University of Southampton

Faculty of Social Sciences

Business School

Studies on High Frequency Financial Markets

by

Thong Minh Dao

Thesis for the degree of Doctor of Philosophy

March 2019

University of Southampton

Abstract

Faculty of Social Sciences

Business School

Thesis for the degree of Doctor of Philosophy

Studies on High Frequency Financial Markets

by

Thong Minh Dao

This thesis examines high frequency financial markets in terms of the relationship among financial instruments. Chapter 2 paves the way for subsequent chapters by providing a detailed delineation of high frequency markets, the main context of this work, and discussing various topics such as the significance and popularity of high frequency trading (HFT) as well as its positive and negative impacts on today's markets. Chapter 3, 4 and 5 are three research papers which focus on multiple aspects of the high frequency relationship among financial assets including correlation, lead-lag effects and volatility transmission. Specifically, chapter 3 studies pairs trading, a popular trading strategy based on correlation and designed to exploit related securities. Among other things, this chapter explains why the literature may have consistently underestimated the level of pairs trading profitability and market inefficiency, and then proposes a new trading rule to correct this bias which outperforms the standard rule used by previous papers. On the other hand, chapter 4 analyses the lead-lag relationship between instruments and identifies an important factor that has an impact on this relationship, namely the rate of information arrival. This chapter has been accepted for publication in Quantitative Finance. Finally, chapter 5 investigates the influence of the Brexit referendum, an important political event, on currency markets. This chapter shows that the event has affected the correlation and volatility spillover among exchange rates in a way that suggests a flight to quality and is consistent with the reduced market integration between the UK and the EU due to the UK's decision to leave the EU.

Table of Contents

Table of Figures

Table of Tables

Chapter 3

Chapter 4

Chapter 5

Research Thesis: Declaration of Authorship

Print name:	THONG MINH DAO
Title of thesis:	Studies on high frequency financial markets

I declare that this thesis and the work presented in it are my own and have been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as:

Dao, T.M., McGroarty, F. and Urquhart, A., 2018. Ultra-high-frequency lead–lag relationship and information arrival. *Quantitative Finance*, 18(5), pp.725-735.
Dao, T.M., McGroarty, F. and Urquhart, A., 2019. The Brexit vote and currency markets. *Journal of International Financial Markets, Institutions and Money*, 59, pp.153-164.

Signature:		Date:	27/03/2019

Acknowledgements

I would like to express my sincere gratitude to Prof. Frank McGroarty and Dr. Andrew Urquhart for their continuous and valuable support during my PhD programme. I would also like to acknowledge the use of the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton, in the completion of this work.

Abbreviations

ADR: American Depositary Receipt

ARCH: Autoregressive Conditional Heteroskedasticity

CME: Chicago Mercantile Exchange

DCC: Dynamic Conditional Correlation

ETF: exchange traded fund

EU: European Union

FTSE: Financial Times Stock Exchange

GARCH: Generalised Autoregressive Conditional Heteroskedasticity

HFT: high frequency trading

IOSCO: International Organisation of Securities Commissions

LIBOR: London Interbank Offered Rate

MIDAS: Market Information Data Analytics System

NYSE: New York Stock Exchange

S&P: Standard and Poor's

SEC: Securities Exchange Commission

TSE: Tokyo Stock Exchange

UK: United Kingdom

US: United States

VAR: Vector Autoregression

Chapter 1

Introduction

1.1. Research context

All of the three studies in this thesis are conducted in the novel context of high frequency financial markets where trading activities of market participants and price movements of securities take place at a very high speed. Under the substantial influence of three main factors, the high frequency world has become a major part of my current system of financial markets over time. The first factor is information technology and infrastructure. Because of radical advances in this area recently, it is now possible for a lot of investors and traders to gather and process a large amount of information and then take action based on their analysis in a way that is much faster, easier and less expensive than ever before (Hagströmer and Nordén, 2013). In addition, algorithmic trading in general and high frequency trading (HFT) in particular have increased in popularity as there are more and more traders and trading firms whose normal trading routine has to rely on the support of a variety of advanced computer programmes, at least to some extent but often quite heavily (Sun et al., 2014).

On the other hand, the second factor is the willingness of exchanges and other trading venues to adopt new technological improvements in their daily operation, which is also of great importance. For the most part, the marketplace has become fully automated and thus is able to handle the high level of trading activities of market participants on a regular basis (Jain, 2005, Brogaard et al., 2014). Moreover, O’Hara (2015) places much emphasis on the essential role of trading venues in the high frequency era as well as their substantial contribution to the impressive growth of HFT, especially in more recent periods. Last but not least, the third factor contributing to the rise of high frequency markets can be traced back to a number of policies and regulations which were introduced to the marketplace in the 2000s in order to promote competition among trading locations (O’Hara, 2015). Over time, these regulatory decisions have led to a gradual increase in the degree of market fragmentation and complexity in trading. As a result of the new fragmented and complex environment, HFT is considered by many to be highly useful and valuable due to its extremely high speed and inherent ability to react quickly to changes in the market, which will certainly help traders to not only deal with difficult challenges but also make use of precious opportunities in a more effective and easier way.

Together, the three factors above have resulted in unprecedented trading intensity in my high frequency markets at the moment. A lot of studies have found that high frequency traders constitute a major group of market participants that accounts for a large part of all financial transactions generated every trading day in (i) both equity and other asset classes such as currency and commodity, (ii) spot markets as well as derivative markets including options and futures, and (iii) not only the US but also many other international markets (e.g. Japan and many countries in Europe) (Kirilenko and Lo, 2013, O'Hara, 2014, Benos and Sagade, 2016, Brogaard et al., 2017). According to Aitken et al. (2015), recent data show that it is not unusual for high frequency traders to be responsible for between 50% and 70% of all the trades in financial markets or even as much as 75% on some occasions. Similarly, O'Hara (2014) and O’Hara (2015) argue that 50% or more of the total trading volume of financial instruments can be attributed to this group of traders.

It is interesting that in addition to their active participation in trading, high frequency traders are also well known for their remarkable commitment and willingness to invest heavily in their business without holding back. To be more specific, O’Hara (2015) points out that the infrastructure upgrade required for a tiny improvement in trading speed (i.e. a few milliseconds) may easily cost at least $500 million, which is clear evidence of how seriously HFT is taken as well as how profitable it can potentially be. Nevertheless, in spite of the apparent significance of HFT, the overall impact of high frequency traders on the system of financial markets as a whole is still currently under ongoing debate since the question remains as to whether their impact is positive or negative in general (Jarnecic and Snape, 2014). As a matter of fact, there is a lot of work to be done to have a better understanding of high frequency markets, which is really necessary and beneficial for participants, exchanges and policy makers alike (Sun et al., 2014).

1.2. Scope and overview

Chapter 2 offers a more detailed discussion about high frequency financial markets in order to further familiarise readers with the main setting of this work. This chapter starts by describing the practice of HFT, its significance and popularity in the market as well as different types of high frequency traders and HFT activities. Then the chapter continues by reviewing the relevant literature on the positive and negative influence of HFT on (i) the marketplace in terms of informational efficiency, volatility and risks of trading, and liquidity and transaction costs, (ii) the interaction of market participants including how high frequency traders interact with their computer programmes or other traders, and (iii) trading venues and market regulators. Finally, the chapter ends by pointing out a few issues of interest related to information and data in the high frequency environment such as properties and patterns of the data, interpretation of the data, and difficulties in dealing with the data. With the background knowledge provided in Chapter 2, the thesis proceeds to the next three chapters which are quantitative in nature and constitute a multifaceted examination of the high frequency relationship among financial instruments.

Chapter 3 investigates the relationship among securities from a trading point of view. More specifically, this chapter focuses on a commonly used trading strategy which aims to take advantage of the high degree of correlation between related assets, namely pairs trading (or also referred to as statistical arbitrage). The chapter puts forward an explanation as to why the standard trading rule employed in earlier studies on this particular strategy is unnecessarily strict and not optimal, which is likely to result in a consistent underestimation of the profit potential of pairs trading and the extent of market inefficiency. Based on this explanation, I propose an alternative rule to correct the bias and then evaluate the new rule against the standard rule in terms of trading performance, with and without an adjustment for risks, in the context of the gold ETF market using ultra-high frequency data (i.e. tick data). To make the results more reliable, I carry out the analysis for both the entire sample period and sub-samples, and adopt an appropriate measure to properly take into account the possibility of a data mining (also called data snooping) bias. If the performance of the alternative rule is significantly better than that of the standard rule, it is safe to say that the former allows more effective utilisation of arbitrage opportunities as well as more accurate reflection of market inefficiency than the latter. Additionally, I also address the issue of whether or not there is an optimal specification for the new rule. Last but not least, if the returns from pairs trading after deducting transaction costs are still higher than the required compensation for risks, this is evidence against the efficiency of gold ETFs at ultra-high frequency.

Chapter 4 studies another important aspect of the relationship among financial instruments, namely their lead-lag effect with one another. The lead-lag effect refers to a situation observed in the market where a security follow earlier movements of another after a certain delay. The analysis in this chapter is conducted in the setting of the equity market, based on a large set of ultra-high frequency data time-stamped to the millisecond of the S&P 500 stock index and related ETFs. Acknowledging the difference between data at ultra-high frequency and data at lower frequencies, the chapter employs a novel model proposed by Hayashi and Yoshida (2005) and then developed further by Hoffmann et al. (2013), which is designed to better accommodate distinct features of ultra-high frequency data while estimating the correlation and cross-correlation for each pair of instruments. After analysing the lead-lag phenomenon, the chapter proceeds to establish a potential connection between this phenomenon and an important factor which is the arrival of information. Using suitable variables as a proxy for the rate of information arrival, I implement a regression analysis to test whether or not this factor indeed has an impact on the lead-lag relationship. Finally, a number of additional tests are also executed as a robustness check in order to make sure that the results are valid and reliable.

Chapter 5 examines how financial assets are related to one another from yet a different angle, namely their dynamic correlation and volatility transmission. To be specific, this chapter analyses how the intraday correlation and volatility spillover among exchange rates have changed in the presence of a significant political event which is the Brexit vote by comparing the quantities of interest before and after the vote. The analysis is based on high frequency data of various major markets including the two currencies directly involved in the Brexit event (i.e. the sterling and the euro) as well as three safe haven instruments (i.e. the Swiss franc, the Japanese yen and gold). The estimation of correlation is derived from the Dynamic Conditional Correlation (DCC) model of Engle (2002) and its improvement by Aielli (2013) while the transmission of volatility is calculated using the generalised Vector Autoregression approach and variance decomposition. The chapter aims to find out (i) whether or not there is any pronounced response to the Brexit vote from the foreign exchange market, and also (ii) whether or not the reaction (if any) is consistent with the decision of the UK to leave the EU. Given such a decision, it is quite reasonable for participants in the market to expect not only an increase in the level of uncertainty but also a decrease in the degree of integration between the UK and the EU.

1.3. Motivation

First and foremost, this thesis is motivated by the important role that high frequency traders play in all areas of financial markets nowadays including different asset classes, spot and derivative markets, and geographically diverse markets. Due to the active participation of high frequency traders in the market and their substantial contribution to trading activities in general (O'Hara, 2014, O'Hara, 2015, Aitken et al., 2015), Sun et al. (2014) do emphasise the necessity and benefits of a better understanding of high frequency markets. On the other hand, the thesis focuses on the relationship among securities which has become increasingly important for the following reasons. Firstly, the universe of available vehicles for trading and investment has been expanding continuously and as a result of this constant development of the existing marketplace, traders and investors have more and more products at their disposal, which means that they are required to pay attention to a lot of assets from many different asset classes together with their derivative instruments. Secondly, the relationship among instruments cannot remain constant and will change over time, which gives rise to the need for regular scrutiny. A good understanding of this relationship is often an essential requirement in order to be able to carry out a proper assessment of risks which can be considered the basis of various activities that market practitioners get involved in frequently including risk and portfolio management, hedging, diversification and trading. As a matter of fact, Chapter 3 studies a popular trading strategy based on the relationship among assets, namely pairs trading, as an example of a practical application of this relationship.

Chapter 3 is motivated by the literature on pairs trading, especially the seminal paper of Gatev et al. (2006). This strategy has received attention from both academics and practitioners alike. For academics, pairs trading is interesting because it can act as a natural mechanism where trading activities of arbitrageurs help eliminate (or at least alleviate) the potential mispricing between two markets and effectively keep their relationship with each other in equilibrium (Kondor, 2009, Alsayed and McGroarty, 2012). Meanwhile for practitioners, pairs trading is attractive since it is a powerful tool to assist them in their pursuit of trading profits. Although this arbitrage strategy is highly popular and widely used by many, the same trading rule is employed by traders in their trading practice as well as by researchers in their research works. This standard rule seems to make sense only because the majority of previous studies on pairs trading have utilised only a single time series of price for every security and hence failed to distinguish the bid price from the ask price, which will lead to the use of an incorrect price for a given transaction (i.e. using the bid price for a buy trade or the ask price for a sell trade) and biased results due to the problem of bid – ask bounce. However, with a more careful treatment of prices by including both the bid and the ask in order to use the correct price for every trade and reflect what actually happens in real markets, it has been revealed that the standard rule may be unnecessarily strict and may result in an underestimation of the full profit potential of pairs trading. As the profitability of trading strategies in general and pairs trading in particular is an important indicator of the level of market efficiency, it is necessary to have a more accurate estimation. Therefore, I put forward a new trading rule which is a more relaxed version of the standard rule used in the literature and should be able to reflect profitability and market efficiency better.

In addition, trading profits have been underestimated even more because of the fact that a large part of the pairs trading literature is based on low frequency data. By default, low frequency data cannot capture as much information as high frequency data so there are certainly many lucrative but short-lived arbitrage opportunities that are visible in high frequency data but invisible in low frequency data. To address this issue, Chapter 3 makes use of data at the highest sampling frequency possible, namely tick data, which will guarantee a successful observation of all available opportunities. On the other hand, the ETF market is chosen to conduct the research because this market is likely to be attractive to many arbitrageurs and generate a lot of arbitrage activities for a variety of reasons including (i) the presence of mispricing instances and potential opportunities for arbitrage (Ackert and Tian, 2000, Engle and Sarkar, 2006), (ii) the low level of risk involved and (iii) the fast growth of and ease of access to this market (Marshall et al., 2013, Kearney et al., 2014). In particular, ETFs based on gold are examined since the economic significance of gold has increased over time (in large part due to its safe haven characteristics) as shown by its price and trading volume as well as the significant demand and attention from investors in general (Pullen et al., 2014, Hauptfleisch et al., 2016). Finally, the issue of market efficiency is also investigated in an attempt to make an empirical contribution to the existing literature from the standpoint of trading, or more specifically, high frequency trading.

