intensity when the fix window is $T = 300$ seconds long from when it is $T = 60$ seconds.

Now for each trade time t_i we sample the price process, $B(t_i)$, and compute the benchmark rate f as the median:

$$f = \text{median}\{B_1, B_2, \dots, B_m\}$$


The average price attained by the hypothetical trader trying to replicate the spot rate is computed in the same manner as Simulation 1, and correspondingly we study the 'tracking errors' e_1, \dots, e_K from a large number K of simulations, for $T = 1$ and $T = 5$.

Results of Simulation

The results of this simulation are reported in Figure 1.10. The red line shows the percentage increase in in RMSE a hypothetical participant that wishes to attain the benchmark experiences under a 60-second window regime to a 300-second window. The mean number of trades in the old regime is 20, and the mean number of trades in the new regime is 20*(the number on the x-axis). The blue line is similar, but now the mean number of trades in the old regime is 30.

Figure 1.10: Simulation of Benchmark Windows with Varying Seconds With Trades

This Figure reports the results of the simulation exercise, which computes tracking error, reported on the y-axis in pips, for the 1-minute and 5-minute window lengths across a varying number of seconds with trades (m), reported on the x-axis. Per-second volatility is calculated assuming a yearly volatility of 0.2.



1..4 Full Regression Tables

Table 1.14: Regression of Representativeness Measure $D_{t,p}$ — Media Event

This table reports coefficient estimates for the window event study for a pooled regression of currency-pair dates. The dependent variable, $D_{t,p}$ is the log ratio of the total volume traded in the fix, versus the control period, for a given a currency-date. *AfterMedia* is a dummy variable, which takes the value of one for the time period after ‘the window event’ on the 12 June 2013. *eurhuf*, *eursek*, *eurusd* and *gbpusd* are dummy variables for each currency pair, representing currency fixed effects. *Vol(level)* is the time-weighted average of the options midpoint price for fix window. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 2pm to the end of the fix.

	Representativeness: $D_{t,p}$
after media	-0.21 -0.97
eurhuf	0.02 0.05
eursek	0.19 0.50
eurusd	0.16 0.43
gbpusd	-0.04 -0.12
vol (level)	0.01 0.09
monthend	2.80*** 4.64
macro	1.37*** 3.68
Weekday Dummy?	Yes
Month Dummy?	Yes
Constant	3.17*** 8.51
Observations	1,035
R ²	0.04

Note: *p<0.1; **p<0.05; ***p<0.01

Table 1.15: Regression of Representativeness Measure $D_{t,p}$ — Window Event

This table reports coefficient estimates for the window event study for a pooled regression of currency-pair dates. The dependent variable, $D_{t,p}$ is the log ratio of the total volume traded in the fix, versus the control period, for a given a currency-date. *AfterWindow* is a dummy variable, which takes the value of one for the time period after ‘the window event’ on the 15 February 2015. *eurhuf*, *eursek*, *eurusd* and *gbpusd* are dummy variables for each currency pair, representing currency fixed effects. *Vol(level)* is the time-weighted average of the options midpoint price for fix window. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 2pm to the end of the fix.

	Representativeness: $D_{t,p}$
after window	-0.06 -0.22
eurhuf	0.73 1.47
eursek	0.41 0.90
eurusd	0.26 0.64
gbpusd	0.07 0.16
vol (level)	0.12 0.81
monthend	-0.65 -0.79
macro	0.08 0.20
Weekday Dummy?	Yes
Month Dummy?	Yes
Constant	3.09*** 6.50
	70
Observations	763
R ²	0.01

Note: *p<0.1; **p<0.05; ***p<0.01

Table 1.16: Regression of Representativeness Measure $D_{t,p}$ — Full Sample Time Trend

This table reports coefficient estimates for a pooled regression of currency-pair dates. The dependent variable, $D_{t,p}$ is the log ratio of the total volume traded in the fix, versus the control period, for a given a currency-date. *datecount* is a time trend variable, which increments one for each date in our sample. *eurhuf*, *eursek*, *eurusd* and *gbpusd* are dummy variables for each currency pair, representing currency fixed effects. *Vol(level)* is the time-weighted average of the options midpoint price for fix window. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 2pm to the end of the fix.

Representativeness: $D_{t,p}$	
datecount	–0.0001 –0.27
eurhuf	0.22 1.16
eursek	0.36* 1.86
eurusd	0.14 0.73
gbpusd	–0.13 –0.68
vol (level)	0.02 0.30
monthend	1.37*** 4.10
macro	0.43** 2.21
Weekday Dummy?	Yes
Month Dummy?	Yes
Constant	2.81*** 10.12
71	
Observations	3,207
R ²	0.02
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 1.17: Regression of Log Ratio of Volume (Fix/Non-Fix) — Media Event

This table reports coefficient estimates for media event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, *Volume* is the log ratio of the total volume traded in the fix, versus the control period, for a given a currency-date. *AfterDummy* is a dummy variable, which takes the value of one for the time period after ‘the media event’ on the 12 June 2013. *Volatility* is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. *Vol(level)* is the time-weighted average of the options midpoint price for the fix window. *Carry(ratio)* and *ShortUSD(ratio)* is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	AUDUSD	EURHUF	Volume EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	0.08	0.35*	0.10	0.18	0.16
	0.65	1.81	0.55	0.83	1.20
Volatility	0.06	0.003	-0.04	0.14	0.14**
	0.92	0.03	-0.47	1.28	2.30
Vol. (level)	-0.03	-0.04	-0.08	0.09	-0.02
	-0.43	-0.57	-1.14	0.75	-0.28
Carry (ratio)	-0.02	-0.16*	0.03	-0.12	-0.01
	-0.39	-1.84	0.42	-1.02	-0.26
Short USD (ratio)	0.02	-0.08	-0.05	-0.06	0.05
	0.37	-0.86	-0.89	-0.56	0.96
Month-end	1.19***	0.92*	0.96***	0.78*	1.22***
	2.82	1.93	4.21	1.93	8.27
Macro	-0.12	0.24	0.03	0.09	-0.10
	-0.98	1.23	0.21	0.37	-0.83
Constant	-1.86***	-3.31***	-1.75***	-4.52***	-1.90***
	-12.91	-12.94	-7.44	-20.80	-13.60
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	124	112	127	94	126
R ²	0.17	0.13	0.08	0.17	0.23

Note:

*p<0.1; **p<0.05; ***p<0.01

Volatility, Carry and Short USD are standardized ratios (fix/nonfix).

Table 1.18: Regression of Log Ratio of Quoted spread (Fix/Non-Fix) — Media Event

This table reports coefficient estimates for the media event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, *Quotedspread* is the log ratio of the time-weighted quoted spread in the fix, versus the control period, for a given a currency-date. *AfterDummy* is a dummy variable, which takes the value of one for the time period after ‘the media event’ on the 12 June 2013. *Volatility* is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. *Vol(level)* is the time-weighted average of the options midpoint price for the fix window. *Carry(ratio)* and *ShortUSD(ratio)* is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	Quoted spread				
	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	−0.10***	−0.01	0.10	0.09**	−0.11***
	−3.85	−0.16	1.39	1.97	−4.48
Volatility	0.01	−0.02	0.04*	0.02	−0.01
	1.24	−0.43	1.80	0.92	−0.48
Vol. (level)	−0.03**	0.03	−0.04	0.09***	−0.02
	−2.45	0.57	−1.45	2.86	−1.60
Carry (ratio)	0.01	−0.02	−0.04	−0.003	0.003
	1.07	−0.38	−1.25	−0.09	0.21
Short USD (ratio)	−0.01	−0.02	0.01	−0.01	0.01
	−1.42	−0.40	0.47	−0.48	0.76
Month-end	0.02	0.22	0.07	−0.07	−0.01
	0.22	0.65	0.68	−0.63	−0.23
Macro	0.02	−0.06	−0.005	0.02	0.01
	0.98	−0.57	−0.08	0.43	0.51
Constant	−0.29***	−0.29**	−0.56***	0.24***	−0.33***
	−12.25	−2.42	−8.61	5.36	−13.09
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	121	130	127	127	126
R ²	0.40	0.03	0.10	0.16	0.22

Note:

*p<0.1; **p<0.05; ***p<0.01

Volatility, Carry and Short USD are standardized ratios (fix/nonfix).

Table 1.19: Regression of Log Ratio of Depth at best (Fix/Non-Fix) — Media Event

This table reports coefficient estimates for the media event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, *Depthatbest* is the log ratio of the time-weighted depth at the best bid and offer in the fix, versus the control period, for a given a currency-date. *AfterDummy* is a dummy variable, which takes the value of one for the time period after ‘the media event’ on the 12 June 2013. *Volatility* is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. *Vol(level)* is the time-weighted average of the options midpoint price for the fix window. *Carry(ratio)* and *ShortUSD(ratio)* is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	Depth at best				
	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	-0.10	-0.003	0.06	-0.02	0.13
	-1.03	-0.04	0.48	-0.39	1.43
Volatility	0.01	0.07	0.02	0.01	-0.04
	0.15	1.58	0.29	0.50	-0.90
Vol. (level)	0.22***	-0.01	-0.06	0.07**	-0.03
	4.59	-0.25	-1.18	2.53	-0.58
Carry (ratio)	0.004	-0.04	0.01	-0.002	-0.02
	0.10	-1.04	0.13	-0.08	-0.46
Short USD (ratio)	0.01	0.04	-0.05	0.03	-0.03
	0.28	0.99	-0.91	1.18	-0.74
Month-end	0.28*	0.02	0.30*	0.08	0.49**
	1.94	0.14	1.85	0.44	2.42
Macro	0.18**	0.06	0.12	0.03	0.15*
	2.08	0.76	1.25	0.64	1.80
Constant	0.84***	-0.32***	0.39***	0.06	0.58***
	8.28	-3.98	3.48	0.88	6.33
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	124	130	127	127	126
R ²	0.25	0.10	0.06	0.11	0.12

Note:

*p<0.1; **p<0.05; ***p<0.01

Volatility, Carry and Short USD are standardized ratios (fix/nonfix).

Table 1.20: Regression of Log Ratio of Depth at Top 10 (Fix/Non-Fix) — Media Event

This table reports coefficient estimates for the media event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, *DepthatTop10* is the log ratio of the time-weighted depth at the best 10 bid and offer price levels in the fix, versus the control period, for a given a currency-date. *AfterDummy* is a dummy variable, which takes the value of one for the time period after ‘the media event’ on the 12 June 2013. *Volatility* is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. *Vol(level)* is the time-weighted average of the options midpoint price for the fix window. *Carry(ratio)* and *ShortUSD(ratio)* is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	Depth at top 10				
	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	–0.01	–0.21**	–0.03	0.01	0.11*
	–0.23	–2.39	–0.31	0.67	1.89
Volatility	0.003	0.07	0.03	–0.004	–0.02
	0.13	1.47	0.54	–0.35	–0.96
Vol. (level)	0.10***	–0.04	–0.08	0.003	–0.01
	3.27	–0.89	–1.59	0.19	–0.25
Carry (ratio)	–0.01	0.02	0.01	–0.01	–0.04
	–0.32	0.43	0.33	–0.92	–1.60
Short USD (ratio)	0.02	0.05	–0.05	0.01	0.02
	0.73	1.38	–1.15	0.94	0.77
Month-end	0.26*	0.05	0.23	0.07	0.32***
	1.77	0.24	1.30	0.96	3.26
Macro	0.10**	0.07	0.14*	–0.02	0.06
	2.14	0.92	1.75	–0.70	1.19
Constant	0.31***	–0.17	0.28***	0.08***	0.19***
	3.72	–1.22	2.84	3.11	3.52
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	124	130	127	127	126
R ²	0.24	0.13	0.10	0.07	0.17

Note:

*p<0.1; **p<0.05; ***p<0.01

Volatility, Carry and Short USD are standardized ratios (fix/nonfix).

Table 1.21: Regression of Log Ratio of Effective Spread (Fix/Non-Fix) — Media Event

This table reports coefficient estimates for the media event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, *Effectivespread* is the log ratio of the volume-weighted spread at the time a trade occurs, versus the control period, for a given a currency-date. *AfterDummy* is a dummy variable, which takes the value of one for the time period after ‘the media event’ on the 12 June 2013. *Volatility* is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. *Vol(level)* is the time-weighted average of the options midpoint price for the fix window. *Carry(ratio)* and *ShortUSD(ratio)* is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	Effective spread				
	AUDUSD (1)	EURHUF (2)	EURSEK (3)	EURUSD (4)	GBPUSD (5)
After dummy	−0.12*** −3.72	−0.05 −0.41	0.18** 2.48	0.01 0.09	−0.10*** −4.07
Volatility	0.01 1.17	0.06 1.05	0.01 0.56	0.10** 2.03	0.02 1.29
Vol. (level)	0.01 0.88	0.12** 2.14	−0.11*** −3.24	0.01 0.14	−0.01 −0.88
Carry (ratio)	−0.01 −0.39	−0.02 −0.36	−0.04 −1.53	−0.07* −1.65	−0.01 −0.96
Short USD (ratio)	−0.003 −0.30	0.05 0.65	0.002 0.08	−0.03 −0.69	0.02 1.40
Month-end	0.12*** 2.81	0.25 0.80	0.30*** 3.57	0.11 0.47	0.15*** 2.70
Macro	−0.03 −1.40	−0.09 −0.65	0.05 0.83	−0.03 −0.36	0.03 1.20
Constant	−0.03 −0.91	−0.40** −2.46	−0.42*** −5.54	0.27** 2.23	−0.09*** −3.46
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	124	112	127	94	126
R ²	0.25	0.07	0.13	0.18	0.21

Note:

*p<0.1; **p<0.05; ***p<0.01

76

Volatility, Carry and Short USD are standardized ratios (fix/nonfix).

Table 1.22: Regression of Log Ratio of Price impact (1ms) — Media Event

This table reports coefficient estimates for the media event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, $Priceimpact(1ms)$ is the log ratio of the volume-weighted price impact of a trade over a 1 millisecond period, versus the control period, for a given a currency-date. $AfterDummy$ is a dummy variable, which takes the value of one for the time period after ‘the media event’ on the 12 June 2013. $Volatility$ is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. $Vol(level)$ is the time-weighted average of the options midpoint price for the fix window. $Carry(ratio)$ and $ShortUSD(ratio)$ is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. $Month-end$ is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. $macro$ is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	Price impact (1ms)				
	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	-0.04	-0.03	0.11	0.08	-0.12*
	-0.47	-0.18	0.94	0.52	-1.83
Volatility	0.04	-0.01	0.02	0.32***	0.03
	1.19	-0.06	0.44	3.64	0.94
Vol. (level)	-0.01	0.10	-0.07	0.06	-0.03
	-0.22	1.21	-1.20	0.61	-0.71
Carry (ratio)	-0.02	-0.04	0.004	-0.01	0.03
	-0.58	-0.47	0.08	-0.12	1.01
Short USD (ratio)	0.01	-0.04	0.06	0.03	0.02
	0.46	-0.43	1.26	0.25	0.61
Month-end	0.26	-0.43	0.21	-0.87**	0.16
	1.03	-0.90	0.84	-2.04	0.57
Macro	-0.06	0.12	-0.07	-0.11	-0.02
	-0.80	0.70	-0.71	-0.61	-0.29
Constant	-0.43***	-0.40*	-0.58***	0.61***	-0.34***
	-4.74	-1.71	-4.97	4.68	-4.40
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	123	106	127	88	126
R ²	0.07	0.07	0.06	0.24	0.09

Note:

*p<0.1; **p<0.05; ***p<0.01

Volatility, Carry and $\sqrt{Short\ USD}$ are standardized ratios (fix/nonfix).

Table 1.23: Regression of Log Ratio of Price impact (1 Second) — Media Event

This table reports coefficient estimates for the media event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, *Priceimpact(1sec)* is the log ratio of the volume-weighted price impact of a trade over a 1 second period, versus the control period, for a given a currency-date. *AfterDummy* is a dummy variable, which takes the value of one for the time period after ‘the media event’ on the 12 June 2013. *Volatility* is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. *Vol(level)* is the time-weighted average of the options midpoint price for the fix window. *Carry(ratio)* and *ShortUSD(ratio)* is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	Price impact (1sec)				
	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	–0.09	0.23	0.10	–0.39*	–0.05
	–1.08	1.00	0.83	–1.92	–0.60
Volatility	0.02	–0.05	0.02	0.30***	0.03
	0.49	–0.41	0.43	2.89	0.59
Vol. (level)	–0.04	0.04	–0.03	–0.07	–0.06
	–0.95	0.47	–0.55	–0.78	–1.26
Carry (ratio)	0.01	0.04	–0.01	0.01	0.04
	0.35	0.40	–0.22	0.08	1.02
Short USD (ratio)	–0.02	–0.06	0.04	–0.02	0.01
	–0.45	–0.57	0.77	–0.24	0.29
Month-end	0.16	–0.60*	0.05	–0.59	0.12
	0.74	–1.75	0.22	–1.08	0.53
Macro	–0.08	0.14	–0.11	–0.15	–0.04
	–1.15	0.72	–1.07	–0.88	–0.48
Constant	–0.64***	–0.40	–0.49***	0.19	–0.47***
	–5.83	–1.45	–4.50	1.01	–4.28
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	123	106	127	82	126
R ²	0.11	0.06	0.04	0.29	0.05

Note:

*p<0.1; **p<0.05; ***p<0.01

Volatility, Carry and Short USD are standardized ratios (fix/nonfix).

Table 1.24: Regression of Log Ratio of Price impact (5 Sec) — Media Event

This table reports coefficient estimates for the media event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, *Priceimpact(5sec)* is the log ratio of the volume-weighted price impact of a trade over a 5 second period, versus the control period, for a given a currency-date. *AfterDummy* is a dummy variable, which takes the value of one for the time period after ‘the media event’ on the 12 June 2013. *Volatility* is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. *Vol(level)* is the time-weighted average of the options midpoint price for the fix window. *Carry(ratio)* and *ShortUSD(ratio)* is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	Price impact (5sec)				
	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	-0.10	0.07	-0.11	0.60	0.04
	-0.88	0.31	-0.54	1.31	0.44
Volatility	-0.03	-0.08	0.06	0.04	0.05
	-0.72	-0.61	1.32	0.18	1.03
Vol. (level)	-0.03	0.11	-0.01	0.30	-0.10**
	-0.52	1.40	-0.12	1.18	-2.14
Carry (ratio)	-0.01	0.04	-0.08	0.005	0.05
	-0.18	0.31	-1.31	0.02	1.23
Short USD (ratio)	-0.02	-0.05	0.01	-0.17	-0.02
	-0.35	-0.50	0.21	-0.74	-0.45
Month-end	0.17	-0.18	0.21	0.25	0.11
	0.68	-0.33	0.93	0.37	0.44
Macro	-0.03	0.05	-0.09	-0.16	-0.04
	-0.31	0.25	-0.59	-0.28	-0.44
Constant	-0.56***	-0.18	-0.30**	0.15	-0.39***
	-3.35	-0.70	-2.01	0.60	-3.55
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	123	101	127	72	124
R ²	0.11	0.05	0.07	0.09	0.09

Note:

79

*p<0.1; **p<0.05; ***p<0.01

Volatility, Carry and Short USD are standardized ratios (fix/nonfix).

Table 1.25: Regression of Log Ratio of Volume (Fix/Non-Fix) — Window Event

This table reports coefficient estimates for the window event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, *Volume* is the log ratio of the total volume traded in the fix, versus the control period, for a given a currency-date. *AfterDummy* is a dummy variable, which takes the value of one for the time period after ‘the window event’ on the 15 February 2015. *Volatility* is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. *Vol(level)* is the time-weighted average of the options midpoint price for the fix window. *Carry(ratio)* and *ShortUSD(ratio)* is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	AUDUSD	EURHUF	Volume EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	0.08	0.34*	0.19	0.61***	0.26**
	0.60	1.76	1.18	3.10	2.01
Volatility	-0.03	-0.15	0.13	0.25***	-0.12**
	-0.30	-1.27	1.58	3.18	-1.97
Vol. (level)	-0.09	-0.03	0.02	-0.15	-0.01
	-1.47	-0.24	0.17	-1.47	-0.27
Carry (ratio)	-0.002	-0.13	0.08	-0.20**	-0.04
	-0.02	-1.46	1.09	-2.27	-0.79
Short USD (ratio)	0.10	0.14	0.03	0.15	0.03
	1.45	1.50	0.39	1.62	0.51
Month-end	2.13***	1.30***	1.19***	0.96**	1.36***
	8.55	2.67	4.01	2.37	13.59
Macro	-0.25*	-0.39*	0.06	-0.16	-0.24**
	-1.72	-1.77	0.36	-0.78	-2.01
Constant	-1.59***	-2.47***	-1.70***	-3.58***	-1.82***
	-8.61	-8.92	-9.36	-15.82	-10.92
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	121	114	120	108	121
R ²	0.34	0.23	0.28	0.23	0.25

Note:

*p<0.1; **p<0.05; ***p<0.01

Volatility, Carry and Short USD are standardized ratios (fix/nonfix).

Table 1.26: Regression of Log Ratio of Quoted spread (Fix/Non-Fix) — Window Event

This table reports coefficient estimates for the window event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, *Quotedspread* is the log ratio of the time-weighted quoted spread in the fix, versus the control period, for a given a currency-date. *AfterDummy* is a dummy variable, which takes the value of one for the time period after ‘the window event’ on the 15 February 2015. *Volatility* is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. *Vol(level)* is the time-weighted average of the options midpoint price for the fix window. *Carry(ratio)* and *ShortUSD(ratio)* is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	Quoted spread				
	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	0.07***	0.33***	−0.02	0.11***	0.06*
	2.77	4.49	−0.42	5.03	1.91
Volatility	0.02	0.03	0.02	0.02	0.03*
	1.35	0.85	0.59	1.37	1.71
Vol. (level)	−0.01	−0.02	−0.02	0.01	0.01
	−1.09	−0.56	−0.83	0.93	0.78
Carry (ratio)	0.01	0.07**	−0.001	0.01	0.02
	0.74	2.25	−0.03	1.34	1.12
Short USD (ratio)	0.02*	0.002	0.04	0.02	0.01
	1.66	0.07	1.22	1.36	0.84
Month-end	−0.04	0.19	0.003	0.01	−0.06
	−0.72	0.61	0.03	0.17	−0.96
Macro	0.005	−0.01	−0.0002	0.04*	0.03
	0.17	−0.16	−0.003	1.72	0.73
Constant	−0.37***	−0.24***	−0.20***	0.01	−0.37***
	−10.77	−3.20	−3.38	0.29	−9.87
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	121	122	120	121	121
R ²	0.16	0.25	0.07	0.32	0.15

Note:

*p<0.1; **p<0.05; ***p<0.01

Volatility, Carry and Short USD are standardized ratios (fix/nonfix).

Table 1.27: Regression of Log Ratio of Depth at best (Fix/Non-Fix) — Window Event

This table reports coefficient estimates for the window event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, *Depthatbest* is the log ratio of the time-weighted depth at the best bid and offer in the fix, versus the control period, for a given a currency-date. *AfterDummy* is a dummy variable, which takes the value of one for the time period after ‘the window event’ on the 15 February 2015. *Volatility* is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. *Vol(level)* is the time-weighted average of the options midpoint price for the fix window. *Carry(ratio)* and *ShortUSD(ratio)* is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	Depth at best				
	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	−0.25***	−0.04	−0.22**	−0.0002	−0.26***
	−4.02	−0.51	−2.46	−0.004	−3.09
Volatility	−0.02	−0.06*	0.06	−0.01	−0.06**
	−0.55	−1.72	0.96	−0.40	−2.00
Vol. (level)	0.004	−0.09*	0.04	0.01	0.02
	0.13	−1.77	0.72	0.47	0.56
Carry (ratio)	−0.04	0.02	0.04	−0.03**	−0.05
	−1.21	0.44	0.74	−2.01	−1.24
Short USD (ratio)	0.10**	−0.05	−0.02	−0.03	0.02
	2.43	−1.18	−0.52	−1.25	0.73
Month-end	0.35***	−0.05	0.30*	0.05	0.29***
	3.08	−0.18	1.95	1.08	3.12
Macro	−0.01	−0.14	0.03	−0.07*	−0.09
	−0.14	−1.45	0.31	−1.78	−1.03
Constant	0.35***	0.002	0.31**	−0.16***	0.51***
	5.56	0.02	2.50	−4.86	5.52
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	121	122	120	121	121
R ²	0.25	0.12	0.18	0.10	0.19

Note:

82

*p<0.1; **p<0.05; ***p<0.01

Volatility, Carry and Short USD are standardized ratios (fix/nonfix).

Table 1.28: Regression of Log Ratio of Depth at Top 10 (Fix/Non-Fix) — Window Event

This table reports coefficient estimates for the window event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, *DepthatTop10* is the log ratio of the time-weighted depth at the best 10 bid and offer price levels in the fix, versus the control period, for a given a currency-date. *AfterDummy* is a dummy variable, which takes the value of one for the time period after ‘the window event’ on the 15 February 2015. *Volatility* is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. *Vol(level)* is the time-weighted average of the options midpoint price for the fix window. *Carry(ratio)* and *ShortUSD(ratio)* is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	Depth at top 10				
	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	–0.20***	0.15	–0.15*	0.06***	–0.07
	–4.08	1.63	–1.79	2.60	–1.31
Volatility	–0.02	–0.05	–0.01	–0.02	–0.05*
	–1.27	–1.50	–0.13	–1.35	–1.67
Vol. (level)	–0.003	–0.06	0.04	–0.01	–0.02
	–0.13	–1.23	0.77	–0.53	–0.66
Carry (ratio)	–0.01	–0.01	0.03	–0.01	–0.001
	–0.75	–0.24	0.66	–1.25	–0.02
Short USD (ratio)	0.05*	–0.10*	–0.01	–0.01	0.03
	1.87	–1.92	–0.24	–1.06	1.08
Month-end	0.27***	–0.20	0.03	–0.03	0.25**
	3.19	–0.78	0.24	–1.14	2.41
Macro	0.01	–0.15	0.10	–0.04*	–0.05
	0.16	–1.46	1.02	–1.87	–0.84
Constant	0.29***	–0.08	0.10	–0.05**	0.20***
	4.03	–0.93	0.81	–2.31	2.64
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	121	122	120	121	121
R ²	0.24	0.10	0.11	0.17	0.11

Note:

Table 1.29: Regression of Log Ratio of Effective Spread (Fix/Non-Fix) — Window Event

This table reports coefficient estimates for the window event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, *Effectivespread* is the log ratio of the volume-weighted spread at the time a trade occurs, versus the control period, for a given a currency-date. *AfterDummy* is a dummy variable, which takes the value of one for the time period after ‘the window event’ on the 15 February 2015. *Volatility* is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. *Vol(level)* is the time-weighted average of the options midpoint price for the fix window. *Carry(ratio)* and *ShortUSD(ratio)* is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm..

	Effective spread				
	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	0.02	0.36***	−0.20***	−0.12*	−0.10***
	0.83	3.94	−3.00	−1.80	−3.28
Volatility	0.002	0.03	0.05	0.01	0.03
	0.15	0.71	1.34	0.22	1.59
Vol. (level)	0.01	0.02	−0.03	0.01	0.001
	0.62	0.30	−0.95	0.32	0.08
Carry (ratio)	0.01	0.07*	0.04	−0.04	0.02
	1.12	1.77	1.28	−1.23	0.98
Short USD (ratio)	0.01	−0.07	0.01	−0.01	0.02
	1.48	−1.38	0.27	−0.16	1.33
Month-end	0.15***	0.10	−0.09	0.20	0.07**
	3.34	0.53	−0.79	1.32	2.10
Macro	−0.03	−0.02	0.001	−0.05	0.02
	−1.08	−0.19	0.01	−0.75	0.50
Constant	−0.10***	−0.25*	0.12*	0.03	−0.11***
	−4.05	−1.76	1.84	0.33	−3.43
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	121	114	120	108	121
R ²	0.15	0.23	0.12	0.09	0.17

Note:

*p<0.1; **p<0.05; ***p<0.01

Volatility, Carry and Short USD are standardized ratios (fix/nonfix).

Table 1.30: Regression of Log Ratio of Price impact (1ms) — Window Event

This table reports coefficient estimates for the window event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, $Priceimpact(1ms)$ is the log ratio of the volume-weighted price impact of a trade over a 1 millisecond period, versus the control period, for a given a currency-date. $AfterDummy$ is a dummy variable, which takes the value of one for the time period after ‘the window event’ on the 15 February 2015. $Volatility$ is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. $Vol(level)$ is the time-weighted average of the options midpoint price for the fix window. $Carry(ratio)$ and $ShortUSD(ratio)$ is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. $Month-end$ is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. $macro$ is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	Price impact (1ms)				
	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	0.13*	-0.12	-0.03	0.25*	0.18**
	1.93	-0.74	-0.32	1.94	2.50
Volatility	-0.04	-0.02	0.22*	0.02	0.09***
	-0.84	-0.32	1.81	0.33	2.71
Vol. (level)	0.07*	0.04	-0.02	0.02	-0.01
	1.72	0.43	-0.30	0.26	-0.43
Carry (ratio)	0.02	0.11	-0.12	0.04	0.06*
	0.69	1.50	-1.17	0.66	1.75
Short USD (ratio)	-0.001	0.06	0.18	0.03	0.03
	-0.04	0.71	1.54	0.48	0.94
Month-end	0.34	0.17	0.20	0.03	0.09
	1.12	0.36	1.17	0.16	0.79
Macro	-0.06	-0.07	-0.08	-0.06	0.11
	-0.70	-0.38	-0.58	-0.54	1.61
Constant	-0.52***	-0.16	-0.16	-0.30**	-0.42***
	-5.31	-0.80	-1.27	-2.30	-6.28
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	121	110	120	104	121
R ²	0.14	0.05	0.25	0.10	0.18

Note:

85

*p<0.1; **p<0.05; ***p<0.01

Volatility, Carry and Short USD are standardized ratios (fix/nonfix).

Table 1.31: Regression of Log Ratio of Price impact (1 Second) — Window Event

This table reports coefficient estimates for the window event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, *Priceimpact(1sec)* is the log ratio of the volume-weighted price impact of a trade over a 1 second period, versus the control period, for a given a currency-date. *AfterDummy* is a dummy variable, which takes the value of one for the time period after ‘the window event’ on the 15 February 2015. *Volatility* is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. *Vol(level)* is the time-weighted average of the options midpoint price for the fix window. *Carry(ratio)* and *ShortUSD(ratio)* is the log ratio (or the log return) of the time- weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	Price impact (1sec)				
	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	0.22***	-0.07	0.02	0.29	0.22***
	3.63	-0.44	0.23	0.73	3.25
Volatility	-0.02	0.04	0.07	-0.31*	0.11***
	-0.41	0.51	1.54	-1.85	3.56
Vol. (level)	0.08**	0.09	0.03	-0.13	0.03
	2.44	1.12	0.83	-0.56	1.02
Carry (ratio)	0.004	0.12*	-0.02	-0.04	0.07**
	0.16	1.80	-0.49	-0.29	2.39
Short USD (ratio)	-0.02	0.02	0.02	-0.05	0.08**
	-0.72	0.20	0.42	-0.31	2.30
Month-end	0.32*	-0.01	0.13	0.50	0.15**
	1.73	-0.01	0.68	1.15	2.14
Macro	-0.12*	0.01	-0.03	-0.13	0.13**
	-1.79	0.06	-0.38	-0.39	1.98
Constant	-0.38***	-0.18	-0.16	-0.57**	-0.48***
	-4.36	-0.91	-1.53	-2.44	-5.72
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	121	110	119	101	121
R ²	0.23	0.06	0.05	0.17	0.30

Note:

86

*p<0.1; **p<0.05; ***p<0.01

Volatility, Carry and Short USD are standardized ratios (fix/nonfix).

Table 1.32: Regression of Log Ratio of Price impact (5 Sec) — Window Event

This table reports coefficient estimates for the window event study using the specification in Formula 1.5.1 for regressions in each currency pair. The dependent variable, *Priceimpact(5sec)* is the log ratio of the volume-weighted price impact of a trade over a 5 second period, versus the control period, for a given a currency-date. *AfterDummy* is a dummy variable, which takes the value of one for the time period after ‘the window event’ on the 15 February 2015. *Volatility* is the log ratio of the time-weighted average of the 1-week options midpoint price for the respective currency in the control period, versus the fix window. *Vol(level)* is the time-weighted average of the options midpoint price for the fix window. *Carry(ratio)* and *ShortUSD(ratio)* is the log ratio (or the log return) of the time-weighted average of these index values in the control period, versus the fix window. Carry is an ETF that employs carry strategies and Short USD is a basket index designed to replicate a TWI of short USD. *Month – end* is a dummy variable that takes the value of one for the last trading day of the month in the respective currency. *macro* is a dummy variable that takes the value of one for macro news events of the highest volatility rating that occur in the period from 9am to 4pm.

	Price impact (5sec)				
	AUDUSD	EURHUF	EURSEK	EURUSD	GBPUSD
	(1)	(2)	(3)	(4)	(5)
After dummy	0.20**	-0.17	0.12	0.08	0.26**
	2.52	-1.00	1.24	0.28	2.37
Volatility	0.05	0.19**	0.21**	-0.06	0.07*
	1.14	2.04	2.49	-0.47	1.67
Vol. (level)	0.09**	0.09	-0.03	0.15	0.01
	2.27	1.00	-0.73	0.64	0.28
Carry (ratio)	0.02	0.11	-0.10	-0.03	0.01
	0.63	1.46	-1.35	-0.20	0.30
Short USD (ratio)	0.01	-0.01	0.14*	-0.09	0.12**
	0.42	-0.12	1.74	-0.56	2.56
Month-end	0.38***	-0.35	0.08	0.38	0.40***
	3.07	-0.78	0.36	1.08	3.10
Macro	-0.11	0.01	-0.03	-0.34	0.14
	-1.18	0.05	-0.29	-1.20	1.26
Constant	-0.25***	-0.51**	-0.26**	-0.11	-0.36***
	-2.65	-1.97	-2.40	-0.52	-4.05
Weekday dummy?	Yes	Yes	Yes	Yes	Yes
Observations	121	106	117	86	121
R ²	0.20	0.13	0.26	0.07	0.18

Note:

87

*p<0.1; **p<0.05; ***p<0.01

Volatility, Carry and Short USD are standardized ratios (fix/nonfix).