Chapter 4 is motivated by earlier studies on the lead-lag relationship in returns of financial assets which do not seem to be consistent with the notion of efficient markets where returns of securities should not be predictable and returns of related securities should be contemporaneously correlated. Being able to understand lead-lag effects may offer more insights into market efficiency and the process of price discovery because the lead-lag relationship is the result of differences in the speed of price discovery among instruments. To understand this relationship better, it is important for me to be aware of its potential determining factors which are likely to have a pronounced impact. This chapter aims to analyse one such factor, namely the arrival of information. This factor in particular is examined in the study since there have been a large number of previous works that emphasise how essential information is to various elements of financial markets including volatility effects (i.e. ARCH and GARCH), the size of transactions, adverse selection and the cost of trading (Gregoriou et al., 2005, Aragó and Nieto, 2005, Frank and Kenneth, 2005).

In order to make Chapter 4 more useful, I decide to focus on closely related instruments (i.e. an equity index together with its tracking ETFs) which constitute a research setting where it is more challenging to analyse the lead-lag relationship. The reason is that these ETFs are designed in a way that allows them to closely follow the index and quickly respond to changes in its value so lead-lag effects (if any) tend to be short-lived and disappear quite fast, which makes it more difficult to capture them. Moreover, as discussed earlier, the ETF market is an important part of financial markets since it is highly attractive to investors as an option for investment due to the high level of accessibility and liquidity. In particular, it has been shown that equity ETFs are better than individual stocks in terms of liquidity (Ruan and Ma, 2012). Among different stock indices, I choose the S&P 500 index for the study in this chapter, which consists of the 500 largest companies in the US, the biggest economy in the world. To be able to observe even tiny lead-lag effects, once again I utilise data sampled at the highest possible frequency (i.e. tick data), unlike many previous studies. In comparison with data at lower sampling frequencies, a major difference of tick data is the irregular arrival of new observations because a new observation may occur at any time and will be recorded as is with no aggregation of data. As a result, the analysis employs a novel and more advanced approach suggested by Hayashi and Yoshida (2005) and enhanced by Hoffmann et al. (2013), which can examine the lead-lag relationship using tick data without modifying the data in any way (e.g. interpolation or resampling at fixed intervals).

Chapter 5 is motivated by the literature on how political and unique events influence financial markets as well as the significance of the Brexit vote, which refers to the decision of the UK to leave the EU. This is an important political event with potential implications for the market due to a few reasons as follows. Firstly, this event is unprecedented in the sense that this is the first time a member of the EU has ever wanted to leave. Secondly, the vote is not only about the UK and the EU and will also have an impact on many other markets around the globe, given the increasingly close linkages among different countries in today's system of markets. Thirdly, even though highly important, the vote in itself is just a part of the whole Brexit situation, which is not over yet because the actual exit deal for the UK is uncertain and still to be negotiated. Therefore, understanding the effect of the vote may contribute to the overall understanding of Brexit when it is over.

On the other hand, studies on the Brexit referendum are somewhat limited since it is a fairly recent event. In addition, a number of previous works have investigated how the vote affects individual markets (e.g. the response of individual stocks to the public announcement of the voting outcome) so I decide to focus my research on the relationship among markets. More specifically, I study two aspects of this relationship, namely correlation and the transmission of volatility. The research context of Chapter 5 is the foreign exchange market which is the largest and most economically significant of all markets in the world. In this market, as much as trillions of dollar are traded on a daily basis and the total daily trading volume of even the biggest equity exchange is still only a tiny fraction of that of the currency market. The analysis covers not only the currencies directly related to the Brexit event (i.e. the British pound and the euro) but also safe haven instruments (i.e. the Swiss franc, the Japanese yen and gold) against the US dollar. These safe haven assets are more important and interesting to examine when an influential event such as the Brexit vote takes place partly due to the uncertainty after the vote with regards to the negotiation between the UK and the EU in the future. The currencies considered in this chapter account for most of the trading activities of participants in the foreign exchange market.

Chapter 2

High frequency financial markets

2.1. Introduction

This introductory section aims to (i) provide an overall view of the modern high frequency financial markets which are the key setting of all the studies in this thesis, (ii) show the significance and popularity of HFT in the current market and (iii) introduce to the readers the different types of high frequency traders.

2.1.1. Overview of the high frequency financial markets

All financial markets operate on a continuous basis when they are open during trading hours and therefore, the information and data from the market are also generated on a continuous basis. However, in the past, it was common to collect datasets at a low frequency, which is mostly due to the high costs associated with the collection of data at higher frequencies (Goodhart and O'Hara, 1997). Fortunately, new improvements in information technology in recent years have made it much easier and cheaper to provide a large number of market participants with fast access to new information. As a direct result of this technological advance, there has been a substantial decrease in the costs of data collection. In addition to collecting information, the subsequent process of analysing that information and acting on it (i.e. by placing trade orders) has also been enhanced to a great extent. Traders have now become capable of carrying out these tasks at increasingly high speed (Hagströmer and Nordén, 2013).

At the same time, the changes on the part of market participants are accompanied by important changes on the part of the exchanges as well. They have shifted to complete automation which allows them to dramatically increase their ability to handle a huge number of transactions on a daily basis (Jain, 2005, Brogaard et al., 2014). All of the above factors combined have made a significant contribution to the extremely high level of trading activities that I observe in the market today. It may be worth pointing out that based on the survey data on spending of large US companies, the financial sector is one of the most prominent industries with regards to the intensity of information technology, namely the proportion of revenue that is spent on information technology (Dewan and Min, 1997, Haferkorn, 2017).

Brogaard et al. (2014) argue that the continuing trend of increased automation has led to an inevitable impact on the traditional role of human traders, many of whom have no other choice but to give way to their electronic counterparts in the current market environment. The use of algorithmic trading has been on the rise, which is evidenced by the fast growing number of market participants who employ a computer algorithm to complete at least some (and in many cases, all) stages of their usual trading process (Sun et al., 2014). The typical process consists of (i) collecting data from the data feed, (ii) analysing the data in order to inform trading decisions, (iii) making decisions based on the results of the analysis, (iv) submitting trade orders to the marketplace through the trading platform, (v) managing these orders and potential trades after submission by taking actions such as hedging or liquidating open positions and finally (vi) as soon as the trades have been closed, reporting the trading results which include the monetary outcome (i.e. profit or loss), the duration of trades as well as a wide range of other important statistics of interest.

HFT is a type of algorithmic trading, which means that it also takes advantage of a computer algorithm to implement trading strategies and carry out the trading process (Hagströmer and Nordén, 2013, Aitken et al., 2015). However, the distinct feature shared by high frequency traders which sets them apart from other algorithmic traders is their great emphasis on high speed and low latency throughout the entire process of trading. Although the term 'high frequency trader' actually refers to a broad group of traders, all of them try their best to be able to analyse their information and then execute their transactions with the highest speed and the lowest latency possible, which will put them ahead of other competing high frequency traders (Sun et al., 2014, O'Hara, 2014). With regards to the latency, it is a term which is generally used to refer to some kind of delay in a system. In the specific context of trading in the financial market, it means the time delay that occurs when the information of trade orders is being sent back and forth between the traders and the marketplace (i.e. the exchanges or other trading venues). Therefore, it is desirable to have a low latency because latency is negatively correlated with speed, or in other words, the system can operate more quickly if it has a lower latency and vice versa (O'Hara, 2014). The constant pursuit of a low latency in HFT is the reason why the increasingly fast interaction between computer algorithms and trading platforms has become one of the most striking feature of many markets in the digital era (Benos et al., 2017).

Regarding the starting point of HFT, O'Hara (2015) shows that even though the technological infrastructure that has paved the way for high frequency traders began its development as early as the 90s, the decisive factor contributing to the introduction of my modern high frequency world was in fact the new changes in market regulations and policies which took place in the 2000s and were originally intended to encourage competition. To be more specific, the Alternative Trading Systems regulation (or Reg. ATS for short) in the US was issued in 2000 in order to allow non-exchange entities to enter the market and compete with the exchanges currently in existence. In addition, another regulation of great importance, namely the National Market System regulation (or Reg. NMS for short), was implemented in 2007 with the purpose of helping shape a common marketplace which consists of a number of different trading venues connected to one another by a set of shared rules in terms of access to the market as well as the priority of trade orders.

As a result of these regulations, the market has become more and more fragmented and this transformation has made the environment for trading activities more complex. For example, the lack of a central trading venue, more often than not, requires market participants to look for available liquidity across a wide range of different locations so that they can optimise the execution of their trade orders. It should be noted that the terms and conditions of trading may vary from place to place and hence, the task of searching for liquidity becomes even more challenging than it already is. To make the most informed decision possible, all of a sudden, it is now necessary for traders to be aware of all the relevant information about their potential trading locations, including but not limited to the structure of rebates, fees and charges (O'Hara, 2015). In the new condition of market fragmentation, HFT has become more valuable than ever because of its inherent ability to quickly identify and gain access to the most beneficial source of liquidity currently available among a lot of options. Another prime example of a complication that came into existence due to the fragmentation is the possibility that at a given time, the price of the same financial asset may be different from one trading venue to another. Consequently, traders need to keep a close eye on those potential discrepancies in the market prices, which may form the basis of profitable arbitrage opportunities, among other things. Again, HFT is highly valuable since it enables traders to keep track of and then take full advantage of these opportunities more easily and effectively.

However, it may be naive to think that high frequency traders are able to make the most of HFT all on their own. I can safely say that in spite of the incredible progress with regards to the technological capabilities, it is not possible for HFT to be as useful as it is now without the much needed external assistance from the exchanges (O'Hara, 2015). More specifically, many of the current exchanges have been offering a special type of paid services which enables traders to trade at a higher speed by (i) gaining early access to information about trading activities on the exchange via its direct feeds which provide traders with more comprehensive data than that observed on the standard consolidated tape freely accessible to all market participants and / or (ii) setting up their own trading systems within close proximity of the matching engine at the exchange (which is also known as co-location) in order to reduce to a minimum the latency in the communication of trade order information between traders and the exchange.

Since the early days of HFT until now, a number of notable studies have pointed out the most common characteristics of high frequency traders (Aitken et al., 2015, O'Hara, 2015, Benos and Sagade, 2016, Brogaard et al., 2017). O'Hara (2015) states that there are three indispensible components in the operation of the typical high frequency traders, namely (i) a well-defined trading strategy, (ii) a reliable computer algorithm used to implement the strategy and finally (iii) an extremely high speed of data analysis and trade order execution to make it more likely for the strategy to succeed. Nevertheless, there is a certain variation in these factors within the large group of high frequency traders. For instance, in terms of the speed, the high frequency world can be roughly divided into two groups, namely the high frequency group (which utilises a very fast speed for its electronic communication with the market) and the ultra-high frequency group (which tries its very best to push the speed to the absolute physical limit humanly possible). At the moment, it is not unusual to record the latency of trade orders in the realm of millisecond (one-thousandth of a second), microsecond (one-millionth of a second) or even as extreme as nanosecond (onebillionth of a second) in some exceptional cases (O'Hara, 2015). To put it in perspective, I need to consider the fact that it takes as long as a few hundred milliseconds for the human eyes to respond to visual stimuli by performing a blink. It is indeed interesting to see that I have come to the point where the blink of an eye is like an eternity in the novel context of ultra-high frequency trading.

The lightning fast speed at the heart of HFT helps me to explain the characteristics and behaviours of high frequency traders. Most obviously, they are able to open as well as close trades very quickly. To take full advantage of their unparalleled speed, high frequency traders tend to make use of trading strategies which are designed to exploit lucrative but relatively short-lived opportunities in the market. Therefore, they end up generating trades with a short holding period during the trading day and no open position is kept overnight (Aitken et al., 2015, Brogaard et al., 2017). Because the trading profit can only be accumulated by movements over time in the price of the financial instrument being traded, the short-lived nature of high frequency trades is likely to impose a restriction on the potential profit to be obtained from each individual trade and thus, the traders have to make up for this limited profit per trade with a sufficiently large amount of trading volume (Hagströmer and Nordén, 2013). Perhaps to serve the purpose of risk management, the large trading volume required is often generated by a large number of small trades instead of a small number of large trades. As a result, high frequency traders tend to enter and exit trades quite frequently throughout a typical trading day, executing as many as thousands of transactions per day (Aitken et al., 2015).

Not only are high frequency traders extremely fast at order execution, they are also equally fast at order cancellation. O'Hara (2015) reports that according to the trading data of the Securities and Exchange Commission (SEC), out of 100 instances of cancellation, 23 to 38 take place within 50 milliseconds or fewer after the initial placement of orders. Another interesting fact is that order cancellation is a much more common occurrence than order execution. To be more specific, nowadays it is not unusual for more than 98% of all orders to be cancelled before having the chance to be executed, which leads to a very high order-to-trade ratio for high frequency traders. Their excessive emphasis on the trading speed and frequency, at the end of the day, is essential for them to be able to compete with other high frequency rivals and significantly increase their chance of survival in a highly competitive environment where the winner takes all (Benos and Sagade, 2016). If one considers the fact that only a few firms at the top take the lion's share of the profits from HFT, it is not difficult to understand the reason why high frequency traders focus so heavily on these vitally important factors.

Sun et al. (2014) argue that a good understanding of HFT is necessary for market participants as well as exchanges and regulators for a couple of reasons. Firstly, market participants who know how HFT works are able to adjust their own trading algorithms for more effective performance and improve their trading results. Secondly, understanding HFT will help exchanges to tailor their current services and possibly introduce new specialised ones to better meet the needs of high frequency traders, which will make them more attractive to HFT clients than other competing trading venues and contribute to their revenue. Thirdly, with the knowledge of HFT, regulators and policy makers can monitor the market and maintain market quality more effectively with appropriate use of the different tools at their disposal such as financial transaction tax or restrictions on certain kinds of trading activities. In the next section, I will take the first step in understanding HFT by taking a close look at the popularity of HFT and the activity of high frequency traders in my current financial markets.