1..5 Miscellaneous Figures and Tables

Figure 1.11: Mean Trading Volume — every 15 Seconds — 3 Months Pre-Post Window Change — GBPUSD

This Figure reports mean trading volume in GBPUSD for 15-second time intervals from 15:55 to 16:05, with means calculated across all time intervals in a 3-month period before and after the window event. Volume is in millions of USD and Time is reported in decimal format.




Figure 1.12: Mean Trading Volume — every 15 Seconds — 3 Months Pre-Post Window Change — AUDUSD

This Figure reports mean trading volume in AUDUSD for 15-second time intervals from 15:55 to 16:05, with means calculated across all time intervals in a 3-month period before and after the window event. Volume is in millions of USD and Time is reported in decimal format.




Figure 1.13: Mean Time-Weighted Quoted Spreads — every 15 Seconds — 3 Months Pre-Post Window Change — GBPUSD

This Figure reports mean time-weighted quoted spreads in GBPUSD for 15-second time intervals from 15:55 to 16:05, with means calculated across all time intervals in a 3-month period before and after the window event. The spread is in absolute values and Time is reported in decimal format.

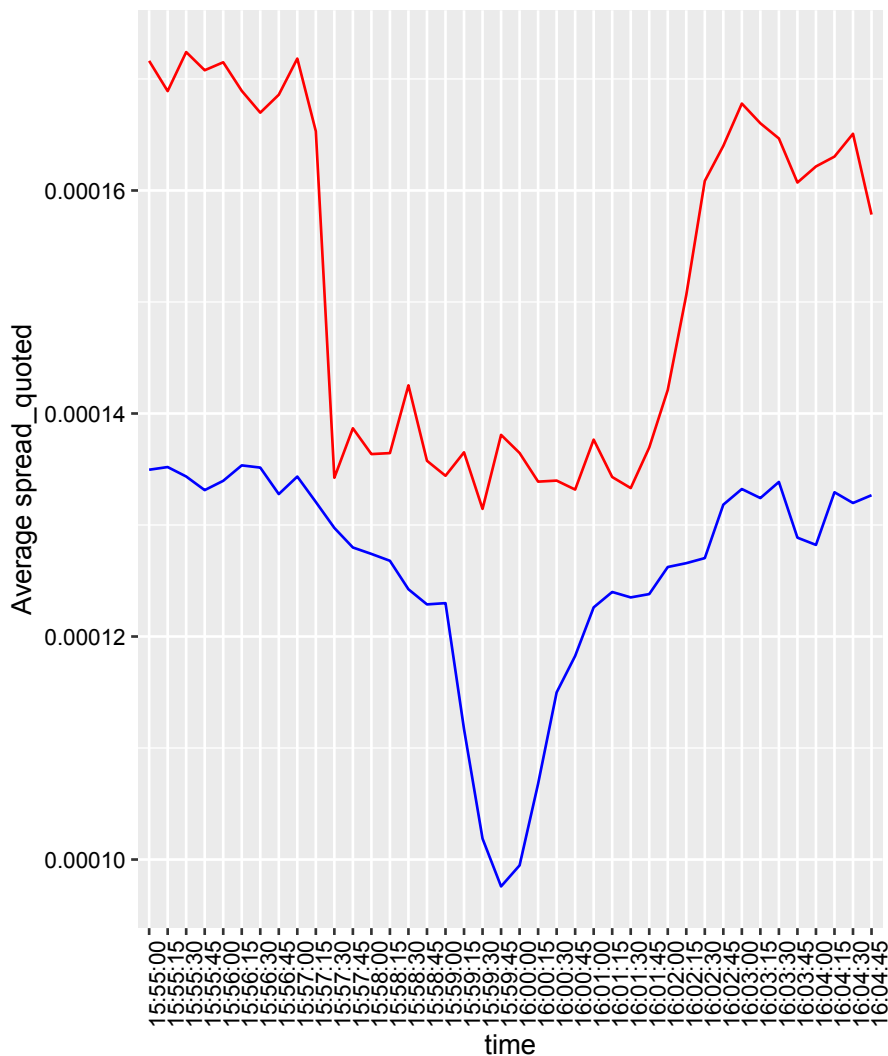


Figure 1.14: Mean Time-Weighted Quoted Spreads — every 15 Seconds — 3 Months Pre-Post Window Change — AUDUSD

This Figure reports mean time-weighted quoted spreads in AUDUSD for 15-second time intervals from 15:55 to 16:05, with means calculated across all time intervals in a 3-month period before and after the window event. The spread is in absolute values and Time is reported in decimal format.




Figure 1.15: Mean Total Volume — Entire Sample — 6am to 10pm — GBPUSD

This Figure reports mean volume in GBPUSD for 30-second time intervals from 6am to 10pm, with means calculated across all dates in our sample.




Figure 1.16: Mean Total Volume — Entire Sample — 6am to 10pm — AUDUSD

This Figure reports mean volume in AUDUSD for 30-second time intervals from 6am to 10pm, with means calculated across all dates in our sample.




Figure 1.17: Mean Time-Weighted Quoted Spreads — Entire Sample — 6am to 10pm — GB-PUSD

This Figure reports time-weighted quoted spreads in GBPUSD for 30-second time intervals from 6am to 10pm, with means calculated across all dates in our sample. The spread is in absolute values.




Figure 1.18: Mean Time-Weighted Quoted Spreads — Entire Sample — 6am to 10pm — AUDUSD

This Figure reports mean time-weighted quoted spreads in AUDUSD for 30-second time intervals from 6am to 10pm, with means calculated across all dates in our sample. The spread is in absolute values.




Figure 1.19: Mean Volume-Weighted Effective Spreads — Entire Sample — 6am to 10pm — GBPUSD

This Figure reports mean volume-weighted effective spreads in GBPUSD for 30-second time intervals from 6am to 10pm, with means calculated across all dates in our sample. The spread is in absolute values.




Figure 1.20: Mean Volume-Weighted Effective Spreads — Entire Sample — 6am to 10pm — AUDUSD

This Figure reports mean volume-weighted effective spreads in AUDUSD for 30-second time intervals from 6am to 10pm, with means calculated across all dates in our sample. The spread is in absolute values.




Figure 1.21: Mean Volume-Weighted 1-Second Price Impact — Entire Sample — 6am to 10pm — GBPUSD

This Figure reports mean volume-weighted price impacts in GBPUSD, calculated over 1 second, for 30-second time intervals from 6am to 10pm, with means calculated across all dates in our sample. The price impact is in absolute values.




Figure 1.22: Mean Volume-Weighted 1-Second Price Impact — Entire Sample — 6am to 10pm — AUDUSD

This Figure reports mean volume-weighted price impacts in AUDUSD, calculated over 1 second, for 30-second time intervals from 6am to 10pm, with means calculated across all dates in our sample. The price impact is in absolute values.




Figure 1.23: Mean Time-Weighted Average Depth at the Best Bid or Offer — Entire Sample — 6am to 10pm — GBPUSD

This Figure reports time-weighted average depths in GBPUSD for 30-second time intervals from 6am to 10pm, with means calculated across all dates in our sample. Depth is reported with millions of USD.




Figure 1.24: Mean Time-Weighted Average Depth at the Best Bid or Offer — Entire Sample — 6am to 10pm — AUDUSD

This Figure reports mean time-weighted average depths in AUDUSD for 30-second time intervals from 6am to 10pm, with means calculated across all dates in our sample. Depth is reported with millions of USD.




Figure 1.25: Mean Time-Weighted Average Depth at the Top 10 Best Bid or Offer Levels — Entire Sample — 6am to 10pm — GBPUSD

This Figure reports time-weighted average depths in GBPUSD for 30-second time intervals from 6am to 10pm, with means calculated across all dates in our sample. Depth is reported with millions of USD.




Figure 1.26: Mean Time-Weighted Average Depth at the Top 10 Best Bid or Offer Levels — Entire Sample — 6am to 10pm — AUDUSD

This Figure reports mean time-weighted average depths in AUDUSD for 30-second time intervals from 6am to 10pm, with means calculated across all dates in our sample. Depth is reported with millions of USD.




Figure 1.27: Volatility — Mean of High Minus Low of Trade Prices in Each 30 Seconds — Entire Sample — 6am to 10pm — GBPUSD

This Figure reports time-weighted average depths in GBPUSD for 30-second time intervals from 6am to 10pm, with means calculated across all dates in our sample. Reported in absolute values.




Figure 1.28: Volatility — Mean of High Minus Low of Trade Prices in Each 30 Seconds — Entire Sample — 6am to 10pm — AUDUSD

This Figure reports the mean, across the entire sample, of the highest minus the lowest trade price in a 30 second interval, from 6am to 10pm. Reported in absolute values.




Figure 1.29: Mean Total Number of Messages Entire Sample — 6am to 10pm — GBPUSD
This Figure reports the total number of messages in each 30 second interval, calculated as a mean across the entire sample. Reported in absolute values.




Figure 1.30: Mean Total Number of Messages Entire Sample — 6am to 10pm — AUDUSD
This Figure reports the total number of messages in each 30 second interval, calculated as a mean across the entire sample. Reported in absolute values.




Figure 1.31: Mean Order-book Slope — Entire Sample — 6am to 10pm — GBPUSD

This Figure reports the total number of messages in each 30 second interval, calculated as a mean across the entire sample. Reported in absolute values.




Figure 1.32: Mean Order-book Slope — Entire Sample — 6am to 10pm — AUDUSD

This Figure reports the mean order-book slope, calculated as the depth at the best bid(ask) less the depth at the top ten buy(sell) levels, divided by 9. This is then calculated as a time-weighted average and then as a mean across the entire sample.




Table 1.33: Summary Statistics During the full trading day from 8am to 8pm — by Currency-Year

Volume is total volume during the trading day from 8am to 8pm. Depth is computed as the average of depth at bid and offer sides of the book (at the best bid and offer and the top 10 levels or all levels). Mean number of messages ('#msg'), quote life ('qlif'), unique TCIDs ('#TCIDs'), number of trades ('#trades') and number of aggressor trades ('#agr.trades') are calculated across all currency-dates. #agr.trades is smaller than #trades because it doesn't include the component orders that make up a trade - of which there are least 2. Quoted spread ('Qtd.Sprd') is time-weighted, effective spreads ('Eff.Sprd.') and price impact (PI) is volume-weighted in basis points.

Pair	Year	# TCIDs	Volume	Depth best	Depth top	Depth total	#msg	q.life	# trades	#agr. trades	Qtd. sprd	Eff. sprd	PI lms	PI ls	PI 5s
audusd	2012	223.2	14950.3	8.1	121.1	283.4	187219.2	245.2	10332.6	4224.4	1.46	1	0.88	0.98	1.07
	2013	222.9	18177.2	5.8	88.7	239.8	191810.9	180.44	12881.1	5511.4	1.76	3.9	0.89	1.18	1.27
	2014	198.5	12218.4	6.3	87	252.6	156363.4	216.48	8556.1	3701.5	1.91	1.3	0.96	1.18	1.28
	2015	183.2	12843.6	5.7	76.5	249	212718.8	146.06	9606.6	4277	2.54	1.6	1.16	1.53	1.6
	2017	138.9	6540.8	6.2	87.4	314.9	186881.7	119.85	4868.5	2107.5	2.04	1.4	1.14	1.34	1.39
	2012	88.2	1001.3	1.6	4.3	32.4	10415.1	944.36	821.7	396.9	6.79	3.7	2.37	2.52	2.73
	2013	88.9	1288.1	2	5.1	34.6	15273	563.57	1048	497.7	6.86	3.5	2.21	2.3	2.48
eurhuf	2014	79.4	1203.8	1.9	5.1	32.8	21162.4	410.13	1013.3	484.7	6	3	2.2	2.1	2.23
	2015	71.1	1215.7	1.9	4.4	30.7	18946.8	422.73	1036.8	500.2	7.89	3.9	2.95	2.66	2.84
	2017	58.1	620.3	2	8.3	41.4	49627.6	88.96	535	252.8	4.41	2.7	1.38	1.65	1.93
eursek	2012	115.5	3674.7	1.9	4.1	35	47752.5	384.83	2805.3	1351.4	2.88	1.5	1.32	1.27	1.4
	2013	107.9	3541.1	2	4.1	33.9	80000.9	170.29	2780.8	1321.6	3.37	1.8	1.48	1.46	1.7
	2014	97.3	2779.7	2	4.3	32.1	77906	128.82	2270.4	1089.8	3.44	1.7	17.22	4.1	9.53
	2015	92.3	3575.4	2.2	4.3	43.6	177314.6	147.86	2912.2	1405	4.88	2.3	1.83	1.92	2.26
	2017	80.7	2337.5	3.1	15.3	106.1	84043.1	50.18	1885.1	889.5	2.7	1.5	0.97	1.37	1.5
	2012	151.2	2047.7	3.4	28.6	48.6	170697.9	101.25	1504.5	724.6	1.28	0.8	0.52	0.91	1.07
	2013	132.8	1678.7	2.4	33.1	57.8	125233.2	78.14	1258.5	616.9	1.62	1	0.53	1.06	1.17
eurusd	2014	114.8	1349.7	2.5	39	61.6	71431.9	120.46	970.4	462.4	1.58	1.1	0.62	1.01	1.11
	2015	113.1	1697.3	2.2	37.5	56.5	167992	141.97	1158.5	562.3	2.22	1.6	0.9	1.71	1.81
	2017	77.8	1196.2	1.3	15.4	55.7	151302.6	38.14	1047.9	489.3	1.28	1	0.66	1.05	1.12
gbpusd	2012	206.3	10615.7	4.5	72.2	135.8	184067	130.59	8114	3609.1	1.16	0.8	0.59	0.74	0.79
	2013	219.8	14557.4	3.6	57	154.2	238280	103.12	10980.5	4927.6	1.38	0.9	0.81	1	1.07
	2014	198.3	10354.9	4	51.2	146.4	205683	150.28	7906.4	3599.1	1.31	0.8	0.69	0.87	0.91
	2015	180.2	11219	3.3	41.5	152.9	251627.8	175.88	8832	4145.6	1.8	1.1	0.78	1.05	1.11
	2017	150	8049.7	3.6	47.1	296.1	222747.3	71.86	6312	2874.6	1.76	1.1	0.9	1.17	1.2

Table 1.34: Media event : Mean volume of aggressive and passive trades, by fix quarter, before and after (*) the event, for GBPUSD and AUDUSD. Volume is first summed across all TCIDs in each participant group, for each currency-date-fix quarter combination, and then averaged across currency pairs and dates. Absolute value ('Tot') in million of base currency, quarterly volume ('Q') as share of total. P-value of two-sample t-test for difference in mean of the ratio (first half)/(second half).

Participant	Before					After					p-value
	Tot	Q1	Q2	Q3	Q4	Tot*	Q1*	Q2*	Q3*	Q4*	
Aggr. Trades:											
Agency Broker	29.4	0.26	0.28	0.20	0.26	26.7	0.29	0.24	0.24	0.23	0.31
Asset Manager	15.0	0.20	0.27	0.31	0.22	26.0	0.12	0.60	0.21	0.08	
Commercial Bank	45.8	0.25	0.29	0.25	0.21	38.2	0.28	0.30	0.23	0.19	0.52
Custodian	24.2	0.29	0.32	0.22	0.18	27.3	0.33	0.26	0.23	0.18	0.13
Dealer	63.4	0.32	0.28	0.24	0.16	65.5	0.29	0.29	0.22	0.20	0.70
Dealer - R	110.5	0.32	0.31	0.22	0.15	98.8	0.31	0.32	0.21	0.17	0.63
Hedge Fund	9.6	0.10	0.27	0.42	0.21	14.9	0.19	0.41	0.13	0.27	0.57
Private Bank	27.5	0.36	0.11	0.36	0.16	8.3	0.24	0.28	0.12	0.36	
Prop Trader	14.4	0.28	0.26	0.23	0.23	17.4	0.30	0.28	0.22	0.20	0.67
Prop Trader - HFT	56.3	0.35	0.28	0.21	0.17	57.8	0.34	0.26	0.22	0.17	0.20
Passive Trades:											
Agency Broker	46.6	0.30	0.24	0.26	0.20	27.0	0.22	0.30	0.27	0.21	0.92
Asset Manager	21.7	0.25	0.31	0.34	0.11	6.3	0.21	0.32	0.16	0.32	
Commercial Bank	41.1	0.25	0.32	0.24	0.19	47.0	0.30	0.29	0.23	0.19	0.77
Custodian	29.3	0.27	0.30	0.25	0.18	23.1	0.28	0.26	0.26	0.20	0.01
Dealer	63.0	0.31	0.31	0.23	0.16	78.5	0.30	0.30	0.23	0.18	0.06
Dealer - R	109.0	0.33	0.31	0.22	0.15	109.5	0.31	0.31	0.21	0.16	0.33
Hedge Fund	7.2	0.20	0.34	0.14	0.33	11.3	0.14	0.19	0.44	0.22	0.66
Private Bank	17.5	0.25	0.37	0.11	0.27	5.9	0.40	0.17	0.17	0.26	
Prop Trader	8.6	0.27	0.28	0.26	0.19	9.4	0.27	0.22	0.22	0.29	0.59
Prop Trader - HFT	37.1	0.38	0.26	0.23	0.14	26.7	0.38	0.25	0.19	0.19	0.18

Table 1.35: Mean number of trades and proportion of seconds with trades during the fix. Calculated as a mean across all seconds in a given year.

Pair	Year	Mean #Trades	% Sec. w Trades
audusd	2012	1.34	0.48
audusd	2013	1.49	0.52
audusd	2014	1.25	0.47
audusd	2015	0.49	0.27
audusd	2017	0.27	0.15
eurhuf	2012	0.07	0.05
eurhuf	2013	0.07	0.06
eurhuf	2014	0.08	0.06
eurhuf	2015	0.04	0.03
eurhuf	2017	0.03	0.03
eursek	2012	0.59	0.36
eursek	2013	0.56	0.33
eursek	2014	0.55	0.29
eursek	2015	0.23	0.14
eursek	2017	0.16	0.11
eurusd	2012	0.05	0.04
eurusd	2013	0.04	0.03
eurusd	2014	0.04	0.03
eurusd	2015	0.03	0.02
eurusd	2017	0.04	0.03
gbpusd	2012	1.19	0.48
gbpusd	2013	1.38	0.51
gbpusd	2014	1.16	0.44
gbpusd	2015	0.50	0.27
gbpusd	2017	0.46	0.23

Table 1.36: Mean number of messages, quote life, unique TCIDs, number of trades and number of aggressor trades, daily by time window (fix or control). The mean is first taken with respect to all trades for a given currency pair-date combination, and then averaged across all currencies.

Period	#msg	q.life	#TCIDs	#trades	#agr.trades
control	25384.2	103.96	79.2	915.3	411.0
fix	1342.5	111.60	29.6	123.7	57.2

Table 1.37: Mean trading volume and depth of orderbook and best prices, top ten levels and total. Depth is computed as average of bid and offer, daily by time window (fix or control). Each measure is first averaged across all observations for a given date-currency pair combination, the aggregated into a single daily mean.

Period	Volume	Depth best	Depth top	Depth total
control	1218.2	3.4	38.6	120.4
fix	213.6	5.4	46.9	129.0

Table 1.38: Mean quoted and effective spreads and price impacts. Quoted spread is time-weighted, effective spreads and price impact is volume-weighted. Unit: basis points. The mean is first taken with respect to all trades for a given currency pair-date combination, and then averaged across all currencies. Daily by time window (fix or control).

Period	Qtd.sprd	Eff.sprd	Pr.impact 1ms	Pr.impact 1s	Pr.impact 5s
control	2.61	1.8	1.20	1.38	1.49
fix	2.15	1.7	1.03	1.12	1.23

Table 1.39: Media event: Mean price impact (5sec), by fix quarter, before and after (*), for GBPUSD and AUDUSD. Basis points. P-value for two-sample t-test of difference in mean price impact across the entire fix.

Participant	Q1	Q2	Q3	Q4	Q1*	Q2*	Q3*	Q4*	p-value
Agency Broker	1.4	0.8	1.2	0.6	0.8	1.0	0.7	-0.4	0.14
Asset Manager	1.0	-1.3	2.1	-1.0	1.0	0.0	0.0	1.3	
Commercial Bank	0.4	0.3	0.6	0.8	0.4	0.7	0.5	0.5	0.44
Custodian	0.6	0.8	0.8	0.8	0.5	0.5	0.8	0.5	0.95
Dealer	0.7	0.7	0.5	0.6	0.6	0.8	0.6	0.9	0.13
Dealer - R	1.2	0.9	0.7	1.0	0.6	0.7	0.6	0.8	0.67
Hedge Fund	0.0	0.4	0.0	1.0	1.1	1.1	0.1	1.1	0.33
Private Bank	-4.1	-0.8	3.1	1.1	1.1	-0.6	0.6	0.0	
Prop Trader	0.4	0.5	0.2	0.6	0.7	0.7	0.8	1.0	0.15
Prop Trader - HFT	0.6	0.6	0.7	0.9	0.7	0.8	1.0	0.7	0.79

Table 1.40: Correlation of flows (net position change) during the control window, for GBPUSD and AUDUSD. Net position change is computed as the sum of signed trade volume across all TCIDs in each category, using trades in the control window of 12pm to 2pm only.

	Broker	Ass.mngr	Cm.bank	Cstd	Dealer	Dealr-R	Hedge	Prop
Agency Broker								
Asset Manager								
Commercial Bank	-0.01	-0.10						
Custodian	-0.01	0.01	-0.12					
Dealer	-0.11	-0.15	-0.10	-0.15				
Dealer - R	-0.12	-0.13	-0.46	-0.10	-0.45			
Hedge Fund	0.05	-0.06	-0.06	-0.01	-0.07	-0.16		
Prop Trader	0.14	0.26	-0.14	-0.07	-0.09	-0.27	0.03	
Prop Trader - HFT	0.12	0.23	-0.18	-0.10	-0.21	-0.35	0.13	0.47

Table 1.41: Pct. negative bid-ask spread (mean, 30 seconds)

year	audusd	eurhuf	eursek	eurusd	gbpusd
2012	0.00	0.00	0.02	0.01	0.00
2013	0.02	0.01	0.02	0.01	0.01
2014	0.03	0.00	0.02	0.02	0.05
2015	0.02	0.00	0.02	0.00	0.02
2017	0.07	0.09	0.00	0.00	0.07

Chapter 2

Trading strategies and information flow around price benchmarks

JO SAAKVITNE, BI NORWEGIAN BUSINESS SCHOOL

Abstract

This paper characterizes equilibrium pricing and trading strategies in a competitive market where a subset of liquidity traders have a preference for executing their trades at a benchmark price. In the model, order flow is at a maximum while price impact is at a minimum when the price benchmark is set. These results are consistent with recent empirical evidence from foreign exchange markets. The market structure in the model gives incentives for the use of manipulative frontrunning strategies, but I show that the presence of a rational market maker partly negates the use of such strategies. This has important implications for benchmark design and understanding benchmark manipulation.

2.1 Introduction

Price benchmarks function as reference points for market participants, primarily for reasons related to reducing asymmetric information regarding the value of an underlying traded financial instrument (Duffie and Stein, 2015b). Price benchmarks often feature explicitly or implicitly in contracts, investment mandates, portfolio valuations and rebalancing rules (see for example Financial Stability Board (2014)). This leads to situations where market participants have incentives to trade at the benchmark price, for example to minimize tracking error against a reference portfolio. An example of such a market structure can be found in foreign exchange (FX), where the most important benchmark is known as the 4pm London fix. The rates from this benchmark are for example used in calculating cross-currency effects common equity and bond indices, in the valuation of exchange-traded funds, in the mark-to-market of assets held in custodian accounts, and in a wide range of other financial contracts (Evans, 2018). We take a closer look at the FX market and its main benchmark in Section 2.1.1.

The general question we are concerned with in this paper is the characterization of equilibrium pricing and trading strategies in a market where a subset of traders have a preference for executing their trades at the benchmark price. A characterization of equilibrium pricing and trading strategies in such a market is important in light of recent major manipulation and collusion scandals related to financial price benchmarks. For example, Ito and Yamada (2015) investigates price and volume data from FX markets looking for evidence of benchmark manipulation. It is useful in such cases to know what patterns one should expect based on a rational and competitive equilibrium model.




Figure 2.1: Mean trade imbalance (left) and 5 second price impact (right) in GBPUSD at the Thomson Reuters Matching trading platform, 15 minutes before, during and 15 minutes after the 4pm fix benchmark. More details in footnote 2.1.

Figure 2.1 show the empirical motivation for this paper, in the form of price impact and trade imbalance in the 15 minutes before, during and 15 minutes after the important FX 4pm benchmark. The plot shows that trade imbalance (defined as net flow divided by total trading volume, averaged across all market participants) peaks during the fix, while the price impact of trading is at a minimum.¹ Marsh et al. (2017) shows that the pattern of Figure 2.1 even prevails when one extends the empirical analysis to cover the entire day. Figure 2.2 shows the corresponding model predictions.

We will highlight three mains results from our model. First, we show that order flow peaks at the time when benchmark price is determined. This result is a consequence of the assumption that a subset of liquidity traders have a preference for transacting at the benchmark price. There is also a clustering of informed trading at the benchmark.

Second, the price impact of order flow reaches a minimum at the benchmark time. Or, equivalently, market depth is at a maximum. Intuitively, the non-informative trading interest of investors who want the benchmark price dominates the market when the benchmark price is set. This result holds whether the size of benchmark orders are correlated with the asset's fair value or not. This mechanism illustrates how a benchmark can serve as a "basin of attraction" for liquidity (Duffie and Stein, 2015b) - the existence of a benchmark generates incentives amongst market participants to transact at the benchmark price, which leads to an agglomeration of uninformed trading at the benchmark and therefore to better liquidity in the sense of lower price impact. The first and second main results are also consistent with the empirical facts in Figure 2.1 as well as Marsh et al. (2017).

Third, the price paths of the model features a temporary downward (upward) price movement at the benchmark time followed by an upward (downward) rebound after the benchmark. Such price patterns surrounding important price benchmarks have been identified empirically (Evans, 2018; Marsh et al., 2017), and earlier theoretical literature

¹Figure 2.1 is based on the working paper on FX benchmarks by Evans, Rime, O'Neill and Saakvitne (2018), which uses data from Thomson Reuters Matching. A complete description of the data and variables can be found there. Trade imbalance is defined for individual market participants, and is calculated as the absolute value of the participant's net position change divided by its total trading volume. An average trade imbalance is computed across all participants for the before, during and after-windows, for each trading day. The variable therefore measures the directionality of trading in these three time periods. Price impact is measured as the change in midpoint price 5 seconds after a trade has occurred. It is calculated for every trade in the sample period, and then volume-weighted into averages for the three time periods shown in the plot.

have associated the pattern with market manipulation (Osler et al., 2016). It is therefore important to note that this pattern can be generated also in a rational competitive equilibrium model, and is as such not alone proof of manipulative strategies being employed.

The model we develop in this paper builds on Kyle (1985). We extend Kyle's model by adding a new agent, called a *dealer*. The dealer has a profit/loss that depends on his transaction prices *relative to the benchmark price*. The dealer profits from selling at an average price higher than the benchmark price (or alternatively buying at a lower price). The dealer has such a profit function because he is executing orders on behalf of his clients, which are not modeled here, and guarantees that these clients receive or pay the benchmark price on their orders. The rest of the model is standard, in the sense that it features a profit-maximizing informed trader, liquidity traders trading for reasons outside of the model, and competitive market makers who observe aggregate order flow and sets market-clearing prices. Our model has three time periods, with the equilibrium price in the middle period being the *benchmark price*. The middle period will be referred to as the *benchmark time*.

2.1.1 The 4pm currency fix and fill-at-fix orders

The market structure we have in mind is reminiscent of the FX market and its main benchmark price. In this market, dealers offer an order type that are explicitly linked to the benchmark price.

FX markets are important in part because they are so large. Trading in FX averaged \$5.1 trillion per day in April 2016 (Bank of International Settlements, 2016b), dwarfing in size any other financial market. FX markets never close, but trading activity is concentrated around European business hours for most currency pairs (Evans, 2018). The most important benchmark price is the London 4 pm fixing rate (the WM/Reuters fix). The benchmark is calculated as a sample average of prices over a 5 minute interval at 4 pm London time². WMR Fixes are used for constructing indices comprising international securities, such as the MSCI indices and Barclays Global Bond Index. They are also routinely used to compute the returns on portfolios that contain foreign currency denominated securities, by for example country tracking funds and exchange traded funds (ETFs), as well as for valuation of foreign securities held in custodial accounts (Evans, 2018).

²Before February 2014 this was a 60 second interval

Trading volumes in major currency crosses spike around the London 4 pm fix, especially in the last day of each month (Melvin and Prins, 2015b). In a survey of fund managers and other investors, Financial Stability Board (2014) clarifies how market participants have incentives to transact at the benchmark rate, thereby creating a self-reinforcing dynamic:

Passively managed funds, including ETFs, [...] use the WMR fix to minimize tracking error and meet mandate transfer requirements. [...] Most investment mandates are benchmarked against global indices that use the WMR 4pm London fixing for FX valuation and transaction purposes. As a result, there is a self-reinforcing dynamic whereby indices are benchmarked versus these fixes, investors tracking those indices seek to minimize their FX risk by transacting directly at those same fixes.

To meet the demand of clients who want to trade at the benchmark price, dealers in FX markets have introduced a special type of order. These orders are called "fill-at-fix", or simply "fix orders". Fix orders are given to dealers *before* the benchmark time (4 pm London time), and a dealer who is accepting a fix order guarantees that the customer will receive the yet-to-be determined benchmark price. The dealer is free to cover his position any way he sees fit, within the constraints imposed by regulation.

Fix orders and the FX fix benchmark are especially interesting because they have been at the center of a very large market manipulation case. In the summer of 2013, news reports began to circulate that regulators were investigating manipulation of the London FX fixings.³ In November 2014, the United Kingdom's Financial Conduct Authority (FCA) imposed fines totaling \$1.7 billion on five of the world's largest banks⁴ for failing to control business practices in their G10 spot foreign exchange trading operations. The FCA determined that the five banks had failed to manage risks around client confidentiality, conflict of interest, and trading conduct. The banks had profited illegally at the expense of their customers and the market from manipulating fixing rates for G10 currencies, and had also used confidential customer order information to collude with other banks⁵. On the same day, the United States Commodity Futures Trading Commission (CFTC) imposed collective fines of \$1.4 billion against the same five banks for attempted manipulation of, and for aiding and abetting other banks' attempts to manipulate, global FX benchmark rates.

³*Traders Said to Rig Currency Rates to Profit Off Clients*, Bloomberg, June 12 2013.

⁴Citibank, HSBC, JPMorgan, RBS and UBS

⁵FCA press release 12 November 2014

The model developed in this paper does not include collusion or illegal information sharing between dealers. A model featuring collusion is developed in [Osler et al. \(2016\)](#). A related model in [Saakvitne \(2016b\)](#) shows how a particular market manipulation strategy can be sustained in a competitive equilibrium. Importantly, neither of these two models include learning market makers, rather they simply assume an exogenous price impact function of trading. In this paper we model dealers executing fix orders in a competitive equilibrium model with rational and learning counterparties.

2.1.2 Related literature

This paper extends the classical model by [Kyle \(1985\)](#). As such it is related to a large body of literature. A survey of the early parts of this literature can be found in [O'Hara \(1995\)](#).

One particular paper that is closely related to ours is [Admati and Pfleiderer \(1988\)](#). They have a market with two types of uninformed liquidity traders. Both types have a exogenously specified amount that they need to trade, but differ in their timing. One class of agents must transact at a specific time, while the other type have some discretion as to when they have to trade. The non-informed discretionary traders are somewhat similar to the dealers in our model. There are however two important differences in how the two models are set up.

Firstly, the dealers of our model have a payoff that depends on their transaction prices relative to the benchmark price. The discretionary traders of Admati and Pleiderer, on the other hand, want to minimize their trading cost and decides their timing based on this criterion. The distinction reflects the idea that we want to model trading strategies and information flows around price benchmark, which is a different research objective than Admati and Pleiderer.

Secondly, information is short-lived in the model of Admati and Pfleiderer, meaning that private information is valuable only for one trading interval. This assumption is a departure from the market model of [Kyle \(1985\)](#), but significantly improves the analytical tractability of their model. Our model features long-lived information. The price we pay is that closed-form analytical expressions are infeasible, at least for the general version of the model, and instead we must solve the model by a numerical algorithm.

Two recent related papers are [Degryse et al. \(2014\)](#) and [Choi et al. \(2017\)](#). Both these papers studies the problem of a large liquidity trader who must trade a fixed amount before a deadline and wishes to minimize the expected cost of trading. The paper by DeGryse et al features short-lived information, while the one by Choi et al has long-lived information. The latter paper extends the methodology of [Foster and Viswanathan \(1996\)](#) to achieve semi-closed form analytical solutions. Our paper differs from both these models in the special incentives of our dealer. While the discretionary traders of earlier models wants to minimize trading cost, the market structure in our model means that our dealer measures his trading profit against the benchmark price. Moreover, the numerical solution algorithm of the current paper allows for more general distributional assumptions than the earlier papers. For example, we do not require noise trades to be neither Gaussian nor iid, although we use this case as a baseline parametrization to allow for better comparison with other models.

This paper is also related to the literature on price benchmarks. An overview of the topic can be found in [Duffie and Stein \(2015b\)](#). [Duffie et al. \(2017\)](#) studies benchmarks in a search-theoretic framework, particularly relevant for over-the-counter (OTC) markets. This is different from our paper, where search costs does not feature in the market structure. Several recent empirical papers documents puzzling regularities surrounding the main benchmark in global foreign exchange markets ([Melvin and Prins, 2015b](#); [Ito and Yamada, 2015](#); [Evans, 2018](#); [Michelberger and Witte, 2016](#); [Marsh et al., 2017](#); [Yamada and Ito, 2017](#)). Some of these findings are explained by models of market manipulation in [Osler et al. \(2016\)](#) and [Saakvitne \(2016b\)](#). These models differs from the one in this paper in that we here endogenize price impact by introducing a learning market maker to our model. We show that this has the crucial effect of nullifying the effectiveness of manipulation strategies by the dealer (see Section [2.2.1](#)).

2.2 Model

The dealer: A key feature of our model is an agent who measures his trading profit relative to a price benchmark. We think of this agent as a dealer who has taken on benchmark orders from customers, akin to the fill-at-fix orders common in FX markets (see Section [2.1.1](#)). Our model has three time periods, and the price in period 2 is the benchmark price.

We let x be the net volume of benchmark orders that the dealer has to fill, defined so that

x is positive when the dealer has to sell. We model x as a random variable, drawn from the distribution F_x . Let ω_n be the dealers net sales in period n . We write the dealer's profit function J as the difference between his sales income and what he gives to his customers,

$$J(\omega) = \sum_{n=1}^3 (\omega_n p_n - c\omega_n^2) - xp_2 \quad (2.1)$$

The term $c\omega_n^2$ represents a (very small) quadratic transaction cost. We discuss this role and implications of this term when we turn to the equilibrium model (Section 2.2.2).

Our dealer is not allowed to take up a speculative position in the asset. In other words, we restrict the dealer to end the day with zero inventory:

$$\omega_1 + \omega_2 + \omega_3 = x \quad (2.2)$$

2.2.1 Exogenous price impact - a model of frontrunning

We first consider the dealer's optimal trading strategy in a market where the price impact of trading is exogenously specified. We shall see that in this much simpler version of the model, the optimal trading strategy of the dealer is to frontrun his client, similar to the model in [Osler et al. \(2016\)](#).

We let λ be a constant price impact from trading, ε_n be a white noise sequence, and model the price p_n as

$$p_n = p_{n-1} + \lambda \omega_n + \varepsilon_n$$

We shall find the trading strategy $(\omega_n, n = 1, 2, 3)$ that solves

$$\max_{(\omega) \in \mathbb{R}^3} \mathbb{E} \left[\sum_{n=1}^3 \omega_n p_n - p_2 x \mid x \right]$$

Subject to the constraint

$$\sum_{n=1}^3 \omega_n = x$$

Note that have for simplicity set the trading cost to zero ($c = 0$).

We can find the optimal strategy by observing that

$$\mathbb{E} \left[\sum_{n=1}^3 \omega_n p_n - p_2 x \mid x \right] = \lambda (\omega_1^2 + \omega_2^2 + \omega_1 \omega_2 - 2\omega_2 x - 2\omega_1 x + x^2)$$

Where we have inserted the price impact function for p_n and eliminated ω_3 by using the constraint. The optimizer is then found by taking the first order conditions with respect to ω_1 and ω_2 . This yields the optimal strategy

$$\omega_1 = \frac{2}{3}x, \quad \omega_2 = \frac{2}{3}x, \quad \omega_3 = -\frac{1}{3}x$$

The optimal trading strategy involves the dealer selling *more* than his total client orders in period 1 and 2, and then buying back assets in period 3. Essentially the dealer is using the price impact caused by his client trades to take a short-lived proprietary position that makes an expected profit. This strategy is called *frontrunning*, and is illegal in most markets⁶. The reason is that the strategy goes against the best interest of the client - the expected profit earned by the dealer is reflected in an poorer execution price for the client than what would prevail if the dealer simply traded the entire client order at time 2. In terms of price dynamics it means that the expected price follows an "up-and-down" path, with the price at the benchmark time being the most extreme observation and then rebounding after the benchmark time. [Osler et al. \(2016\)](#) show that this price pattern is even more pronounced when there are several dealers engaging in collusion.

A related case is studied in [Saakvitne \(2016b\)](#). In that model the price impact of trading is taken to be transitory rather than permanent. The resulting trading strategy is termed "banging the close", which is a known manipulation strategy, illegal in many markets.

2.2.2 Endogenous price impact

In the model of Section 2.2.1, price impact of trading was exogenously assumed. In this section we are interested in studying how the trading strategy of the dealer is affected when we introduce a learning market into the model. To that end we extend the classic model of [Kyle \(1985\)](#).

There are four classes of agents in the model: an informed trader, a group of liquidity (noise) traders, a market maker, and a dealer executing benchmark orders. The dealer can be thought of as representing a subset of liquidity traders who wants to transact at

⁶Frontrunning is commonly defined as entering into a trade with advance knowledge of a large transaction that will influence the price of the underlying security to capitalize on the trade. This behavior is explicitly forbidden by most market regulators, see for example FINRA rule 5270, *Front Running of Block Transactions*. More details can be found in [Harris \(2003, ch. 7.5\)](#)

the benchmark price, and shall model this agent as we did in Section 2.2.1.

The informed trader knows the fair value of the asset, as in the standard Kyle model, but he does not know the size of benchmark orders received by the dealer. The dealer, on the other hand, knows the size of his benchmark order, but he does not know the fair value of the asset. The objective of both the informed trader and the dealer is to maximize their respective profits, however the two profit functions look different. The informed trader profits by buying the asset when it is undervalued and selling it when it is overvalued. The dealer profits by buying the asset at a lower price than the benchmark price, and selling it a higher price.

In contrast to the original Kyle model we have in our model a small quadratic cost associated with trading, meaning that to trade a volume of z the agent pays a cost cz^2 , where c is constant. The numerical value of the cost c can be taken to be arbitrarily small if one so wishes, but that it is non-zero serves an important role in the numerical solution of the model. Specifically, it ensures that the fixed-point iteration scheme used to solve the model converges to a unique equilibrium. Without the cost component there exists an infinite number of possible equilibria where the informed trader and the dealer trades arbitrarily large amounts back and forth with each other. Although the trading cost serves an important computational role, its small size means that it is economically insignificant.⁷

The informed trader:

We let θ_n denote the asset sales of the informed trader in trading period n . The informed trader knows the value v of the asset, which we model as drawn from the distribution F_v . The objective of the informed trader is to profit from his information, and his profit is expressed as

$$V(\theta) = \sum_{n=1}^3 (\theta_n(p_n - v) - c\theta_n^2) \tag{2.3}$$

The liquidity traders:

In each trading period n , liquidity (or "noise") traders submit net orders of u_n . We model the triplet $(u_n, n = 1, 2, 3)$ as drawn from the distribution F_u .

⁷Numerical experiments using the baseline parameterization of Section 2.3.1 shows that for the informed trader, the ratio of trading costs to profit has a mean of $2.58 * 10^{-7}$ and a standard deviation of $4.26 * 10^{-6}$. For the dealer this ratio has a mean of $1.87 * 10^{-7}$ and $2.12 * 10^{-5}$.

The market maker:

The market maker functions as in the original Kyle model - he observes total order flow y_n and sets the price p_n . Total order flow is just the sum of orders coming from the dealer, the informed trader and the noise traders:

$$y_n = \omega_n + \theta_n + u_n \quad (2.4)$$

The market maker is operating in a competitive environment, which leads to the efficiency condition that the market price p_n must equal the expected fair value of the asset, conditional on the public information set \mathcal{F}_n^M :

$$p_n = \mathbb{E}[v \mid \mathcal{F}_n^M] \quad (2.5)$$

The publicly available information set \mathcal{F}_n^m is defined in equation (2.6).

To summarize the information structure of the model: the dealer knows the size of benchmark orders (x) at the start of the model, and observes total order flow. The informed trader knows the fair value of the asset, and also observes total order flow. The market maker only observes total order flow. We formalize this structure in the filtrations $(\mathcal{F}_n^d, \mathcal{F}_n^i, \mathcal{F}_n^m)_{n=1}^3$ for the dealer, informed trader and market maker respectively.⁸ We define these filtrations for $n = 1, 2, 3$ by

$$\begin{aligned} \mathcal{F}_n^i &= \sigma(v, y_n, \dots, y_1) \\ \mathcal{F}_n^d &= \sigma(x, y_n, \dots, y_1) \\ \mathcal{F}_n^m &= \sigma(y_n, \dots, y_1) \end{aligned} \quad (2.6)$$

We are interested in characterizing *trading strategies* of the informed trader and the dealer. We say that a trading strategy is admissible if it is a function adapted to that agent's filtration. Intuitively this means the trading strategy respects the agents information set.

An equilibrium in this model is defined as

- An admissible informed strategy (θ_n^*) that maximize $\mathbb{E}[V(\theta) \mid v]$ subject to the pricing rule (p_n) and the dealer's strategy (ω_n^*) .
- An admissible dealer strategy (ω_n^*) that maximize $\mathbb{E}[J(\omega) \mid x]$ subject to the no-inventory constraint (2.2), the pricing rule (p_n) and the informed strategy (θ_n^*) .

⁸The notation $\sigma(z)$ here means the sigma-algebra generated by the random variable z . It should not be confused with the variances of various distribution, which will later be denoted by the constant σ .

- A process (p_n) that satisfies the pricing rule (2.5) under the trading strategies (θ_n^*) and (ω_n^*) .

We solve the model for a **linear equilibrium** with **rational expectations**, meaning that we require each agent's trading strategies to be a linear function of his own state variable and the conditional expectation of the unknown state variable:

$$\begin{aligned} \omega_1 &= \alpha_1 x & \theta_1 &= \beta_1 (v - p_0) \\ \omega_2 &= \alpha_2 x + \alpha_3 \mathbb{E}[v \mid \mathcal{F}_2^d] & \theta_2 &= \beta_2 (v - p_1) + \beta_3 \mathbb{E}[x \mid \mathcal{F}_2^i] \\ \omega_3 &= x(1 - \alpha_2 - \alpha_1) & \theta_3 &= \beta_4 (v - p_2) + \beta_5 \mathbb{E}[x \mid \mathcal{F}_3^i] \end{aligned}$$

That each agent conditions on the state variable unknown to him carries economic significance; the dealer will over the course of the model become informed, since by virtue of knowing his own order flow he can make a more precise estimate of the fair value v than what the market maker can form from public information alone. Similarly, since the informed trader knows his own orders, he can estimate the size of benchmark orders and use that estimate in his own strategy.

We solve the model numerically. Monte Carlo sample averaging is combined with a projection approach to compute the conditional expectations involved, and a fixed point iteration scheme is combined with numerical optimization to find equilibrium coefficients. The details of the solution algorithm is in Appendix 1.⁹

Note that, in contrast to Kyle's original model, we *do not* require the pricing rule to be linear. Nor do we require the pricing rule or the trading strategies to be Markovian, as the agents' expectations are formed based on their entire information set. Our equilibrium definition does however rule out states that encode non-payoff relevant information, which excludes the use of signals, negotiation and cooperation between agents (for example "cheap talk" or contracts).

⁹Some readers may find it uncomfortable that we proceed with numerical analysis without a formal equilibrium existence proof. Judd (1998) argues the case for proceeding: First, if we did have an existence proof, or if we used some other appropriate approximate equilibrium definitions, we could proceed as we do below. Second, if an equilibrium does not exist, we would expect our solution method not to work. Thirdly, any candidate equilibria that we identify with our solution method are ϵ -equilibria for small ϵ , and if there are no pure equilibria then these provide natural alternative solutions where agents make small optimization errors (see e.g. Shoham and Leyton-Brown (2008, Section 3.4.7)) for more details.

2.3 Results

2.3.1 Baseline parameters

The baseline parametrization of the model is chosen to match the assumptions of the classical Kyle-model and later extensions. In particular we take all random variables involved to be independent zero-mean Gaussian, and the sequence of noise trades to be iid. This structure is useful when one wants to derive analytical results, but is not necessarily the most realistic. Table 2.1 summarize the baseline parametrization.

$F_x = \mathcal{N}(0, \sigma_x)$	$\sigma_x = 10$
$F_v = \mathcal{N}(0, \sigma_v)$ i.i.d	$\sigma_v = 1$
$F_u = \mathcal{N}(0, \sigma_u)$	$\sigma_u = 1$
$c = 10^{-4}$	$K = 5000$

Table 2.1: Baseline parameter values

We define *expected absolute order flow* as $\mathbb{E}[|y_n|]$ and *price impact of order flow* λ_n as

$$\lambda_n := \frac{\text{cov}(y_n, \Delta p_n)}{\text{var}(y_n)}$$

Both expected order flow and price impact are by-products of our solution algorithm and can therefore be computed through sample averaging.

	t=1	t=2	t=3
Mean	-0.01	0.00	0.00
St.dev.	0.67	0.25	0.39
Skewness	-0.01	-0.33	-0.12
Kurtosis	2.86	3.62	2.94

Table 2.2: Price change - first four moments (K=5000)

Figure 2.2 shows several of the key results from the model. First, as shown in the top-left panel, expected absolute order flow peaks at the benchmark time ($t = 2$). Price impact is nearly zero at the benchmark time however, as shown in the top-middle panel.

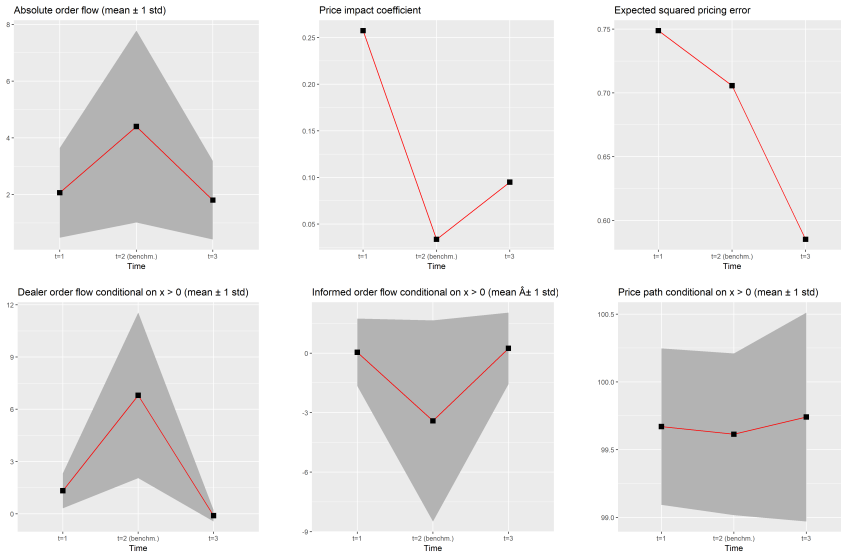


Figure 2.2: Baseline results

	(t_1, t_2)	(t_2, t_3)	(t_1, t_3)
Orderflow correlation	0.67	0.78	0.76

Table 2.3: Orderflow correlation

$\alpha_1 = 0.17$	$\beta_1 = 1.68$
$\alpha_2 = 0.78$	$\beta_2 = 6.04$
$\alpha_3 = 0.01$	$\beta_3 = -0.24$
	$\beta_4 = 3.51$
	$\beta_5 = 0.20$

Table 2.4: Equilibrium coefficients - baseline parameters

The reason is that the market is dominated by the benchmark trades done by the dealer in this period. Since the dealer does not trade for reasons concerning the fair value of the asset, these trades are uninformative, and hence price impact is low.

The market price does converge to the fair value, however. The top-right panel show that the expected squared pricing error decrease each trading round, with by far the largest decrease in the final trading round.

The bottom-left panel provides the explanation for the two previous results; the dealer is executing a large share of his total order in the fixing round, thereby driving up the expected absolute order flow in this trading round. It is also interesting to notice that the dealer does not trade in the opposite direction of his client orders; this distinguish the dealer's execution strategy from the "frontrunning" that takes place when price impact is exogenous (Section 2.2.1 and Osler et al. (2016)).

Unlike the classic Kyle model, here the informed trader trades also for reasons not related to the fundamental value of the asset; at the benchmark time he trades in the opposite direction of the dealer, in order to profit from any non-efficient price impact caused by the dealer's trades. This can be seen in the bottom-middle panel.

The bottom-right panel show the price path conditional on benchmark orders being positive. There is a slight up-and-down movement in the mean price path, reflecting the price disturbance caused by the dealer executing benchmark trades. The benchmark trades are uninformative but still have a price impact. The reason is that market maker cannot separate between informed and uninformed trade volume, he only observes the aggregate order flow. Empirically, Evans (2018) and Osler et al. (2016) document a similar pattern in foreign exchange rate around the 4pm London fix.

Orderflow in this model is autocorrelated (Table 2.3), unlike in the original Kyle model. This autocorrelation stems from the benchmark orders that the dealer has to execute, and allows the informed trader to form an estimate of the size of benchmark orders. The informed trader use this estimate to trade in the opposite direction of the dealer during the benchmark, signified by the negative value of β_3 in Table 2.4. Similarly, the dealer is able to form a more precise estimate of the fair asset value than what the market maker can form from observing public information (orderflow) only. In equilibrium, however the dealer does *not* condition on this information when submitting his trades, as signified by the coefficient α_3 in Table 2.4.

2.3.2 Sensitivity to variances

In this section we study how various key features of the model depends on the variances of the underlying random variables. We shall let these variances range over the values

$$\sigma_x = [6, 8, 10, 12, 14]$$

$$\sigma_v = [0.6, 0.8, 1.0, 1.2, 1.4]$$

$$\sigma_u = [0.6, 0.8, 1.0, 1.2, 1.4]$$

We change one variance at the time, and let all other parameters be as in the benchmark case.

First we study how the profits of both the dealer and the informed trader depend on the parameters of the model. We do that by computing equilibrium expected profits for the range of variances under consideration. Figure 2.3 show expected profits, normalized by the baseline level of profits. The profits of the informed trader, signified by the blue line, increase in the variance of benchmark orders (σ_x), while the profits of the dealer (red line) decrease in this variance.

The profit of both the dealer and the informed trader are increasing the variance of the fair value (σ_v), as shown in the middle panel. The right-hand panel shows that the profit of the dealer is much more sensitive to the level of noise (σ_u) than the profit of the informed trader.

Summing up, a high uncertainty in the fair value and the amount of noise are beneficial for the dealer, while the variation in the direction and size of benchmark orders is detrimental. For the informed trader, returns on information are higher in the variation of both benchmark orders and the fair value, while variation in the level of noise is only very slightly beneficial.

Next, we study how benchmark price volatility and price deviation from the fair value depends on the variances of the underlying random variables x, v, u . We compute equilibrium price standard deviation and the root mean square error between the price and the fair value. Note that the standard deviations are "cross-sectional", in the sense that we compute the standard deviation of p_n across K random samples for each $n = 1, 2, 3$, and similarly for the RMSE of the price measured against the fair value.




Figure 2.3: Profits as function of σ_x (left), σ_v (middle) and σ_u (right). Y-axis show percent of baseline value. Red circles are dealer profit, blue triangles are informed trader's profit.

Figure 2.4 shows the results for price standard deviation. The left-hand panel shows that pre-benchmark price volatility is initially increasing in the variance of benchmark orders, and the volatility across all periods decrease for high levels of variation in benchmark orders. The variance of the fair value, on the other hand, has an equal impact on price volatility across all time periods, as shown in the middle panel of Figure 2.4. Increased noise trading, shown in the right-hand panel, *decrease* the volatility of the pre-benchmark price. An interpretation of this is that when noise goes up, the price impact of trading goes down, and in particular the price impact of the dealer trading at the first trading round goes down. This second-order effect dominates the first-order effect that increased noise has on the price, for all but very low values of noise variance.




Figure 2.4: Price standard deviation as function of σ_x (left), σ_v (middle) and σ_u (right). Y-axis show percent of baseline value. Red circles are the pre-benchmark period ($t=1$), green triangles are the benchmark time ($t=2$), and blue squares are the post-benchmark period ($t=3$).

Figure 2.5 shows how the price deviation from fair value, measured as RMSE between

the two quantities, depends on the underlying variances. First of all, we see that the price deviation decrease for each trading round, reflecting the convergence of the price to the fair value. For the highest level of variance in benchmark orders, the price deviation increase across all time periods. Price deviation as a function of the variance in fair value is U-shaped - both a higher and a lower variance of the fair value than in the baseline increases the price deviation.




Figure 2.5: Price deviation from fair value as function of σ_x (left), σ_v (middle) and σ_u (right). Y-axis show percent of baseline value. Red circles are the pre-benchmark period ($t=1$), green triangles are the benchmark time ($t=2$), and blue squares are the post-benchmark period ($t=3$).

Lastly, we study how price impact depends on the parameters of the model. As in the other experiments, we compute equilibrium price impacts for the range of variances under consideration. The results are shown in Figure 2.6. We see that the relation between price impact and amount of uncertainty is somewhat complicated. Price impact is decreasing in the variance of benchmark orders, which means that the higher the average absolute size of benchmark orders, the lower is price impact. The middle plot shows that price impact at the benchmark is increasing in the variance of the fair value. This is intuitive: higher σ_v entails that the informed trader is on average "more informed", which increase price impact. Price impact is first decreasing and then increasing in the variance of uninformed trades (noise). Intuitively, when there is more noise, the market maker is more uncertain as to whether order flow at the benchmark time is coming from the informed trader or the uninformed dealer. This is an interesting contrast to [Admati and Pfleiderer \(1988\)](#), where price impact is only decreasing in the variance of noise.




Figure 2.6: Price impact as function of σ_x (left), σ_v (middle) and σ_u (right). Y-axis show percent of baseline value. Red circles are the pre-benchmark period ($t=1$), green triangles are the benchmark time ($t=2$), and blue squares are the post-benchmark period ($t=3$).

2.4 Discussion

Price benchmarks has a tendency to work as basins of attraction for liquidity, as pointed out by [Duffie and Dworczak \(2014\)](#). They can also be accompanied by strong incentives for some market participants to transact at the benchmark price ([Financial Stability Board, 2014](#)). Both these economic forces can exhibit self-reinforcing dynamics. The resulting market structure is modeled in this paper, and we show testable implications for both order flow and price impact. The implications have empirical support in recent evidence from FX markets. An open question is whether similar dynamics can be found in other financial markets with important benchmarks, such as money markets, commodity derivatives or interest rate swaps. In equity markets closing prices plays a similar role as the fix does in FX markets.

[Osler et al. \(2016\)](#) show how a model where dealers are frontrunning their customers results in a "convex" price path, both in a competitive situation and in a situation with collusion. The model in this paper has no frontrunning and features a competitive market, yet still produce a convex price path.

We have also shown that without a learning market maker in the model, a manipulative trading strategy would be optimal for the dealer. Put simply, the learning market maker will in equilibrium understand the execution strategy of the dealer and adjust the price impact of trading downwards. Since there is an informed trader present in the market, and the market maker only observes aggregate order flow, the price impact of trading at the benchmark is not exactly zero, which leaves some profit to the dealer, and a certain

amount of distortion of the benchmark price. At a general level, we show how the introduction of an endogenous price impact function significantly affects optimal execution strategies. This last point also carries important implications for the literature on optimal order execution, where exogeneous price impact functions are commonplace.

2.A Numerical solution method

We solve the model numerically via a combination of a sample average approximation, projection approximations of conditional expectations, numerical optimization and fixed point iteration.

Specifically, we first fix a large sample S of realizations of the random variables x, v, u and arbitrary trading strategies (α_n) and (β_n) . Given these, we can compute order flow, and then approximate the conditional expectation of agents by regressing the unknown state variable onto polynomial powers of the variables in each agent's information set.

For example, in period 2 the pricing rule (2.5) is estimated by regressing v on $(y_1, y_2, y_1^2, y_2^2, y_1 y_2, y_1^3, \dots)$.

Next, we use the given trading strategies (α_n, β_n) and regression coefficients to approximate expected profits $\mathbb{E}[V]$ and $\mathbb{E}[J]$ by their sample average.

The last step of the algorithm is a fixed point-iteration: take as given one trading strategy and numerically optimize the other, proceeding back and forth between the strategies. Since the sample S is fixed, these are deterministic optimizations. When and if the coefficients converge, we have reached the assumed equilibrium functions (α_n^*, β_n^*) . The full algorithm is described in Box 1.

Box 1: Numerical Solution Algorithm

1. Specify the distributions F_v, F_u, F_x and draw K realizations of v, x, u_1, u_2, u_3 .
2. Fix some initial values of the strategy coefficients (α_n) and (β_n) for $n = 1, 2, 3$, satisfying $\alpha_1 + \alpha_2 + \alpha_3 = 1$.
3. Approximate $\mathbb{E}[v | \mathcal{F}_n^m]$ by the least squares prediction \hat{v}_n made by regressing our K realizations of v onto a sequence of polynomial powers of $(y_j, j \leq n)$.

4. Similarly, regress x on polynomial powers of $(y_j, j \leq n)$ and v to form $\hat{x}_n \approx \mathbb{E}[x | \mathcal{F}_n^i]$, and regress v on polynomial powers of $(y_j, j \leq n)$ and v to form $\hat{v}_n^d \approx \mathbb{E}[v | \mathcal{F}_n^d]$.
5. For each $k = 1, \dots, K$, use the strategy coefficients (α_n, β_n) to compute profits of the informed trader and the dealer, V_k and J_k .
6. Approximate expected profits by averaging over the K simulations,

$$\bar{J} = K^{-1} \sum_{k=1}^K J_k$$

$$\bar{V} = K^{-1} \sum_{k=1}^K V_k$$

7. Update (α_n) to the optimizer of \bar{J} by repeating step 3-6.
8. Update (β_n) to the optimizer of \bar{V} by repeating step 3-6.
9. Repeat step 7 and 8 until (α_n) and (β_n) converges.

The complete source code of the algorithm is available online¹⁰. In the next section we discuss the properties of the algorithm in more detail. We first offer some general observations.

In a wide range of numerical experiments, we have seen the algorithm outlined above to converge to the same parameters over several runs with different randomization seeds. We have thus not experienced instabilities related to non-uniqueness of equilibria. The trading cost (c) is crucial in this respect, as previously noted. When the algorithm converge, the projection theorem and the law of large numbers ensures that, given mild

¹⁰sites.google.com/view/jo-saakvitne-academic-homepage

regularity conditions and large enough M and K , we have

$$p_n \approx \mathbb{E}[v \mid \mathcal{F}_n^m]$$

$$\bar{J}(\omega_n^*) = \max \bar{J}(\omega_n) \approx \max \mathbb{E}[J(\omega_n) \mid \mathcal{F}_n^d]$$

$$\bar{V}(\theta_n^*) = \max \bar{V}(\theta_n) \approx \max \mathbb{E}[V(\theta_n) \mid \mathcal{F}_n^i]$$

We use third-degree polynomials of $(y_j, j \leq n)$ as basis functions in our approximation of p_n . In an iid Gaussian setting, for example the classical Kyle model, it would in fact be sufficient to use only the first power of y_n (a linear regression). Other types of polynomials can be used, the important thing is that the polynomials forms an orthogonal basis for v in L_2 space. We have verified that our results do not change if we use other common types of orthogonal polynomials. We have used a sample size of five thousand draws per simulation run ($K = 5000$). As convergence criteria in the fixed point iteration we have required the Euclidean norm of the change in strategy coefficients to be below 10^{-2} . In our baseline parametrization of the model we have found that the rate of convergence is quite fast, and that 30 simulation runs (step 9) are often sufficient for a reasonable approximation to the equilibrium values. Our experiments indicate that 200 simulation runs are usually sufficient to achieve a convergence tolerance of $< 10^{-4}$.

An advantage of the solution method used here is that we can free ourselves from the reliance on the assumption of joint normality of the random variables involved. As noted by [O'Hara \(1995\)](#), this assumption is problematic because it often directly affects a model's predictions regarding factors such as optimal trade size and the effects of volume. We employ a baseline parametrization that involves jointly normal random variables, but in our sensitivity analysis we loosen this assumption (Section [2.B.2](#)).

2.A.1 Convergence results for the numerical solution algorithm

The algorithm consists of several steps: solving two stochastic optimizations via sample averaging, finding an equilibrium via a fixed-point iteration, and approximating a conditional expectation via a basis function regression. Providing a general convergence result for the algorithm is difficult, since one needs to consider limits for the number of basis functions, sample size and fixed-point iterations.

One standard result that can be applied to our setting concerns the convergence of the Monte Carlo sample averages to the true expected values. If we constrain the vector of

strategy coefficients $\mathbf{b} = [b_1, \dots, b_j]$ to take values in a compact set $\mathcal{B} \subset \mathbb{R}^j$, we have the following result.

Proposition 1. *Take as given an informed strategy (ω_n) and a pricing rule $p_n(y_1, \dots, y_n)$. Suppose the price rule is jointly continuous. Let \mathcal{B} be a compact subset of \mathbb{R}^j . Consider the stochastic optimization problem*

$$\max_{\mathbf{b} \in \mathcal{B}} \mathbb{E}[J(\mathbf{b}) | x] \quad s.t \quad \sum_{n=1}^3 b_n = 1$$

Suppose the regularity condition $\mathbb{E}[\sup_{\mathbf{b} \in \mathcal{B}} |J_k(\mathbf{b})|] < \infty$ holds. Then, as $K \rightarrow \infty$,

i)

$$K^{-1} \sum_{k=1}^K J_k \rightarrow \mathbb{E}[J(\mathbf{b}) | x] \text{ uniformly, with probability } 1$$

Furthermore, suppose the stochastic problem is uniquely maximized at \mathbf{b}^* , and let $\bar{\mathbf{b}}_K^*$ be the optimizer of the corresponding deterministic problem with sample size K . Then, as $K \rightarrow \infty$,

ii)

$$\bar{\mathbf{b}}_K^* \xrightarrow{P} \mathbf{b}^*$$

Proof. Part i) follows from the Uniform Law of Large Numbers, see [Newey and McFadden \(1994, Lemma 2.4\)](#). Part ii) follows from [Newey and McFadden \(1994, Theorem 2.1\)](#). \square

An equivalent result holds for the informed trading strategy (ω_n) . One could furthermore build on standard asymptotic theory and apply the central limit theorem to derive distributional properties of the Monte Carlo estimators, see [Homem-de Mello and Bayraktan \(2014\)](#).

In the statement of [Proposition 1](#) we assumed that the pricing rule $p_n(\cdot)$ is jointly continuous in order flow (y_n) . In the general case where p_n is the conditional expectation of v , this condition is hard to verify without imposing additional restrictions on the distributions F_x, F_v, F_ε . When p_n is a polynomial function of $(y_j, j \leq n)$ ¹¹ however, as is the case in our algorithm, joint continuity is immediate. Approximating the conditional expectation by truncating the basis function expansion to the first M terms introduce its own error term. However, when the conditional expectation is square integrable it is possible to derive bounds on this truncation error - see [Zanger \(2013, 2018\)](#) for more details.

¹¹For example $p_n = \gamma_1 y_1 + \gamma_2 y_2 + \gamma_3 y_1^2 + \gamma_4 y_2^2 + \gamma_5 y_1 y_2$

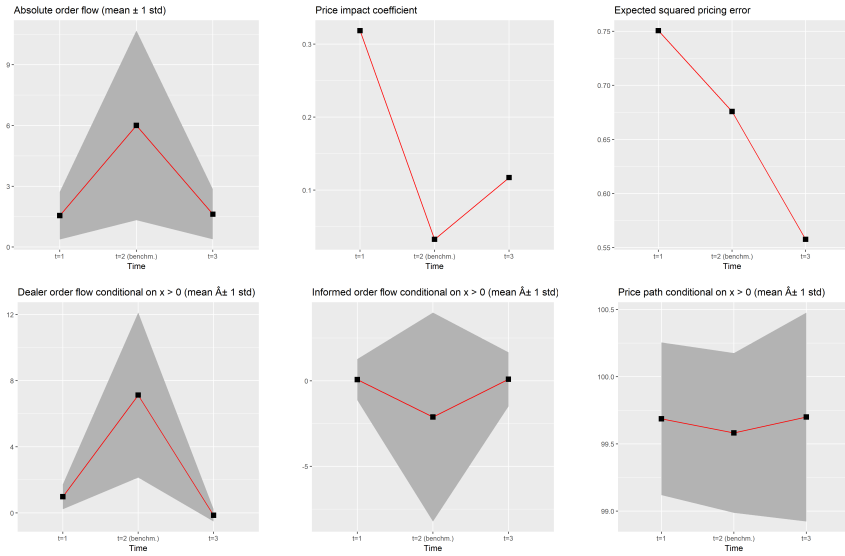


Figure 2.7: Correlation fair value and benchmark orders

2.B Additional model sensitivity checks

In this section we report the output of two additional sensitivity checks - introducing correlation between benchmark orders and the fair value, and allowing for time-varying and autocorrelated liquidity trading.

2.B.1 Correlation between benchmark orders (x) and the fair value (v)

What happens if the size and direction of benchmark orders are correlated with the fair value of the asset? To answer this question, we change the baseline parameters such that benchmark orders x and the fair value v are multivariate normal random variables with correlation coefficient 0.5. All other parameters are as in the baseline scenario.

Figure 2.7 shows that correlation of benchmark orders and the asset's fair value does not change any of the baseline results substantially. The price impact in the trading round after the benchmark is somewhat higher (top-middle panel). The order flow coming from the informed trader now reflects the informed trader tends to trade in the same

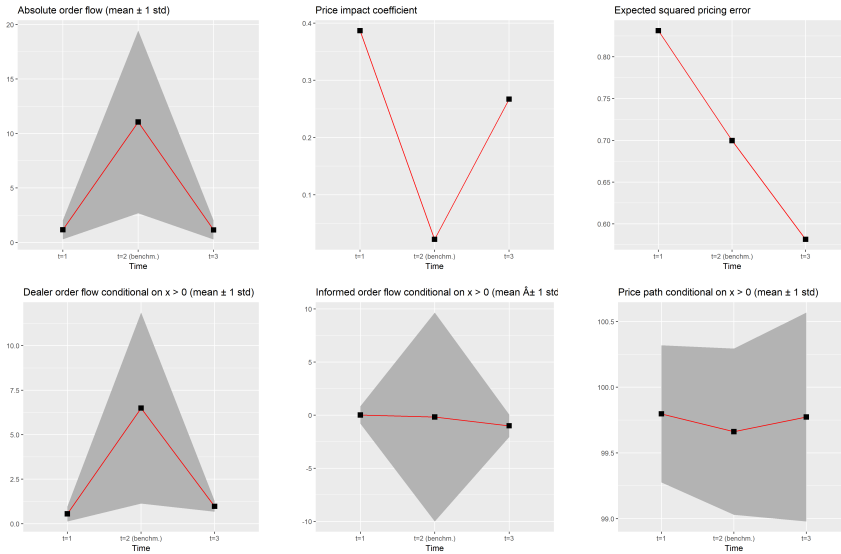


Figure 2.8: Time-varying and autocorrelated noise

direction as the dealer (bottom-middle panel).