2.1.2. Popularity and activity of HFT

As discussed in the previous section, Hagströmer and Nordén (2013) and Aitken et al. (2015) point out that due to the relatively short holding period of typical high frequency trades and the resultant restriction on the profit potential of each individual trade, high frequency traders need to enhance their overall profits by generating a large amount of trading volume. This is one of the main reasons why an increasing number of studies on both the US and non-US markets have shown that high frequency traders are responsible for a substantial part of all trading activities and they are among major participants not only in equity markets but also in many other types of financial markets (Kirilenko and Lo, 2013, Hagströmer and Nordén, 2013, Brogaard et al., 2017, O'Hara, 2014, O'Hara, 2015, Benos and Sagade, 2016). In addition, the widespread popularity of HFT can be observed not only in spot markets but in derivative markets as well such as futures and options (Benos and Sagade, 2016). As a result, it is not surprising that HFT has become one of the key components of the day-to-day operation of most major hedge funds and there are also a number of proprietary trading firms which focus solely on HFT strategies (O'Hara, 2014).

With regards to the estimation of the participation rate of high frequency traders in the market, there is a variation among the estimated figures from different sources due to certain difficulties in identifying which trades are high frequency trades and which ones are not. Aitken et al. (2015) report that based on recent data, high frequency traders may normally account for 50 – 70% of all the transactions in equity markets in the US and sometimes as much as 75% on the NASDAQ stock exchange. Consistent with the figures reported by Aitken et al. (2015), O'Hara (2014) and O’Hara (2015) state that these traders generate at least 50% of the total trading volume in stock markets and even more in other markets such as the foreign exchange market. A similar figure has also been provided by Easley et al. (2011) and Hendershott et al. (2011), who examine the US market and find that the proportion of trading activities related to HFT increased gradually to 73% on the NYSE in the last decade.

In addition to the US market, HFT activities have also been on the rise in other markets all over the world (including Japan, many European countries and more) and have made a great contribution to their overall trading volume (O'Hara, 2014). In the Japanese market, the advent of HFT was marked by the introduction of a novel trading platform called Arrowhead in 2010 on the Tokyo Stock Exchange (TSE), which is the largest equity exchange located outside the US. The annual report of the TSE in 2011 pointed out that after only a short time since this important event (i.e. by April 2011), the new platform, together with related co-location services, resulted in a decrease in order latency and a substantial increase in the share of HFT in all market activities from zero to 36% (Jain et al., 2016). Moreover, after Arrowhead was launched, the quote-to-trade ratio increased by more than 100%, which is consistent with the fact that high frequency traders tend to cancel a large part of their trade orders prior to execution. Regarding European countries, Jacob Leal et al. (2016) report that there was a sharp increase in the popularity of HFT in European markets during the past ten years and it is considered one of the major innovations in the financial sector. Quantitatively, the activities of high frequency traders contribute (i) between 24% and 43% of all executed transactions in Europe according to the data of the European Securities and Markets Authority in 2014 (Haferkorn, 2017) and (ii) 35% in the UK in particular (Aitken et al., 2015). Similar to Japan as well as Europe, HFT is also popular in Canada, making up as much as 40% of all trading activities in the Canadian market (Aitken et al., 2015).

The significance of HFT is evidenced not only by its increasing popularity but also by the huge amount of capital that the market is willing to invest in its supporting technological infrastructure. One of the most prominent examples of this strong financial commitment is the joint venture between the Chicago Mercantile Exchange (CME) and the NASDAQ stock exchange to enhance the connection between the computers of the former in Chicago and the servers of the latter in New Jersey. As a result of this co-operation, traders can transmit their trade order information marginally faster (i.e. 4 milliseconds faster, to be exact) than it normally takes between those two locations (O'Hara, 2014). This seems to be a tiny reduction in order latency if I recall a previously mentioned fact that human beings need a few hundred milliseconds to recognise visual stimuli and blink the eyes in response to such stimuli. However, it may be surprising how expensive such a seemingly trivial improvement can be and even more surprising that someone is ready to pay for it. O’Hara (2015) suggests that this technological upgrade could easily have cost more than $500 million, which clearly shows that high frequency traders are prepared to go to extreme length in order to make the most of HFT. Now that I have examined the popularity of HFT, I am going to investigate different types of high frequency traders in the next section.

2.1.3. Types of high frequency traders

The term ‘high frequency traders’ is used to refer to a large and diverse group of market participants. In the ongoing debate among regulators with regards to HFT, great emphasis has been placed on the recognition of and differentiation between a number of commonly used trading strategies which constitute HFT (Hagströmer and Nordén, 2013). In addition, it is also highly important for regulators to understand the potential influence of these strategies on many aspects of market quality such as market efficiency, volatility, liquidity and more. According to the International Organisation of Securities Commissions (IOSCO) as well as the SEC, such knowledge is vital to the design of effective market regulations. O’Hara (2015) points out that generally speaking, HFT activities can be divided into beneficial ones and harmful ones which are also called predatory activities. As a result of this distinction, market regulators and trading exchanges have considered and introduced new changes in the market design in a continuous attempt to attract good high frequency traders and, at the same time, prevent bad high frequency traders from exploiting other participants.

In terms of the beneficial HFT strategies, they consist of such activities as arbitrage or market making which make a contribution to the improvement of market quality (James and Douglas, 2013). For example, let me consider arbitrage activities. In essence, arbitrageurs simultaneously trade two closely related financial instruments which are based on the same underlying but temporarily mispriced, buying the underpriced instrument and selling the overpriced one to lock in the profit which will be realised when their prices eventually converge and the mispricing disappears. Because of this process, arbitrage can help me to enforce the Law of One Price and maintain informational efficiency across multiple markets. Therefore, traders are more confident that all available information is actually reflected in current prices of financial assets which represent their true fundamental value. This confidence makes traders feel more comfortable with placing trade orders, increasing the level of trading activities in general. On the other hand, market making activities contribute directly to the enhancement of not only overall liquidity but also the distribution of liquidity among markets since market makers tend to move available liquidity to where it is highly valuable (i.e. where there is a substantial demand for it) (O'Hara, 2014). Improved liquidity will lead to a reduction in the costs of liquidity search and trade execution. It is clear that these beneficial activities of high frequency traders need to be supported and encouraged by regulators, policy makers and exchanges alike (Robert and Philip, 2012).

In stark contrast to the good side of HFT presented above, HFT also has a bad side where predatory algorithms are used to carry out trading strategies which are manipulative in nature (James and Douglas, 2013). These algorithms often aim to take advantage of a weakness in the microstructure of the market by taking certain actions in order to trigger a microstructure mechanism which is likely to result in a predictable and unfavourable outcome for other market participants (O'Hara, 2014). Quote stuffing is a common example of this kind of malicious behaviours. Basically, quote stuffing refers to the constant submission of a massive number of trade orders which are then cancelled within a very short time after being placed. The purpose of this practice is not to have orders executed but to overwhelm trading exchanges with a huge amount of incoming messages, significantly increase the burden on the processing of orders and slow down trading activities of competing traders (O'Hara, 2015). Even though such a harmful strategy is obviously not legal, the problem is that it is not easy to identify and catch high frequency traders who engage in quote stuffing. As a matter of fact, it may be reasonable for traders to submit many trader orders but cancel them afterwards. The reason is that in the search for available liquidity, it is often necessary to send orders to multiple trading venues to have an idea where liquidity currently exists. Before the intended trade has a chance to be executed, the trader may need to keep updating his / her orders to reflect new information. After the trade has been executed at a certain location, all of the remaining orders previously submitted to other places will be cancelled. In short, the challenge of correctly distinguishing quote stuffing from a legitimate search for liquidity is what makes it difficult to catch quote stuffers.

Interestingly, quote stuffing is only one example to illustrate that the line between normal and bad HFT activities cannot be defined clearly at times. According to Hagströmer and Nordén (2013), harmful trading strategies, such as the quote stuffing described above, belong to a more diverse group of strategies which is known as opportunistic HFT. They also suggest that studies in the future should break down this group to evaluate the prevalence of malicious strategies because not all opportunistic strategies are necessarily malicious. For instance, Robert and Philip (2012) argue that many high frequency traders make their trading decisions based on market signals and their HFT technology allows them to trade fast enough to make good use of potentially short-lived signals before they are fully incorporated into current market prices and thus lose their value. Although the nature of this strategy is opportunistic, it is safe to say that it is acceptable since it does not attempt to manipulate the marketplace with a malicious intent. On the contrary, manipulative strategies have to be discouraged and prevented without a doubt.

With regards to the significance of sub-groups of high frequency traders in relation to one another, Hagströmer and Nordén (2013) study market making HFT as well as other types of HFT and find that market makers account for the majority of HFT activities, generating 63 – 72% of the total trading volume and as much as 81 – 86%

of all limit trade orders. Furthermore, market makers also have a lower level of inventory and a higher contribution to the supply of liquidity compared to other types of traders. Given the overwhelming presence of market makers in the HFT world at the moment, it is not surprising that they will experience a more pronounced impact than others from any policy or regulation aimed at high frequency traders as a whole.

In another research paper on the different kinds of HFT, Benos and Sagade (2016) investigate equity markets in the UK and evaluate not only (i) the contribution of high frequency traders to price discovery in general but also (ii) in particular, the variation in the respective contribution of each component in the cross-section of traders. These authors divide all of the traders into two groups based on whether they demand or supply liquidity. Those who apply an aggressive trading strategy are likely to require liquidity whereas those who use a more passive strategy such as market making tend to provide liquidity. Responsible for two-thirds of the contribution of HFT to price discovery, aggressive traders have been found to contribute more to price discovery than passive ones, which suggests that they are more informed. Nevertheless, Benos and Sagade (2016) also show that it may not be easy to increase the scale of aggressive strategies. Consistent with the finding that aggressive trades have a higher level of information content than passive ones, changes in the order flow of aggressive traders are often in the same direction as changes in the future price while those of passive traders are unrelated or move in the opposite direction of future price changes. On the other hand, the way these groups manage their respective inventories may be similar or different depending on the trading horizon. Even though their inventories are managed in very different ways over a short horizon (i.e. measured in seconds), all of them exhibit not only a tendency of mean reversion but the same pattern of serial correlation as well over a long horizon (i.e. measured in hours). Having studied the types of high frequency traders, I am going to analyse in more details their various impacts on many aspects of financial markets in the following section.

2.2. Impact of HFT on the system of financial markets

2.2.1. Impact on the market

Jarnecic and Snape (2014) state that there is a complex relationship between the behaviours of investors (as shown by the flow of their submitted trade orders) and several key factors in the market such as the incorporation of new information into the marketplace, the price dynamics of financial instruments, the nature of liquidity and the cost of transactions, among other things. This inherently complex linkage has become even more complicated since the arrival of high frequency traders who are able to handle the process of order execution (including the submission and possibly update or cancellation of orders) at a much higher speed than many other market participants. Furthermore, O'Hara (2014) also argues that since these low latency traders joined the market, there have been some fundamental changes in its structure and trading activities in general have indeed become faster and faster. However, in spite of the importance of HFT, the role of high frequency traders is still not completely understood. The debate is going on at present with regards to the contribution of HFT to financial markets because it is not entirely clear whether this contribution is positive or negative (Jarnecic and Snape, 2014). The reason is that high frequency traders have been found to be related to a wide variety of common phenomena documented in the existing finance literature (e.g. temporary shocks to market prices, an increase in the rate of order cancellation, a decrease in bid-ask spreads and trade sizes, to name a few) and their impact on financial markets is not always straightforward. In this section, I investigate the effect of HFT on three important areas of the market, namely (i) informational efficiency, (ii) volatility and trading risks, and (iii) liquidity and transaction costs.

2.2.1.1. Impact on market efficiency

The role of HFT in price discovery is essential (Aitken et al., 2015). Manahov et al. (2014) examine exchange rates of many major currencies against the US dollar sampled every minute from August 2012 to March 2013 and find that the presence of high frequency traders indeed contributes positively to the process of price discovery. As a result, an increasing number of both theoretical and empirical studies agree that HFT in general, and HFT market making in particular, helps improve the informational efficiency of financial markets which is often considered one of the main determinants of market quality (Carrion, 2013, Brogaard et al., 2014, Manahov et al., 2014, Aitken et al., 2015). The positive impact of high frequency traders on market efficiency may result from the high level of information content in their transactions, as evidenced by the relationship between the direction of their trade orders and price changes of financial instruments. Investigating transaction data of a large group of high frequency traders in 120 stocks listed on two major US equity exchanges (i.e. NYSE and NASDAQ) in 2008 and 2009, Brogaard et al. (2014) document that on average days as well as highly volatile days, high frequency trade orders are usually in the same direction with permanent changes in price but in the opposite direction of temporary changes. Therefore, one can use these orders to forecast future price changes over a short trading horizon of a few seconds. It has also been observed that in addition to changes in the price of individual securities, high frequency trades are correlated with several other important factors such as market-wide movements in price, publicly available information (e.g. releases of macroeconomic news) and imbalances in the limit order book.

However, it should be mentioned that there are exceptions to the general agreement on the beneficial effect of HFT on price discovery and market efficiency. A notable example is the study of Robert and Philip (2012), which shows that actions of individual high frequency traders are often triggered by the appearance of a certain signal so when many of these traders observe a common signal, their independent actions become co-ordinated and they may collectively create an instance of mispricing. Potential deviations of security prices from their intrinsic values due to such a mechanism are one of the reasons why some commentators and the media have shown their concerns with regards to the possibility of HFT being exploited for harmful and unethical trading activities with the sole purpose of market manipulation (Aitken et al., 2015).