2.B.2 Time-varying autocorrelated noise

The classical Kyle model and many later models typically assume that liquidity (“noise”) trading is iid white noise. An advantage of the numerical solution algorithm employed in this paper is that it allows us to study the impact of a departure from this assumption.

For this experiment we let noise trading be described by time-varying variance and time-series correlation. Specifically, we let the noise trades u be three-dimensional jointly normal with correlation matrix Σ_u where

$$\Sigma_u = \begin{bmatrix} 1 & 0.3 & 0 \\ 0.3 & 5 & 0.3 \\ 0 & 0.3 & 1 \end{bmatrix}$$

Figure 2.8 shows that qualitatively the main results of the model are unchanged under this more complicated stochastic structure.

	t=1	t=2	t=3
Mean	-0.01	-0.00	0.00
St.dev.	0.56	0.45	0.39
Skewness	-0.03	-0.11	0.09
Kurtosis	3.18	2.91	2.88

Table 2.5: Price change - first four moments (K=5000) - autocorrelated noise

	(t_1, t_2)	(t_2, t_3)	(t_1, t_3)
Orderflow correlation	0.73	-0.42	-0.39

Table 2.6: Orderflow correlation - autocorrelated noise

$\alpha_1 = 0.07$	$\beta_1 = 0.81$
$\alpha_2 = 0.88$	$\beta_2 = 13.83$
$\alpha_3 = 0.01$	$\beta_3 = 0.96$
	$\beta_4 = 0.16$
	$\beta_5 = -0.15$

Table 2.7: Equilibrium coefficients - autocorrelated noise

Chapter 3

A note on optimal trade times for financial quotes with a last look

GIULIA DI NUNNO, UNIVERSITY OF OSLO

JO SAAKVITNE, BI NORWEGIAN BUSINESS SCHOOL

Abstract

We model the option value embedded in "last look"-quotes, the structure of which is similar to an American barrier option. To account for the realities of high-frequency financial prices we introduce the time-changed discrete Levy process as a model for price dynamics, and show how the optimal stopping problem associated to the quote can be solved via Least Squares Monte Carlo. For various special cases we provide explicit formulae. We also solve the optimal stopping problem for the cases where the price process follow a Brownian motion, and a Skellam process.

3.1 Introduction

Our goal is to study the risk involved when a trader makes a binding fixed-price offer to buy or sell a stock or other security in the financial market. Such an offer is called a quote. A trader who provides a quote assumes a risk, since he takes a loss if the value of the quoted security has a detrimental development before the quote expires. This risk is often referred to as the risk of "being picked off". To protect themselves against this risk, traders sometimes provides quotes with a feature known as "last look". The essence of the last look is that the quote is automatically cancelled if the price of the security cross a certain threshold, even though the quote hasn't expired yet. Quotes with last look or similar features are used both on bilateral trading venues and on limit order book markets. In this paper we are chiefly concerned with the bilateral setting.

We approach last look-quotes by noting an analogy from option pricing. The payoff from a last look-quote has the same form as the payoff from an American barrier option, and by analogy we propose to put a monetary measure on the risk of being picked off by evaluating the largest possible expected payoff under a risk-neutral measure. In finding this expectation we also characterize the optimal stopping time associated with the quote. The time scales involved in quotes are very different from that of traditional option pricing however, which leads us to revisit common assumptions about financial price modelling, and to propose a price dynamic that builds on a time-changed discrete-valued Levy process.

The main elements of novelty in our paper are the following. First, by taking up and extending the idea of an analogy between quotes and options, we show how the special last look-feature of a quote can be modelled as an exotic option. Second, we propose a new model for the dynamics of high-frequency financial prices, built on a time-changed discrete Levy process. Third, for certain special cases of the price dynamics we provide analytical solutions to the optimal stopping problem associated with last look-quotes. Fourth, for the general form price dynamics we provide a convergence result on the solution to a discretized form of the model. We illustrate the discretization approach via a Least Square Monte Carlo simulation exercise.

The results in this paper can be useful not only for practical risk-management purposes, but also for regulators and others engaging in market design, as the picking off-risk inherent in posting quotes is closely tied to the cost of providing liquidity, and hence to transaction costs and overall functioning of market venues.

The idea that a financial quote can be regarded as an option was first introduced by Copeland and Galai (1983), who compared observed bid-ask spreads with the predictions of the Black-Scholes formula for European options. A more detailed comparison of the similarities and differences between options and quotes was done by Frino et al. (2006), who performed an empirical investigation into the option value of the limit order book. Lehmann (2006) built on the analogy between options and quotes to study limit orders and their relation to Arrow-Debreu state prices. In contrast to these earlier papers we shall explicitly model the early exercise (American) feature of quotes. Moreover we show how the special last look-feature of certain quotes maps into the structure of an American barrier option. Last look-quotes are also studied by Cartea et al. (2015) and Oomen (2016), although in models with different emphasis than ours.

The fact that quotes involve an option-like structure is easily seen through an example: say a trader makes an offer to buy a specific share for 100 dollars, and the offer is binding for the next 5 seconds. If the market value of the share drops to 99 dollars after 3 seconds, the counterpart can profit from accepting the offer. In order to evaluate the risk taken on by someone who makes a binding fixed-price offer, it is necessary to study the decision making problem of the trader receiving the offer: at what time should he or she trade on the quote? The execution time that maximizes the expected payoff is the *optimal trade time*. The downside risk that the quote provider takes on in this example is often called the risk of being *picked off*, and we quantify this as the value of the option that the receiver of the quote can claim.

In our model, a quote is said to have a last look-feature if it is automatically cancelled if the price moves past a specified barrier. The last look-feature puts a cap on the potential loss for the supplier of the quote, thereby limiting its risk. Last look-quotes are common in foreign exchange markets, where there has also been a controversy concerning the design and use of last look-features. This makes a mathematical model for last look-quotes of particular practical and regulatory interest ¹.

Even though our model for the risk evaluation of financial quotes has many similarities to option pricing models, the time scale involved when studying quotes is typically seconds or milliseconds, which is orders of magnitude shorter than that of a traditional options. For this reason we need to revisit the standard assumptions of financial price

¹See for example "Currency Trading's Last Look' Rules Are Changing, BOE Says", Bloomberg news, 28th July 2016.

modelling. In particular we discuss the role of the discounting numéraire, and in view of the extremely short time scales involved we disregard it in the modelling. Moreover, real-world prices are only allowed to take discrete values, a fact that also becomes visually quite apparent when one studies prices over very short time intervals (see also e.g. [Harris \(1994\)](#) , [Angel \(1997\)](#) and [Werner et al. \(2015\)](#)).

In this work we consider three types of underlying price models: the Brownian motion, the Skellam process, and a time-changed discrete Levy process. The Brownian motion is of interest since it is a traditional benchmark model, popular for practical applications. The Skellam process however, is better suited to describe the underlying security due to its discrete values, see [Barndorff-Nielsen et al. \(2012\)](#). The model that has the best potential of the three to capture all the stylized facts of high-frequency financial prices is the one based on a time-changed discrete Levy process, which we to the best of our knowledge are the first to introduce in this context.

In fact, financial prices also exhibit a high-frequency clustering effect, which we account for by introducing a random time change in the construction of the price process. The main idea of change-of-time methods is to find a simple representation for a stochastic process with a complicated structure, using some known process and a change-of-time process. A review of time change-methods in stochastic calculus can be found in [Barndorff-Nielsen and Shiryaev \(2015\)](#). For applications to finance see [Swishchuk \(2016\)](#). In this paper we use an integrated Cox-Ingersoll-Ross process as our time change, an approach introduced by [Carr et al. \(2003\)](#). The paper of [Gong and Zhuang \(2017\)](#) is related to ours in that they study the pricing of traditional American options under a time-changed price process. The optimal stopping time of the traditional American option differs from ours because the last look-problem studied in the present work involves also a knock-out boundary.

The structure of this paper is the following. In the background section we discuss the economic function of last look-quotes and the institutional details surrounding this mechanism, and review four stylized facts of high-frequency financial prices. In the next section we proceed with formulating quotes and the associated optimal stopping problem mathematically, and we introduce the time-changed discrete Levy process. We then turn to solving the model. We provide an analytical solution to the optimal stopping problem for a special case of our price process, the time-changed Skellam. The general case is solved numerically, and we provide a convergence result for the discrete approximating solution. In the final section we study last look quotes under two alternative price processes, the Brownian motion and the (ordinary) Skellam process. In both cases

we provide an explicit solution of the optimal stopping problem and derive formulae for the value of the quote, the distribution of the optimal stopping time and the probability that the last look comes into effect.

3.1.1 Background

Last look quotes

The last look feature in a quote is associated with a threshold level of prices, so that when the price of the security being quoted crosses the boundary, the quote becomes invalid. A formal definition is given later on. Last look quotes are perhaps best known from foreign exchange markets, where their use has been the subject of a recent controversy, but the feature is also used in fixed income trading (for example on the MTS BondVision bond trading platform). Also in equity markets there are order types with features similar to last look (for example the "Discretionary peg" order offered on the IEX stock exchange). Last look-features are used in both limit order books and in bilateral trading systems.

The model in the present paper is applicable to quotes in a bilateral trading setting more than a limit order book. The reason is that in a limit order book there will be many market participants that can trade on a given quote. Competition between participants will drive the execution time of the quote to the first time it is profitable to trade (in our model this would be when the fair value crosses the execution price of the quote). The optimal stopping problem is trivial in this case, although several other interesting problems arise - see for example [Goettler et al. \(2005\)](#) and [Roşu \(2009\)](#) for equilibrium models of the limit order book, and [Foucault et al. \(2016\)](#) and [Budish et al. \(2015\)](#) for issues related to latency arbitrage.

In a bilateral setting, on the other hand, the optimal stopping problem associated with the last look is not trivial. And if one wants to value the option embedded in the quote, one needs to tackle this problem. Such a valuation is economically important because it determines trading costs - in conditions where the option value is high, liquidity providers have to cover in the option value through higher spreads. Participants with the technology and strategy to take advantage of stale quotes can possibly regain the higher spreads through executing quotes with the optimal policy described here. Participants without the technology to do so have to bear the higher trading cost. Deeper analysis along these lines are performed by both [Cartea et al. \(2015\)](#) and [Oomen \(2016\)](#).

Bilateral provision of quotes can take place as streaming to customers, common in FX, or in response to a customer's request for quote, common in bond trading. In either case the last look-feature is used to reduce the downside risk taken by the liquidity provider. In particular liquidity providers want to reduce the losses incurred on them by faster traders engaging in cross-platform arbitrage. Unfortunately, to the best of our knowledge there are currently no empirical studies on last look-quotes at the present. Conversations with traders and analysts in FX markets suggests that the frequency of last look being activated is highly dependent on the type of customer receiving the quote - one participant stated that high frequency traders and hedge funds commonly saw their executions stopped by last look in 10-20 % of all trades, and for certain participants it could be as high as 50 %. For other types of customer the last look feature is hardly ever exercised. Publicly available information on the use of last look quotes by the major FX dealing banks shows that the parameters of the last look feature are often set individually for each customer.

Given that last look protects liquidity suppliers against the risk of being picked off and hence reduce their risk of providing liquidity, competitive forces will normally ensure that the cost reduction of suppliers is reflected in lower trading costs for customers. Put differently, we would expect "slow traders", such as institutional investors, to benefit from a trading mechanism that limits high frequency arbitrage. Whether this holds in practice is an empirical question that has yet to be answered. However, liquidity providers have also been accused of misusing the feature, to the detriment of their customers. One major dealing bank went beyond employing last look as a purely "defensive measure", and instead used it as "general filter to reject unprofitable orders", according to the New York Department of Financial Services, who fined the bank in question \$ 150 million in 2015². In the wake of this case, regulators and market participants agreed on new best practices regarding the use of last look, stating that "Market Participants employing last look should be transparent regarding its use and provide appropriate disclosures to Clients" ([Global Foreign Exchange Committee, 2017](#), Principle 17).

Last look is particular relevant in connection with minimum resting times on quotes. A minimum resting time is a requirement laid down by the trading platform or regulators saying that a quote cannot be cancelled until a minimum amount of time has passed. The attention surrounding minimum resting times has increased, as regulators

²[New York Department of Financial Services press release November 15th 2015](#)

and market participants have become increasingly concerned about algorithmic trading causing “phantom liquidity”, a phenomenon in which quotes are submitted and subsequently cancelled within a very short time frame (Blocher et al., 2016). There is an ongoing debate on whether this is detrimental to the quality of markets, see for example Hendershott et al. (2011); Budish et al. (2015); Foucault et al. (2016). Several regulatory responses have been proposed, among them a rule requiring all quotes to have a minimum resting time (Jorgensen et al., 2016). A question of interest to regulators is therefore how the introduction of a minimum resting time will actually affect trading costs (Brogaard, 2011; Furse et al., 2011). It seems intuitive that a minimum resting time entails a larger risk of being picked off, because the supplier of the quote can not adjust the quoted prices in reaction to new information until the minimum resting time has expired. The model presented in the present work allows one to quantify how the risk of being picked off depends on the minimum life length of the quote (T).

Connection to option pricing theory

A trader that has provided a buy quote to another trader has the obligation, but not the right, to sell at the specified price at any time until the quote expires, even if the fair value of the asset moves to the suppliers disfavour. In this sense the provider of the quote has taken up a risk. Then we see that the position of someone who has supplied a quote is in fact very much like that of someone who is short an (American) put option, an observation first made by Copeland and Galai (1983)³. The quoted price takes the role as the strike price of the option, and the validity time of the quote is the expiration time of the option. The receiver of the quote is in a position as if he were long an American put option. Clearly, if the quote is an offer to sell, the situation would be reversed and the provider of the quote is short an American call option. The risk taken up by the provider of the quote is effective if the quote is executed against. The *monetary value* of this risk corresponds to the value, or “fair price”, of the corresponding option.

In the framework of option pricing introduced by Black and Scholes (1973), the price of a traded option is found using hedging and the principle of no arbitrage. The price is unique, since the options (and all financial claims) are perfectly replicable. The fundamental building block of this theory is the so-called replicating portfolio: a self-financing portfolio of assets that exactly replicates the payoff from the option. The mar-

³Copeland and Galai did not account for the possibility of early exercise however, but instead treated the quoted as an European option

ket model is called complete. Prices of financial claims determined by the no-arbitrage principle are called "fair". The no arbitrage pricing principle provides an interval for fair prices also when the market model is incomplete, that is when not all financial claims are perfectly replicable.

However, in the present framework of quotes, it is not clear how we should apply the idea of no arbitrage pricing and replicating portfolios in the current setting: this portfolio would normally involve a position in the underlying asset, but the process of acquiring such a position is exactly what we are modelling in the first place.

Non arbitrage prices can also be found using a so-called risk-neutral evaluation. That is, the fair price is given by the expected discounted payoff under a risk-neutral probability measure. If the market is complete, the measure is unique and it provides the unique non-arbitrage price for any traded payoff, as the expectation of the discounted payoff under the risk neutral probability measure (Jeanblanc et al., 2009). If the market model is not complete, the interval of arbitrage prices corresponds to all the risk-neutral measures of the incomplete model.

In the present work, we apply the *idea* of risk neutral pricing of the option embedded in the quote to evaluate the risk of posting firm quotes. The expectation of the discounted payoff of the quote, under a given martingale measure, represents a monetary measure of the risk involved in supplying a firm quote.

To see why the expectation under a risk neutral measure can be a useful benchmark, imagine a hypothetical complete market where one is able to continuously trade in the quoted asset. In this market the quote can be replicated by an American option, and the initial value of the hedging portfolio of the American option can be found via its risk-neutral expectation.

No discounting

We have argued that financial quotes are in some respects very much like American options. Unlike common option pricing models however, we will disregard discounting. There are two reasons for this.

First, the time scales involved in modeling financial quotes are very different from those of traditional option pricing models. The expiration time of a quote in modern mar-

kets will typically be measured in seconds or milliseconds, while American stock options have expiry dates measured in months or years. Cash flows occurring months or years into the future are significantly affected by discounting, and we therefore cannot abstract it away without fundamentally changing the structure of decision making problems. For cash flows that are seconds or milliseconds into the future things are different. The empirical literature on intraday interest rates finds evidence of rates in the range of 0.1 to 0.9 basis points per hour (Furfine, 2001; Kraenzlin and Nellen, 2010; Jurgilas and Žikeš, 2014). These numbers implies per-second interest rates in the range of 2.8×10^{-9} to 2.5×10^{-7} per cent. On the other hand, a volatility coefficient of 20% per year, commonly found for stocks indices, translates into a per-second volatility of $\frac{.2}{(252*8*60*60)^{1/2}} \approx 7 \times 10^{-5}$, several orders of magnitude larger than any realistic discount rate. We see that discounting is simply not of any real significance over very short time horizons.

Secondly, it is not clear that per-second or per-millisecond discount rates are even meaningful economic concepts. The hourly discount rates mentioned earlier are thought to be due to structural properties of payment systems. There is as far as we know no evidence that these structural properties carry over to much shorter time intervals.

3.1.2 Stylized facts of high-frequency prices

The randomness in our model will stem from the fair price of an asset. We will model this fair price as an integer-valued Levy process subject to a random time change. In this section we review the arguments behind this choice, in the form of four stylized facts about financial prices relevant for time horizons of seconds or less.

Prices take values on a discrete set

Financial prices take values on a discrete set of points called *ticks*, see Harris (1994); Angel (1997); Werner et al. (2015) for more details on the tick grid.

In mathematical finance it is common practice to model prices as random processes taking values on the real line. In some way the continuous line is an "approximation" of a very dense price grid. This is fine when the model deals with horizons of days, months and weeks. However it has different impact when the model is to be applied for time spans of seconds or milliseconds.

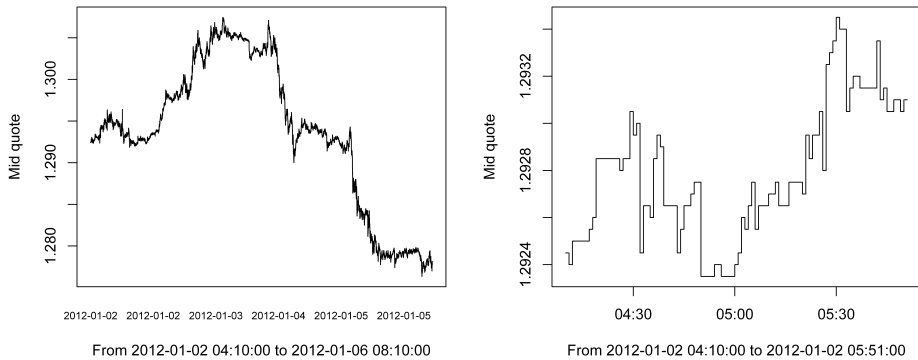


Figure 3.1: EURUSD exchange rate over different time intervals. The data is aggregated into one-minute intervals. Source: Thomson Reuters.

Figure 3.1 illustrates how the granularity of the price grid becomes even visually apparent at short time intervals.

Prices change value in continuous time

Most trading systems treat time as continuous, in the sense that orders are processed continuously on a first come–first served basis and recorded in chronological order (Budish et al., 2015). An accurate model for financial prices should therefore have a time index which takes values on the real line.

Note that trading systems give recorded trades a digital time-stamp. This means that although trading takes place in continuous time, transaction *data* have discrete-valued time indices. For the applications in this paper it is the actual events that are important, not how they are stored in digital systems.

Prices sometimes change by more than one tick

Price changes occur when new quotes enter the market or old quotes are cancelled or traded against. A price change of more than one tick is commonly seen in practice. The frequency distribution of the increment size is however rapidly decreasing - Barndorff-Nielsen et al. (2012) analyse a data set consisting of equity futures, Treasury bonds, commodity futures and money market futures, and show that price increments of more than five ticks are very rare in these securities.

Price changes are clustered in time

Price changes are clustered in time for different reasons and in different manners.

First, prices change more frequently at certain times of the day. The frequency of price changes is usually lower during night time than day time, and it is higher around the fixed release times for important information such as company reports and macroeconomic indicators. If the model for prices refers to horizons of a day or more, this form of clustering is important. However, it is not important when the time horizon is seconds or less; either the model is dealing with a daytime-second or a night time-second, and it is either dealing with a second when important information is scheduled to be released or it is not. We will therefore disregard any cyclical time-of-day clustering effect during the short time intervals we are considering, although the model could easily be extended to account for such effects (see [Borovkova and Schmeck \(2017\)](#) for an example of this in the context of electricity markets).

There is however a second clustering effect that is present also in very short time horizons: we are more likely to see price change in the next few seconds if another price change has just occurred. A variety of modeling approaches has been suggested for dealing with this effect, see [Bauwens and Hautsch \(2009\)](#) for more details. The economic mechanisms behind this clustering effect is still being debated, possible explanations include heterogeneous information arrivals ([Andersen and Bollerslev, 1997](#)), investor learning ([Banerjee and Green, 2015](#)) and behavioral models ([Cont, 2007](#)). We shall account for this second clustering effect by introducing a random time change in our asset price model. To the best of our knowledge, it is the first time that such an approach is taken when studying quotes.

To illustrate the time clustering of price changes we consider intraday data on the exchange rate between US Dollars and Euro (EURUSD) on June 4th 2018. The data is sourced from Bloomberg, and consists of 1440 observations, one for each minute of the day. Associated with each observation is the mid price at the start of the minute, and the number of "ticks" in the one-minute interval. A "tick" corresponds to a change in the bid price, ask price, last trade price or volume. There is an average of 780.6 ticks in each interval, with a minimum of 422 and a maximum of 2564. The tick grid for the EURUSD exchange rate is in 0.0001 Euros. [Figure 3.2](#) show the empirical distribution of price changes over these one minute-intervals, meaning a count of how often the difference in open prices between two consecutive minutes take given values. In around 40 per cent of the intervals there is a price change of one tick up or down. Price changes of




Figure 3.2: Histogram of EURUSD price changes over one-minute time intervals on June 4th 2018, in pips. Source: Bloomberg.

two ticks or more occur in around 20 percent of the intervals, with the maximum price change observed this day being 7 ticks.

Figure 3.3 shows the estimated autocorrelation function of price changes and squared price changes in the intraday EURUSD data. As expected there is little evidence of autocorrelation in the price changes, but significant evidence of autocorrelation in the squared price changes. The latter is a well-known symptom of the high-frequency clustering effect - large price changes tend to be clustered in time.

3.2 Model

3.2.1 The last look-quote

We assume a filtered probability space $(\Omega, \mathcal{F}, P, \mathbb{F})$ satisfying the usual conditions, a fixed time horizon $T < \infty$, and an adapted random process (S_t) called the **price process**, taking values in the measurable space (E, \mathcal{E}) , where $E \subseteq \mathbb{Z}$. The price process repre-




Figure 3.3: Autocorrelation of EURUSD price changes (left) and squared changes (right) over one-minute time intervals on June 4th 2018. Source: Bloomberg.

sents the “fair value” of the asset being quoted⁴. We shall assume throughout that (S_t) has the martingale property - we take a closer look at what this assumption entails when we discuss our price process in the next section.

The last look quote is associated with a boundary $B \in E$ such that the quote becomes invalid (“knocked out”) if the fair price ever crosses B .

Definition 1. A quote with last look is represented by a triplet $(K, T, B) \in E \times \mathbb{R}_+ \times E$

The quote is a *sell quote* if $S_0 < K < B$.

The quote is a *buy quote* if $B < K < S_0$.

We shall use the notation (\bar{K}, T, \bar{B}) for a sell quote and $(\underline{K}, T, \underline{B})$ for a buy quote.

⁴Exactly what should be understood by the word “fair” depends on the particular context. In the stock market-example discussed earlier, one could take “fair value” to mean the net present value of expected future dividend payments discounted at the relevant risk-adjusted discount rate. In general, (S_t) represents the current value of the asset, against which all quotes are compared. We assume that (S_t) is common knowledge to all participants in the market.

For a given constant $B \in E$ and a stochastic process (S_t) , we let T_B denote the **first hitting time**,

$$T_B := \inf\{t \geq 0 \mid S_t = B\}$$

Note that our notion of hitting time is the time when the process (S_t) *takes the value* B . Many books on stochastic calculus use a similar notation to denote the first time the process is greater than or equal to B , a time we will here refer to as T_B^+ :

$$T_B^+ := \begin{cases} \inf\{t \geq 0 \mid S_t \geq B\} & , \text{ if } B > S_0 \\ \inf\{t \geq 0 \mid S_t \leq B\} & , \text{ if } B < S_0 \end{cases}$$

The two random times T_B and T_B^+ are equal almost surely if the process (S_t) has continuous paths. If the process has jumps, however, these two times can differ. The distinction between T_B and T_B^+ is in fact crucial for many of the problems and arguments in this paper.

For a given a stochastic process (S_t) and constant $B \in E$ we write the **knockout time** T_{\dagger} as

$$T_{\dagger}(B) := \begin{cases} \inf\{t \geq 0 \mid S_t > B\} & , \text{ if } B > S_0 \\ \inf\{t \geq 0 \mid S_t < B\} & , \text{ if } B < S_0 \end{cases}$$

We write the **running maximum** (M_t) as

$$M_t := \sup_{u \in [0, t]} S_u$$

We write the **running minimum** (m_t) as

$$m_t := \inf_{u \in [0, t]} S_u$$

We say that a quote is *executed* when, and if, the receiver of a quote decides to trade. A sell quote executed at time $\tau \in [0, T]$ gives payoff $(S_\tau - \bar{K})$ to receiver of the quote. Similarly, the supplier loses $-(S_\tau - \bar{K})$. Following classical assumptions on rationality of traders, it is natural to assume that the receiver executes the quote only if he does not lose relative to the fair value. So we write the payoff of the receiver as $(S_\tau - \bar{K})_+$, and of the supplier as $-(S_\tau - \bar{K})_+$.

Definition 2. The *payoff* associated to a quote executed at stopping time τ is denoted $\psi(S_\tau)$.

If the quote is a sell quote, we have $\psi_{sell} = (S_\tau - \bar{K})_+$

If the quote is a buy quote, we have $\psi_{buy} = (\underline{K} - S_\tau)_+$

For both buys and sells, ψ is a non-negative convex function. This fact will be used repeatedly. Hereafter we introduce the last look feature in a quote.

Definition 3. The *payoff* to the receiver from a **last look sell quote** executed at stopping time τ is

$$\mathbf{I}_{\{M_\tau \leq B\}} \psi_{sell}(S_\tau)$$

The *payoff* to the receiver from a **last look buy quote** executed at time τ is

$$\mathbf{I}_{\{m_\tau \geq B\}} \psi_{buy}(S_\tau)$$

We note that the payoff functions from last look-quotes are the same as American barrier options; up-and-out calls in the case of sell quotes, down-and-out puts in the case of buy quotes.

Definition 4. We define the risk from supplying a last look-quote by

$$\begin{aligned} V_{sell}(\bar{K}, T, \bar{B}) &= \sup_{\tau \in \mathcal{T}} \mathbb{E} \left[\mathbf{I}_{\{M_\tau \leq \bar{B}\}} \psi_{sell}(S_\tau) \right] \\ V_{buy}(\underline{K}, T, \underline{B}) &= \sup_{\tau \in \mathcal{T}} \mathbb{E} \left[\mathbf{I}_{\{m_\tau \geq \underline{B}\}} \psi_{buy}(S_\tau) \right] \end{aligned} \tag{3.1}$$

Where \mathcal{T} is the set of all stopping times taking values in $[0, T]$ almost surely.

Any stopping time that attains the supremum in (3.1) is called an **optimal stopping time**. An optimal stopping time may not be unique. In such cases we will be particularly interested in the first optimal stopping time.⁵


⁵Let $\tau^* \in \mathcal{T}$ be an optimal stopping time. We say τ^* is the **first optimal stopping time** if, for any optimal stopping time $\sigma^* \in \mathcal{T}$, we have $P(\tau^* \leq \sigma^*) = 1$.

To keep the notation simple we shall mostly describe sell quotes in the remainder of this paper. Moreover, we shall omit the bar in \bar{K} and \bar{B} , and simply refer to sell quotes as (K, T) and (K, T, B) , keeping in mind that these quotes satisfy

$$S_0 < K < B$$

The problem we solve is thus to characterize the risk V associated with last look quotes, and the optimal trading time τ^* . Figure 3.4 illustrates.

Figure 3.4: Illustration of the problem



3.2.2 The price process

In this section we construct a price process that behaves in accordance with the stylized facts laid out in Section 3.1.2: a discrete-valued continuous-time process exhibiting changes that are clustered in time. We achieve this by starting out from a general integer-valued Levy process, which we then subject to a random time change. To achieve clustering we shall use the Cox-Ingersoll-Ross process as our time change. This process has become a standard model for this purpose, and has the additional benefit of producing a tractable final model. It is the mean-reverting property of the CIR process that introduce time-clustering of price changes. To our knowledge a similar effect has not been demonstrated using Levy subordinators, the main alternative approach in time-change methods.

Let (N_t^+) and (N_t^-) be two independent \mathbb{F} -adapted Poisson processes with intensities λ^+ and λ^- . Let $(u_n)_{n \in \mathbb{N}}$ and $(d_n)_{n \in \mathbb{N}}$ be two sequences of iid random variables taking values in \mathbb{N} , with $(u_n), (d_n), (N_t^+)$ and (N_t^-) all independent from each other. Let $(T_n, n = 1, 2, \dots)$ be the arrival times for (N_t^+) and (N_t^-) . We also require that u_n and d_n are \mathcal{F}_{T_n} -measurable for all n , and assume that $\mathbb{E}[u_1] < \infty$ and $\mathbb{E}[d_1] < \infty$.

We define the **up** and **down** processes as

$$U_t := \sum_{n=1}^{N_t^+} u_n \quad (3.2)$$

$$D_t := \sum_{n=1}^{N_t^-} d_n \quad (3.3)$$

The representation of an integer-valued Levy process as a compound Poisson is completely general: The Levy-Ito decomposition makes it clear that any integer-valued Levy process can be written as a Poisson integral (Applebaum, 2009, Theorem 2.4.16), since the drift, the Brownian integral and the compensated Poisson integral are necessarily zero for the process to take on only integer values. Since any Poisson integral can be written as a compound Poisson process (Applebaum, 2009, Theorem 2.3.9), it can (by independence) also be written as a sum of an "up" and "down" process as we have done here.

We write the **integer-valued Levy process** (L_t) as

$$L_t := L_0 + U_t - D_t \quad (3.4)$$

Barndorff-Nielsen et al. (2012) refer to this class of processes as "type P processes", and show how they relate to a subordinated simple Poisson process.

Proposition 2. *The process (L_t) is an \mathbb{F} -martingale if and only if*

$$\frac{\lambda^+}{\lambda^-} = \frac{\mathbb{E}[d_1]}{\mathbb{E}[u_1]}$$

Throughout the paper we assume that the process (L_t) is a martingale.

The characteristic function of L_t is given by

$$\mathbb{E}[e^{iuL_t}] = \exp[t\gamma_L(u)]$$

Where

$$\gamma_L(u) = \lambda \sum_{j=-\infty}^{\infty} (e^{iu_j} - 1)\mu(j)$$

and μ is the Levy measure associated with (L_t) .

Let $(\theta, \eta, \sigma, y_0)$ be given constants, and let (Y_t) be the integrated CIR process,

$$dy_t = \theta(\eta - y_t)dt + \sigma\sqrt{y_t}dW_t$$

$$Y_t := \int_0^t y_s ds$$

The process Y_t will be our random trading clock, as in Carr et al. (2003). The characteristic function for Y_t is given by

$$\begin{aligned} \mathbb{E}[e^{iuY_t}] &= \phi(u, t, y_0; \kappa, \eta, \lambda) \\ &= A(t, u)e^{B(t, u)y_0} \end{aligned}$$

Where

$$\begin{aligned} A(t, u) &= \frac{\exp(\frac{\kappa^2 \eta t}{\lambda^2})}{\left(\cosh(\frac{\gamma t}{2}) + \frac{\kappa}{\gamma} \sinh(\frac{\gamma t}{2})\right)^{2\kappa\eta/\lambda^2}} \\ B(t, u) &= \frac{2iu}{\kappa + \gamma \coth(\frac{\gamma t}{2})} \\ \gamma &= \sqrt{\kappa^2 - 2\lambda^2 iu} \end{aligned}$$

We construct the price process, the **time-changed discrete Levy** (TCDL) as

$$S_t := L_{Y_t} \tag{3.5}$$

The process (S_t) does in a manner inherit the martingale property from (L_t) , as (S_t) has the martingale property with respect to the filtration generated by the time changed process and the time change (Tankov and Voltchkova, 2009, Lemma 15.2).


The characteristic function of the TCDL is given by (see Carr et al. (2003))

$$\mathbb{E}[e^{iuS_t}] = \phi(-i\gamma_L(u), t, y_0, \theta, \eta, \sigma)$$

The TCDL process differs from the class of integer-valued Levy processes defined by [Barndorff-Nielsen et al. \(2012\)](#). They show how one can construct an integer-valued Levy process through subordination of a simple Poisson process, and show that the subordinated process always can be written as a compound Poisson. Subordination, which is also a form of time change-technique, ensures that the process remains within the class of Levy processes. Here, we start from a compound Poisson, which we then subject to a time-change through composition with the integrated CIR process, with the aim of generating volatility-clustering. The random time change also means that we go outside the Levy setting.

Figure 3.5 shows a sample path of the TCDL process, along with the time-change process (y_t). A simulation algorithm is provided in the appendix.

Figure 3.5: A realization of the time-changed discrete Levy (upper panel), and the corresponding rate of time-change (lower panel).



3.3 Solution to the optimal stopping-problem

Below we solve two special cases analytically. For the general case we resort to numerical solutions.