Contrary to the concerns about manipulation raised by the media above, Aitken et al. (2015) provide evidence that the active participation of high frequency traders plays a significant part in enhancing the price discovery process and reducing the dislocation of market prices at the end of the trading day in terms of not only the magnitude but the frequency as well. Their conclusion is based on an extensive set of transaction data on 22 equity exchanges in a wide range of developed and emerging markets (including the US, UK and a lot of other countries) during the time period from January 2003 to June 2011. Interestingly, they find that this positive influence of HFT on price discovery is more significant under certain conditions including (i) when the cause of price dislocation is likely to be manipulative behaviours, (ii) expiration dates of options and (iii) the end of the month. Furthermore, they also find that price dislocation is reduced by the presence of high frequency traders more than by trading regulations, market surveillance and legal enforcement. More importantly, these findings still hold when using different proxies for the presence of HFT such as co-location, trade size or the cancellation rate of trade orders.

In addition to providing their empirical results, Aitken et al. (2015) point out two potential reasons which may help explain why it is not surprising that HFT activities generally contribute to the reduction of market manipulation attempts. The first reason is that most of the monitoring systems used on trading exchanges are often built with more emphasis on identifying suspicious trading patterns than capturing single instances of manipulation. Meanwhile, HFT activities are carried out by computers based on certain algorithms and as a result, they are bound to generate some kind of observable patterns and malicious trading behaviours are likely to raise a red flag to the surveillance systems. On the other hand, the second reason is that a number of previous studies in the literature have shown that HFT tends to have a beneficial effect on price discovery and informational efficiency (Carrion, 2013, Manahov et al., 2014, Brogaard et al., 2014). When financial markets are more efficient, they are more difficult to be manipulated successfully. In summary, this section has studied the impact of HFT on price discovery and efficiency. The next section is going to examine how high frequency traders influence volatility and trading risks which are also highly important factors to be considered by participants in the market.

2.2.1.2. Impact on the volatility and market risks

In terms of the level of market volatility, some studies show that when there is an increase in the trading activities of high frequency traders, volatility is usually reduced (Hagströmer and Nordén, 2013, Sun et al., 2014). The finding of Hagströmer and Nordén (2013) is based on a comprehensive dataset of all high frequency trades and quotes (timestamped to the nanosecond) of 30 Swedish stocks on the Stockholm stock exchange with the highest trading volume in 2011. On the other hand, Sun et al. (2014) investigate the equity market in the US using data at the tick level in three years from 2008 to 2010. A possible explanation for the mitigating effect of HFT on volatility is that more often than not, high frequency traders have been found to make a positive contribution to the process of price discovery and market efficiency. To be more specific, HFT allows financial transactions to take place at a very fast pace so new information can be incorporated into market prices through trading activities much more quickly and frequently, which makes it less common to observe large price shocks due to information accumulated over some period of time.

Because HFT tends to reduce the level of volatility, one can expect a lower level of volatility risks as a result of HFT. Moreover, as pointed out in the previous section, HFT has also been shown to contribute to the reduction of price dislocation and market manipulation, even though in theory it is possible for high frequency traders to take advantage of their advanced HFT technology and infrastructure to tamper with markets for their own gains in a harmful and malicious way. In addition to the risks related to market volatility and manipulation, HFT can also influence another type of risks, namely systematic risks. In a study on HFT in the foreign exchange market which is the largest financial market in the world to date, Chaboud et al. (2014) provide some evidence that trading behaviours of machine traders are less diverse than those of human traders. As a direct result of this lack of diversity, it is often the case that trades generated by computer algorithms have a higher level of correlation with one another compared to trades generated by human traders. Similar to Chaboud et al. (2014), Aitken et al. (2015) analyse the equity market in many countries and argue that if a lot of high frequency traders (who rely heavily on computer programmes) choose to employ similar trading strategies, their resultant transactions should be highly correlated, which is likely to lead to an increase in systematic risks. It is worth noting that unlike these machine traders, even if many traditional human traders use the same strategy, their trades may still be fairly different from one another due to the involvement of subjective elements (i.e. personal emotions and discretionary judgements) in the process of making a trade decision. These factors, which vary from trader to trader, do not play a role in the case of automatic algorithms. In the next section, I am going to study how HFT affects liquidity and transaction costs in financial markets.

2.2.1.3. Impact on market liquidity and trading costs

A large number of previous studies on HFT have reported that high frequency traders tend to enhance liquidity of the market and at the same time reduce bid – ask spreads (i.e. an important component of trading costs) (Manahov et al., 2014, Brogaard et al., 2014, Carrion, 2013, Hendershott et al., 2011, Sun et al., 2014, Chaboud et al., 2014). Transaction costs have been decreasing during the last three decades, especially in recent times, and there has been empirical evidence that the arrival of high frequency traders in the Canadian market has led to a decrease in trading costs (O'Hara, 2015). Using a high frequency dataset on trading and quoting activities of constituent stocks of the FTSE 100 index (which is usually considered to be representative of the equity market in the UK), Jarnecic and Snape (2014) investigate how high frequency traders supply additional liquidity to the market by analysing their strategies for submitting trade orders to the limit order book. According to the findings of this study, high frequency participants provide liquidity on a continuous basis and hence contribute to resolving the issue of liquidity imbalance in the limit order book, which makes markets more liquid. It should be noted that this observation still holds in different market conditions (i.e. slow and stable or fast and volatile), which lends further support to the positive role of HFT in enhancing liquidity. Jarnecic and Snape (2014) attribute the willingness of high frequency traders to take part in the provision of liquidity to the fact that their computer algorithms are able to help them lower the risks of inventory management and improve their market making activities. By supplying more liquidity when financial markets are not sufficiently liquid, HFT offers investors and traders a highly valuable and useful service, particularly in times of uncertainty.

On the other hand, despite the positive contributions of HFT confirmed by the earlier research shown above, a number of concerns have been raised with regards to its potential negative impact on liquidity. First of all, it has been argued that the presence of high frequency traders in the market making scene will make the current competition among market makers for market share become much more fierce. If existing traditional market makers cannot compete effectively with their high frequency rivals or do not find it worthwhile to do so, they may decide to cut back on the amount of capital committed to their business or even worse they may put an end to their operation and leave the market altogether, which will have a negative effect on market making and liquidity in general without a doubt (Jarnecic and Snape, 2014).

Secondly, a strategy for order submission commonly used by high frequency traders requires them to submit large volume of trade orders with small sizes and then often cancel these orders quickly after submission but before execution. Because of such a strategy, it is possible that even though HFT provides liquidity continuously, not much is actually added to the market. Moreover, the ability of high frequency traders to place and cancel their orders very fast and the short-lived nature of those orders as a consequence of this behaviour may reduce the quality of liquidity since it becomes more and more difficult to know exactly where to find liquidity (if there is any at all) among a myriad of potential trading venues. The quality of liquidity may deteriorate further due to some significant differences between high frequency market makers and their traditional counterparts. O'Hara (2015) argues that although high frequency traders supply liquidity with the frequent use of limit orders in the same way as traditional market makers, they only place their trade orders on one side of the limit order book (i.e. either buy or sell) for a given traded asset. Even more importantly, one needs to remember that these high frequency traders do not have any commitment or obligation to provide liquidity on a regular basis. These features of high frequency market makers have raised some concerns, and rightfully so, that high frequency market making can potentially cause illiquidity and hence instability in the market from time to time (Ananth, 2012).

Thirdly, the appearance of HFT also has several notable implications for the transparency of liquidity. O'Hara (2015) shows that because it is easier for predatory high frequency traders to locate and take advantage of other market participants whose transactions are large enough in size, cautious investors who are afraid of and want to avoid the risk of being exploited by these predatory traders often try to hide their activities by reducing the size of their trades if possible. That is one of the reasons why there is a substantial decrease in trade sizes in the US stock market with the average size being only more than 200 shares and as much as 20% or more of all transactions make use of odd lots (i.e. order sizes which are smaller than the normal unit of trading for a specific financial instrument) (O'Hara et al., 2014). In addition, investors can also hide their trades by gaining access to liquidity in dark pools which constitute an alternative system for trading outside traditional exchanges where participating investors have the opportunity to place and execute their trade orders, which are usually large, with no requirement to disclose any of their intentions and activities to the general public while they are waiting for their orders to be executed. On the other hand, one has to be aware of not only the advantages but also the disadvantages of dark pools due to the inherent lack of transparency such as potential conflicts of interest between pool members and pool owners or the possibility that participants may not receive the best available price. Despite all of these pitfalls, in the current state of financial markets where more and more high frequency traders have joined the market over time, dark pools have become an increasingly important and necessary tool at the disposal of many investors which can be used to help them adapt to new changes in their trading environment and optimise the results of their trading endeavours (O'Hara, 2015).

Interestingly, although the existence of HFT may influence the quality of market liquidity in a negative way as discussed above, the findings of Jarnecic and Snape (2014) suggest that high frequency traders do not really do any harm to markets which are already sufficiently liquid. Throughout section 2.1 of this chapter, I have considered the impact of HFT on three main aspects of market quality, namely (i) informational efficiency, (ii) volatility and risks, and (iii) liquidity and transaction costs. In the following section, I am going to shift my focus from how HFT affects market quality to how HFT affects participants in the market.

2.2.2. Impact on market participants

2.2.2.1. Relationship between high frequency traders and their trading machines

Since HFT activities are too fast for a human trader to handle manually, it goes without saying that the only way for high frequency traders to implement their HFT strategies is to rely on advanced computer algorithms whose operation is fully automatic. This requirement certainly has significant implications for human traders in terms of their role in the trading process as well as their relationship with the machines. An obvious advantage of computer programmes over humans is that they are able to capture and process new information from the market and then react to potential signals in this information by placing, updating or cancelling trade orders at a much higher speed than any human being ever could. Additionally, not only can they surpass the limits of human cognitive abilities, they can also make up for the imperfect rationality of human traders (Aitken et al., 2015), which is especially necessary and important in times of stress when it is very easy to make incorrect judgements and potentially costly mistakes which could have been prevented. To be more specific, while it is definitely not possible for me to trade in a rational manner 100% of the times, machines can always be completely logical and objective in analysing a situation and making a decision no matter what happens in the market simply because they do not have to suffer from the harmful interference of negative and undesirable feelings and emotions (i.e. mainly fear and greed) associated with humans (Borch and Lange, 2017).

Having said that, one has to bear in mind that automatic algorithms also have their own risks compared to human traders. For example, if there is a mistake in a trading programme but its user is not careful enough, it may take a long time to be found and the user may only recognise the mistake after the resultant damage accumulated over time has become considerable and clearly visible. In comparison, it is often easier and faster for human traders to realise mistakes in their own trading because they personally and manually go through the entire trading process of every transaction by themselves step by step. Also, as already discussed in an earlier section, automated trading strategies have the potential to be highly correlated with one another and hence increase the level of systematic risks due to the lack of human elements which are unique for every individual trader.

For the same reason, automated trading by definition is not flexible. On the other hand, market conditions in general tend to vary over time. Therefore, in order to make sure that algorithmic strategies can adapt to new environments and maintain the best trading performance possible, it is still an important requirement for them to have a certain level of human monitoring and intervention from time to time when necessary in response to the possibility of unexpected changes in market conditions. In fact, based on a number of interviews with and observations of high frequency traders, Borch and Lange (2017) argue that the exponential growth of the high frequency world does not actually make human traders become redundant and irrelevant but instead it has led to interesting changes with regards to the role of humans and how they interact with trading machines in the novel paradigm of HFT. High frequency traders still need to keep their emotions under control but unlike manual traders who need to prevent their own trading from being interfered with by their negative emotions, the challenge of algorithm users is to avoid letting those emotions make them impulsively and unnecessarily tamper with their trading programmes at the wrong time. Furthermore, it has also been shown that algorithmic traders attempt to make themselves more disciplined using a set of techniques and strict rules to achieve the goal of completely avoiding or at the very least minimising the likelihood of emotional attachment to and interference with their algorithms.

2.2.2.2. Relationship between high frequency traders and the rest of the market

The continuous rise of HFT and the resultant shifts in the dynamics of current financial markets all around the globe have not only introduced fundamental changes in the interaction between high frequency traders and their trading computers but also presented profound implications for the interaction between high frequency traders and other traders in the marketplace. A great deal of concerns have been raised regarding the issue of fairness which is without a doubt one of the most essential issues to take into account in financial markets in general as well as in high frequency markets in particular, especially from the point of view of trading exchanges, regulators and policy makers (James and Douglas, 2013). Fairness is highly important because it is directly related to whether markets are sustainable or not. If a certain group of investors do not find the market fair to them and there is nothing or little they can do about their circumstances, they are likely to stop their trading endeavours and leave the market altogether, which will cause many undesirable consequences for the rest of the market in one way or another, such as a reduction in the level of market activities, liquidity and interest of the general public in the market. Fairness is an even more important consideration in the high frequency context since O'Hara (2015) states that it is the high level of overall complexity involved in HFT that makes me even more concerned about the question of whether markets are unfair to some subset of participants or not.

High frequency traders invest quite heavily in extremely advanced trading equipment and technology to make sure that their costly day-to-day HFT operation is consistently profitable and worthwhile so in comparison with others, they have access to remarkably superior tools (and also information in many cases). If there is in fact a lack of fairness in the market, other traders are highly likely to be the victim. In order to examine this potential lack of fairness, one needs to understand how high frequency traders interact with and possibly have some negative externalities (intentional or otherwise) on other investors. In terms of the intentional impact, it is the result of deliberate behaviours of predatory high frequency traders who are more than willing to obtain their own financial gains at the expense of others in a malicious manner. There are a few ways they can try to manipulate markets and exploit other investors. Firstly, as described in one of the previous sections, they can use a technique known as 'quote stuffing' (Aitken et al., 2015) where they take advantage of their superior trading speed to submit a large amount of limit orders which are not actually intended for execution and will be cancelled very quickly after submission. The sheer purpose of this practice is simply to massively increase the load of trading information and data which need to be processed by exchanges with the hope of slowing down the processing speed of exchanges and thus the trading speed of other investors.