3.3.1 Time-changed Skellam Process

When the price process only moves in increments of ± 1 , it becomes a *time-changed Skellam process*.

Proposition 3. *Assume the size of up- and downticks is constant one, $u_i = d_i = 1$ for all i . Then, the optimal stopping time τ^* is given by*

$$\tau^* = T \wedge T_B$$

Proposition 3 nests the cases $T = \infty$ and $B = \infty$; in the former case $\tau^* = T_B$ while in the latter case $\tau^* = T$.

Proposition 3 says that when the underlying price follows a (time-changed) Skellam process, the receiver of a last look quote can expect to do no better than to wait until either the quote expires, or the price process hits the boundary B . The reason is the following: either the sample space realization is such that the quote is going to be killed by the last look-feature, in which case one cannot expect to do better than wait until the process (S_t) hits the boundary B . Or, the quote is not going to be killed, in which case one cannot expect to do better than wait until the quote expires at time T . In either case, stopping before $T_B \wedge T$ gives a lower expected payoff than continuing. The key assumption for this argument to hold is, heuristically speaking, that the underlying process never "jumps past" the barrier ($P(M_{\tau^*} > B) = 0$), meaning that it is safe to wait until the exact moment when $S_t = B$. This condition is clearly satisfied when the process only jumps in size one.⁶ The formal proof is given in the appendix.

Proposition 4 (Symmetry of buy and sell quotes). *Let (S_t) be a time-changed Skellam process. Suppose we have the sell and buy quotes (\bar{K}, T, \bar{B}) and $(\underline{K}, T, \underline{B})$ satisfying*

$$\begin{aligned}\bar{K} &= -\underline{K} \\ \bar{B} &= -\underline{B}\end{aligned}$$

⁶Recall that there is zero probability that a Poisson process makes two jumps at the same time

If the price process is symmetric and satisfies $S_0 = 0$, meaning that we have the equality of law

$$S_t \stackrel{\text{law}}{=} -S_t \quad \text{all } t \geq 0$$

Then,

$$V_{\text{sell}}(\bar{K}, T, \bar{B}) = V_{\text{buy}}(\underline{K}, T, \underline{B})$$

3.3.2 Optimal stopping with infinite horizon and no time change

If there is no time change and the last look-problem has an infinite time horizon, the free-boundary problem associated with optimal stopping becomes a system of linear equations, which can be solved easily. We provide a heuristic overview of the argument.

Specifically, let $\sigma = 0$ and $y_0 = \eta$ (meaning there is no random time change). Also let $T = \infty$, meaning that the quote never expires. Then, the process (S_t) is a compound Poisson process with infinitesimal generator

$$\mathcal{A}f(y) = \lambda \sum_{x \in E} p(x)[f(x+y) - f(y)] \quad (3.6)$$

Where $p(0) = 0$ by convention.

Let $V(x)$ be the value function associated with (3.1). Recall the free boundary problem associated with optimal stopping (Peskir and Shiryaev, 2006):

$$\begin{aligned} \mathcal{A}V &= 0 \text{ in the continuation region} \\ \mathcal{A}V &= \psi \text{ in the stopping region} \end{aligned}$$

The problem is known as a free boundary problem because we need to simultaneously determine the stopping boundary *and* the value function V . In our case, let $q \in E$ denote the lower boundary of the stopping region, and define the chopped-off payoff function as

$$\bar{\psi}(x) = \begin{cases} 0 & \text{if } x > B \\ \psi(x) & \text{if } x \leq B \end{cases}$$

The free boundary-problem then becomes

$$V(y) = \bar{\psi}(y) \quad \text{for } y \geq q \quad (3.7)$$

$$V(y) = \sum_{x \in E} p(x)V(x+y) \quad \text{for } y < q \quad (3.8)$$

Since $\bar{\psi}(x) = 0$ for $x > B$ and $\lim_{x \rightarrow -\infty} V(x) = 0$, we can approximate the free boundary problem by a finite system of linear equations:

$$\mathbf{V} = (\mathbf{1} - \mathbf{P})^{-1} \Psi \quad (3.9)$$

where Ψ is a vector with $\Psi_n = \bar{\psi}(n)$ for $n \geq q$ and 0 otherwise.

For any $q \leq B$ the system (3.9) is fully determined and can be solved. Denote the corresponding value vector by $\mathbf{V}_k, k = B, B-1, \dots$. We can identify the optimal q^* by observing that $\mathbf{V}_{q^*} \gg \mathbf{V}_k$ for all k .

The general case differs from the special case considered here in two important respects. First, when the time horizon is not infinite, the value function V becomes also a function of the time remaining $(T - t)$, which introduce a time derivate term in the system of equations (3.7) (Øksendal and Sulem, 2005). Second, when we apply a random time change, we also change the infinitesimal generator of (S_t) .

3.3.3 General case

We rely on numerical methods for the general case. As is usual for models based on random time change we do not have closed-form expressions for probability densities. Therefore it seems a standard backward induction scheme can not be used. However, there are several other possibilities that are still open.

First, one could try to find the infinitesimal generator of the time-changed compound Poisson and write down the free boundary-problem associated with the optimal stopping rule (Peskir and Shiryaev, 2006). One can then hope to solve this problem via numerical PDE methods. Although we think this possibility is a promising avenue for future research, we have not pursued it further here.

Second, one can exploit the fact that there is a one-to-one correspondence between the characteristic function and the probability density of a random variable. Since we know the characteristic function of the TCDL process one can use Fast Fourier Transform or Fourier-Cosine series to implement an efficient numerical solution (Fang and Oosterlee, 2009). This method has the advantage of being very fast, which would be a considerable advantage in practical use. For details on implementing Fourier transform methods we refer to Gong and Zhuang (2017), who use Fourier-Cosine series to price American options under a time-changed Levy process, and Ding et al. (2012) who applies the same

method to Bermudan barrier options.

Third, we can simulate the TCDL process and use Monte Carlo-techniques. In the next section we provide implementation details the use of the Least Squares Monte Carlo algorithm to approximate the continuation value $c(x, t_n)$, as in [Longstaff and Schwartz \(2001\)](#). An algorithm for simulation of the TCDL process can be found in the appendix.

All three solution methods listed here require a discretization of the problem (3.1). In the next section we provide a general proof that the optimal stopping time for (3.1) can be approximated arbitrarily well by the solution to a discrete optimal stopping problem.

Discretization scheme

We shall reformulate the last look-problem into a Markovian structure by adding the *graveyard state* $\{\dagger\}$. Formally, let $\bar{E} := \{x \in E : x \leq B\} \cup \{\dagger\}$ and extend the payoff function ψ (see Definition 2) by setting $\psi(\dagger) = 0$. Define the *extended price process* by

$$\bar{S}_t = \begin{cases} S_t & \text{if } t < T_{\dagger} \\ \dagger & \text{if } t \geq T_{\dagger} \end{cases}$$

For any fixed integer N let the *time grid* \mathcal{G}_N be the set of $(1 + 2^N)$ equally spaced points defined by

$$\{0 = t_0 < t_1 < \dots < t_{1+2^N} = T \mid t_n - t_{n-1} = c \text{ all } n \geq 1\}$$

Let \mathcal{M}_N be the set of stopping times taking values in \mathcal{G}_N almost surely.

Define the *discrete stopping problem* with an extended state-space,

$$\hat{V}_N = \sup_{\sigma \in \mathcal{M}_N} \mathbb{E} [\psi(\bar{S}_{\sigma})] \quad (3.10)$$

Theorem 5. *Let τ^* be the first optimal stopping time of the last look-problem (3.1). Let σ_N^* be the first optimal stopping time of the discrete problem (3.10). Then,*

$$\sigma_N^* \xrightarrow{P} \tau^*$$

Lemma 6. *(\bar{S}_t) is a Markovian process.*

Define the **Bellman equations** as

$$v_N(x) := \psi(x) \quad (3.11)$$

$$c_n(x) := \sum_{y \in \bar{E}} v_{n+1}(y) p(y | x), \quad (3.12)$$

$$v_n(x) := \max\{\psi(x), c_n(x)\}, \quad n = 0, 1, \dots, N-1 \quad (3.13)$$

We will refer to c_n as the *continuation value*, and note that it is a conditional expectation.

Lemma 7. *The process*

$$Z_n := v_n(X_n), n = 0, 1, \dots, N$$

is a supermartingale.

Proposition 8. *Under the assumptions of this section the stopping time*

$$\sigma^* := T \wedge \inf\{n \in \mathcal{G}_N \mid \psi(S_n) = v_n\}$$

is optimal for (3.10).

For any given grid \mathcal{G}_N one can use Proposition 8 to solve the discrete optimal stopping problem, if one can compute the continuation value c_n ,

$$c_n(x) = \mathbb{E}[v_{n+1} \mid \bar{S}_n = x]$$

Unfortunately, for the class of processes considered here the probability densities $p(x | y)$ are not known in closed form. This makes it challenging to compute conditional expectations.

The least squares Monte Carlo method of [Longstaff and Schwartz \(2001\)](#) approximates the conditional expectation via simulation. Specifically, the approximation is done by regressing future realized payoffs on polynomial functions of the price process. Since one can choose polynomials that form a basis for the Hilbert space of square-integrable functions, one can under the right conditions approximate the conditional expectation arbitrarily well this way.

In Figure 3.6 we have plotted the optimal stopping boundary. Figure 3.7 shows the approximate stopping probabilities in the time interval $[0, T]$.

Figure 3.6: Execution boundary of the last look-problem





Figure 3.7: Stopping probabilities over time.



3.4 Last look under alternative price processes

3.4.1 Brownian motion

The Brownian motion is a standard building block in mathematical finance and is very common in practical applications. It is therefore of interest to study the last look-problem when the price process follows a Brownian motion. Hence, in this section we let (S_t) be a Brownian motion with $S_0 = x_0$.

The optimal stopping for the last look-problem (3.1) is given by $\tau^* = T \wedge T_B$. The proof of this claim is exactly the same as the proof of Proposition 3. The critical property used in this proof is that there is zero probability of the price process "jumping past" the barrier, $P(M_{T \wedge T_B} > B) = 0$, a property that the Brownian motion satisfies by virtue of its continuous paths.

For the Brownian motion we can also find an explicit formula for the expected value of the quote attained by following the optimal rule, $\mathbb{E}[\psi(S_{T \wedge T_B})]$. For this we need the density of the stopped process, $S_{T \wedge T_B}$, which can be found through the reflection principle. For easy reference we state the reflection principle first. See [Jeanblanc et al. \(2009, Chapter 3\)](#) for a proof.

Proposition 9 (Reflection principle for Brownian motion). *Let W_t be a standard Brownian motion. Let x and y be given real numbers satisfying $y \geq 0$ and $x \leq y$. Let (M_t) be the running supremum of the (W_t) , meaning that $M_t = \sup_{u \in [0, t]} W_u$. The following holds:*

$$P(W_t \leq x, M_t \geq y) = P(W_t \geq 2y - x)$$

Corollary 10. *For a fixed t , we have the equality of law*

$$M_t \stackrel{\text{law}}{=} |W_t|$$

Lemma 11. *Let (S_t) and T_B be as defined above, and let $\phi(\cdot)$ be the standard normal density. The density f of the stopped process $S_{T \wedge T_B}$ is given by*

$$f(u) = \phi\left(\frac{u - x_0}{\sigma\sqrt{T}}\right) - \phi\left(\frac{u + x_0 - 2B}{\sigma\sqrt{T}}\right)$$

It turns out that the expected value attained by following the optimal stopping rule can be expressed as a linear combination of quotes *without* the last look feature⁷

⁷This result has an analogy in option pricing theory - The value of a European barrier option can be decomposed into two vanilla European options. Here, the "option feature" of the quote is in practice European, since early exercise is never optimal.

Proposition 12. Let (S_T) and T_B be as defined above, and let $\mathbb{E}_r[f(S_t)]$ denote the expectation of $f(S_t)$ when one changes the starting point of (S_t) from x_0 to r . We have

$$\sup_{\tau \in \mathcal{T}} \mathbb{E}_{x_0} [\mathbf{I}_{\{M_\tau \leq B\}} \Psi(S_\tau)] = \mathbb{E}_{x_0} [\Psi(S_T)] - \mathbb{E}_{2B-x_0} [\Psi(S_T)]$$

We can also compute the probability that the last look-feature is activated:

Proposition 13. Let (S_t) and T_B be as defined above, and let $\Phi(\cdot)$ be the standard normal CDF. Then,

i) The hitting time T_B has the scaled inverse chi-square distribution,

$$T_B \stackrel{\text{law}}{=} \frac{(B-x_0)^2}{\sigma^2 Z^2}, \quad Z \sim \mathcal{N}(0, 1)$$

ii) The probability of the last look-feature coming into effect is given by

$$P(T_B \leq T) = 2 - 2\Phi\left(\frac{B-x_0}{\sigma\sqrt{T}}\right)$$

3.4.2 Skellam process

Proposition 3 stated that the optimal stopping time when the price process is a time-changed Skellam process is $\tau^* = T \wedge T_B$. For the special case without a time-change we can find explicit formulas, which it turns out has a structure mirroring those found for the Brownian motion. Throughout this section we assume that (S_t) is a martingale and a Skellam process (meaning the the intensity of up and down-jumps are the same). We start by characterising the probability mass function of the stopped process, by adapting the reflection principle.

Proposition 14. The probability mass function of the stopped process $S_{T \wedge T_B}$ is given by

$$P(S_{T \wedge T_B} = x) = \begin{cases} 0 & \text{if } x > B \\ P(N_T = B) & \text{if } x = B \\ P(S_T = x) - P(S_T = 2B - x) & \text{if } x < B \end{cases}$$

We can now compute the value of the last look-quote.

Proposition 15. Let $I_k(x)$ be the modified Bessel function of the first kind. The value of (3.1) is given by

$$\sup_{\tau \in \mathcal{T}} \mathbb{E} [\mathbf{I}_{\{M_\tau \leq B\}} \Psi(S_\tau)] = e^{-\lambda T} \left(\sum_{x < B} \Psi(x) [I_{|x|}(\lambda T) - I_{|2B-x|}(\lambda T)] + \Psi(B) \frac{(\lambda T)^B}{B!} \right)$$

We can also compute the probability that the last look comes into effect, as for the case when the price follows a Brownian motion.

Proposition 16. *The probability of the last look coming into effect, $P(T_B \leq T)$, is given by*

$$\begin{aligned} P(T_B \leq T) &= 1 - P(N_T \leq B - 1) \\ &= 1 - e^{-\lambda T} \sum_{n=0}^{B-1} \frac{(\lambda T)^n}{n!} \end{aligned}$$

3.5 Conclusions

In this paper we modelled the optimal stopping problem associated with last look-quotes, and showed how the mathematical structure of the problem is similar to that involved in pricing an American barrier option. Furthermore, we introduced the time-changed discrete Levy process in order to account for important stylized facts about high-frequency prices. The optimal stopping problem must in general be solved with numerical methods, and we provide a discretization scheme and a convergence result. Some special cases involving the time-changed discrete Levy are solved explicitly: when the price process only jumps in size 1 (the time-changed Skellam process), and when there is an infinite horizon and no time change. We also solve the optimal stopping problem for Brownian motions and ordinary Skellam processes, and provide formulas for quote value and stopping probabilities.

3.A Proofs

Proof of Proposition 2. Integrability follows from applying Wald's equation:

$$\begin{aligned} \mathbb{E}[L_t] &= \mathbb{E}[U_t] + \mathbb{E}[D_t] \\ &= (\lambda^+ \mathbb{E}[u_1] + \lambda^- \mathbb{E}[d_1])t < \infty \end{aligned}$$

The process (L_t) is \mathbb{F} -adapted because (N_t^+) , (N_t^-) , (u_n) and (d_n) are all assumed \mathbb{F} -adapted.

For the martingale property, observe that

$$\begin{aligned}\mathbb{E}[L_t \mid \mathcal{F}_j] &= \mathbb{E}[U_t - D_t \mid \mathcal{F}_j] \\ &= U_j + (t-j)\lambda^+ \mathbb{E}[u_1] - D_j - (t-j)\lambda^- \mathbb{E}[d_1] \\ &= S_j + (t-j)(\lambda^+ \mathbb{E}[u_1] - \lambda^- \mathbb{E}[d_1])\end{aligned}$$

We see that the martingale property holds if and only if

$$\lambda^+ \mathbb{E}[u_1] - \lambda^- \mathbb{E}[d_1] = 0$$

Which is equivalent to the stated claim. \square

Proposition 3. We prove the case $S_0 < B$ (the sell quote). Take any stopping time $\tau \in \mathcal{T}$. We shall first show that

$$\mathbb{E}[\psi(S_\tau) \mathbf{1}_{\{M_\tau \leq B\}}] \leq \mathbb{E}[\psi(S_{\tau^*}) \mathbf{1}_{\{M_{\tau^*} \leq B\}}]$$

Let $A := \{\omega \in \Omega \mid M_{T \wedge T_B} \leq B\}$.

We have that

$$\begin{aligned}\psi(S_\tau) \mathbf{1}_{\{M_\tau \leq B\}} &= \underbrace{\psi(S_\tau) \mathbf{1}_{\{M_\tau \leq B\}} \mathbf{1}_{\{\tau < T \wedge T_B\}} \mathbf{1}_A}_{\text{I}} \\ &\quad + \underbrace{\psi(S_\tau) \mathbf{1}_{\{M_\tau \leq B\}} \mathbf{1}_{\{\tau \geq T \wedge T_B\}} \mathbf{1}_A}_{\text{II}} \\ &\quad + \underbrace{\psi(S_\tau) \mathbf{1}_{\{M_\tau \leq B\}} \mathbf{1}_{A^c}}_{\text{III}}\end{aligned}$$

We consider the terms I and II separately. We shall see that term III vanishes in expectations, since A^c is a null set.

First consider term I. On the set $\{\tau < T \wedge T_B\}$ we have $\mathcal{F}_\tau \subseteq \mathcal{F}_{T \wedge T_B}$, and hence we can use Doob's optional sampling theorem on the martingale (S_t) . Furthermore we apply Jensen's inequality for conditional expectations, and write

$$\begin{aligned}\psi(S_\tau) \mathbf{1}_{\{M_\tau \leq B\}} \mathbf{1}_{\{\tau < T \wedge T_B\}} \mathbf{1}_A &= \psi(\mathbb{E}[S_{T \wedge T_B} \mid \mathcal{F}_\tau]) \mathbf{1}_{\{M_\tau \leq B\}} \mathbf{1}_{\{\tau < T \wedge T_B\}} \mathbf{1}_A \\ &\leq \mathbb{E}[\psi(S_{T \wedge T_B}) \mid \mathcal{F}_\tau] \mathbf{1}_{\{M_\tau \leq B\}} \mathbf{1}_{\{\tau < T \wedge T_B\}} \mathbf{1}_A \\ &= \mathbb{E}[\psi(S_{T \wedge T_B}) \mid \mathcal{F}_\tau] \mathbf{1}_{\{M_{T \wedge T_B} \leq B\}} \mathbf{1}_{\{\tau < T \wedge T_B\}} \mathbf{1}_A\end{aligned}$$

The last equality use that on the set $\{\tau < T \wedge T_B\} \cap A$ we have $\mathbf{1}_{\{M_\tau \leq B\}} = \mathbf{1}_{\{M_{T \wedge T_B} \leq B\}} = 1$.

The above inequality also show that it is never optimal to stop before $T \wedge T_B$. Hence, if $T \wedge T_B$ is indeed optimal, it must also be the first optimal stopping time.

Now consider term II. By the definition of T_B and since ψ is monotonically increasing we have $\psi(S_{T_B}) = \psi(B) > \psi(x)$ for all $x < B$. The assumption $\tau \in \mathcal{F}$ means that $\tau \leq T$. Therefore on the set $\{\tau \geq T \wedge T_B\} \cap A$ it must be the case that $T_B \leq T$. On the set A we also have $\mathbf{1}_{\{M_{T_B} \leq B\}} = 1 \geq \mathbf{1}_{\{M_\tau \leq B\}}$. Therefore,

$$\begin{aligned} \mathbf{1}_{\{\tau \geq T_B\}} \mathbf{1}_{\{M_\tau \leq B\}} \psi(S_\tau) \mathbf{1}_{\{A\}} &\leq \mathbf{1}_{\{\tau \geq T_B\}} \mathbf{1}_{\{M_\tau \leq B\}} \psi(S_{T_B}) \mathbf{1}_{\{A\}} \\ &\leq \mathbf{1}_{\{\tau \geq T_B\}} \mathbf{1}_{\{M_{T_B} \leq B\}} \psi(S_{T_B}) \mathbf{1}_{\{A\}} \\ &= \mathbf{1}_{\{\tau \geq T_B \wedge T\}} \mathbf{1}_{\{M_{T_B \wedge T} \leq B\}} \psi(S_{T_B \wedge T}) \mathbf{1}_{\{A\}} \\ &= \mathbf{1}_{\{\tau \geq T_B \wedge T\}} \mathbf{1}_{\{M_{T_B \wedge T} \leq B\}} \psi(\mathbb{E}[S_{T_B \wedge T} \mid \mathcal{F}_\tau]) \mathbf{1}_{\{A\}} \end{aligned}$$

The last line use that $\mathcal{F}_{T \wedge T_B} \subseteq \mathcal{F}_\tau$ for $\{\omega \in \Omega \mid \tau \geq T \wedge T_B\}$, and thus $S_{T_B \wedge T}$ is \mathcal{F}_τ -measurable.

Combining our considerations for term I and term II, we get

$$\begin{aligned} \psi(S_\tau) \mathbf{1}_{\{M_\tau \leq B\}} &\leq \\ &\mathbb{E}[\psi(S_{T \wedge T_B}) \mid \mathcal{F}_\tau] \mathbf{1}_{\{M_{T \wedge T_B} \leq B\}} \mathbf{1}_{\{\tau < T \wedge T_B\}} \mathbf{1}_{\{A\}} + \\ &\mathbb{E}[\psi(S_{T \wedge T_B}) \mid \mathcal{F}_\tau] \mathbf{1}_{\{M_{T \wedge T_B} \leq B\}} \mathbf{1}_{\{\tau \geq T \wedge T_B\}} \mathbf{1}_{\{A\}} + \\ &\psi(S_\tau) \mathbf{1}_{\{M_\tau \leq B\}} \mathbf{1}_{\{A^c\}} \\ &= \mathbb{E}[\psi(S_{T \wedge T_B}) \mid \mathcal{F}_\tau] \mathbf{1}_{\{M_{T \wedge T_B} \leq B\}} \mathbf{1}_{\{A\}} + \psi(S_\tau) \mathbf{1}_{\{M_\tau \leq B\}} \mathbf{1}_{\{A^c\}} \end{aligned}$$

We shall apply the expectation operator on the preceding inequality. Note that on $\{A\}$ we have $\mathbf{1}_{\{M_{T \wedge T_B}\}} = 1$. Moreover since $P(A) = 1$ by assumption we have $P(A^c) = 0$ and hence $\mathbb{E}[Y \mathbf{1}_{\{A^c\}}] = 0$ for any random variable Y by the properties of the Lebesgue

integral. Therefore, by the Tower property of conditional expectations,

$$\begin{aligned}
\mathbb{E} [\psi(S_\tau) \mathbf{1}_{\{M_\tau \leq B\}}] &\leq \mathbb{E} \left[\mathbb{E}[\psi(S_{T \wedge T_B}) \mid \mathcal{F}_\tau] \mathbf{1}_{\{M_{T \wedge T_B} \leq B\}} \mathbf{1}_{\{A\}} \right] \\
&= \mathbb{E} [\mathbb{E}[\psi(S_{T \wedge T_B}) \mid \mathcal{F}_\tau]] \\
&= \mathbb{E} [\psi(S_{T \wedge T_B})] \\
&= \mathbb{E} \left[\psi(S_{T \wedge T_B}) \mathbf{1}_{\{M_{T \wedge T_B} \leq B\}} \right]
\end{aligned}$$

Since $T \wedge T_B \in \mathcal{T}$ and τ was arbitrary, we have proved that $T \wedge T_B$ is the first optimal stopping time. The case $S_0 > B$ (the buy quote) follows the same steps. \square

Proposition 4. Take as given the last look-quotes (\bar{K}, T, \bar{B}) and $(\underline{K}, T, \underline{B})$. First note that, for any $t \in [0, T]$, we have

$$\begin{aligned}
P(S_t > \bar{B}) &= P(S_t > -\underline{B}) \\
&= P(-S_t < \underline{B}) \\
&= P(S_t < \underline{B})
\end{aligned}$$

Which implies that

$$\begin{aligned}
T_{\dagger}(\bar{B}) &:= \inf\{t \geq 0 \mid S_t > \bar{B}\} \\
&\stackrel{law}{=} \inf\{t \geq 0 \mid S_t < \underline{B}\} \\
&=: T_{\dagger}(\underline{B})
\end{aligned}$$

Also note that for any $\tau \in \mathcal{T}$,

$$\{M_\tau \leq \bar{B}\} = \{\tau < T_{\dagger}(\bar{B})\}$$

and

$$\{m_\tau \geq \underline{B}\} = \{\tau < T_{\dagger}(\underline{B})\}$$

Therefore we have that

$$\begin{aligned}
P(M_\tau \leq \bar{B}) &= P(\tau < T_{\dagger}(\bar{B})) \\
&= P(\tau < T_{\dagger}(\underline{B})) \\
&= P(m_\tau \geq \underline{B})
\end{aligned}$$

Using the above we get that

$$\begin{aligned}
\mathbb{E} \left[\psi_{\text{sell}}(S_\tau) \mathbf{1}_{\{M_\tau \leq \bar{B}\}} \right] &= \mathbb{E}[\max(0, S_\tau - \bar{K}) \mid \tau < T_\dagger(\bar{B})] P(\tau < T_\dagger(\bar{B})) \\
&= \mathbb{E}[\max(0, (\underline{K} - S_\tau)) \mid \tau < T_\dagger(\underline{B})] P(\tau < T_\dagger(\underline{B})) \\
&= \mathbb{E}[\psi_{\text{buy}}(S_\tau) \mathbf{1}_{\{m_\tau \geq \underline{B}\}}]
\end{aligned}$$

By assumption we have that

$$S_{T_{\bar{B}}} = \bar{B} = -\underline{B} = -S_{T_{\underline{B}}}$$

and since S_t is symmetric we have $T_{\bar{B}} \stackrel{\text{law}}{=} T_{\underline{B}}$. This means that

$$\begin{aligned}
P(S_{T \wedge T_{\bar{B}}} \leq x) &= P(S_T \leq x \mid T < T_{\bar{B}}) + P(S_{T_{\bar{B}}} \leq x \mid T \geq T_{\bar{B}}) \\
&= P(-S_T \leq x \mid T < T_{\underline{B}}) + P(S_{T_{\underline{B}}} \leq x \mid T \geq T_{\underline{B}}) \\
&= P(-S_{T \wedge T_{\underline{B}}} \leq x)
\end{aligned}$$

Showing that the we have symmetry of distribution also for the random variables $(S_{T \wedge T_{\bar{B}}}, S_{T \wedge T_{\underline{B}}})$. Using that $T \wedge T_{\bar{B}}$ is optimal (Proposition 3), we have that

$$\begin{aligned}
V_{\text{sell}}(\bar{K}, T, \bar{B}) &= \sup_{\tau \in \mathcal{T}} \mathbb{E} \left[\psi_{\text{sell}}(S_\tau) \mathbf{1}_{\{M_\tau \leq \bar{B}\}} \right] \\
&= \sup_{\tau \in \mathcal{T}} \mathbb{E}[\psi_{\text{buy}}(S_\tau) \mathbf{1}_{\{m_\tau \geq \underline{B}\}}] \\
&= V_{\text{buy}}(\underline{K}, T, \underline{B})
\end{aligned}$$

And the proof is complete. \square

Proof of Theorem 5. Take as given $\delta > 0$ and $\varepsilon > 0$. We shall prove that there exists an $N \in \mathbb{N}$ such that

$$P(|\tau^* - \sigma_N^*| > \varepsilon) < \delta$$

Let \mathcal{J} be the set of jump times of \bar{S}_t , meaning the set of random variables

$$\mathcal{J} := \{j : \Omega \rightarrow \mathbb{R}_+ : |\bar{S}_j - \bar{S}_{j-}| > 0\}$$

Define the constant u as

$$u := \sup \left\{ t \in \mathbb{R}_+ \mid P \left(\mathcal{J} \cap (\tau^*, \tau^* + t) \neq \emptyset \right) < \delta \right\}$$

Note that for all $t \leq u$,

$$P(\bar{\mathcal{S}}_{\tau^*} \neq \bar{\mathcal{S}}_{\tau^*+t}) < \delta$$

Let $N \in \mathbb{N}$ be such that

$$2^{-N} < \min\{u, \varepsilon\}$$

Now define the random variable n^* as the first grid point coming after the optimal stopping time τ^* ,

$$n^* = \min\{n \in \mathcal{G}_N \mid n > \tau^*\}$$

It is clear that $n^* \in \mathcal{M}_N$. Moreover, the time grid is constructed so that

$$\tau^* < n^* < \tau^* + u$$

Next we show that n^* is an optimal stopping time for the discrete problem. Define the set A as the region of the sample space where there are no jumps in $(\tau^*, \tau^* + u)$,

$$A := \left\{ \omega \in \Omega \mid \mathcal{J} \cap (\tau^*, \tau^* + u) = \emptyset \right\}$$

Observe that

$$\begin{aligned} \mathbf{1}_A \mathbb{E}[\psi(\bar{\mathcal{S}}_{n^*})] &= \mathbf{1}_A \mathbb{E}[\psi(\bar{\mathcal{S}}_{\tau^*})] \\ &= \mathbf{1}_A \left(\sup_{\tau \in \mathcal{T}} \mathbb{E}[\psi(\bar{\mathcal{S}}_{\tau})] \right) \\ &\geq \mathbf{1}_A \left(\sup_{\sigma \in \mathcal{M}_N} \mathbb{E}[\psi(\bar{\mathcal{S}}_{\sigma})] \right) \end{aligned}$$

Where the last inequality follows from observing that $\mathcal{M}_N \subseteq \mathcal{T}$. This shows that on the set A , $\sigma_N^* = n^*$. Therefore,

$$P(\sigma_N^* \neq n^*) \leq P(A^c) < \delta$$

Since $\tau^* - n^* < \varepsilon$ it follows that

$$P(|\tau^* - \sigma_N^*| > \varepsilon) < \delta$$

Which concludes the proof. \square

Lemma 6. We shall prove that for all $0 \leq s \leq t$ and any bounded Borel-measurable function $f : E \rightarrow \mathbb{R}$, we have

$$\mathbb{E}[f(\bar{\mathcal{S}}_t) \mid \mathcal{F}_u] = \mathbb{E}[f(\bar{\mathcal{S}}_t) \mid \bar{\mathcal{S}}_u]$$

Where the notation $\mathbb{E}[\cdot | X]$ is understood to mean $\mathbb{E}[\cdot | \sigma(X)]$. Note that the expectations involved are well-defined by our standing assumptions. We have that

$$\mathbb{E}[f(\tilde{S}_t) | \mathcal{F}_u] = \underbrace{\mathbb{E}[f(\tilde{S}_t)\mathbf{1}_{\{\tilde{S}_u \neq \dagger\}} | \mathcal{F}_u]}_{(*)} + \underbrace{\mathbb{E}[f(\tilde{S}_t)\mathbf{1}_{\{\tilde{S}_u = \dagger\}} | \mathcal{F}_u]}_{(**)}$$

For the term $(*)$ we have

$$\begin{aligned} \mathbb{E}[f(\tilde{S}_t)\mathbf{1}_{\{\tilde{S}_u \neq \dagger\}} | \mathcal{F}_u] &= \left(\mathbb{E}[f(\tilde{S}_t)\mathbf{1}_{\{\tilde{S}_t = \dagger\}} | \mathcal{F}_u] + \mathbb{E}[f(\tilde{S}_t)\mathbf{1}_{\{\tilde{S}_t \neq \dagger\}} | \mathcal{F}_u] \right) \mathbf{1}_{\{M_u \leq B\}} \\ &= \left(\mathbb{E}[\mathbf{1}_{\{\sup_{j \in [u,t]} S_j > B\}} | \mathcal{F}_u] f(\dagger) + \mathbb{E}[f(S_t)\mathbf{1}_{\{\sup_{j \in [u,t]} S_j \leq B\}} | \mathcal{F}_u] \right) \mathbf{1}_{\{M_u \leq B\}} \\ &= \left(\mathbb{E}[\mathbf{1}_{\{\sup_{j \in [u,t]} S_j > B\}} | S_u] f(\dagger) + \mathbb{E}[f(S_t)\mathbf{1}_{\{\sup_{j \in [u,t]} S_j \leq B\}} | S_u] \right) \mathbf{1}_{\{M_u \leq B\}} \\ &= \left(\mathbb{E}[\mathbf{1}_{\{\sup_{j \in [u,t]} S_j > B\}} | \tilde{S}_u] f(\dagger) + \mathbb{E}[f(S_t)\mathbf{1}_{\{\sup_{j \in [u,t]} S_j \leq B\}} | \tilde{S}_u] \right) \mathbf{1}_{\{M_u \leq B\}} \\ &= \left(\mathbb{E}[\mathbf{1}_{\{\sup_{j \in [u,t]} S_j > B\}} | \tilde{S}_u] f(\dagger) + \mathbb{E}[f(\tilde{S}_t)\mathbf{1}_{\{\sup_{j \in [u,t]} S_j \leq B\}} | \tilde{S}_u] \right) \mathbf{1}_{\{M_u \leq B\}} \\ &= \mathbb{E}[f(\tilde{S}_t)\mathbf{1}_{\{\tilde{S}_u \neq \dagger\}} | \tilde{S}_u] \end{aligned}$$

For the term $(**)$ we have

$$\begin{aligned} \mathbb{E}[f(\tilde{S}_t)\mathbf{1}_{\{\tilde{S}_u = \dagger\}} | \mathcal{F}_u] &= f(\dagger) \mathbb{E}[\mathbf{1}_{\{\tilde{S}_u = \dagger\}} | \mathcal{F}_u] \\ &= f(\dagger) \mathbb{E}[\mathbf{1}_{\{\tilde{S}_u = \dagger\}} | \tilde{S}_u] \\ &= \mathbb{E}[f(\tilde{S}_t)\mathbf{1}_{\{\tilde{S}_u = \dagger\}} | \tilde{S}_u] \end{aligned}$$

Combining $(*)$ and $(**)$ gives

$$\begin{aligned} \mathbb{E}[f(\tilde{S}_t) | \mathcal{F}_u] &= \mathbb{E}[f(\tilde{S}_t)\mathbf{1}_{\{\tilde{S}_u \neq \dagger\}} | \tilde{S}_u] + \mathbb{E}[f(\tilde{S}_t)\mathbf{1}_{\{\tilde{S}_u = \dagger\}} | \tilde{S}_u] \\ &= \mathbb{E}[f(\tilde{S}_t) | \tilde{S}_u] \end{aligned}$$

And the proof is complete. □

Lemma 7. We have

$$\begin{aligned} \mathbb{E}[Z_{n+1} | \mathcal{F}_n] &= \mathbb{E}[Z_{n+1} | X_n] \\ &= \mathbb{E}[v_{n+1}(X_{n+1}) | X_n] \\ &\leq v_n(X_n) \\ &= Z_n \end{aligned}$$

Which proves the claim. □

Proposition 8. Let the process (Z_n) be as defined in Lemma 7. From that result and from the optional stopping-theorem for supermartingales, we have that, for any stopping time $\sigma \in \mathcal{M}$,

$$\mathbb{E}[Z_\sigma] \leq \mathbb{E}[Z_0] = v_0(x)$$

We use this and the definition of v_n to conclude that for any n and any stopping time $\sigma \in \mathcal{M}$,

$$v_0(x) \geq \mathbb{E}[Z_\sigma] \geq \mathbb{E}[\psi(X_\sigma)]$$

We shall show that inequality holds with equality for σ^* , which then proves our claim.

For this purpose we consider, for $n = 0, 1, \dots, N-1$, the stopping time $\sigma^* \wedge (n+1)$ and the stopped process $Z_{\sigma^* \wedge (n+1)}$. Note that the random variables $\mathbf{1}_{\{\sigma^* \leq n\}}$ and $\mathbf{1}_{\{\sigma^* > n\}}$ are both \mathcal{F}_n -measurable. We can therefore write

$$\begin{aligned} \mathbb{E}[Z_{\sigma^* \wedge (n+1)} | X_n] &= \mathbb{E}[\mathbf{1}_{\{\sigma^* \leq n\}} Z_{\sigma^*} + \mathbf{1}_{\{\sigma^* > n\}} Z_{\sigma^*} | X_n] \\ &= \mathbf{1}_{\{\sigma^* \leq n\}} Z_{\sigma^*} + \mathbb{E}[\mathbf{1}_{\{\sigma^* > n\}} Z_{\sigma^*} | X_n] \\ &= \mathbf{1}_{\{\sigma^* \leq n\}} Z_{\sigma^*} + \mathbb{E}[\mathbf{1}_{\{\sigma^* > n\}} V_n | X_n] \\ &= \mathbf{1}_{\{\sigma^* \leq n\}} Z_{\sigma^*} + \mathbf{1}_{\{\sigma^* > n\}} V_n \\ &= Z_{\sigma^* \wedge n} \end{aligned}$$

Therefore,

$$\begin{aligned} v_0(X_0) &= \mathbb{E}[Z_0] = \mathbb{E}[Z_{\sigma^* \wedge 0}] \mathbb{E}[Z_{\sigma^* \wedge 1}] = \dots = \mathbb{E}[Z_{\sigma^* \wedge N}] = \mathbb{E}[Z_{\sigma^*}] \\ &= \mathbb{E}[v_{\sigma^*}(X_{\sigma^*})] \\ &= \mathbb{E}[\psi(X_{\sigma^*})] \end{aligned}$$

Which concludes the proof. □

Proof of Lemma 11. The proof is a consequence of the reflection principle for the Brownian motion. Let (W_t) be the standard Brownian motion, and let here (Y_t) be the running supremum

$$(Y_t) := \left(\sup_{u \leq t} W_u, t \geq 0 \right)$$

Take first the joint probability law of the terminal value W_T of the Brownian motion and the running supremum over $[0, T]$:

$$\begin{aligned}
 P(W_T \leq x, Y_T \leq y) &= P(W_T \leq x) - P(W_T \leq x, Y_T \geq y) \\
 &= \Phi\left(\frac{x}{\sqrt{T}}\right) - P(W_T \leq x, Y_T \geq y) \\
 &= \Phi\left(\frac{x}{\sqrt{T}}\right) - P(W_T \leq x - 2y) \\
 &= \Phi\left(\frac{x}{\sqrt{T}}\right) - \Phi\left(\frac{x - 2y}{\sqrt{T}}\right)
 \end{aligned}$$

Where the first line use the law of total probability and the second-to-last line use the reflection principle. Set $\hat{x} = \frac{x-x_0}{\sigma}$ and $\hat{y} = \frac{y-x_0}{\sigma}$, and note that by the properties of the Brownian motion and the Normal distribution we have

$$P(S_T \leq x) = P(W_T \leq \hat{x}) \quad \text{and} \quad P\left(\sup_{t \leq T} S_t > y\right) = P\left(\sup_{t \leq T} W_t > \hat{y}\right)$$

Hence, if we let $F(x, y)$ denote the joint distribution of $(S_T, \sup_{t \leq T} S_t)$, we get

$$\begin{aligned}
 F(x, y) &:= P(S_T \leq x, \sup S_T \leq y) \\
 &= P(W_T \leq \hat{x}, Y_T \leq \hat{y}) \\
 &= \Phi\left(\frac{x - x_0}{\sigma\sqrt{T}}\right) - \Phi\left(\frac{x + x_0 - 2y}{\sigma\sqrt{T}}\right)
 \end{aligned}$$

Or, equivalently in terms of the joint density function,

$$f(u, y) = \phi\left(\frac{u - x_0}{\sigma\sqrt{T}}\right) - \phi\left(\frac{u + x_0 - 2y}{\sigma\sqrt{T}}\right)$$

The lemma now follows from setting $y = B$. □

Proof of Proposition 12. From the proof of Proposition 3 we know the optimal stopping time to be $\tau^* = T_B \wedge T$, so we must evaluate $\mathbb{E}_{x_0}[\psi(S_{T_B \wedge T})]$. This is straightforward when we use the density of the stopped process from Lemma 11 :

$$\mathbb{E}_{x_0}[\psi(S_{T_B \wedge T})] = \int_{\mathbb{R}} \psi(u) \phi\left(\frac{u - x_0}{\sigma\sqrt{T}}\right) du - \int_{\mathbb{R}} \psi(u) \phi\left(\frac{u + x_0 - 2B}{\sigma\sqrt{T}}\right) du$$

The first integral is the expectation of $\psi(S_T)$ when the process (S_t) starts in x_0 . The second integral can also be seen as the expectation of $\psi(S_T)$, but now the starting point of the process (S_t) has been shifted to $2B - x$. □

Proof of Proposition 13. We first prove i) for the sell quote, meaning that $B > S_0$. We have

$$\begin{aligned} M_T &:= \sup_{t \in [0, T]} S_t \\ &= \sup_{t \in [0, T]} \frac{W_T - x_0}{\sigma} =: Y_T \end{aligned}$$

Define

$$y := \frac{B - x_0}{\sigma}$$

We apply the following corollary to the reflection principle for the Brownian motion:

$$P(M_T \leq B) = P(|W_T| \leq B)$$

Hence, we have

$$\begin{aligned} P(T_B \leq T) &= P(B \leq M_T) \\ &= P(y \leq Y_T) \\ &= P(y \leq |W_T|) \\ &= P(y \leq |Z|\sqrt{T}) \\ &= P\left(\frac{y}{\sqrt{T}} \leq |Z|\right) \end{aligned}$$

Which implies that the hitting time T_B has the scaled inverse chi-square law,

$$T_B \stackrel{\text{law}}{=} \frac{(B - x_0)^2}{\sigma^2 Z^2}$$

To prove ii), we could use the CDF of the scaled inverse chi-square distribution, properties of the incomplete gamma function and its relation to the normal CDF. However, the claim can also be derived using only the reflection principle and the symmetry of the normal distribution:

$$\begin{aligned}
P(T_B \leq T) &= P(M_T > B) \\
&= P(Y_T > y) \\
&= 1 - P(Y_T \leq y) \\
&= 1 - (P(W_T \leq y) - P(W_T \leq -y)) \\
&= 1 - P(W_T \leq y) + 1 - P(W_T \geq -y) \\
&= 2 - P(W_T \leq y) - P(-W_T \leq y) \\
&= 2 - P(W_T \leq y) - P(W_T \leq y) \\
&= 2 - 2\Phi\left(\frac{y}{\sqrt{T}}\right) \\
&= 2 - 2\Phi\left(\frac{B - x_0}{\sigma\sqrt{T}}\right)
\end{aligned}$$

This completes the proof. \square

Proof of Proposition 14. The case $x > B$ follows from noting that $P(M_{T_B} > B) = 0$. Therefore $P(S_{T \wedge T_B} > B) = 0$.

For the case $x < B$ we first note that

$$\begin{aligned}
P(S_{T \wedge T_B} = x) &= P(S_T = x, M_T < B) \\
&= P(S_T = x) - P(S_T = x, M_T \geq B)
\end{aligned}$$

Where the second equality follows from the law of total probability.

We shall derive an adaptation of the reflection principle to the Skellam process in order to turn the expression $P(S_T = x, M_T \geq B)$ into one not involving M_T .

Since $P(M_{T_B} > B) = 0$ we can conclude that $S_{T_B} = B$ almost surely. Recall that the Skellam process is a Levy process. Moreover, the standing assumption that (S_t) is a martingale implies that $\lambda^+ = \lambda^-$ and hence that the process is symmetric, meaning that for any t ,

$$S_t \stackrel{\text{law}}{=} -S_t$$

Using these properties we deduce that

$$\begin{aligned}
P(S_T = x, M_T \geq B) &= P(S_T = x, T_B \leq T) \\
&= P(S_T = x \mid T_B \leq T)P(T_B \leq T) \\
&= P(S_T - S_{T_B} = x - B \mid T_B \leq T)P(T_B \leq T) \\
&= P(S_{T - T_B} = x - B \mid T_B \leq T)P(T_B \leq T) \\
&= P(-S_{T - T_B} = x - B \mid T_B \leq T)P(T_B \leq T) \\
&= P(S_{T - T_B} = B - x \mid T_B \leq T)P(T_B \leq T) \\
&= P(S_T - S_{T_B} = B - x \mid T_B \leq T)P(T_B \leq T) \\
&= P(S_T = 2B - x \mid T_B \leq T)P(T_B \leq T) \\
&= P(S_T = 2B - x, M_T \geq B) \\
&= P(S_T = 2B - x)
\end{aligned}$$

The last equality use that since $x < B$ we have $2B - x > B$, and thus on the set $\{S_t > 2B - x\}$ we have $M_t > B$ almost surely. Hence, we have

$$P(S_{T \wedge T_B} = x) = P(S_T = x) - P(S_T = 2B - x)$$

For the case $x = B$ we first note that

$$P(S_{T \wedge T_B} = x) = P(M_T = B)$$

We shall now use another adaptation of the reflection principle to prove that

$$M_t \stackrel{\text{law}}{=} N_t, \text{ any } t \geq 0$$

Following the same steps as in the previous case, we have that

$$\begin{aligned}
P(M_T \leq B) &= P(S_T \leq B, M_T \leq B) \\
&= P(S_T \leq B) - P(S_T \leq B, M_T > B) \\
&= P(S_T \leq B) - P(S_T \leq B \mid T_B < T)P(T_B < T) \\
&= \dots \\
&= P(S_T \leq B) - P(S_T \geq 2B - x, M_T > B) \\
&= P(S_T \leq B) - P(S_T \geq 2B - x) \\
&= P(S_T \leq B) - P(S_T \geq 2B - B) \\
&= P(S_T \leq B) - P(S_T \geq B) \\
&= P(S_T \leq B) - P(-S_T \geq B) \\
&= P(S_T \leq B) - P(S_T \leq -B) \\
&= P(|S_T| \leq B) \\
&= P(|U_T + D_T| \leq B) \\
&= P(N_T \leq B)
\end{aligned}$$

Equality of laws implies that $P(M_T = B) = P(N_T = B)$.

□

Proof of Proposition 15. We know from Proposition 3 that the optimal stopping time is $T \wedge T_B$. Hence, we get

$$\begin{aligned}
\sup_{\tau \in \mathcal{F}} \mathbb{E} [\mathbf{1}_{\{M_\tau \leq B\}} \psi(S_\tau)] &= \mathbb{E} [\psi(S_{T \wedge T_B})] \\
&= \sum_{x \in \mathbb{Z}} \psi(x) P(S_{T \wedge T_B} = x)
\end{aligned}$$

The claim now follows from applying the density of the stopped process (Lemma 14).

□

Proof of Proposition 16. We have

$$\begin{aligned}
 P(T_B \leq T) &= P(M_T \geq B) \\
 &= 1 - P(M_T \leq B - 1) \\
 &= 1 - P(N_T \leq B - 1) \\
 &= 1 - e^{-\lambda T} \sum_{n=0}^{B-1} \frac{(\lambda T)^n}{n!}
 \end{aligned}$$

And the proof is complete. □

3.B Algorithms

We here provide an algorithm for simulation the TCDL process, and for the numerical solution of the last look-optimal stopping problem.

A couple of things should be noted regarding the simulation of the TCDL process. First, we use an Euler-Maruyama scheme to discretize the CIR SDE. There is a general issue with how to avoid negative values in the square-root term of the discretized SDE. We use what is known as "full truncation", which means that we disregard realizations with negative values (see for example [Andersen \(2007\)](#) for more details). Next we simulate the jump times of the time-changed Levy process. To this end we observe that the time changed Poisson processes driving (S_t) are in fact Cox processes, with intensities proportional to the time-change (Y_t) ([Borovkova and Schmeck, 2017](#)). Given realizations of (Y_t) , we can therefore sample the jump times of the price process by sampling Cox processes with intensities proportional to (Y_t) . This can be done via thinning a dominating Poisson process ([Burnecki et al., 2004](#)). Finally, we simulate the jump sizes (u_i) and (d_i) corresponding to the jump times, which gives us the paths of the process S_t .

Simulation algorithm for the TCDL process:

Output: Price paths of the TCDL process.

1. For each discrete time step, compute increments of CIR process. Recompute any negative increments until no negative increments remain.

2. For each path of the resulting CIR process, draw jump times of the dominating Poisson process.
3. Thin the dominating Poisson process using the criterion in [Burnecki et al. \(2004\)](#), with intensities equal to the state of the CIR process.
4. Assign the jumps of the thinned process as either *up* or *down*, using thinning and the relative intensities of the up and down processes.
5. Draw jump sizes from the specified distributions.
6. Compute the TCDL path as the cumulative sum of jump sizes occurring at the simulated jump times.

The Least Squares Monte Carlo Algorithm was introduced by [Longstaff and Schwartz \(2001\)](#), and a body of later literature exists that extends the original algorithm and studies its properties. We shall therefore not go into these details here, but for the sake of completeness we outline the form of the algorithm that we employ.

The Least Squares Monte Carlo Algorithm:

Input: simulated paths of the TCDL process, and a payoff function
Output: expected value of American barrier option on TCDL process

1. For all price paths that cross the barrier, set the extended price process to zero at and after the knockout time. Use extended price process for rest of algorithm.
2. At the final time point, compute payoff from stopping, set $STOP = 1$ for paths where payoff is positive, and set $CASHFLOW =$ stopping payoff.
3. Move backwards to next time point, and
 - Compute payoff from stopping.
 - For paths where payoff is positive, compute expected continuation value by regressing $CASHFLOW$ next period on polynomial powers of current state of TCDL process.

- set $STOP = 1$ for paths where payoff from stopping is larger than expected continuation value.
 - set $CASHFLOW$ to continuation value where $STOP = 0$, to stopping payoff where $STOP = 1$.
4. Move backwards to next time point, and repeat until first time point is reached.
 5. Compute expected value as average $CASHFLOW$ at the first time point where $STOP = 1$.

Chapter 4

Order anticipation and large traders - evidence from FX markets

GEIR BJØNNES, BI NORWEGIAN BUSINESS SCHOOL

NEOPHYTOS KATHITZIOTIS, UNIVERSITY OF HAMBURG

JO SAAKVITNE, BI NORWEGIAN BUSINESS SCHOOL

Abstract

We provide novel evidence on how a major FX dealer bank adjusts its inventory *before* particularly large customer orders are executed. This is to our knowledge the first time such evidence is presented. We also study pre-trade price dynamics, and show that the observed price patterns differ significantly before large customer trades than at other times, in a manner consistent with pre-trade inventory adjustments. We frame pre-trade inventory adjustments in the context of a simple model, showing how such adjustments follow naturally from profit-maximizing behavior of dealers. Our empirical results indicates that pre-trade price impact is a significant source of indirect trading costs for large dealer-intermediated trades in OTC markets.

4.1 Introduction

We examine how a global foreign exchange liquidity provider adjusts its inventory immediately *before* certain large customer orders are booked. Such trading behaviour by liquidity providers is sometimes labelled "frontrunning" or "pre-hedging". Anticipatory trading strategies are forbidden in many jurisdictions and marketplaces.¹ FX markets are however subject to much less regulation and oversight than many other asset markets, which makes them an ideal setting for a study of anticipatory trading strategies. We use a unique data set from a global foreign exchange dealer bank, and we are able to provide direct evidence on the existence and extent of anticipatory trading practices. This is to our knowledge the first time such evidence is presented. We also provide a novel method to validate the time stamps of financial trade data.

Our analysis consists of two main parts. First, we examine the dealer bank's inventory adjustment in the time period surrounding trades that are especially susceptible to frontrunning. These are particularly large trades performed in a manner involving direct interaction between dealer and customer. We find that, for such trades, there is a strongly significant *pre-trade* inventory adjustment, up to 30 minutes before the trade time. We do not find this kind of adjustment for ordinary, non-susceptible trades.

Second, we examine price dynamics before susceptible trades. We find that price dynamics differ significantly before susceptible trades than before other trades. These extraordinary price dynamics are consistent with a price impact stemming from the bank's observed inventory adjustments.

We illustrate the simple economic intuition of anticipatory trading in a small model. The model shows how profit maximization by the bank leads to anticipatory inventory adjustments and corresponding pattern in pre-trade price dynamics.

Figure 4.1 shows our main result on anticipatory inventory adjustment. The figure shows the mean accumulated inventory change of the bank in the 30 minutes before and after the booking times of trades that are particularly susceptible to frontrunning. The "susceptible trades" in this figure, and throughout our study, are trades with a size larger than

¹For example, in 2018 a US court sentenced the former head of FX cash trading at HSBC to two years in prison for his involvement in the frontrunning of a large currency order placed by an oil company - see e.g. *Ex-HSBC FX Trader Sentenced to 2 Years, Sent Directly to Prison*, Bloomberg Markets, April 26th 2018.

25 million EUR (approximately 30 million USD)² that are done either directly between the bank and the customer (e.g. phone, chat, email, fax) or on the bank's proprietary single bank trading platform (SBP). These are the trades that most likely present a with both an opportunity and the incentive to engage in anticipatory trading, something that we discuss in more detail below. The red bars of the figure shows how the bank's inventory adjusts *before* the susceptible trade is booked - in the 30 minutes before the bank books a large sale to (purchase from) a customer, we on average see the bank buying (selling) around 12 million EUR. We do not find evidence of a similar pattern for ordinary non-susceptible trades. A statistical analysis of this pattern is performed in Section 4.5.

The rest of this paper is organized as follows. In the remainder of this section we describe and discuss anticipatory trading strategies in more detail, and place the current paper in the context of existing literature. In Section 4.2 we convey the economic intuition of these trading practices, and why they are forbidden in many markets, through a simple sequential trading model. We then describe our unique data set (1.3), our empirical method (4.4) and our results (4.5). The Appendix contains descriptive statistics.

4.1.1 Anticipatory trading


We define anticipatory trading as any trading strategy that relies on certain or uncertain knowledge about other participant's future trading.

It is important to note that this definition does not necessarily entail a manipulative or otherwise illegal trading practice; if for example one market participant is able to infer from publicly available market data that another participant is in the middle of executing a large parent order via trade splitting, it would be rational but not illegal for the detecting participant to engage in anticipatory trading (see [Van Kervel and Menkveld \(2017\)](#) for an example).

Anticipatory trading by a dealer would however typically be considered illegal if it involves the use of privileged client information about a future trade to engage in trading

²Market practitioners have indicated that sizes of 25m and above are considered as large trades requiring the use of targeted risk management activities as for example the manual intervention by dealers. While we use 25m for our analysis, our results are robust with other potentially relevant thresholds too, as shown in Tables ?? and ?? in the Appendix.

Figure 4.1: The figure shows accumulated inventory change from 30 minutes before a susceptible trade is booked to 30 minutes after. Inventory change is computed as the cumulative sum of signed traded quantity, excluding all susceptible trades. To reconcile buys and sells in the same plot, we sign trades as follows: when the susceptible trade is a buy, all other sell trades are signed negative and other buy trades are signed positive. When the susceptible trade is a sell, all other sell trades are signed positive and other buy trades are signed negative.



intended to benefit the dealer. For example, traders at HSBC were found guilty of fraud in a US court after using the knowledge about an upcoming large customer order to engage in self-serving trading at the expense of their client.³ Before conducting the large customer order, traders at HSBC took up a proprietary position that benefited from the price impact of the customer trade. This benefit came at the expense of the customer, an oil company, as it got a worse price on its trade than what it would have gotten if HSBC

³Former Global Head of HSBC's Foreign Exchange Cash-Trading Found Guilty of Orchestrating Multimillion-Dollar Front-Running Scheme, US Department of Justice, October 23rd 2017.

had refrained from frontrunning.

We can not in any certain manner discern between legal and illegal anticipatory trading using our data. As such, we do not make any claims as to the legality of trading patterns uncovered in our analysis. The trading patterns we document can very well be caused by a rational and presumably legal positioning by the bank in anticipation of predictable events, not involving any misuse of confidential customer information. Nor is there always a clear line separating nefarious trading strategies from the ordinary ones; it is for example common for traders to exchange gossip and rumours about what is going on in the market, and trading in order to position one's inventory based on such rumours may or may not be illegal, depending on the specific circumstances. The distinction between legal and illegal forms of anticipatory trading in the context of US legislation is explored in e.g. [Scopino \(2014\)](#) and [Markham \(1988\)](#).

Anticipatory trading is not a risk-free trading strategy, whether it is based on certain or uncertain knowledge about the future actions of customers or other participants. These trading strategies are based on taking up a position in order to benefit from the price impact of an future trade, but price impact can often be several orders of magnitude lower than price volatility, meaning that any profit generated by an anticipatory strategy can easily be wiped out by non-related movements in the market price.

In our analysis, we focus on anticipatory trading around certain trades, termed "susceptible trades". These are trades of at least 25 million EUR that takes place either via direct bilateral communication with the customer (e.g. chat, phone, fax or email) on through the bank's single bank platform (SBP). We choose to focus on large trades because the incentive for engaging in anticipatory trading is much stronger than before smaller trades; the total price impact generated by a trade is increasing in the size of the trade, and the expected profit generated by anticipatory trading comes from this price impact. Our results are robust to choosing other thresholds than 25 million EUR. We focus on direct and SBP trades because we believe the dealer has more opportunity to frontrun these trades: An order that is negotiated over phone clearly provides more opportunity for frontrunning, as does an order that is placed well in advance via email. Since large orders placed on the SBP triggers human interaction and bilateral negotiation, we deem also these trades to be susceptible to frontrunning.

Although there is a large existing literature on frontrunning by intermediaries, there is little in the way of direct evidence. Early literature focused on the behaviour of so-called

“dual traders” at the Chicago Mercantile Exchange: [Fishman and Longstaff \(1992\)](#) find evidence of frontrunning, but [Chakravarty and Li \(2003\)](#) use CFTC audit trail data and arrive at the opposite conclusion. [Cai \(2003\)](#) use audit trail transaction data to examine the trading behaviour of market makers in the Treasury bond futures market around the collapse of Long-Term Capital Management in 1998. He finds evidence that market makers engaged in frontrunning, trading for their own accounts 1-2 minutes *before* executing customer orders originating from LTCM. [Comerton-Forde and Tang \(2007\)](#) find some evidence of frontrunning by a subset of brokers at the Toronto Stock Exchange, but [Anand and Subrahmanyam \(2008\)](#) investigates the same market and concludes there is no evidence in favour of front running. [Chaturvedula et al. \(2015\)](#) finds strong evidence of frontrunning before bulk trades in the Indian stock market.

The theoretical literature on frontrunning have focused on welfare implications. [Danthine and Moresi \(1998\)](#) studies frontrunning in an extension of the model from [Kyle \(1985\)](#), and concludes that frontrunning has either no or positive consequences for welfare. [Bernhardt and Taub \(2008\)](#) shows how frontrunning can introduce serial correlation in order flow. [Sannikov and Skrzypacz \(2014\)](#) develops an equilibrium trading model where liquidity providers engage in frontrunning of large investors, and finds that this in general is welfare improving.

[Harris \(2003\)](#) points out that when dealers are informed about the common trading intentions of their clients they may enter into profitable strategies by trading ahead of client orders or by moving prices against the anticipated trade. The ability of dealers to anticipate orders is dependent on market structure and more common in markets where dealers operate under a dual role [Röell \(1990\)](#); [Fishman and Longstaff \(1992\)](#). Firstly, in their capacity as market makers, dealers will receive a fee to intermediate trades between end-customers and secondly they may trade on their own account. Through the intermediation of client trading dealers are better positioned to infer the future direction of a client order. This in turn creates an incentive for them to free-ride on order flow information ([Pagano and Röell, 1992](#)). Because trading is viewed as a zero sum game, the trading costs of market participants will be the trading gains of dealers. Order anticipation raises transaction costs for market participants ([Friederich and Payne, 2014](#)) for two reasons. Firstly when a dealer buys (sells) in anticipation of a client order in the same direction, he will move prices upwards (downwards) causing prices to overshoot. Secondly, this trading activity by the dealer withdraws liquidity from the market when particular market participants need it the most.

Understanding why these issues are relevant for our study requires some perspective of how OTC markets operate. On most exchanges (like the one for equities), liquidity providers move first, when they submit quotes for all market participants to see, which means that the counterparty has not been determined when those quotes are disseminated. In OTC markets in contrast, the liquidity demander moves first by requesting quotes from a dealer. The difference in sequence creates a critical asymmetry in pre-trade anonymity: when quoting prices, market makers in auction markets do not know, and OTC dealers do know, the identity and the common trading intentions of their counterparties. Such information is important for several reasons. Knowledge of a trader's identity may help a dealer induce information about future trades i.e. the size and direction of order flow. Consider for example the rational behavior of a corporate customer that requires liquidity in a certain currency pair at pre-defined frequency (perhaps every Friday at 2pm for weekly liquidity / financial reporting purposes). Similar to corporate clients, institutional investors tend to execute large transactions that relate to portfolio rebalancing activities at known frequencies (i.e. month or quarter end). Because in these instances, clients will require larger amounts of liquidity than usual, they will be exposed to higher transaction costs if the dealer decides to engage in anticipatory trading.

4.2 Model

We model a situation where a dealer bank has knowledge of a customer trade before the actual transaction is formalized. This sort of situation arises for example when trades are negotiated over phone or electronic chat. Another common example is customers that follow a fixed or highly predictable trading pattern (for example a customer who buys a large amount of currency every Friday at 1 pm). In our model, the dealer has certain knowledge of the coming trade. In actuality it seems reasonable to assume that the dealer has only a forecast or an informative signal about the size and directions of future customer trades. The focus of our model is solely to convey a simple intuition, and we therefore abstract from this uncertain.

The time line of the model is the following. We refer to the time when the dealer first learns of the coming trade as the *arrival time*. The time when the trade is actually done is the *transaction time*. We refer to the trading activity done by the dealer in the time interval between arrival and transaction as *anticipatory trading* or *frontrunning*, and the

trading activity done after the transaction time as *inventory management*. Figure 4.2 illustrates.

Figure 4.2: Timeline



We model this situation in an n -period sequential trade-model. In each trading period the dealer submits a given order volume that he wishes to trade, and market clearing price results. We take as given that the market clearing price is a downwards sloping linear function of the net volume that our dealer wants to transact, meaning that is a price impact from the dealer's trades. The model we present here is therefore not an equilibrium model - we only model one side of the market (the dealer), and the price impact of trading is taken as exogenous. A more involved theoretical analysis can be found in [Saakvitne \(2018\)](#), in which a partial equilibrium model is presented where the market learns about the dealers actions, and price impact is endogenous, essentially extending the classic model by [Kyle \(1985\)](#).

The dealer knows beforehand which trading round the large customer trade will occur, as well as the size and direction of the trade. This removes an important source of uncertainty likely present in real-world considerations. We do however model uncertainty relating to exogenous movements in the market price. We do not allow the dealer to take up a long-lived speculative position, meaning that our dealer bank starts and ends the model with zero inventory.⁴

The following example illustrates the simple intuition of the model.

⁴ [Osler et al. \(2016\)](#) analyze a similar model, where the dealer banks is allowed to take up a position of a certain maximum size.

Example. Suppose that the initial price of asset A is \$100, and that the model covers three trading rounds. The dealer knows that a customer will sell him 9 units of asset A in trading round two, and that therefore that he will pay his customer the round 2-market price (he does not know what the round 2-market price will actually be). We denote the volume traded by the dealer in the first trading round by x_1 , and assume a linear pricing rule with a price impact coefficient of 0.1:

$$p_1 = 100 + .1x_1$$

Similarly, we have $p_2 = p_1 + .1x_2$ and $p_3 = p_2 + .1x_3$.

Since the dealer is constrained to end period three with zero inventory, and will buy 9 units from his customer in period 2, he must therefore sell a net amount of 9 units over the course of the three trading rounds.

Consider first the strategy where he does no trading in round one and three, but immediately resell the 9 units the he buys from the customer in trading round 2. The market clearing price in round two will be $100 - .1*9 = \$99.10$. The dealer receives this price from his selling in the interdealer market, and the customer receives this price from the dealer. Thus, the dealer makes no profit.

Consider now instead the strategy where the dealer sells 3 units in trading round one. He receives $100 - .1*3 = \$99.70$ on these units. In round two he sells 6 units, for a price of $99.70 - .1*6 = 99.10$. The customer receives $\$99.10 * 9 = \891.90 from the dealer, while the dealer earns $99.70*3 + 99.10*6 = \$893.70$. We see that the dealer makes a profit of \$ 1.80.

When the dealer makes a profit, logic requires that someone else must have made a loss. It is the customer who loses whatever the dealer profits, because the customer receives a worse price on his trade than what he would have done if the dealer did not frontrun him. Here we see the rationale for why frontrunning is forbidden in many markets - the dealer is profiting from his information about his customer's future trading interest, at the expense of the customer. In certain cases this would be deemed a misuse of inside information by market regulators.

Let us now take the perspective of the dealer, and ask whether we can do even better. It turns out that there exists a unique optimal trading strategy, which involves selling *more* than the total customer order in periods one and two, and buying back the remainder in period three: the optimal thing to do is to sell 6 units in period 1, sell 6 units in period 2, and buy 3 units in period three.

We now derive a formal model, building on the intuition of the previous example. Suppose the price in period t can be written as a linear function of the previous price, the volume traded and a white noise-term ε_t ,

$$p_t = p_{t-1} + \lambda x_t + \varepsilon_t \quad (4.1)$$

There are a total for N trading rounds, and the dealer knows that in period k he will buy y units from a customer. The model extends to the case where the customer is buying by allowing y to take negative values.

Denote the inventory of the dealer by X_t . The dealer's trades are denoted $(x_n, n = 1, \dots, N)$. Since the dealer is not allowed to take up a long-lived speculative position, we have $X_0 = 0$ and the constraint

$$\sum_{n=1}^N x_n = y \quad (4.2)$$

The profit of the dealer is simply the income he earns on his own trades in the interdealer market, less the price he pays to his customer:

$$\pi(x_1, \dots, x_N) = \sum_{n=1}^N p_n x_n - p_k y \quad (4.3)$$

The dealer in our model maximizes a mean-variance type objective function:

$$U(\pi) = \mathbb{E}[\pi | y] - \gamma \text{var}(\pi | y) \quad (4.4)$$


The general model is solved by numerical methods, but the simple case with a risk neutral dealer and three trading periods ($N = 3, k = 2, \gamma = 0$) is straightforward to solve analytically, and we do so in the next section.

4.2.1 Risk neutral dealer in a three-period model

When the dealer is risk neutral ($\gamma = 0$), he maximizes expected profit. When in addition there are only three trading periods, one before and one after the transaction time, the model is straightforward to solve: one inserts the pricing function and the inventory assumption into the expected profit function, and solves the first order condition of a maximum:

$$\frac{\partial \mathbb{E}[\pi]}{\partial x_1} = 0, \quad \frac{\partial \mathbb{E}[\pi]}{\partial x_2} = 0$$

Figure 4.3: Timeline of the three-period model



Solving the resulting system of two equations yields

$$x_1^* = -\frac{2}{3}y, \quad x_2^* = -\frac{2}{3}y, \quad x_3^* = +\frac{1}{3}y$$

This optimal strategy has three notable features, the first two of which also carries over to the general model.

First, it is optimal for the dealer to start his trading before the large customer transaction is finalized (i.e. before the transaction time). One way to interpret this is that the dealer becomes informed once he learns about the coming customer order - because there is a price impact in this market, the dealer has a non-zero expectation about future short-term price changes. Trading on such information is the very definition of frontrunning, as discussed in Section 4.1.

Second, we see from the fact that $|x_1| + |x_2| > y$ that the dealer is *overtrading*, meaning that he trades more than the total customer order in the two first periods taken together. He then offloads his remaining inventory in the final trading period.