Secondly, in addition to quote stuffing which has an effect on the market as a whole, high frequency traders may also pose another type of risks to specific participants in the form of adverse selection (Manahov et al., 2014, Aitken et al., 2015, Brogaard et al., 2017). The term 'adverse selection' refers to a situation that takes place when there is an asymmetry of information between two parties (i.e. one of them has access to a crucial piece of information which the other is not aware of). In that case, the party with the informational advantage will usually make use of it by attempting to initiate a transaction which is favourable for them but not for the other party. Needless to say, such a transaction would probably not be agreed upon by the other party if they knew the same information as well. Brogaard et al. (2017) provide a good example of adverse selection in the context of spot and futures markets which is based on the relationship between these two closely related types of markets. When there is an increase or decrease in the price of equity index futures, the corresponding constituent stocks in the spot market are often expected to follow the change in the futures market with a short delay. Because high frequency traders are able to capture and process the new piece of information from the futures faster than others, they can exploit it to buy or sell the constituents against limit orders of less informed investors who are not yet ready to act on the same information in time and hence cannot delete or adjust the price of their orders in a timely manner. Brogaard et al. (2017) report that a particular kind of high frequency trades, namely short selling trades, has a pronounced effect of adverse selection on existing limit orders. Their study is based on a large transaction dataset of 758 stocks in 2008 provided by the NASDAQ equity exchange and every trade comes with a timestamp up to the millisecond as well as an identifier which clearly shows whether the two parties taking part in the trade are high frequency traders or not.

Thirdly, another way high frequency traders can exploit low frequency traders is that the former can take action based on the new market information revealed by trading activities of the latter (O'Hara, 2014, Benos and Sagade, 2016, Jacob Leal et al., 2016). More specifically, as shown by Benos and Sagade (2016), HFT firms are capable of accurately forecasting the flow of incoming order submission of non-HFT firms in the future which is driven by their information. This may explain why many high frequency traders can get their hands on and freely utilise the information obtained by the efforts of other traders in the market without putting in similar efforts to acquire that information on their own. Furthermore, Jacob Leal et al. (2016) show that after gaining some information from low frequency traders, high frequency traders will apply their trading strategies to take advantage of it. Generally speaking, information free-riding, together with the questionable practice of quote stuffing and adverse selection mentioned above and possibly many other shady techniques of predatory high frequency traders, has had the combined effect of making markets become a less favourable and more difficult environment for other investors to maintain their operation. As an inevitable consequence of these problems, those investors may decide to decrease their participation in the market, which in turn will reduce the provision of liquidity and ultimately the overall quality of financial markets (Manahov et al., 2014, Aitken et al., 2015, Brogaard et al., 2017).

Market participants are influenced not only by the intentional impact of high frequency traders (i.e. predatory trading) but also by their unintentional impact which results from normal HFT activities with no malicious intent. Because more often than not, high frequency traders place their trades based on market inefficiencies and trading opportunities which are short-lived in nature, the short duration of their typical trading signals requires them to use aggressive strategies with regards to the submission of their trade orders to increase their chance of getting the orders filled and not missing out on potentially profitable transactions (Jarnecic and Snape, 2014). As discussed earlier, the aggressive order submission of high frequency traders (i.e. a large number of small orders) followed by equally aggressive cancellation or resubmission has altered properties of the limit order book in more ways than one. According to the Securities and Exchange Commission in the US, due to these important changes in the limit order book, there have been a number of criticisms and complaints especially from non-high frequency traders about the increasingly high level of complexity and difficulty related to trading and order execution (Jarnecic and Snape, 2014, Manahov et al., 2014, Sun et al., 2014, Brogaard et al., 2014, Hoffmann, 2014). Specifically, HFT has brought about flashing quotes and thus uncertain liquidity (Sun et al., 2014) which may be an unwanted result of a legitimate search for available liquidity on the part of the high frequency traders themselves and in fact not necessarily due to some manipulative practice such as quote stuffing. Either way, it has become more and more challenging for investors in general and large investors in particular to have their trade orders executed effectively, which will eventually make them incur more costs and affect their trading negatively (Chordia et al., 2013, Brogaard et al., 2014, Hoffmann, 2014).

Even though non-high frequency traders may have to suffer from both the intentional and unintentional influence of high frequency traders, unfortunately it is not easy for them to do something about it and change their current situation for the better. The reason is that HFT operation is very expensive and requires a lot of capital investment to run, particularly the fixed costs incurred during the initial installation, so it is really hard to be able to compete with high frequency traders without being committed to a similar level of financial investment which is simply not feasible for many traders (Aitken et al., 2015). Therefore, O'Hara (2014) points out that in order to survive and thrive in the new trading environment with the presence of HFT, the most practical thing that low frequency traders can do is to try to make their activities more difficult to identify and reduce the likelihood of being on the radar of predatory high frequency algorithms. This goal can be achieved by using more suitable and smarter strategies for the execution of trade orders such as gaining access to the liquidity in dark pools and / or dividing an original parent order into a number of smaller child orders, as mentioned previously.

Other than the interaction between high frequency traders and the rest of the market as well as the effect of the former on the latter, the issue of fairness in high frequency financial markets can also be examined from different perspectives, namely the principle of equal outcomes or equal opportunities (James and Douglas, 2013). Equal opportunities mean that everyone has to receive the same treatment and is subject to the same regulatory procedure and rules with regards to their trading, and no one is allowed to receive any unjustified preferential treatment from exchanges or policy makers. Nevertheless, O’Hara (2015) gives a warning that it is not a simple task to reach the final conclusive answer for the question of whether the market is indeed unfair to a certain kind of investors or not since it is not easy to come up with an exact definition of fairness and even more challenging to provide an accurate measurement of fairness in the setting of financial markets. As a result, the author suggests that a viable solution to this problem for researchers may be to indirectly evaluate fairness after the fact by using the willingness of traders to take part in the market as a proxy for fairness. The rationale behind this recommendation is that the level of fairness is expected to be directly and positively correlated with how willing investors are to join the market so the willingness to trade (or lack thereof) is likely to be an accurate and reliable reflection of how fair (or unfair) the market is to its participants.

Regardless of the approach and proxy used for the assessment of fairness, James and Douglas (2013) place strong emphasis on making a clear distinction between a tool and its user. It is important to note that as a matter of fact, some HFT strategies are beneficial to the market and can improve its quality (i.e. market making, for example) so it is not fair to consider the whole practice of HFT unfair. To make a correct and well-informed judgement on fairness, one has to realise that fairness is not necessarily about the high speed of trading but, more importantly, about what high frequency traders use that speed for. In other words, both high and low frequency traders may choose to engage in legitimate or predatory activities, which is what truly matters for fairness. Up to this point in the chapter, I have already considered how HFT can affect the quality of markets and daily activities of market participants in a wide range of both positive and negative ways. In the next section below, I am going to investigate the influence of HFT on another integral part in the system of financial markets, namely trading exchanges and market regulators, and how these entities have to adapt in response to many new changes brought about by high frequency traders.

2.2.3. Impact on trading venues and regulations

2.2.3.1. Impact on trading venues

Generally speaking, the competition among rival trading venues to be able to attract new customers and also at the same time retain existing ones has become more and more intense over the years due to a few reasons. Firstly, the higher level of sophistication of advanced trading algorithms developed and used in recent times has allowed many traders to better cope with the current condition of market fragmentation and hopefully improve their order execution and trading results by actively looking for and quickly gaining access to the most desirable liquidity available among an increasing number of alternative potential trading locations. The second reason for the increase in competition is the fact that several market regulations of the SEC so far including Reg. ATS (the Regulation of Alternative Trading System) and Reg. NMS (the Regulation of National Market System) have opened the door for non-exchange entities to enter the market and compete directly with existing trading exchanges (O'Hara, 2015). As a result, each and every exchange has to try its hardest in an attempt to outperform its competitors (i.e. not only other exchanges but also other types of trading platforms such as dark pools) and be the best choice for traders to meet their trading needs, particularly when it comes to liquidity (Aitken et al., 2015).

Even though all exchanges want to rapidly expand their customer base as much as they can in order to boost their chance of success in this highly competitive environment, there is a dilemma that they need to face with regards to the participation of high frequency traders. At first glance, it seems obvious that they are valuable clients and their presence should be greatly appreciated by exchanges because of a simple reason. It is a well-known fact that high frequency traders account for a large part of trading volume in the market and make a substantial contribution to the high level of activities observed on numerous financial exchanges on a day-to-day basis not only in the US but also in various countries around the world, especially in recent periods. However, despite the huge potential for additional revenue, exchanges also have to take into account the potential loss of revenue from the rest of their client base if they allow too many high frequency traders to trade on their platform together with all of the other traders (O'Hara, 2015). As long as other traders do not find it favourable for them to co-exist with high frequency traders in the same ecosystem and do not want to run the risk of becoming an unfortunate victim of certain HFT practices, they will probably be discouraged from using trading venues which are too friendly to HFT and thus will be more likely to choose other locations to carry out their usual trading activities.

Due to the possibility of a conflict of interest between high frequency traders and everyone else in the market, it is necessary for exchanges to carefully and thoroughly consider a number of concerns related to the issue of fairness which has already been discussed in details. To ensure that the market is not unfair for anyone in the presence of HFT and that exchanges can manage to strike a balance between making money and sustaining the marketplace, indeed there is not an easy answer to the key question of which level of services should be offered to high frequency traders. Whether exchanges provide them with the same services as everyone else or superior services if they can afford (e.g. more specialised types of trade orders, faster communication with the order matching engine with the use of co-location, earlier access to essential information, among other things) will have significant implications and make a big difference for other market participants. O'Hara (2015) shows that in order to meet the diverse demands of different kinds of traders, new trading venues tend to adopt novel market designs with special and attractive features which are tailored to satisfy specific requirements of certain investors. On the other hand, to reach the same goal of being able to respond to a wide variety of trading needs, many existing venues decide to create markets within markets with separate sets of microstructure properties which are intended to accommodate and benefit particular subsets of their client base. As a result of these innovative moves, trading in general has become fragmented but still streamlined at the same time.

In terms of markets which aim to support HFT and enable high frequency traders to utilise their speed advantage to a great extent, notable features of the trading mechanism include (i) sophisticated order types to help improve trade execution while also optimising the cost of trading, (ii) extremely fast access to information in the marketplace and (iii) equally fast communication between traders and their trading platform. For instance, at the moment it is possible to transmit data and information from the NASDAQ stock exchange to the Chicago Mercantile Exchange in only 4.13 milliseconds and the NYSE has plans to reduce this tiny period of delay even further (O'Hara, 2015). These factors are among the most important things for high frequency traders to keep in mind before they make the final decision about where to trade. In the meantime, exchanges must strive to offer the best trading conditions possible and use those as their competitive edge and selling point to attract and retain high frequency traders who, in turn, will help exchanges to survive the fierce competition with their rivals.

In stark contrast to HFT-friendly trading venues, other markets are designed in such a way that can limit HFT activities and discourage high frequency traders from joining the marketplace so that non-high frequency traders can be protected from predatory

HFT practices and conduct their trading with more confidence. One of the methods that can be used for the purpose of mitigating the speed advantage of high frequency traders is to adopt a new rule regarding the priority of order execution which favours brokers (i.e. if there are multiple trade orders submitted at the same price by both a broker and a high frequency trader, the order of the broker will be prioritised and executed prior to the order of the high frequency trader). Another approach to create a similar effect is to add a certain delay to orders coming from high frequency traders to slow them down (O'Hara, 2015). Together with trading venues, market regulations have also undergone many substantial changes as a result of HFT which are going to be studied in more details in the following section.

2.2.3.2. Impact on regulations

Market regulators and policy makers have paid a lot of attention to HFT due to its tremendous and rapid growth in trading volume and activities as well as its numerous potential effects on vital components of the financial system, namely financial markets, market participants and trading venues (Sun et al., 2014, Aitken et al., 2015). In addition, there are even more reasons to keep a close eye on HFT because of the fact that high frequency traders have been found to contribute to and exacerbate the Flash Crash, which refers to a remarkable event that occurred in the US on May 6, 2010 and led to incredible instability and turmoil in not only financial markets but also the entire economy at large. This event was aptly named the Flash Crash to correctly reflect its nature which is a massive depreciation in market value within a very short amount of time. Easley et al. (2011) conduct research on the Flash Crash and show that it has resulted in the largest price drop in a trading day ever seen in the long history of the Dow Jones Industrial Average index. To be more specific, on the day of this tragic event, the index lost up to 1000 points and as much as $1 trillion in market value disappeared into thin air in a matter of minutes.

More importantly, Jain et al. (2016) argue that the high degree of correlation among trading activities of high frequency traders (which has been examined in a previous section) may be the cause of an increase in the level of systematic risk and may lead to extreme conditions of the market such as the Flash Crash. In spite of the great importance of this mechanism, it has not yet received sufficient attention and deserves more scrutiny from scholars as well as market regulators who still place too much emphasis on individual stocks instead of their correlation and systematic risks when it comes to rules and regulations. Those authors stress that it is definitely necessary for the government to pay special attention to the correlation among financial assets and consider this fundamental factor more seriously when monitoring and regulating markets in order to prevent similar events to the Flash Crash from happening again in the future. Their research paper uses an extensive dataset of a developed equity market outside the US, namely Japan, which consists of quote and trade data of stocks on the Tokyo Stock Exchange from 2008 to 2011. In a related study, Benos et al. (2017) utilise unique high frequency transaction data in the electronic order book of the London Stock Exchange in 2012 which include various information about executed trades of all of the constituent stocks in the FTSE 100 index (i.e. a common proxy for the UK equity market), with timestamps to the nearest second and identifiers to clearly distinguish HFT firms from other traders in the market. Their findings show that a number of key indicators of trading activities (such as the flow of trade orders, total volume and net positions for example) exhibit a remarkably higher degree of similarity among high frequency traders in comparison with another group of market participants, namely investment banks.

In addition to the potential to add to systematic risks, there has been an ongoing and heated debate with regards to the issue of social welfare in relation to the practice of HFT as a whole, in other words, the overall impact and contribution of HFT not only in the context of financial markets but also in the broader setting of society (Biais et al., 2015, Budish et al., 2015). The tough competition among high frequency traders to be the fastest has naturally led them, whether they want to or not, to enter a technological arms race all the time in order to be able to keep up with and hopefully surpass one another in terms of trading speed. Nevertheless, the constant escalation of investment in the most advanced and sophisticated infrastructure available to be used for the sole purpose of communication and trading has amounted to a huge sum of money which could have been more productive if spent on more meaningful endeavours for the sake of financial markets in particular as well as the economy and society in general. The total amount of capital invested in this endless arms race is so massive that it raises the question of whether or not HFT adds value and has a positive impact at the end of the day (Chordia et al., 2013).

Given the various concerns about HFT, policy makers in many countries have considered different measures whose purpose is to neutralise or at least alleviate as much as possible the detrimental effects of HFT, especially predatory trading practices (Robert and Philip, 2012, Benos and Sagade, 2016). In order to regulate HFT activities more effectively, the first important step for the governing body is to view the market from the same point of view as that of high frequency traders. In this aspect, the SEC has made a big improvement in catching up with the rapid evolution of HFT by setting up MIDAS (short for Market Information Data Analytics System). MIDAS is a software package specifically built for HFT by one of the leading HFT firms to provide market regulators with the same set of trading tools often used by high frequency traders and hence put them in the shoes of the type of traders that they try to keep under their surveillance (O'Hara, 2014, O'Hara, 2015). Although modern applications in technology such as the MIDAS software are a fundamental and significant achievement in and of themselves, they are just the first step towards the ultimate goal of a well regulated system of financial markets. What is more important is how these applications can inform and build a solid foundation for the next step where the government have to issue appropriate regulations or changes to existing ones based on the information and data obtained from them.

On the other hand, Manahov et al. (2014) state that even though there has been a great deal of research on HFT, unfortunately the available evidence and findings are still not conclusive in terms of the impact of high frequency traders. Therefore, a heated debate is still in progress for regulatory bodies all over the world regarding whether or not (and how) to impose some kind of limits on HFT or even to take a very extreme measure such as banning it completely. However, in practice, a number of rules and requirements have already been introduced to the market in an attempt to address the issue of HFT. For instance, now high frequency traders are usually asked to (i) submit firm quotes at competitive prices to provide a reliable and regular source of liquidity as market makers, (ii) pay taxes on their transactions, (iii) refrain from excessive cancellation of their trade orders and excessive trading based on the use of computer algorithms to minimise malicious and harmful trading behaviours such as the technique of quote stuffing, and / or (iv) describe their automatic trading programme in details and perform all of the necessary tests before actually applying them to the market in order to avoid regrettable situations such as flash crashes which could have been prevented in the first place (Sun et al., 2014, O'Hara, 2015, Benos and Sagade, 2016).

Regarding the rules, O'Hara (2015) is concerned that while many of them are reasonable, some are not really suitable for the current situation of financial markets (e.g. transaction taxes or restriction on order cancellation and the use of algorithms). Furthermore, since it has been found that a large percentage of high frequency traders plays the important role of market makers who are beneficial to the market and hence should be encouraged, Hagströmer and Nordén (2013) argue that policies may potentially do more harm than good if they are applied indiscriminately to everyone. As a result, to maximise the positive effects and minimise the negative effects of regulatory efforts on markets, it is advisable for regulators to separate predatory HFT from normal HFT in one way or another and then target their regulations at a specific group based on this distinction rather than the whole population of high frequency traders in general. Now that I have investigated the influence of high frequency traders on regulatory issues, my analysis of how HFT can have an impact on the system of financial markets is finally complete. In the next section, I am going to examine the last issue in this chapter, namely the evolution of financial information and data in the high frequency era.

2.3. Some issues of high frequency data and information

Compared to traditional low frequency data sampled at long intervals, high frequency data sampled at short intervals generally have some notable differences in three main aspects including (i) properties and patterns found in the data, (ii) interpretation of the data, and (iii) difficulties and challenges in dealing with the data.

2.3.1. Properties and patterns in the data

In terms of the properties of the data, Andersen and Bollerslev (1998) find that forecasts of volatility using high frequency data tend to be more accurate and reliable than those using low frequency data over daily or longer horizons, based on their study on daily and five-minute exchange rates of two of the most actively traded currencies worldwide (i.e. the German mark and the Japanese yen) against the US dollar during a six-year period from October 1987 to September 1993. In a follow-up research paper, Andersen et al. (1999) analyse the German mark – US dollar exchange rate and extend the length of the dataset in the previous study to cover as much as a decade (i.e. between December 1986 and November 1996) which also includes the original time period of study. Again, they confirm their earlier findings that it is indeed possible to enhance forecasts of volatility substantially in practice with the use of high frequency returns.

Another example to illustrate the differences in data patterns and properties at different sampling frequencies is the issue of market efficiency and the profitability of technical trading rules. Taylor and Allen (1992) examine the foreign exchange market and report that technical analysis becomes more applicable when there is an increase in the frequency of trading. In other words, technical analysis plays a more important role in high frequency data than in low frequency data and it may be a good idea for traders to incorporate at least some technical rules into their overall trading plan when the sampling frequency of their data is high. A possible reason for the increased significance and effectiveness of trading rules at high frequencies may be that by using data sampled at these frequencies, traders are able to identify complex but lucrative patterns and dynamics which cannot be observed in low frequency data and then take advantage of this potential lack of efficiency in financial markets to earn positive excess returns (Goodhart and O'Hara, 1997). In addition, as markets become more and more mature and efficient, the duration of trading opportunities such as arbitrage is expected to be shorter and they are likely to disappear more quickly so it may be the case that these trade signals are impossible to capture and act upon without the use of high frequency data.

In line with this explanation, Narayan et al. (2015) conduct an empirical study on the currency market based on a dataset of 5-minute exchange rate observations and provide evidence to show that it is in fact profitable to trade the foreign exchange market and there is a considerable amount of money to be made. However, Hudson et al. (2017) take one step further and argue that not all of the technical trading rules are created equal because there are two types of rules, namely those based on price momentum and those based on price reversal, which may or may not be able to generate a favourable trading performance depending on the frequency currently used to sample the data. More specifically, when the frequency increases, the performance of the former rules tends to deteriorate whereas that of the latter rules tends to improve and vice versa when the frequency decreases, which may have something to do with the high level of noise often contained in high frequency data due to HFT activities. The presence of a lot of noise means that the market price will usually fluctuate around a certain area instead of following a strong and smooth trend in one direction. As a result, such a condition will probably be more suitable for trading strategies based on mean reversion so they are likely to perform well and generate positive returns on investment while those based on momentum or trend following are not.

2.3.2. Interpretation of the data

There are interesting differences between high and low frequencies in terms of not only the behaviours of market data mentioned above but also the way that data can be interpreted or in other words, the information content of data. Regarding the information content of high frequency data as opposed to low frequency data, there are several important and noteworthy changes in (i) the significance of certain types of information, (ii) the nature of information and (iii) the visibility of information. Firstly, in terms of the significance of information, the traditional view puts more emphasis on fundamental information which is information directly related to the intrinsic value of financial assets (e.g. financial statements of companies in the case of stocks) or the overall condition of markets and the economy (e.g. macroeconomic data such as interest rates and unemployment rates), and less emphasis on microstructure information which is technical details about trading activities of securities in the market (e.g. the flow of orders in the order book of trading venues). On the contrary, in current high frequency financial markets, microstructure information is just as necessary as fundamental information or perhaps even more so (O'Hara, 2015).

Secondly, with regards to the nature of information, Easley et al. (2016) report that at some trading venues, market participants who can afford it can now purchase the right to gain early access to essential market information such as the order flow a short time before it becomes publicly available to others. As a result, for a very brief period of time ranging from a few seconds to merely a few milliseconds, the information which is supposed to be public information and accessible to everyone has turned into private information that can only be seen by a limited group of elite trading firms who will try to exploit their informational advantage as much as possible to make a profit, probably at the expense of others. However, one should note that it may still not be possible to completely prevent or eliminate this potentially unfair advantage even if exchanges offer the same access to information to each and every investor in order to properly address various concerns about the issue of fairness and hopefully avoid being accused of favouritism and discriminatory treatment. The reason is that the inherent speed advantage of high frequency traders due to their use of state-of-the-art technology and possibly co-location can still allow them to get their hands on new information more quickly than other traders even if this information is released to everybody at the same time.

Thirdly, regarding the visibility of information, the previous literature on the microstructure of financial markets is mostly interested in the active side of a transaction in the form of a market order executed at the price quoted by a specialist or against a limit order in the order book, because it is believed that market orders are often driven by information and thus they represent the underlying information (O’Hara, 2015). The commonly adopted rationale is that investors in possession of a certain piece of information can only make use of their information by actively initiating a transaction via market orders, specifically, those informed of good news need to actively place a buy order to make a profit while those informed of bad news have to actively place a sell order to profit from their news. Nevertheless, traders and their trading algorithms in the high frequency environment usually operate in a way that is more complex and not so straightforward. In order to optimise the execution of orders and hide their trading activities from predatory computer programmes at the same time, traders now have the tendency to divide the original orders (also called parent orders) that they want to execute, especially large ones, into a number of smaller orders (also called child orders) which should be easier to fill. When some or all of these child orders eventually get filled and turn into actual transactions, those resultant trades are supposed to be related to one another since the child orders themselves come from the same parent order so they are not entirely independent. Therefore, series of trades become informative and quite possibly even more so than individual trades, which makes it more difficult and complicated to fully capture and study the elusive information that forms the basis of trading decisions of participants in the market.

On the other hand, the practice of splitting orders can also help shed some light on why more attention should be paid to the use of odd lots in trading which has become more common and played a more essential role in the context of HFT today. To have a general idea about the popularity of odd lots, one can refer to the work of O'Hara et al. (2014), which points out that on average, out of every five transactions in equity markets, one is made using odd lots while for some stocks, at least 50% of all of the trades are made using them. In addition, it is also shown in the same study that not only are odd lots used on a regular basis in trading but they generally have high information content as well, which is consistent with the idea that many informed traders take advantage of odd lots to be able to execute their transactions while simultaneously still managing to keep their information and intention hidden from other traders as much as possible. Information in high frequency markets has become less and less visible due to not only the fragmentation of trade orders but also the reduced use of market orders to take action on new information even though this type of orders has traditionally been considered to be a valuable and intuitive signal to the underlying information as explained above. Because of the fee structure in use at many trading venues, investors have found that in general using market orders does not seem to be a cost-effective way to trade, which is the reason why they have switched from market orders to other less costly alternatives such as limit orders in their effort to reduce the cost of trading.

O’Hara (2015) provides a good example of how strategies for trade execution in the high frequency world may mask the information behind observed transactions and even make market data become misleading, resulting in the possibility of an incorrect interpretation. Let us consider a scenario as follows: a trader has just received favourable information about a certain company and expects its stock price to go up as a result so he / she wants to make money from this particular piece of information by placing a large buy order to purchase a lot of shares in this company prior to its highly likely appreciation in the future. Previously, in most cases such an order would simply translate to one or maybe just a few market buy transactions. Conversely, computer algorithms used nowadays tend to (i) divide this order into many buy limit orders, (ii) submit them to one or more order books on multiple exchanges, and then (iii) update or cancel them later if necessary. When new sell market orders from other traders are submitted to the order book and matched against those buy limit orders submitted earlier, the resultant transaction will be registered on records as a sell trade because the active side of this trade (i.e. market order) wants to sell. Due to this mechanism, in the end the original intention to buy which is motivated by important information will ultimately show up as a sell trade in the historical transaction data of the exchange.

2.3.3. Difficulties and challenges with the data

Given the current landscape of financial markets with an incredibly huge number of financial assets, market participants and trading venues, it is really hard to imagine the vast amount of data and information generated by the whole system in each and every trading day. More importantly, innovations in technology over time have allowed this massive volume of data to be distributed more widely and hence an increasing number of new and more detailed datasets have become accessible to practitioners as well as academics for trading and research purposes. However, the sheer magnitude of available data also poses a number of challenges. These days datasets, particularly high frequency ones, are not only more expensive to gain access to but also more difficult to store and carry out analyses on, even basic ones (O'Hara, 2015).

In addition to the quantity, the nature of high frequency data may also create some problems for users during the process of sampling and analysis. To understand this point more clearly, let me take a close look at the highest frequency of data possible, namely tick data. Tick data are data in their natural form, which means that a new data point is recorded as soon as there is a new instance of the type of events that I am interested in so every point in the data represents and corresponds to a unique instance that has happened. For example, if I am interested in financial transactions, I will have a new observation for every trade that has been made while if I am interested in quoting activities, I will have a new observation for every quote that has arrived at the market. In other words, tick data capture everything that has occurred and it is just not possible to have a higher frequency for data sampling than the tick frequency. Therefore, tick data are the building blocks of data at lower frequencies which are sampled by selecting tick observations at regular intervals and ignoring the remaining tick observations so every data point at lower frequencies often represents more than one instance of the event of interest. Because of the procedure of data sampling described above, one can say that tick data are natural and lower frequency data are artificial in a sense. Also, unlike time series sampled at lower frequencies which are temporally aligned with one another, tick observations by definition do not appear at regular intervals, which raises an interesting but challenging issue in terms of the treatment of time. Specifically, traditional techniques for the analysis of financial time series are designed to handle temporally aligned data so now the question is how previous statistical techniques can be adapted and / or novel alternative solutions can be introduced to provide insights into the contemporaneous dynamics of temporally misaligned high frequency time series.

Unfortunately, the treatment of time is not the only potential problem with high frequency data that needs to be addressed as O'Hara (2015) is concerned that there may also be other issues with quote data due to a number of common techniques currently used to conduct quoting activities. More often than not, the submission of quotes will soon be followed by a lot of revisions, cancellations and resubmissions, which leads to the widespread phenomenon of flickering quotes and uncertainty with regards to the actual price level of financial instruments at the moment and thus poses a new challenge for the empirical analysis of quote data. On the other hand, the findings of Hasbrouck and Saar (2013) show that this instability of price quotes will result in several negative and undesirable consequences such as a reduction in the information content of quotes and a higher level of risk associated with order execution for investors, among other things. It has also been emphasised that the volatility caused by trading at extremely short horizons is much higher than that caused by trading at longer horizons based on fundamental information (either public or private), which may explain why it is usually the case that market data, especially price data, behave in a more erratic way at short trading horizons compared to long ones. According to the study of O'Hara (2015), problems and difficulties with the data on stock markets are the main reason why the SEC took a fairly long time (i.e. nearly six months, to be specific) to carry out their investigation of the Flash Crash on May 6th, 2010.