Third, the optimal trading strategy of the dealer does not depend on the price impact parameter λ . It is however straightforward to verify that his expected profit increase linearly in λ . It is the existence of a price impact from trading that enables the frontrunning strategy. It follows that the dealer has stronger incentives to frontrun when the price impact is high. For the same reason the dealer would ideally like that the price impact is high when he has the opportunity to frontrun his customer.

4.2.2 The general model

We solve the general n -period model with risk aversion by applying a sample average approximation to attain the dealer's optimal solution.^{5 6}

Figure 4.4 show model-predicted inventory for the dealer in the trading periods before and after the transaction time k , for a unit customer order. The effect of the customer order itself is not included in these figures, in order to make them comparable to the corresponding empirical plot in Figure 4.1.

As shown in Figure 4.4a, when the dealer is risk neutral and has 30 trading periods, he uses the trading periods before the transaction time to gradually build up an inventory of seven times the size of the actual customer order. His inventory reaches a maximum at the transaction time, and then linearly declines until he satisfy the zero-net position constraint when the model ends.

Figures 4.4b-4.4d show the three effects of risk aversion:

First, the amount of overtrading done by the dealer declines rapidly in the degree of risk aversion. When $\gamma = 0.01$ the dealer builds up a total inventory around twice the size of the customer order, while for $\gamma = 0.1$ and $\gamma = 1$ the total position is just barely larger than the customer order.

Secondly, the rate of inventory accumulation before the transaction time is also rapidly decreasing in the degree risk of risk aversion. When risk aversion is high ($\gamma = 1$), the


⁵ The main idea of SAA methods is to turn a stochastic optimization into a deterministic one. This is achieved by first fixing a large sample S of the random variables involved and approximating theoretical moments such as expectation and variance by their sample counterparts. One then performs a numerical optimization routine on this fixed sample, and notes that the specific solution to the deterministic problem converges to the solution of the stochastic problem as the sample size of S increases. See [Homem-de Mello and Bayraktan \(2014\)](#) for a review of the method. In the solution of the general model laid out in Section 4.2.2, we have set the parameters to $N = 30$ trading rounds, where the transaction time k is in round 15. The price impact parameter λ is set to 0.3 and the white noise-term in the pricing equation is Gaussian iid with unit variance. We use a sample size of 5000 draws of the random sequence $(\epsilon_n, n = 1, 2, \dots, N)$

⁶Semi-closed form solutions to a related model is given in [Osler et al. \(2016\)](#). If one is willing to use a continuous-time formulation, the model can be made analytically tractable by an application of Ito's formula, and one can then proceed to solve via deterministic optimal control theory - please contact the author for more details on this.

dealer engages in just a tiny amount of frontrunning, and only in the trading period right before the transaction time.

The third effect of risk aversion relates to the unwinding of the dealers proprietary position, meaning how he reduces his position back to neutral and satisfies his zero net-position constraint. The larger the risk aversion coefficient, the faster is the unwinding.

Figure 4.4: Model-predicted inventory change for various degrees of risk aversion. $N = 30, k = 15, \varepsilon \sim N(0, 1), y = 1$.



4.3 Data

Our data come from the largest OTC market, spot FX, where trading averages USD 1.7 trillion per day (Bank of International Settlements, 2016). The dataset contains all spot deals of a top-10 forex dealer bank in EUR/USD (Euromoney FX survey, 2012) over more than three months of trading, from 2 January 2012 to 20 April 2012 (68 trading days). The sample includes 471,844 transactions. EUR/USD is the most liquid currency pair with a daily turnover of roughly USD 400 billion (Bank of International Settlements, 2016).

Our data set is unique since it combines several customer and interdealer trading venues at high frequency. We are not aware of any other data set in the literature that combines all these features. For each transaction, the data provides the following information: currency pair, date and time stamp of the trade (to the second), transaction price, quantity traded, sign of trade (buy or sell), initiating party, portfolio within the bank to which each trade is assigned, counterparty ID, trading venue (e.g. Hotspot), prime broker (if any), and markup. Customer markups are measured in pips (\$0.0001 per euro) relative to the prevailing core price. The core price reflects information on price and depth from several trading platforms such as EBS, Reuters and Hotspot. Our bank is a major FX market maker so it sets prices in the vast majority trades. The data contains trades with all types of customers such as retail brokers, hedge funds, real asset managers, corporates, and small banks regarded as customers (ranked below top 50 by Euromoney FX Survey, 2012). In addition, the data also includes all trades with FX trading banks (ranked top 50).

These data are well suited for examining frontrunning. Most data sets focus exclusively on dealer-to-customer trades (D2C) or dealer-to-dealer trades (D2D), and provide much less details than our dataset. For instance, Evans and Lyons (2002) and Bjonnes and Rime (2005) focus only on the D2D trades, while Evans and Lyons (2005) and Froot and Ramadorai (2005) focus on D2C trades. The OTC datasets that do include D2C trades often provide no information about customers (Green et al., 2006), or they group customers into a few broad types, such as financial and commercial (Evans and Lyons, 2005). Our data, by contrast, provides individual customer identifiers and we can divide customers into several customer categories. With data covering all trades, including both the customer and interdealer market, we can also track inventory changes. In order to study frontrunning, this is an important feature of the dataset.

The data are also well-suited for analysing today's FX market since they identify over 20 trading venues; most other FX dataset are limited to interdealer platforms (Evans and Lyons, 2002) or they predate the market's fragmentation (Osler et al., 2011). Our empirical analysis is carried out on a trade-by-trade basis and customer-by-customer basis and fully exploits the customer identities in our data. Other features of the same data set are explored in Bjonnes et al. (2017).

Further descriptive statistics on the main data set is reported in Section 4.A of the Appendix.

The second source of information in our analysis is provided by Reuters and contains quotes, depth and transaction data for the time period covered by the proprietary dataset of our bank (frequency is 1/10 sec.). We calculate midquotes (and depth) using the data from Reuters in order to estimate trading costs at different times.

4.4 Methodology

We report two main pieces of analysis in this paper. The analysis examines how the dealer bank adjusts its inventory before and after susceptible customer trades, with a particular emphasis on the anticipatory adjustment. The second analysis examines price dynamics before the same trades. In addition, in Section 4.6 we describe and report a third piece of analysis, aimed at verifying the validity of the time stamps in our data. This analysis can be seen as a robustness check.

The model of Section 4.2 has predictions on both pre-trade inventory adjustments and pre-trade price dynamics. If the bank is engaging in anticipatory trading, we would expect to see a significant build-up of inventory before certain susceptible trades are booked. As a result of this build-up, market prices will be driven down before a customer sells, and up before a customer buys, for these particular trades.

4.4.1 Inventory analysis

The aim of the inventory analysis is to examine changes to the banks inventory position after, and more importantly, *before* the transaction times a trade susceptible to frontrunning is booked.

In Section 4.1 we discussed which characteristics of a trade that gives a liquidity provider opportunity and incentives to engage in anticipatory trading. For the purpose of the inventory analysis we operationalize these characteristics as the set of trades satisfying the following two criteria:

1. Trade size of 25 million Euro or more
2. Trade occurs either on the bank's single bank platform (SBP) or via direct bilateral communications with the customer (e.g. phone, electronic chat, email etc.)

The set of trades satisfying these criteria are referred to as "susceptible trades". The inventory analysis is performed as follows. Let $f(t)$ be the net inventory change of the bank over the time period $(t, t + 1)$, where t is a time stamp given in whole minutes (for example 2012-01-01 14:15:00). We denote the **pre-trade inventory change** at time lead x as $y_x(t)$,

$$y_x(t) = \begin{cases} (+1) \sum_{n=1}^x f(t-n) & \text{if customer buys} \\ (-1) \sum_{n=1}^x f(t-n) & \text{if customer sells} \end{cases} \quad (4.5)$$

Defined in this way, $y_0(t)$ is the inventory change over the minute where susceptible trade took place, $y_5(t)$ is the inventory change over the 5 minutes before the trade took place, and so on. The sign change for trades where the customer sold to the bank ensures that we can treat buys and sells in the same manner later in our analysis.

Let $T = (t_1, t_2, \dots, t_N)$ be the trade times for the N susceptible trades, rounded down to the nearest whole minute. If the bank is adjusting its inventory in anticipation of large customer trades, we expect $y_x(t)$ to be positive for the susceptible trade times $t \in T$. In Table 4.1, we report the mean \bar{y}_x for all susceptible trades, meaning

$$\bar{y}_x = N^{-1} \sum_{t=t_1}^{t_N} y_x(t) \quad (4.6)$$

To test the statistical significance of average pre-trade inventory changes \bar{y}_x , we seek to compare the inventory change for susceptible trades with the corresponding change for ordinary trades. To this end, we compute a bootstrap distribution of average pre-trade inventory changes for *any* trade, not only the susceptible trades.

The bootstrap distribution is built as follows. Let $H = (h_1, h_2, \dots, h_K)$ be the time stamps of all trades. Table 4.4 shows there are $K = 275,251$ trades in our sample. One run of

the bootstrap procedure involves drawing $N = 74$ trade times at random from H , and computing \bar{y}_x per equations (4.5) and (4.6) for these trades. We perform a large number of such runs to build up a bootstrap distribution of \bar{y}_x under the null hypothesis that pre-inventory changes for susceptible trades are no different than from any other trade. A bootstrap p-value can then be found by comparing \bar{y}_x computed on the susceptible trade set against the quantiles of the distribution $F_{\bar{y}_x}^*$.

4.4.2 Pre-trade price dynamics

If the bank is engaging in frontrunning, we expect there to be a pattern to market prices in the period preceding especially large trades, caused by the price impact of the bank trading in anticipation of the large trade. We investigate the presence of such a pattern by measuring changes to the midpoint price in the Reuters reference data.

Let $m(t)$ be the reference midpoint price at time t . For a given trade occurring at time t , we define the **pre-trade price impact** at time lead x , $z_x(t)$ as

$$z_x(t) = \begin{cases} m(t) - m(t-x) & \text{if customer buys} \\ m(t-x) - m(t) & \text{if customer sells} \end{cases}$$

In our empirical investigation we consider time leads x of 0, 1, 5, 10, 15, 20, 25 and 30 minutes.

The pre-trade price impact z is used as dependent variable in a regression analysis. As independent variables of interest, we consider trade volume (unsigned), and a dummy variable indicating whether the customer is a broker/bank or not ⁷ In addition we use daily volatility as a control variable. Daily volatility has been standardized for inter-pretability. For a given time lag x we estimate the regression equation

$$z_x(t) = \alpha + \beta_1 \text{VOLUME}(t) + \beta_2 \text{VOLUME}^2(t) + \beta_3 D_{NONBANK}(t) + \beta_4 \text{VOL} + e_t \quad (4.7)$$

We estimate these regressions using both ordinary least squares and, as a robustness check, using quantile regressions. The regressions are estimated time lags (x) of 5, 10 and 30 minutes.

⁷The non-banks in this sample would be real money investors, hedge funds and large corporations.

We expect frontrunning to be more pronounced for large trades, since these trades provide a stronger incentive for anticipatory trading. We also expect frontrunning to be more pronounced for trades with clients that have a more predictable trading pattern, such as corporation and funds. We shall therefore test whether the coefficients β_1 and β_3 are significantly different from zero, and interpret rejection of the null hypothesis as evidence for frontrunning.

4.5 Results

Both the inventory analysis and price change analysis provide evidence consistent with frontrunning, and in particular the model in Section 4.2.

4.5.1 Pre-trade inventory effects

There is a pre-trade inventory adjustment before large trades involving human interaction that is significant at the 1% level, for time leads of 10, 15 and 30 minutes. This is seen in Table 4.1. The red bars of Figure 4.1 illustrates the build-up of this inventory adjustment over the 30 minutes before large trades. Tables 4.11 and 4.12 contains the output of the same bootstrapping procedure used to produce Table 4.1, but for size thresholds of 20 and 30 million EUR respectively. These tables show that our results are robust to changing the size threshold of 25 million EUR used throughout this paper.

The pre-trade inventory adjustments seen from Table 4.1 and Figure 4.1 are consistent with both economic intuition and the frontrunning model presented in Section 4.2. We do however not find robust evidence for the overtrading and subsequent inventory reversal predicted by the model. The lack of such evidence can be seen visually in Figure 4.1.

There can be several reasons for the lack of evidence in favour of overtrading and subsequent inventory reversal. One possibility is that the bank does not engage in this aspect of frontrunning, perhaps because the strategy is deemed as involving unacceptable levels of reputational risk, or as being non-compliant with internal or external regulations. There would also be market risk associated with the overtrading predicted by our simple model - the bank would be exposed to adverse movements in the spot rate above and beyond what is implied by ordinary market making. The theoretical model in fact shows

Percentile	Time lead			
	5 min	10 min	15 min	30 min
1%	-1.78	-2.02	-2.06	-2.72
5%	-1.12	-1.37	-1.49	-2.15
50%	-0.04	0.16	0.47	0.87
95%	1.42	1.86	2.62	4.96
99%	1.87	2.42	2.65	6.07
\bar{y}_x	0.46	2.73	5.25	7.79

Table 4.1: Mean pre-trade inventory changes and quantiles of corresponding bootstrap distribution. The table shows the pre-trade inventory adjustments, as defined by equations (4.5) and (4.6), for all trades (percentiles) and susceptible trades (\bar{y}_x), at various time leads (5-30 minutes). At leads of 10 minutes and above, the pre-trade inventory adjustments before susceptible trades ($\bar{y}_{10}, \bar{y}_{15}, \bar{y}_{30}$) are outside the 99% bootstrap percentile observed for non-susceptible trades, which is interpretable as statistical significance at the 1% level. Units: million EUR.

(Figure 4.4) that the extent of overtrading decrease in aversion to market risk.


Another possibility for why we are not seeing overtrading is that significant parts of the bank's inventory adjustment is done through either other currency pairs or through internal netting. There are facts supportive of this interpretation: the mean size of the susceptible trades in our sample is 72 million Euro, while the mean *post-trade* inventory adjustment hovers around 8 to 10 million Euro for various time lags (shown in the blue bars of Figure 4.1).

4.5.2 Pre-trade price effects

Figure 4.5 shows average and median trading costs in two samples - small trades (4.5a) and large trades (4.5b) - and illustrates our results on pre-trade price effects. Small and large trades are those with trade size below and above 25 million EUR respectively.

The figure shows two interesting things. First, we see that the market midpoint typically changes in the customers disfavor in the 30 minutes preceding a large trade. If the average customer is benchmarking his trading costs against the prevailing market midpoint at the time of the trade ($x = 0$), his average trading cost would be a little less

Figure 4.5: Customer spread versus market midprice: the plots show the mean and median difference between the actual price paid by the customer and the market midpoint price at the Reuters trading platform. The market midpoint price is measured at the time of the trade (right-most part of x-axis) and for various time leads up to 30 minutes before the actual trade (left-most part of x-axis). Units: pips.



than 6 pips. If, on the other hand, the customer benchmarks his trading costs against the prevailing market midpoint 20 minutes before the actual trade ($x = 20$), he would find his trading costs to be a little over 8 pips. The difference of 2 pips can be defined as the frontrunning costs.


A second observation from figure 4.5 is that the average customer seems to be timing his trade in a particular manner, evident from the upwards slope of costs for the full sample of trades. The average buy trade is placed after a period of falling market (mid)price, the average sell trade is placed after a period of raising market (mid)price. This pattern is intuitive – when the market price goes up it triggers profitable trading opportunities for market participants waiting to sell (for example related to hedging motives, information, speculation and so on), while when the market price goes down it triggers profitable trading opportunities for market participants waiting to buy. This pattern is interesting because it implies that for the average trade, pre-trade price impact is negative, which makes the positive pre-trade price impact for large trades all the more striking.

Table 4.2 summarize the results from estimating the regression equation (4.7). We see that there is a significant size effect to pre-trade price impact, of approximately 0.15 pips per million EUR. The coefficient on squared volume means that the relation between size and pre-trade price impact is concave; the relation is plotted in Figure 4.6. The positive coefficient on the dummy "Non-liq.provider" indicates that there is significantly higher pre-trade price impact when the customer is not a broker or a bank, which in our data means that it is classified as a real money investor, a hedge fund or a multinational corporation.

The coefficient on hourly price volatility has a negative sign, and the effect is stronger the longer the time lead. This is consistent with the model we have presented, and we would argue also with economic intuition; when price volatility is higher, the market risk involved in taking up an anticipatory position is also higher. A risk averse trader will therefore take up a smaller anticipatory position, all else being equal, and he will begin his anticipatory trading closer to the expected trade time.

The estimated parameters in the quantile regression, shown in Table 4.3, have the same signs and are of a similar magnitude as the OLS parameters. This rules out that our results are heavily influenced by outliers.

Figure 4.6: Estimated relation between trade size and pre-trade price impact, in pips.



	30 minutes	10 minutes	5 minutes
Intercept	-3.4458*** (0.0621)	-2.3224*** (0.0388)	-1.7427*** (0.0283)
Size	0.1474*** (0.0264)	0.1814*** (0.0165)	0.1728*** (0.0120)
Size squared	-0.0006* (0.0003)	-0.0011*** (0.0002)	-0.0011*** (0.0001)
Non-liq.provider	2.0084*** (0.5163)	1.6117*** (0.3227)	1.1780*** (0.2347)
Hourly vola.	-2.2424*** (0.0612)	-1.2644*** (0.0382)	-0.7880*** (0.0278)
Num. obs.	84852	84849	84854

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 4.2: Pre-trade price impact regression models - OLS. The dependent variable is the difference between price paid by the customer and the market midpoint price at the Reuters trading platform, where the midpoint has been measured at 5,10 and 30 minutes. The unit of the dependent variable is pips, which is the fourth decimal of the exchange rate. "Size" and "Size squared" are continuous variables containing the size of the trade in million EUR. "Non-liq.provider" is a dummy variable encoding whether the customer is a broker/bank or not (the non-bank customers are real money investors, hedge funds and large corporations).

4.6 Robustness - are the timestamps reliable?

4.6.1 Time lag analysis - A method for validating time stamps in trade data

The time lag analysis is a methodology developed to ensure that the time stamps in our data are correct; and in particular that large trades are not booked with a time lag. This methodology can also be applied in other contexts than the current paper, so we find it useful to formulate our problem in more general terms.

The general problem looks as follows. Suppose we have records of N financial trades

	30 minutes	10 minutes	5 minutes
Intercept	-4.3901*** (0.0906)	-2.5215*** (0.0486)	-1.6797*** (0.0306)
Size	0.1170** (0.0386)	0.1746*** (0.0199)	0.1553*** (0.0118)
Size squared	-0.0003 (0.0003)	-0.0010*** (0.0002)	-0.0010*** (0.0001)
Non-liq.provider	2.2420*** (0.5318)	1.7520*** (0.2599)	0.9848*** (0.2139)
Hourly vola.	-3.0928*** (0.0936)	-1.4405*** (0.0557)	-0.8053*** (0.0356)
Num. obs.	84852	84849	84854

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 4.3: Pre-trade price impact regression models - Quantile regression (median). The dependent variable is the difference between price paid by the customer and the market midpoint price at the Reuters trading platform, where the midpoint has been measured at 5,10 and 30 minutes. The unit of the dependent variable is pips, which is the fourth decimal of the exchange rate. "Size" and "Size squared" are continuous variables containing the size of the trade in million EUR. "Non-liq.provider" is a dummy variable encoding whether the customer is a broker/bank or not (the non-bank customers are real money investors, hedge funds and large corporations).

of a particular currency pair. Sells and buys are distributed evenly amongst the recorded trades. Each record i consists of a transaction price v_i and a time stamp t_i . Each trade is timestamped when it is entered into the electronic bookkeeping system, a process called "booking" the trade. However, the researcher suspects that some time goes by from when a trade is actually done (the transaction time of Figure 4.2) until it is booked. In other words, the trades might be booked with a time lag.

Suppose further that we have available a time series $m(t)$ of market mid point prices of the currency pair in question. The transactions in the trade records are thought to take place at the market mid price plus/minus a markup. The markup may vary from trade to trade, but is believed to be zero on average, as there is a roughly equal number of buys

and sells in the transaction records.

We model this situation by assuming that each trade i actually took place at time $t_i - \theta$, where θ is a non-negative parameter to be estimated from the data. We also assume that each transaction takes place at the prevailing market price plus a Gaussian white noise component (the markup):

$$v_i = m(t_i - \theta) + \varepsilon_i$$

$$\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$$

Further, we model the market mid price of the currency pair $m(t)$ by a Brownian motion $B(t)$:

$$m(t) = \sigma B(t)$$

To estimate θ , define for each i the **pricing error** $\Delta_i(x)$:

$$\begin{aligned} \Delta_i(x) &= v_i - m(t_i - x) \\ &= m(t_i - \theta) + \varepsilon - m(t_i - x) \\ &= \sigma[B(t_i - \theta) - B(t_i - x)] + \varepsilon \\ &\sim \mathcal{N}(0, \sigma^2|\theta - x| + \sigma_\varepsilon^2) \end{aligned}$$

It is clear that $\mathbb{E}[\Delta_i(x)] = 0$ for all x . It is also clear that the variance of $\Delta_i(x)$ is at a minimum when $x = \theta$, so that

$$\theta = \arg \min_x \mathbb{E} [(\Delta_i(x))^2]$$

We approximate the expectation by the sample mean, which gives the estimator $\hat{\theta}$:

$$\hat{\theta} := \arg \min_x \frac{1}{N} \sum_{i=1}^N \Delta_i(x)^2 \quad (4.8)$$


The estimator $\hat{\theta}$ is an extremum estimator, a class of estimators which are known to be asymptotically consistent also under much weaker conditions than what we have assumed here (Newey and McFadden, 1994).

Figure 4.7a maps a simulation of the root mean squared error (RMSE) $\sqrt{N^{-1} \sum_{i=1}^N (\Delta_i(x))^2}$ as a function of the time lag x . The figure shows the mean of 1000 simulations, where the true time lag θ has been set to 15 units of time (minutes). Notice that the RMSE equals the standard deviation of the markup (σ_ε) when the minimum is attained. Figure 4.7b shows the RMSE when the true time lag is zero ($\theta = 0$).


The sum of N squared standard normal random variables is known to be chi-squared distributed with N degrees of freedom (Abramowitz and Stegun, 1964, p. 940). It follows that for a fixed x , the sum $\sum_{i=1}^N \Delta_i^2$ is a scaled chi-square:

$$\sum_{i=1}^N \Delta_i^2 = (\sigma^2 |\theta - x| + \sigma_\varepsilon^2) \chi_N^2 \quad (4.9)$$

(a) Theoretical RMSE - 15 minutes lag
($\sigma = .8, \sigma_\varepsilon = 1, \theta = 15$)



(b) Theoretical RMSE - zero time lag
($\sigma = .8, \sigma_\varepsilon = 2, \theta = 0$)



4.6.2 Results of time lag analysis

Applying the time lag estimation method to our data, we find strong evidence in favor of an accurate time stamp on the average trade. For the full sample, the minimizer of equation 4.8 is zero, and the RMSE increase in accordance with the theoretical distribution laid out in equation 4.9 both backward and forward in time. Figure 4.8a show the


empirical RMSE for all $\sim 257,000$ trades in the sample, for time lags and leads ranging from 1 to 30 minutes.

For the sample of trades involving human interaction (i.e. marked as "direct"), the evidence is in favor of a time lag between 0 and 3 minutes, with a point estimate (minimizer of equation 4.8) of 2 minutes. Figure 4.8b show the empirical RMSE for these $\sim 4,300$ trades.

In our analyses of both inventory change and price dynamics, we find evidence of pre-trade effects on a 30 minute horizon. The time lag analysis of this section show that it is highly unlikely that these effects are caused by a time lag in the bank's booking of trades in its internal systems.

(a) Empirical RMSE - all trades
($n = 257,241$)

(b) Empirical RMSE - direct trades only
($n = 4,316$)



4.A Descriptive statistics and additional robustness

Table 4.4: Absolute volume by trade size. Unit: Thousand EUR.

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.	N
All trades	0.00	23	120	500	500	180000	257241
Small trades	0.00	20	100	270	500	1000	239314
Medium trades	1000.00	1500	2000	3200	4000	25000	17800
Large trades	25000.00	30000	39000	58000	57000	180000	127

Table 4.5: Number of trades by venue.

	Direct	SBP	MBP 1	MBP 2	API
All trades	4316	75329	60796	8755	108045
Small trades	3223	68854	55979	6782	104476
Medium trades	1019	6442	4814	1966	3559
Large trades	74	33	3	7	10

Table 4.6: Mean trade size by venue. Unit: Thousand EUR.

	Direct	SBP	MBP 1	MBP 2	API
All trades	2848.65	647.65	523.35	1254.76	227.84
Small trades	232.04	344.11	379.81	459.36	152.43
Medium trades	6078.35	3704.57	2170.32	3850.44	2353.51
Large trades	72338.63	37238.70	36141.54	42868.17	31592.00

Table 4.7: Customer spread by trade size. Unit: Pips.

	1st Qu.	Median	Mean	3rd Qu.	N
All trades	-1.20	0.20	0.14	1.50	257241
Small trades	-1.20	0.20	0.14	1.50	239314
Medium trades	-1.50	0.10	0.00	1.70	17800
Large trades	-5.00	0.30	5.30	6.30	127

Table 4.8: Price change ($p_{t-30} - p_t$) from t=-30 to 0 minutes by trade size. Unit: Pips.

	1st Qu.	Median	Mean	3rd Qu.	N
All trades	-12.50	-2.00	-2.43	7.50	257241
Small trades	-13.00	-2.50	-2.50	7.50	239314
Medium trades	-11.00	-1.00	-0.93	10.00	17800
Large trades	-6.50	3.00	2.10	9.30	127

Table 4.9: Customer spread by venue. Unit: Pips.

	1st Qu.	Median	Mean	3rd Qu.	N
All	-1.20	0.20	0.14	1.50	257241
Direct	-1.50	0.50	-0.20	2.50	75329
SBP	-1.00	0.20	0.26	1.50	60796
MBP 1	-1.30	0.10	0.04	1.50	108045
MBP 2	-1.30	0.30	0.14	1.70	8755
API	-1.30	0.10	0.12	1.50	4316

Table 4.10: Price change ($p_{t-30} - p_t$) from t=-30 to 0 minutes by venue. Unit: Pips.

	1st Qu.	Median	Mean	3rd Qu.	N
All	-12.50	-2.00	-2.43	7.50	257241
Direct	-10.00	-0.50	-0.93	8.50	75329
SBP	-13.00	-3.00	-3.40	6.50	60796
MBP 1	-12.00	-1.50	-1.30	9.50	108045
MBP 2	-10.00	-1.50	-2.20	6.50	8755
API	-13.00	-2.00	-2.40	7.50	4316

Percentile	Time lead			
	5 min	10 min	15 min	30 min
1%	-1.96	-3.21	-3.22	-3.86
5%	-1.56	-1.74	-2.44	-2.95
50%	0.08	0.41	0.26	0.29
95%	1.80	1.97	2.98	3.76
99%	2.48	2.38	3.51	4.15
\bar{y}_x	6.78	7.97	7.26	10.85

Table 4.11: Mean pre-trade inventory changes and quantiles of corresponding bootstrap distribution when the size threshold for "susceptible trades" is set to 20 million EUR. The table shows the pre-trade inventory adjustments, as defined by equations (4.5) and (4.6), for all trades (percentiles) and susceptible trades (\bar{y}_x), at various time leads (5-30 minutes). At leads of 10 minutes and above, the pre-trade inventory adjustments before susceptible trades ($\bar{y}_{10}, \bar{y}_{15}, \bar{y}_{30}$) are outside the 99% bootstrap percentile observed for non-susceptible trades, which is interpretable as statistical significance at the 1% level. Units: million EUR.

Percentile	Time lead			
	5 min	10 min	15 min	30 min
1%	-2.83	-3.88	-4.30	-5.92
5%	-2.61	-2.68	-3.35	-4.49
50%	-0.32	0.22	0.24	-0.11
95%	1.82	2.45	3.47	4.89
99%	2.57	2.78	4.11	6.75
\bar{y}_x	11.06	12.63	12.36	16.85

Table 4.12: Mean pre-trade inventory changes and quantiles of corresponding bootstrap distribution when the size threshold for "susceptible trades" is set to 30 million EUR. The table shows the pre-trade inventory adjustments, as defined by equations (4.5) and (4.6), for all trades (percentiles) and susceptible trades (\bar{y}_x), at various time leads (5-30 minutes). At leads of 10 minutes and above, the pre-trade inventory adjustments before susceptible trades ($\bar{y}_{10}, \bar{y}_{15}, \bar{y}_{30}$) are outside the 99% bootstrap percentile observed for non-susceptible trades, which is interpretable as statistical significance at the 1% level. Units: million EUR.

Bibliography

- Abramowitz, Milton and Irene A Stegun**, *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*, Vol. 55, Courier Corporation, 1964.
- Abrantes-Metz, Rosa M., Michael Kraten, Albert D. Metz, and Gim S. Seow**, “Liber manipulation?,” *Journal of Banking & Finance*, 2012, 36 (1), 136–150.
- Admati, Anat R and Paul Pfleiderer**, “A theory of intraday patterns: Volume and price variability,” *The Review of Financial Studies*, 1988, 1 (1), 3–40.
- Anand, Amber and Avanidhar Subrahmanyam**, “Information and the intermediary: Are market intermediaries informed traders in electronic markets?,” *Journal of Financial and Quantitative Analysis*, 2008, 43 (1), 1–28.
- Andersen, Leif BG**, “Efficient simulation of the Heston stochastic volatility model,” 2007.
- Andersen, Torben G and Tim Bollerslev**, “Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns,” *The Journal of Finance*, 1997, 52 (3), 975–1005.
- Angel, James J.**, “Tick size, share prices, and stock splits,” *The Journal of Finance*, 1997, 52 (2), 655–681.
- Applebaum, David**, *Lévy processes and stochastic calculus*, Cambridge university press, 2009.
- Aquilina, Matteo, Ibikunle Gbenga, Vito Mollica, and Tom Steffen**, “Benchmark Regulation and Market Quality,” 2017, 27.
- Arnold, Martin and Daniel Schaefer**, “Banks speed up shift to forex automation,” *Financial Times*, June 2014.

- Banerjee, Snehal and Brett Green**, “Signal or noise? Uncertainty and learning about whether other traders are informed,” *Journal of Financial Economics*, 2015, 117 (2), 398–423.
- Bank of International Settlements**, “Triennial Central Bank Survey of foreign exchange and OTC derivatives markets in 2016,” Technical Report, Bank of International Settlements 2016.
- , “Triennial Central Bank Survey of foreign exchange and OTC derivatives markets in 2016,” Technical Report, Bank of International Settlements 2016.
- Bardorff-Nielsen, Ole E and Albert Shiryaev**, *Change of time and change of measure*, Vol. 21, World Scientific Publishing Company, 2015.
- , **David G Pollard, and Neil Shephard**, “Integer-valued Lévy processes and low latency financial econometrics,” *Quantitative Finance*, 2012, 12 (4), 587–605.
- Baron, Matthew D., Jonathan Brogaard, Björn Hagströmer, and Andrei A. Kirilenko**, “Risk and Return in High-Frequency Trading,” SSRN Scholarly Paper ID 2433118, Social Science Research Network, Rochester, NY 2017.
- Bauwens, Luc and Nikolaus Hautsch**, *Modelling financial high frequency data using point processes*, Springer, 2009.
- Bernhardt, Dan and Bart Taub**, “Front-running dynamics,” *Journal of Economic Theory*, 2008, 138 (1), 288 – 296.
- Bjonnes, Geir and Dagfinn Rime**, “Dealer behavior and trading systems in foreign exchange markets,” *Journal of Financial Economics*, 2005, 75 (3), 571–605.
- Black, Fischer and Myron Scholes**, “The pricing of options and corporate liabilities,” *The Journal of Political Economy*, 1973, pp. 637–654.
- Blocher, Jesse, Rick Cooper, Jonathan Seddon, and Ben Van Vliet**, “Phantom Liquidity and High-Frequency Quoting,” *The Journal of Trading*, 2016, 11 (3), 6–15.
- Borovkova, Svetlana and Maren Diane Schmeck**, “Electricity price modeling with stochastic time change,” *Energy Economics*, 2017, 63, 51–65.
- Breedon, Francis and Paolo Vitale**, “An empirical study of portfolio-balance and information effects of order flow on exchange rates,” *Journal of International Money and Finance*, April 2010, 29 (3), 504–524.

- Brogaard, Jonathan A**, “Minimum quote life and maximum order message-to-trade ratio,” *Techreport, Foresight, Government Office for Science, UK*, 2011, 25.
- Budish, Eric, Peter Cramton, and John Shim**, “The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response,” *The Quarterly Journal of Economics*, 2015, 130 (4), 1547–1621.
- Burnecki, K, W Härdle, and R Weron**, “An introduction to simulation of risk processes, in Encyclopedia of Actuarial Science, eds. J. Teugels, B. Sundt,” 2004.
- Cai, Fang**, “Was there front running during the LTCM crisis?,” 2003.
- Carr, Peter, Hélyette Geman, Dilip B Madan, and Marc Yor**, “Stochastic volatility for Lévy processes,” *Mathematical Finance*, 2003, 13 (3), 345–382.
- Cartea, Álvaro, Sebastian Jaimungal, and Jamie Walton**, “Foreign exchange markets with last look,” 2015.
- Chaboud, Alain, Benjamin Chiquoine, Erik Hjalmarsson, and Clara Vega**, “Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market,” *The Journal of Finance*, 2014, 69 (5), 2045–2084.
- Chakravarty, Sugato and Kai Li**, “An examination of own account trading by dual traders in futures markets,” *Journal of Financial Economics*, 2003, 69 (2), 375–397.
- Chaturvedula, Chakrapani, Nupur Pavan Bang, Nikhil Rastogi, and Satish Kumar**, “Price manipulation, front running and bulk trades: Evidence from India,” *Emerging Markets Review*, 2015, 23 (Supplement C), 26 – 45.
- Chochrane, Jim**, “The New WM/Reuters fixing methodology doesn’t fix everything,” Technical Report, ITG 2015.
- Choi, Jin, Kasper Larsen, and Duane Seppi**, “Information and trading targets in a dynamic market equilibrium,” 2017.
- Comerton-Forde, Carole and Kar Mei Tang**, “Anonymity, frontrunning and market integrity,” *The Journal of Trading*, 2007, 2 (4), 101–118.
- Cont, Rama**, “Volatility clustering in financial markets: empirical facts and agent-based models,” in “Long memory in economics,” Springer, 2007, pp. 289–309.