2.4. Chapter summary

The current chapter has introduced readers to high frequency financial markets by (i) explaining the practice of HFT, (ii) demonstrating how popular and significant HFT is in markets at the moment, and (iii) describing different kinds of HFT activities. After providing this overview and background knowledge, the chapter has proceeded to examine the various ways that high frequency traders can have an impact on the system of financial markets including the markets themselves, participants in these markets as well as trading venues and regulatory bodies. More specifically, with regards to markets, I have studied the influence of HFT on three key components of market quality, namely (i) market efficiency, (ii) volatility and risks, and (iii) liquidity and the cost of trading. Meanwhile, in terms of market participants, I have discussed the interaction (i) between high frequency traders and their trading algorithms, and (ii) between high frequency traders and other traders. Finally, following the discussion about the effect of HFT, I have presented some interesting issues about data and information at high frequencies and pointed out a few notable differences between them and their low frequency counterpart regarding (i) characteristics of and patterns in the data, (ii) making sense of the data and (iii) challenges and difficulties with the data. In general, the arguments are mixed and the evidence is not conclusive with regards to the overall impact of HFT since high frequency traders have been found to affect the ecosystem of markets positively in some ways but negatively in others. However, it is universally agreed that one thing is for certain: high frequency traders have become a major force and played a significant role in the marketplace, which has gradually attracted more and more special attention from market regulators and policy makers. The presence of HFT has introduced many fundamental changes in not only how financial markets work but also how market data and information should be treated.

Chapter 3

High frequency pairs trading in gold ETFs

Abstract

Based on a large dataset of gold ETFs, I find arbitrage opportunities in the gold ETF market which can be exploited by high frequency traders. To my knowledge, this is the first paper to study pairs trading of gold ETFs using tick data. Able to execute their orders with minimal delay and take advantage of potentially short-lived opportunities, high frequency traders can make a positive excess return after including transaction costs. Consistent with Grossman and Stiglitz (1976) and Grossman and Stiglitz (1980), this profitability may be compensation for arbitrage efforts and incentivise arbitrageurs to eliminate mispricing. I also explain why the trade exit rule of full convergence used in previous studies may not be optimal and propose a rule based on partial convergence which outperforms the standard full-convergence rule. Specifically, changing the exit rule from full convergence to partial convergence can increase pairs trading returns and also enhance the risk-adjusted performance. Therefore, partial convergence enables better exploitation of arbitrage opportunities and more accurate reflection of market inefficiency than full convergence.

3.1. Introduction

Arbitrage and the Law of One Price are among the most important and fundamental principles in financial markets and thus they have been studied to a great extent in the previous finance literature. On one hand, a number of research papers have documented the existence of mispricing instances which may actually be exploited to make a profit (e.g. Froot and Dabora, 1999, Gatev et al., 2006, Gagnon and Karolyi, 2010, Do and Faff, 2012, Marshall et al., 2013). On the other hand, many studies have emphasised limits to arbitrage (e.g. De Long et al., 1990, Shleifer and Vishny, 1997, Abreu and Brunnermeier, 2002, D'Avolio, 2002, Grossmann et al., 2007). In particular, pairs trading, which is also commonly referred to as statistical arbitrage, is one of the most popular arbitrage strategies and it is widely used by many large financial institutions such as hedge funds and investment banks, among others. Interestingly, the basic underlying idea of this well-known trading strategy is in fact fairly intuitive and quite easy to understand. The entire process of pairs trading can be quickly explained and summarised in a few simple and straightforward steps as follows – (i) the first and most crucial step is to find a pair of closely related stocks (or other financial assets of interest) which tend to move together in the same direction with similar magnitude on a regular basis, (ii) then the second step is simply to wait for a sufficient divergence in their prices (where they become mispriced in relation to each other), which is likely to happen sooner or later for a variety of reasons, and (iii) the third and also last step is to take advantage of the highly likely subsequent convergence (because the price divergence is only temporary in most cases) by buying the underpriced stock and selling the overpriced one at the exact same time (Gatev et al., 2006).

In essence, the reason why pairs trading in general has usually been found to work and be able to generate a positive return on investment is that even though similar instruments may sometimes diverge in price due to the imperfect efficiency of financial markets, in the end this state of disequilibrium will go away and these instruments will eventually resume their normal behaviours and relationship with each other as a result of the combined trading activities of arbitrageurs (Do and Faff, 2010). Pairs trading can help maintain and enhance the level of market efficiency since arbitrage activities will lead to a reduction in not only the frequency but also the magnitude of mispricing instances between securities (Kondor, 2009). Similar to

Kondor (2009), Alsayed and McGroarty (2012) examine the Law of One Price in the context of the ADR [1] market and show that pairs trading indeed plays a major role in maintaining the price parity between stocks and their corresponding ADRs. It should be pointed out that these findings are consistent with the argument of Grossman and Stiglitz (1980) that (i) markets are not automatically efficient on their own from the start and they are also not perfectly efficient at all times, (ii) there must be some inefficiencies and potentially lucrative trading opportunities as a kind of incentive and reward to motivate traders to take part in the market and profit from these opportunities, and (iii) because of their trading endeavours, market inefficiencies will gradually be reduced and markets will become more efficient over time.

The motivation of my study is the fact that pairs trading has generally been shown to be profitable in the literature (e.g. Gatev et al., 2006, Do and Faff, 2010, Do and Faff, 2012, Jacobs and Weber, 2015). In one of the seminal papers on pairs trading, Gatev et al. (2006) find that this strategy is able to generate a return of as much as 11% on an annual basis while more recently, Jacobs and Weber (2015) provide evidence that it is possible for traders to consistently earn at least 12% per annum in profits with pairs trading. However, the application of pairs trading based on high frequency and even ultra-high frequency data has not received much attention in earlier works so my research will focus on the analysis of pairs trading using data at the highest sampling frequency possible, namely tick data.

Research on high frequency data has become more important and interesting in the current setting of our financial markets where information is gathered and actions are taken at an increasingly high speed (see Goldstein et al., 2014 for an overview of high frequency trading which has also been presented in details in the preceding chapter). In today's market, even being just marginally faster than other competitors can make a big difference to one's trading results and decide whether their transactions are going to make money or not, especially during time periods when markets are likely to be more sensitive than usual with a high level of overall volatility (e.g. public announcements of significant news and data with regards to the macroeconomic situation in general or a certain industry or company in particular). As a result of how

[1] ADR stands for American Depositary Receipt.

essential speed is to their trading performance, many high frequency trading firms have been doing whatever it takes to gain and maintain their valuable competitive advantage over their rivals in terms of speed, from investing more and more heavily in advanced technology and equipment to co-locating their trading systems in close proximity to the matching engine inside the trading venue. Hasbrouck and Saar (2013) find that because of this constant arms race, the reaction time of traders to information and signals from the market has been dramatically reduced to no more than a few milliseconds, which is hundreds of times faster than the average time that it normally takes a person to perform an eye blink. High frequency traders are the type of traders that generates the most trading activities on a daily basis and thus, it is not difficult to see why they are among the most important participants in financial markets at the moment. I find that the strategy of pairs trading is indeed profitable for high frequency traders. The excess returns from this trading strategy may be considered an incentive to encourage arbitrage activities as well as compensation for a number of risks and costs usually involved in this kind of trading (Grossman and Stiglitz, 1976, Grossman and Stiglitz, 1980).

Most of the previous studies on pairs trading in the literature (e.g. Gatev et al., 2006, Do and Faff, 2010, Do and Faff, 2012) often rely on the criterion of full convergence (i.e. the two financial instruments are required to converge completely after their temporary divergence in price) to exit currently open pairs trades. I explain in details why the commonly used rule for liquidation of positions based on full convergence may not be optimal and may lead to consistent underestimation of the profitability of pairs trading and the level of market inefficiency. As a result, I propose a better trade exit rule which is based on partial convergence as a superior alternative to help correct the existing bias and improve the standard mechanism of pairs trading. Partial convergence is achieved when the two stocks (or other securities) converge to some extent but not completely and thus it is a more flexible and achievable condition than the strict condition of full convergence. Then I show that the partial convergence rule performs more favourably than the full convergence rule in different sample periods, enabling more effective and profitable exploitation of available pairs trading opportunities as well as more accurate reflection of market inefficiency.

My study on pairs trading in high frequency markets is conducted in the context of the ETF market, which has been growing at a fast speed (Shin and Soydemir, 2010, Caginalp et al., 2014, Kearney et al., 2014). First of all, my choice of market is motivated by earlier findings of mispricing instances of ETFs in the literature (e.g. Ackert and Tian, 2000, Engle and Sarkar, 2006). If mispricing does in fact occur between ETFs, then the ETF market is an interesting and potentially profitable environment for pairs trading which is a strategy designed specifically to take advantage of the relative mispricing between financial assets. Moreover, one can say that this setting is also appropriate for an arbitrage study because in general ETFs are quite easily accessible to most traders and less risky to apply arbitrage strategies to than in other settings (Marshall et al., 2013). More specifically, divergence risk (i.e. the risk that the two securities do not converge or they take a really long time to do so) is fairly low for two reasons as follows. Firstly, investors put their money in ETFs in order to continuously track the current performance of the underlying index or asset so as a way to attract the interest and attention of many investors, the issuing firms of ETFs always need to try to the best of their ability to minimise any potential tracking errors and make sure that their ETFs are actually able to do what they are supposed to (i.e. follow the underlying as closely as they possibly can). Secondly, if an instance of mispricing between ETFs and their underlying happens for one reason or another, it is possible for investors to exchange their shares in the ETFs for the underlying asset or index constituent stocks so that they can make some money from the mispricing, which should act as a natural mechanism to effectively correct the mispricing and keep prices in equilibrium (Engle and Sarkar, 2006).

Among the different types of ETFs available in the market at the moment, I place my focus on gold ETFs in particular in this study. As a result of the significant appreciation of gold in the first decade after 2000 (Pullen et al., 2014), interest in and attention to gold and gold investment have been on the rise in the academic literature (see O'Connor et al., 2015 for an informative overview). Even though the price of gold was only $250 per ounce in 2001, after a period of slightly more than ten years, it has increased six times to as much as $1500 or more per ounce in 2012 (Blose and Gondhalekar, 2013). More importantly, over the years gold has continued to reinforce its solid position as an essential asset and commodity in investment portfolios that a lot of investors often cannot do without, especially in times when overall market conditions and environment are tough to trade, because of its safe haven properties which have been well documented in previous research (Baur and Lucey, 2010, Baur and McDermott, 2010, Bredin et al., 2015). In addition, the World Gold Council also reported that there was a substantial increase in the demand for gold in general and gold ETFs in particular as viable instruments for investment during the period from July 2008 to March 2009, and this demand has continued to grow (Pullen et al., 2014). The turnover of gold around the globe was estimated to be about 4000 metric tons on a day-to-day basis, with a total trading value which can even be compared to that of all of the stock exchanges in the world combined [2]. In fact, gold is often considered to be a part of the foreign exchange market which is the most enormous market to have ever existed in history by any metrics, and the turnover of gold is only lower than that of the four most actively traded currency pairs (Hauptfleisch et al., 2016). Given the importance of gold and gold ETFs, my study on pairs trading of gold ETFs may have valuable implications for investors. I hope to provide insights into ETF pairs trading beyond the equity ETFs in Marshall et al. (2013). Because (i) pairs trading exploits the mispricing of ETFs which relates to their ability to track the underlying index or asset and (ii) this tracking ability differs among ETFs, pairs trading performance may differ between different types of ETFs. Indeed, I show that gold ETFs may be inefficient while Marshall et al. (2013) point out that their results do not necessarily support market inefficiency for equity ETFs.

I collect bid-ask quotes (time-stamped to milliseconds) of the two most liquid US gold ETFs, namely SPDR Gold Shares and iShares Gold Trust from 2005 to 2010. This is the first study to examine pairs trading of gold ETFs using tick data, which can be useful in several ways. Firstly, the use of ultra-high frequency data allows me to capture short-lived mispricings unobservable in low frequency data (e.g. daily) which is the data frequency often used in the literature. In fact, while I find very few arbitrage opportunities in the daily data of gold ETFs, there are a number of opportunities in the tick data. Following Marshall et al. (2013), I only consider relative mispricing of at least 0.2% and find a comparable number of arbitrage opportunities as in equity ETFs found in Marshall et al. (2013) on an annual basis. Secondly, I enhance the standard

[2] Bank of International Settlement's report on the global foreign exchange market activity. Available at http://www.bis.org/publ/rpfxf10t.pdf

pairs trading mechanism used in the literature by proposing a better trading rule based on partial convergence which is supported by both ex-ante justification and ex-post results. Thirdly, I show that pairs trading outperforms the buy-and-hold strategy on a risk-adjusted basis, which suggests that the gold ETF market may be inefficient.

This study proceeds as follows. In section 3.2, I review the literature on pairs trading performance and the explanation for such performance, and propose a classification scheme for pairs trading studies. Section 3.3 describes my data. Section 3.4 presents the methodology of my pairs trading analysis based on both full convergence and partial convergence. Section 3.5 reports and discusses the results. Section 3.6 provides a conclusion.

3.2. Literature review

3.2.1. Pairs trading performance

Gatev et al. (2006), testing pairs trading in four decades (1962 – 2002) with daily data, find that self-financing pairs trades generate up to 11% of annualised excess returns on average. Pairs trading profits are robust to transaction costs, short-selling costs and short recalls even in out-of-sample test. They also find that a large part of profits comes from short positions. However, their zero-investment assumption (i.e. short sale proceeds are used to finance the long position in a trade so no upfront capital is required) may not be realistic and hence some authors assume a 50% margin requirement (e.g. Mitchell et al., 2002, Marshall et al., 2013). In another paper, Alsayed and McGroarty (2012) indicate that arbitrage profitability does not correlate with broad market performance.

Do and Faff (2010), extending the study of Gatev et al. (2006) to 2009, confirm that although pairs trading performed well before 1990 (even during the 1987 crash), it has deteriorated since the 1990s (evidenced by regular unprofitable months). This deterioration may be caused by the reduced mispricing frequency in recent periods as a result of increasingly popular algorithmic trading (Akram et al., 2009). Nevertheless, Marshall et al. (2013), who study high frequency pairs trading between US ETFs tracking the S&P 500 index from 2001 to 2010, report greater profit magnitude from recent mispricings. Moreover, despite decreasing profitability, pairs trading shows favourable performance in tough times (Do and Faff, 2010) and more importantly, its risk-adjusted returns remain stable over time (Gatev et al., 2006).