- Copeland, Thomas E. and Dan Galai**, “Information effects on the bid-ask spread,” *Journal of Finance*, 1983, 38 (5), 1457–1469.
- Danthine, Jean-Pierre and Serge Moresi**, “Front-running by mutual fund managers: a mixed bag,” *Review of Finance*, 1998, 2 (1), 29–56.
- de Mello, Tito Homem and Güzin Bayraksan**, “Monte Carlo sampling-based methods for stochastic optimization,” *Surveys in Operations Research and Management Science*, 2014, 19 (1), 56–85.
- Degryse, Hans, Frank de Jong, and Vincent van Kervel**, “Does order splitting signal uninformed order flow,” Technical Report, working paper 2014.
- Ding, Deng, Ning ying Huang, and Jing ya Zhao**, “An efficient algorithm for Bermudan barrier option pricing,” *Applied Mathematics-A Journal of Chinese Universities*, 2012, 27 (1), 49–58.
- DuCharme, Michael**, “Does Trading at the Fix fix FX?,” May 2013.
- Duffie, Darrell and Jeremy C. Stein**, “Reforming LIBOR and Other Financial Market Benchmarks,” *Journal of Economic Perspectives*, May 2015, 29 (2), 191–212.
- **and Jeremy C Stein**, “Reforming LIBOR and other financial market benchmarks,” *Journal of Economic Perspectives*, 2015, 29 (2), 191–212.
- **and Piotr Dworzak**, “Robust benchmark design,” Technical Report, National Bureau of Economic Research 2014.
- **and _**, “Robust benchmark design,” resreport NBER Working Paper No. w20540, National Bureau of Economic Research 2018.
- , – , **and Haoxiang Zhu**, “Benchmarks in search markets,” *The Journal of Finance*, 2017, 72 (5), 1983–2044.
- Evans, Martin DD**, “Forex trading and the WMR fix,” *Journal of Banking & Finance*, 2017.
- , “Forex trading and the WMR Fix,” *Journal of Banking & Finance*, 2018, 87, 233–247.
- Fang, Fang and Cornelis W Oosterlee**, “Pricing early-exercise and discrete barrier options by Fourier-cosine series expansions,” *Numerische Mathematik*, 2009, 114 (1), 27.

- FEMR**, “Fair and Effective Markets Review - Final Report,” June 2015.
- Financial Stability Board**, “Foreign Exchange Benchmarks,” Technical Report, Foreign Exchange Benchmark Group, Financial Stability Board 2014.
- Fishman, Michael J and Francis A Longstaff**, “Dual trading in futures markets,” *The Journal of Finance*, 1992, 47 (2), 643–671.
- Foster, F Douglas and S Viswanathan**, “Strategic trading when agents forecast the forecasts of others,” *The Journal of Finance*, 1996, 51 (4), 1437–1478.
- Foucault, Thierry, Roman Kozhan, and Wing Wah Tham**, “Toxic arbitrage,” *Review of Financial Studies*, 2016, p. hhw103.
- Friederich, Sylvain and Richard Payne**, “Trading anonymity and order anticipation,” *Journal of Financial Markets*, 2014, 21, 1–24.
- Frino, Alex, Elvis Jarnecic, and Thomas H McInish**, “The Option Value of the Limit Order Book,” in “Advances In Quantitative Analysis Of Finance And Accounting: Essays in Microstructure in Honor of David K Whitcomb,” World Scientific, 2006, pp. 57–71.
- FSB**, “Foreign Exchange Benchmarks - Final Report,” September 2014.
- , “Report on progress in implementing the September 2014 recommendations,” October 2015.
- Furfine, Craig H**, “Banks as monitors of other banks: Evidence from the overnight federal funds market,” *The Journal of Business*, 2001, 74 (1), 33–57.
- Furse, C, A Haldane, C Goodhart, D Cliff, JP Zigrand, K Houstoun, O Linton, and P Bond**, “The future of computer trading in financial markets,” Technical Report, Working paper, Foresight, Government Office for Science, London (UK). URL <http://www.bis.gov.uk/foresight/our-work/projects/current-projects/computer-trading> 2011.
- Global Foreign Exchange Committee**, “FX Global Code,” Technical Report, Global Foreign Exchange Committee 2017.
- Goettler, Ronald L, Christine A Parlour, and Uday Rajan**, “Equilibrium in a dynamic limit order market,” *The Journal of Finance*, 2005, 60 (5), 2149–2192.

- Gong, Xiaoli and Xintian Zhuang**, “American option valuation under time changed tempered stable Lévy processes,” *Physica A: Statistical Mechanics and its Applications*, 2017, 466, 57–68.
- Harris, Larry**, *Trading and exchanges: Market microstructure for practitioners*, Oxford University Press, USA, 2003.
- Harris, Lawrence E**, “Minimum price variations, discrete bid–ask spreads, and quotation sizes,” *The Review of Financial Studies*, 1994, 7 (1), 149–178.
- Hendershott, Terrence, Charles M Jones, and Albert J Menkveld**, “Does algorithmic trading improve liquidity?,” *The Journal of Finance*, 2011, 66 (1), 1–33.
- ITG**, “Examining the WM/Reuters London Close through the Prism of Foreign Exchange Transaction Cost Analysis,” April 2014.
- Ito, Takatoshi and Masahiro Yamada**, “Was the Forex Fixing Fixed?,” Technical Report, National Bureau of Economic Research 2015.
- **and** —, “Did the Reform Fix the London Fix Problem?,” *Journal of International Money and Finance*, 2017.
- , **Richard K. Lyons, and Michael T. Melvin**, “Is there private information in the FX market? The Tokyo experiment,” *The Journal of Finance*, 1998, 53 (3), 1111–1130.
- Jeanblanc, Monique, Marc Yor, and Marc Chesney**, *Mathematical methods for financial markets*, Springer Science & Business Media, 2009.
- Jorgensen, Kjell, Johannes Atle Skjeltop, and Bernt Arne Ødegaard**, “Throttling hyperactive robots - Order to trade ratios at the Oslo Stock Exchange,” 2016.
- Judd, Kenneth L**, *Numerical methods in economics*, MIT press, 1998.
- Jurgilas, Marius and Filip Žikeš**, “Implicit intraday interest rate in the UK unsecured overnight money market,” *Journal of Financial Intermediation*, 2014, 23 (2), 232–254.
- Kervel, Vincent Van and Albert J Menkveld**, “High-frequency trading around large institutional orders,” 2017.
- Killeen, William P, Richard K. Lyons, and Michael J. Moore**, “Fixed versus flexible: Lessons from EMS order flow,” *Journal of International Money and Finance*, 2006, 25 (4), 551–579.

- King, Michael R., Carol L. Osler, and Dagfinn Rime**, “The market microstructure approach to foreign exchange: Looking back and looking forward,” *Journal of International Money and Finance*, November 2013, 38, 95–119.
- Kraenzlin, Sébastien and Thomas Nellen**, “Daytime is money,” *Journal of Money, Credit and Banking*, 2010, 42 (8), 1689–1702.
- Kyle, Albert S.**, “Continuous auctions and insider trading,” *Econometrica*, 1985, pp. 1315–1335.
- Lehmann, Bruce N.**, “Notes for a contingent claims theory of limit order markets,” in “Advances In Quantitative Analysis Of Finance And Accounting: Essays in Microstructure in Honor of David K Whitcomb,” World Scientific, 2006, pp. 39–56.
- Leifeld, Philip**, “texreg: Conversion of Statistical Model Output in R to LATEX and HTML Tables,” *Journal of Statistical Software, Articles*, 2013, 55 (8), 1–24.
- Longstaff, Francis A and Eduardo S Schwartz**, “Valuing American options by simulation: a simple least-squares approach,” *Review of Financial studies*, 2001, 14 (1), 113–147.
- Mancini, Loriano, Angelo Ranaldo, and Jan Wrampelmeyer**, “Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums,” *The Journal of Finance*, 2013, 68 (5), 1805–1841.
- Markham, Jerry W.**, “Front-Running-Insider Trading under the Commodity Exchange Act,” *Cath. UL Rev.*, 1988, 38, 69.
- Marsh, Ian W, Panagiotis Panagiotou, and Richard Payne**, “The WMR Fix and its Impact on Currency Markets,” 2017.
- Melvin, Michael and John Prins**, “Equity hedging and exchange rates at the London 4p.m. fix,” *Journal of Financial Markets*, January 2015, 22, 50–72.
- **and** – , “Equity hedging and exchange rates at the London 4p.m. fix,” *Journal of Financial Markets*, 2015, 22 (C), 50–72.
- **and Xixi Yin**, “Public information arrival, exchange rate volatility, and quote frequency,” *The Economic Journal*, 2000, 110 (465), 644–661.
- Mende, Alexander**, “09/11 on the USD/EUR foreign exchange market,” *Applied Financial Economics*, 2006, 16 (3), 213–222.

- Menkhoff, Lukas**, “The noise trading approach—questionnaire evidence from foreign exchange,” *Journal of International Money and Finance*, 1998, 17 (3), 547–564.
- Menkveld, Albert J.**, “High frequency trading and the new market makers,” *Journal of Financial Markets*, 2013, 16 (4), 712–740.
- Michelberger, Patrick Steffen and Jan Hendrik Witte**, “Foreign exchange market microstructure and the WM/Reuters 4 pm fix,” *The Journal of Finance and Data Science*, 2016, 2 (1), 26–41.
- Mooney, Attracta**, “Passive funds grow 230% to \$6tn,” *Financial Times*, May 2016.
- MSCI**, “MSCI Index Calculation Methodology - MSCI Equity Indexes,” April 2018.
- Newey, Whitney K and Daniel McFadden**, “Large sample estimation and hypothesis testing,” *Handbook of Econometrics*, 1994, 4, 2111–2245.
- O’Hara, Maureen**, *Market microstructure theory*, Vol. 108, Blackwell Publishers Cambridge, MA, 1995.
- Øksendal, Bernt Karsten and Agnes Sulem**, *Applied stochastic control of jump diffusions*, Vol. 498, Springer, 2005.
- Oomen, Roel**, “Last look,” *Quantitative Finance*, 2016, pp. 1–14.
- Osler, Carol, Alasdair Turnbull et al.**, “Dealer trading at the fix,” Technical Report 2016.
- Osler, Carol L, Alexander Mende, and Lukas Menkhoff**, “Price discovery in currency markets,” *Journal of International Money and Finance*, 2011, 30 (8), 1696–1718.
- Pagano, Marco and Ailsa Röell**, “Auction and dealership markets: what is the difference?,” *European Economic Review*, 1992, 36 (2-3), 613–623.
- Peiers, Bettina**, “Informed traders, intervention, and price leadership: A deeper view of the microstructure of the foreign exchange market,” *The Journal of Finance*, 1997, 52 (4), 1589–1614.
- Peskir, Goran and Albert Shiryaev**, *Optimal stopping and free-boundary problems*, Springer, 2006.

- Pragma**, “Trading the 4pm FX Fix: A Window into Market Impact,” August 2015.
- Reuters, Thomson**, “WM/Reuters FX Benchmarks - Spot & Forward Rates Methodology Guide,” November 2017.
- Röell, Ailsa**, “Dual-capacity trading and the quality of the market,” *Journal of Financial Intermediation*, 1990, 1 (2), 105–124.
- Roşu, Ioanid**, “A dynamic model of the limit order book,” *The Review of Financial Studies*, 2009, 22 (11), 4601–4641.
- Saakvitne, Jo**, “‘Banging the Close’: Price Manipulation or Optimal Execution?,” SSRN Scholarly Paper ID 2753080, Social Science Research Network, Rochester, NY September 2016.
- , “‘Banging the Close’: Price Manipulation or Optimal Execution?,” 2016.
- , “Trading strategies and information flow around price benchmarks,” Technical Report, BI Norwegian Business School 2018.
- Sannikov, Yuliy and Andrzej Skrzypacz**, “Dynamic trading: Price inertia, front-running and relationship banking,” Technical Report, Working paper 2014.
- Schaumburg, Ernst**, “Has Automated Trading Promoted Efficiency in the FX Spot Market?,” Technical Report, New York Federal Reserve March 2014.
- Scopino, Gregory**, “The (Questionable) Legality of High-Speed ‘Pinging’ and ‘Front Running’ in the Futures Markets,” 2014.
- Shoham, Yoav and Kevin Leyton-Brown**, *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*, Cambridge University Press, 2008.
- Swishchuk, Anatoliy**, *Change of time methods in quantitative finance*, Vol. 32, Springer, 2016.
- Tankov, Peter and Ekaterina Voltchkova**, “Jump-diffusion models: a practitioners guide,” *Banque et Marchés*, 2009, 99 (1), 24.
- Vaughan, Gavin Finch Liam and Ambereen Choudhury**, “Traders Said to Rig Currency Rates to Profit Off Clients,” *Bloomberg.com*, June 2013.
- Werner, Ingrid M, Yuanji Wen, Barbara Rindi, Francesco Consonni, and Sabrina Buti**, “Tick size: theory and evidence,” 2015.

Yamada, Masahiro and Takatoshi Ito, “The forex fixing reform and its impact on cost and risk of forex trading banks,” *Finance Research Letters*, 2017, 21, 157–162.

Zanger, Daniel, “Quantitative error estimates for a least-squares Monte Carlo algorithm for American option pricing,” *Finance and Stochastics*, 07 2013, 17.

Zanger, Daniel Z, “Convergence of a least-squares Monte Carlo algorithm for American option pricing with dependent sample data,” *Mathematical Finance*, 2018, 28 (1), 447–479.

Series of Dissertations

2019

2/2019 *Jo Albertsen Saakvitne*
Essays on Market Microstructure

1/2019 *Per-Magnus Moe Thompson*
All you need is love?: Investigating leadership from leaders' attachment experiences in close relationships

2018

15/2018 *Stefania Sardo*
Questioning normality: A study of obduracy and innovation in the O&G industry

14/2018 *Ingvild Müller Seljeseth*
The Inevitable Thucydides's Trap?: How Hierarchical Instability and Threat Influences Leaders' Openness to Inputs from Others

13/2018 *Daniela Carmen Cristian*
The Utility of Hedonics Beyond Pleasure: Positive Outcomes of Hedonic Consumption

12/2018 *Daniel Øgård Kinn*
Essays in Econometrics

11/2018 *Ragnar Enger Juelsrud*
Essays in Macro-Finance

10/2018 *Bisrat Agegnehu Misganaw*
On entrepreneurial teams and their formation in science-based industries

9/2018 *Martin Blomhoff Holm*
Consumption

8/2018 *Helene Lie Røhr*
Essays on decision making dynamics in politics and consumer choice

7/2018 *Hoang Ho*
Are Human Resource Management (HRM) Systems Good or Bad for Employee Well-Being?: An Investigation of the Well-being Paradox from the Mutual Gains and Critical Perspectives

6/2018 *Rannveig Røste*
Innovation in Public Services: Wicked Problems and Multi-layered Solutions

5/2018 *Mathias Hansson*
The Complex Simplicity of Patterns of Action: Organizational Routines as Source of Adaptation and Performance

4/2018 *Mariia Koval*
Antecedents and Consequences of Unplanned Alliance Terminations

3/2018 *Maximilian Rohrer*
Three Essays in Financial Economics

2/2018 *Viacheslav Lurkov*
The role of alliance networks for firms' geographic scope and performance: A structural embeddedness perspective

1/2018 *Huy-Vu Nguyen*
Wealth Inequality

2017

14/2017 *Mirha Suangic*
Aspirations and Daring Confrontations: Investigating the Relationship between Employees' Aspirational Role-Identities and Problem-Oriented Voice

13/2017 *Beniamino Callegari*
A heterodox theoretical interpretation

12/2017 *Sepideh Khayati Zahiri*
Essays on Predictive Ability of Macroeconomic Variables and Commodity Prices

11/2017 *Tonje Hungnes*
Reorganizing healthcare services: Sensemaking and organizing innovation

10/2017 *Eileen Fugelsnes*
From backstage to consensus: A study of the Norwegian pension reform process

9/2017 *Knut-Eric Neset Joslin*
Experimental Markets with Frictions

8/2017 *Sumaya AlBalooshi*
Switching between Resources: Psychosocial Resources as Moderators of the Impact of Powerlessness on Cognition and Behavior

7/2017 *Zongwei Lu*
Three essays on auction theory and contest theory

6/2017 *Frode Martin Nordvik*
Oil Extraction and The Macroeconomy

5/2017 *Elizabeth Solberg*
Adapting to Changing Job Demands: A Broadcast Approach to Understanding Self-Regulated Adaptive Performance and Cultivating it in Situated Work Settings

4/2017 *Natalia Bodrug*
Essays on economic choices and cultural values

3/2017 *Vegard Høghaug Larsen*
Drivers of the business cycle: Oil, news and uncertainty

2/2017 *Nikolay Georgiev*
Use of Word Categories with Psychological Relevance Predicts Prominence in Online Social Networks

1/2017 *Sigmund Valaker*
Breakdown of Team Cognition and Team Performance: Examining the influence of media and overconfidence on mutual understanding, shared situation awareness and contextualization

2016

- 12/2016 *Xiaobei Wang*
Strategizing in Logistics Networks: The role of Logistics Service Providers as mediators and network facilitators
- 11/2016 *Prosper Ameh Kwei-Narh*
A mid-range theory of monitoring behaviors, shared task mental models, and team performance within dynamic settings
- 10/2016 *Anton Diachenko*
Ownership type, performance, and context: Study of institutional and industrial owners
- 9/2016 *Ranvir S. Rai*
Innovation in Practice: A Practice-Theoretical Exploration of Discontinuous Service Innovations
- 8/2016 *Gordana Abramovic*
Effective Diversity Management on the Line - Who and How? On the role of line managers in organisations with a diverse workforce
- 7/2016 *Mohammad Ejaz*
Why do some service innovations succeed while others fail? A comparative case study of managing innovation processes of two internet-based aggregation financial services at Finn.no (Penger.no)
- 6/2016 *Asmund Rygh*
Corporate governance and international business: Essays on multinational enterprises, ownership, finance and institutions
- 5/2016 *Chen Qu*
Three Essays on Game Theory and Patent-Pool Formation
- 4/2016 *Arash Aloosh*
Three Essays in International Finance
- 3/2016 *John-Erik Mathisen*
Elaborating and Implemental Mindsets in Business-related Behavior: An Investigation of the Interface between Cognition and Action. How goals and plans emerge and relate to cognition and action in business.
- 2/2016 *Øyvind Nilsen Aas*
Essays in industrial organization and search theory
- 1/2016 *Dominique Kost*
Understanding Transactive Memory Systems in Virtual Teams: The role of integration and differentiation, task dependencies and routines
- 2015**
- 8/2015 *Andreea Mitrache*
Macroeconomic Factors and Asset Prices – An Empirical Investigation
- 7/2015 *Lene Pettersen*
Working in Tandem: A Longitudinal Study of the Interplay of Working Practices and Social Enterprise Platforms in the Multinational Workplace
- 6/2015 *Di Cui*
Three Essays on Investor Recognition and Mergers & Acquisitions

- 5/2015 *Katja Maria Hydle*
Cross border practices: Transnational Practices in Professional Service Firms
- 4/2014 *Ieva Martinkenaite-Pujanauskienė*
Evolutionary and power perspectives on headquarters-subsiary knowledge transfer: The role of disseminative and absorptive capacities
- 3/2015 *Tarje Gaustad*
The Perils of Self-Brand Connections: Consumer Response to Changes in Brand Image
- 2/2015 *Jakob Utgård*
Organizational form, local market structure and corporate social performance in retail
- 1/2015 *Drago Bergholt*
Shocks and Transmission Channels in Multi-Sector Small Open Economies
- 2014**
- 13/2014 *Nam Huong Dau*
Asset Pricing with Heterogeneous Beliefs and Portfolio Constraints
- 12/2014 *Vegard Kolbjørnsrud*
On governance in collaborative communities
- 11/2014 *Ignacio García de Olalla López*
Essays in Corporate Finance
- 10/2014 *Sebastiano Lombardo*
Client-consultant interaction practices: Sources of ingenuity, value creation and strategizing
- 9/2014 *Leif Anders Thorsrud*
International business cycles and oil market dynamics
- 8/2014 *Ide Katrine Birkeland*
Fire Walk with Me: Exploring the Role of Harmonious and Obsessive Passion in Well-being and Performance at Work
- 7/2014 *Sinem Acar-Burkay*
Essays on relational outcomes in mixed-motive situations
- 6/2014 *Susanne H.G. Poulsson*
On Experiences as Economic Offerings
- 5/2014 *Eric Lawrence Wiik*
Functional Method as a Heuristic and Research Perspective: A Study in Systems Theory
- 4/2014 *Christian Enger Gimsø*
Narcissus and Leadership Potential - The measurement and implications of narcissism in leadership selection processes
- 3/2014 *Mehrad Moeini-Jazani*
When Power Has Its Pants Down: Social Power Increases Sensitivity to Internal Desires
- 2/2014 *Yuriy Zhovtobryukh*
The role of technology, ownership and origin in M&A performance
- 1/2014 *Siv Staubo*
Regulation and Corporate Board Composition

2013

- 9/2013 *Bjørn Tallak Bakken*
Intuition and analysis in decision making: On the relationships between cognitive style, cognitive processing, decision behaviour, and task performance in a simulated crisis management context
- 8/2013 *Karl Joachim Breunig*
Realizing Reticulation: A Comparative Study of Capability Dynamics in two International Professional Service Firms over 10 years
- 7/2013 *Junhua Zhong*
Three Essays on Empirical Asset Pricing
- 6/2013 *Ren Lu*
Cluster Networks and Cluster Innovations: An Empirical Study of Norwegian Centres of Expertise
- 5/2013 *Therese Dille*
Inter-institutional projects in time: a conceptual framework and empirical investigation
- 4/2013 *Thai Binh Phan*
Network Service Innovations: Users' Benefits from Improving Network Structure
- 3/2013 *Terje Gaustad*
Creating the Image: A Transaction Cost Analysis of Joint Value Creation in the Motion Picture Industry
- 2/2013 *Anna Swärd*
Trust processes in fixed-duration alliances: A multi-level, multi-dimensional, and temporal view on trust
- 1/2013 *Sut I Wong Humborstad*
Congruence in empowerment expectations: On subordinates' responses to disconfirmed experiences and to leaders' unawareness of their empowerment expectations
- 2012**
- 9/2012 *Robert Buch*
Interdependent Social Exchange Relationships: Exploring the socially embedded nature of social exchange relationships in organizations
- 8/2012 *Ali Faraji-Rad*
When the message feels right: Investigating how source similarity enhances message persuasiveness
- 7/2012 *Marit Anti*
Commercial friendship from a customer perspective: Exploring social norm of altruism in consumer relationships and self-interest-seeking behavior
- 6/2012 *Birgit Helene Jevnaker*
Vestiges of Design-Creation: An inquiry into the advent of designer and enterprise relations
- 5/2012 *Erik Aadland*
Status decoupling and signaling boundaries: Rival market category emergence in the Norwegian advertising field, 2000-2010
- 4/2012 *Ulaş Burkay*
The Rise of Mediating Firms: The Adoption of Digital Mediating Technologies and the Consequent Re-organization of Industries

- 3/2012 *Tale Skjølsvik*
Beyond the 'trusted advisor': The impact of client-professional relationships on the client's selection of professional service firms
- 2/2012 *Karoline Hofslett Kopperud*
Well-Being at Work: On concepts, measurement, and leadership influence
- 1/2012 *Christina G. L. Nerstad*
In Pursuit of Success at Work: An Empirical Examination of the Perceived Motivational Climate, Its Outcomes and Antecedents

2011

- 12/2011 *Kjell Jørgensen*
Three Articles on Market Microstructure at the Oslo Stock Exchange (OSE)
- 11/2011 *Siri Valseth*
Essays on the information content in bond market order flow
- 10/2011 *Elisabet Sørfjorddal Hauge*
How do metal musicians become entrepreneurial? A phenomenological investigation on opportunity recognition
- 9/2011 *Sturla Lyngnes Fjesme*
Initial Public Offering Allocations
- 8/2011 *Gard Paulsen*
Betwixt and between: Software in telecommunications and the programming language Chill, 1974-1999
- 7/2011 *Morten G. Josefsen*
Three essays on corporate control
- 6/2011 *Christopher Wales*
Demands, designs and decisions about evaluation: On the evaluation of postgraduate programmes for school leadership development in Norway and England
- 5/2011 *Limei Che*
Investors' performance and trading behavior on the Norwegian stock market
- 4/2011 *Caroline D Ditlev-Simonsen*
Five Perspectives on Corporate Social Responsibility (CSR): An empirical analysis
- 3/2011 *Atle Raa*
Fra instrumentell rasjonalitet til tvetydighet: En analyse av utviklingen av Statskonsults tilnærming til standarden Mål- og resultatstyring (MRS) 1987-2004
- 2/2011 *Anne Louise Koefoed*
Hydrogen in the making - how an energy company organises under uncertainty
- 1/2011 *Lars Erling Olsen*
Broad vs. Narrow Brand Strategies: The Effects of Association Accessibility on Brand Performance

2010

- 8/2010 *Anne Berit Swanberg*
Learning with Style: The relationships among approaches to learning, personality, group climate and academic performance
- 7/2010 *Asle Fagerstrøm*
Implications of motivating operations for understanding the point-of-online-purchase: Using functional analysis to predict and control consumer purchasing behavior
- 6/2010 *Carl J. Hatteland*
Ports as Actors in Industrial Networks
- 5/2010 *Radu-Mihai Dimitriu*
Extending where? How consumers' perception of the extension category affects brand extension evaluation
- 4/2010 *Svanhild E. Haugnes*
Consumers in Industrial Networks: a study of the Norwegian-Portuguese bacalhau network
- 3/2010 *Stine Ludvigsen*
State Ownership and Corporate Governance: Empirical Evidence from Norway and Sweden
- 2/2010 *Anders Dysvik*
An inside story – is self-determination the key? Intrinsic motivation as mediator and moderator between work and individual motivational sources and employee outcomes. Four essays
- 1/2010 *Etty Ragnhild Nilsen*
Opportunities for learning and knowledge creation in practice
- 2009**
- 8/2009 *Erna Senkina Engebretsen*
Transportation Mode Selection in Supply Chain Planning
- 7/2009 *Stein Bjørnstad*
Shipshaped: Kongsberg industry and innovations in deepwater technology, 1975-2007
- 6/2009 *Thomas Hoholm*
The Contrary Forces of Innovation: An Ethnography of Innovation Processes *in the Food Industry*.
- 5/2009 *Christian Heyerdahl-Larsen*
Asset Pricing with Multiple Assets and Goods
- 4/2009 *Leif-Magnus Jensen*
The Role of Intermediaries in Evolving Distribution Contexts: A Study of Car Distribution
- 3/2009 *Andreas Brekke*
A Bumper? An Empirical Investigation of the Relationship between the Economy and the Environment
- 2/2009 *Monica Skjøld Johansen*
Mellom profesjon og reform: Om fremveksten og implementeringen av enhetlig ledelse i norsk sykehusvesen

1/2009

Mona Kristin Solvoll

Televised sport: Exploring the structuration of producing change and stability in a public service institution

2008

7/2008

Helene Loe Colman

Organizational Identity and Value Creation in Post-Acquisition Integration: The Spiralling Interaction of the Target's Contributive and the Acquirer's Absorptive Capacities

6/2008

Fahad Awaleh

Interacting Strategically within Dyadic Business Relationships: A case study from the Norwegian Electronics Industry

5/2008

Dijana Tiplic

Managing Organizational Change during Institutional Upheaval: Bosnia-Herzegovina's Higher Education in Transition

4/2008

Jan Merok Paulsen

Managing Adaptive Learning from the Middle

3/2008

Pingying Zhang Wenstøp

Effective Board Task Performance. Searching for Understanding into Board Failure and Success

2/2008

Gunhild J. Ecklund

Creating a new role for an old central bank: The Bank of Norway 1945-1954

1/2008

Øystein Strøm

Three essays on corporate boards

2007

6/2007

Martha Kold Bakkevig

The Capability to Commercialize Network Products in Telecommunication

5/2007

Siv Marita Fosstenløkken

Enhancing Intangible Resources in Professional Service Firms. A Comparative Study of How Competence Development Takes Place in Four Firms

4/2007

Gro Alteren

Does Cultural Sensitivity Matter to the Maintaining of Business Relationships in the Export Markets? An empirical investigation in the Norwegian seafood industry

3/2007

Lars C. Monkerud

Organizing Local Democracy: The Norwegian Experience

2/2007

Siv Marina Flø Karlsen

The Born Global – Redefined. On the Determinants of SMEs Pace of Internationalization

1/2007

Per Engelseth

The Role of the Package as an Information Resource in the Supply Chain. A case study of distributing fresh foods to retailers in Norway

2006

10/2006

Anne Live Vaagaasar

From Tool to Actor - How a project came to orchestrate its own life and that of others

- 9/2006 *Kjell Brynjulf Hjertø*
The Relationship Between Intragroup Conflict, Group Size and Work Effectiveness
- 8/2006 *Taran Thune*
Formation of research collaborations between universities and firms: Towards an integrated framework of tie formation motives, processes and experiences
- 7/2006 *Lena E. Bygballe*
Learning Across Firm Boundaries. The Role of Organisational Routines
- 6/2006 *Hans Solli-Sæther*
Transplants' role stress and work performance in IT outsourcing relationships
- 5/2006 *Bjørn Hansen*
Facility based competition in telecommunications – Three essays on two-way access and one essay on three-way access
- 4/2006 *Knut Boge*
Votes Count but the Number of Seats Decides: A comparative historical case study of 20th century Danish, Swedish and Norwegian road policy
- 3/2006 *Birgitte Grøgaard*
Strategy, structure and the environment. Essays on international strategies and subsidiary roles
- 2/2006 *Sverre A. Christensen*
Switching Relations - The rise and fall of the Norwegian telecom industry
- 1/2006 *Nina Veflen Olsen*
Incremental Product Development. Four essays on activities, resources, and actors
- 2005**
- 6/2005 *Jon Erland Bonde Lervik*
Managing Matters: Transferring Organizational Practices within Multinational Companies
- 5/2005 *Tore Mysen*
Balancing Controls When Governing Agents in Established Relationships: The Influence of Performance Ambiguity
- 4/2005 *Anne Flagstad*
How Reforms Influence Organisational Practices: The Cases of Public Roads and Electricity Supply Organisations in Norway
- 3/2005 *Erlend Kvaal*
Topics in accounting for impairment of fixed asset
- 2/2005 *Amir Sasson*
On Mediation and Affiliation. A Study of Information Mediated Network Effects in The Banking Industry
- 1/2005 *Elin Kubberød*
Not just a Matter of Taste – Disgust in the Food Domain

2004

- 10/2004 *Sverre Tomassen*
The Effects of Transaction Costs on the Performance of Foreign Direct Investments - An empirical investigation
- 9/2004 *Catherine Børve Monsen:*
Regulation, Ownership and Company Strategies. The Case of European Incumbent Telecommunications Operators
- 8/2004 *Johannes A. Skjeltorp*
Trading in Equity Markets: A study of Individual, Institutional and Corporate Trading Decision.
- 7/2004 *Frank Elter*
Strategizing in Complex Contexts
- 6/2004 *Qinglei Dai*
Essays on International Banking and Tax-Motivated Trading
- 5/2004 *Arne Morten Ulvnes*
Communication Strategies and the Costs of Adapting to Opportunism in an Interfirm Marketing System
- 4/2004 *Gisle Henden*
Intuition and its Role in Strategic Thinking
- 3/2004 *Haakon O. Aa. Solheim*
Essays on volatility in the foreign exchange market
- 2/2004 *Xiaoling Yao*
From Village Election to National Democratisation. An Economic-Political Microcosm Approach to Chinese Transformation
- 1/2004 *Ragnhild Silkoset*
Collective Market Orientation in Co-producing Networks
- 2003**
- 2/2003 *Egil Marstein*
The influence of stakeholder groups on organizational decision-making in public hospitals
- 1/2003 *Joyce Hartog McHenry*
Management of Knowledge in Practice. Learning to visualise competence
- 2002**
- 6/2002 *Gay Bjercke*
Business Landscapes and Mindscapes in People's Republic of China. A Study of a Sino-Scandinavian Joint Venture
- 5/2002 *Thorvald Hærem*
Task Complexity and Expertise as Determinants of Task Perceptions and Performance: Why Technology-Structure Research has been unreliable and inconclusive
- 4/2002 *Norman T. Sheehan*
Reputation as a Driver in Knowledge-Intensive Service Firms:
An exploratory study of the relationship between reputation and value creation in petroleum exploration units
- 3/2002 *Line Lervik Olsen*
Modeling Equity, Satisfaction and Loyalty in Business-to-Consumer Markets

- 2/2002 *Fred H. Strønen*
Strategy Formation from a Loosely Coupled System Perspective. The Case of Fjordland
- 1/2002 *Terje I. Våland*
Emergence of conflicts in complex projects. The role of informal versus formal governance mechanisms in understanding interorganizational conflicts in the oil industry
- 2001**
- 6/2001 *Kenneth H. Wathne*
Relationship Governance in a Vertical Network Context
- 5/2001 *Ming Li*
Value-Focused Data Envelopment Analysis
- 4/2001 *Lin Jiang*
An Integrated Methodology for Environmental Policy Analysis
- 3/2001 *Geir Høidal Bjønnes*
Four Essays on the Market Microstructure of Financial Markets
- 2/2001 *Dagfinn Rime*
Trading in Foreign Exchange Markets. Four Essays on the Microstructure of Foreign Exchange
- 1/2001 *Ragnhild Kvålshaugen*
The Antecedents of Management Competence. The Role of Educational Background and Type of Work Experience
- 2000**
- 1/2000 *Per Ingvar Olsen*
Transforming Economies. The Case of the Norwegian Electricity Market Reform