Trade duration depends on the timeframe of pairs trading and data frequency. At one end, the average duration of trades on the daily timeframe is 3.75 months (Gatev et al., 2006). At the other end, on the tick timeframe, the strong mean reversion of pairs reduces the arbitrage length to minutes (Alsayed and McGroarty, 2012, Marshall et al., 2013).

3.2.2. Explanation for pairs trading profitability

The documented profitability of pairs trading results from the relationship between the two stocks in a pair, compensation for arbitrage efforts and microstructure effects. The relationship in a pair refers to both fundamental and technical relationship. Fundamentally, the two pair components are close substitutes when they belong to the same industry (Gatev et al., 2006). Do and Faff (2010) find that industry homogeneity has statistically significant impacts on pairs trading returns and increased granularity of industry classification can improve performance. Specifically, the 48-industry categorisation of Fama and French (1997) performs better than the four-group classification (i.e. Financials, Industrials, Transportation and Utilities). There is also evidence of industry-specific profitability (i.e. Financials and Utilities pairs outperform Industrials and Transportation pairs) because of industry-specific levels of company homogeneity.

The technical relationship refers to price behaviours of the two stocks in a pair. Marshall et al. (2013) find that price divergences are often followed by fast convergences. If the initial divergences are considered overreaction, the subsequent convergences might be reversals of overreaction. Furthermore, the past convergence frequency influences returns. Stocks whose prices often intersect in the formation period (i.e. when pairs are selected based on minimum sum of squared differences between normalised price series) are likely to be profitable pairs in the testing period. Adding the number of price intersections to pairs selection criteria increases mean excess returns (Do and Faff, 2010).

From another viewpoint, mispricing incentivises the enforcement of the Law of One Price [3] (Alsayed and McGroarty, 2012) and pairs trading profits may compensate for non-convergence risk, information cost and other risks and costs (Marshall et al., 2013). Nevertheless, only a part of the profitability is explained by exposure to five factors (i.e. market risk, firm size, firm value, momentum and reversal). Additionally, firm-specific volatility affects arbitrage performance whilst systematic volatility does not (Do and Faff, 2010). Finally, due to the contrarian nature of pairs trading, microstructure effects might bias its observed profitability upward (Conrad and Kaul, 1989, Jegadeesh and Titman, 1995). If trade prices are used instead of quotes, contrarian trades are likely to be taken on the wrong (and more favourable) side of the spread. If the current trend is up (down), the observed trade price is more likely to be at the ask (bid) so a contrarian strategy will mistakenly sell at the ask and buy at the bid. I use quote data to address this issue.

With regards to the downward trend in pairs trading profitability over time, Do and Faff (2010) conclude that it is attributable to an increase in market efficiency and arbitrage risks. To reach this conclusion, they use winning trades and losing trades to capture the market efficiency effect and arbitrage risk effect respectively. If the efficiency effect is significant, profits will be smaller and less frequent over time; if the risk effect is significant, losses will be increasingly large and regular. Although decreasing transaction costs have increased market efficiency by attracting more pairs trading, especially from hedge funds since 1989 (Gatev et al., 2006); only 30% of the performance deterioration is because of improved efficiency while the rest is caused by higher arbitrage risks (Do and Faff, 2010). Figure 3.1 demonstrates the profitability explanation.

[3] The Law of One Price implies that in efficient markets, financial instruments with the same cash flows should trade at the same price, regardless of their creation methods (Akram et al., 2009). More generally, two securities whose pay-offs are close to each other should have prices which are equally close to each other (Chen and Knez, 1995).

Figure 3.1. Profitability explanation. This figure explains pairs trading profitability, in general and over time. The stock relationship includes fundamental and technical relationship. The compensation covers arbitrage risks and costs. The arrows show the direction of explanation (i.e. the causes point to the effects). 'Profitability trend' points to 'overall profitability' because the trend is a part of the overall performance.

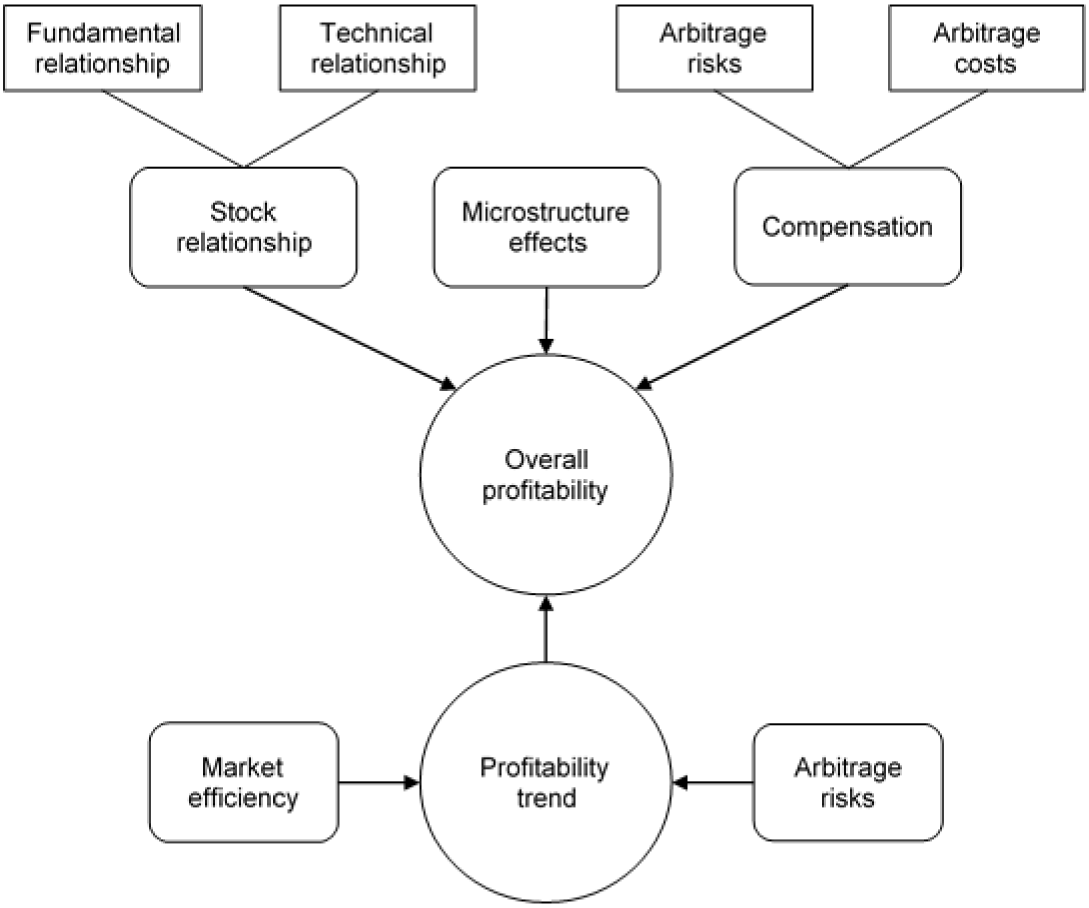

3.2.3. Limits to arbitrage

This section will discuss some important limits to arbitrage which traders have to face when attempting to exploit arbitrage opportunities. The most obvious impediment to arbitrageurs is transaction costs (i.e. bid-ask spreads and commissions) which they need to overcome before making any profit. In general, even when a trader manages to make accurate forecasts in terms of the market's direction, he / she may still lose money if the returns from his / her trades are not sufficient to cover transaction costs. Moreover, trading costs are an even bigger concern in the high-frequency context because on average a high-frequency trade tends to generate a lower return than a low-frequency trade (e.g. on the daily timeframe) due to a shorter holding period to accumulate profits. I address this issue by including costs in my analysis (i.e. bid-ask spread and commission).

In addition to costs, arbitrageurs must also take into account the risks associated with their activities, one of which is liquidity risk. Liquidity refers to the availability of financial instruments in the market at a given time or during a certain period, so liquidity risk is the risk of an asset not having enough readily available trading volume. As a result, it is not easy or even not possible for traders to buy or sell their desired amount of the security in question. More specifically, they will either (i) not place any order at all, (ii) place a partial order, (iii) place a full order but at a less favourable price or (iv) in the worst case, place a partial order at a less favourable price. All of these situations are not optimal for their trading strategy. My empirical analysis examines the issue of liquidity at the level of overall market as well as individual trades.

Another highly important risk in arbitrage is divergence risk. Since arbitrageurs can only make money if the market restores its equilibrium and mispriced instruments revert to their fundamental values, they run the risk of divergence where mispricing persists and the diverging assets do not converge at all or they only do so after a long period of time. Alsayed and McGroarty (2012) show that arbitrage activities are discouraged to a certain extent by this uncertainty in the duration of trades which may lead to a number of undesirable consequences. Firstly, institutional traders often have to report their performance on a regular basis so if a trade is still ongoing and not closed in time, they cannot realise the profit and show it in their report. Secondly, as long as a trade is open, it is exposed to adverse price movements and thus a winning trade at the moment may eventually turn into a loss. Thirdly, the longer a trade is taking place, the longer the trading capital is required for that transaction and hence not ready to finance another trade. However, divergence risk is fairly low for the ETF market in my study for two reasons. Firstly, the purpose of ETFs is to track the underlying so their issuers always try to minimise their tracking errors in order to attract investors. Secondly, traders can arbitrage away potential mispricing by simply exchanging their ETF shares for the underlying (Engle and Sarkar, 2006).

With regards to the explanation for persistent mispricing, Schultz and Shive (2010) suggests that sometimes mispricing is not corrected simply because it is not actually mispricing but rather a reflection of subtle differences in the characteristics of largely similar securities. For example, different classes of shares in the same company may seem identical at first glance but in fact offer different voting rights, and the class with better rights should trade at a premium. Nevertheless, the ETFs in my paper are identical because they are based on the same physical asset (i.e. gold). On the other hand, Abreu and Brunnermeier (2002) argue that mispricing may exist for a long time due to the sub-optimal co-ordination among arbitrageurs in eliminating market disequilibrium, or in other words, because they do not take advantage of the trading opportunity at the same time. This mismatched timing is the result of some traders (i) not having sufficient funds for a new trade due to their current commitment to other trades, or (ii) wanting to wait for an even more significant trading signal (i.e. larger mispricing). As arbitrageurs do not take action simultaneously, the price gap will disappear gradually rather than immediately. According to Kondor (2009), when the gap gets smaller and smaller, the profit potential decreases so new traders become less interested in exploiting it, which results in a lower level of trading activities and slower convergence towards the end of the process.

In addition to the interaction among arbitrageurs, the interaction between arbitrageurs and irrational market participants may also contribute to the persistence of mispricing and divergence risk (Black, 1986, De Long et al., 1990). By definition compared to arbitrageurs, irrational traders have the opposite expectation about future developments of market prices and will act on their beliefs by trading in the opposite direction. Therefore, the competition between these two groups determines whether future prices will be able to reflect the intrinsic value of securities. If arbitrageurs are stronger, prices will be brought back to equilibrium levels and mispricing will be removed sooner or later. Conversely if they are weaker, prices will be driven further away from fundamentals and mispricing will become more substantial. Although harmful to arbitrageurs, irrational investors only trade based on their own expectations and they unintentionally cause negative impacts on arbitrageurs. Arbitrageurs also have to consider the possibility of being exploited intentionally by predatory traders (Brunnermeier and Pedersen, 2005) who aim to make a profit at the expense of other market participants. To be more specific, if predatory traders can identify arbitrage transactions, they will trade in the direction of irrational traders to make the mispricing worse for arbitrageurs. Then if an arbitrageur cuts his / her losses (due to risk management or a margin call), he / she must exit the trade at a worse price than the entry price while the predators stand to benefit.

3.2.4. Pairs trading classification

In order to better position this research in the literature, I propose a classification scheme for pairs trading studies. Based on the strength of the fundamental relationship between the two stocks in a pair, these studies can be divided into three types, namely the loose form, semi-strict form and strict form. In my loose form, the relationship is purely statistical and the two stocks have no fundamental reason to move together. For example, two stocks in two unrelated industries (e.g. Starbucks and IBM) are likely to be influenced by different factors so they may behave differently and their observed relationship is only statistical. In my semi-strict form, the two stocks belong to the same industry or economic sector and they are likely to move together because both of them are affected by common factors (e.g. industry-specific regulation, supply and demand). For instance, if the demand for banking products is increasing and there are two reasonably comparable banks, it is expected that the demand for both banks will increase. As a result, their stocks should be affected in a similar manner and show similar movements. From an economic viewpoint, the products of the two firms in a semi-strict pair can be substitutes for each other to some extent. Finally, in my strict form, the two stocks have an inherent relationship which should force them to move together if markets are efficient. Specifically, they represent the same index or asset; in other words, they are different covers of the same content and close substitutes for each other (e.g. ETFs tracking the S&P 500 index).

As I move from the loose form to semi-strict form to strict form, the requirement becomes increasingly stringent and the number of eligible pairs decreases. My three types of pairs trading have a nested relationship as shown in Figure 3.2. Loose form pairs include semi-strict form pairs because pairs in the same industry (i.e. semi-strict form) constitute a subset of all pairs in the economy (i.e. loose form); and similarly, semi-strict form pairs include strict form pairs. Some examples of loose form and semi-strict form studies are Gatev et al. (2006) and Do and Faff (2010) where they match pairs of US stocks based on statistical criteria at first and then add the criterion of industry homogeneity by matching only stocks in the same category (i.e. Financials, Industrials, Transportation or Utilities). Strict form studies include Alsayed and

McGroarty (2012) and Marshall et al. (2013) who examine UK stock-ADR pairs and a US ETF pair tracking the S&P 500 index respectively. I classify the literature in Table 3.1. Based on this classification, my research on gold ETFs belongs to the strict form because these ETFs track the same asset so they have an inherent relationship.

Figure 3.2. Three types of pairs trading. This figure shows my three types of pairs trading and their nested relationship. Google-Ford pair belongs to the loose form since they operate in two unrelated industries (i.e. Internet information and automobile). Google-Yahoo pair belongs to the semi-strict form as they are in the same industry and offer fairly similar products and services.

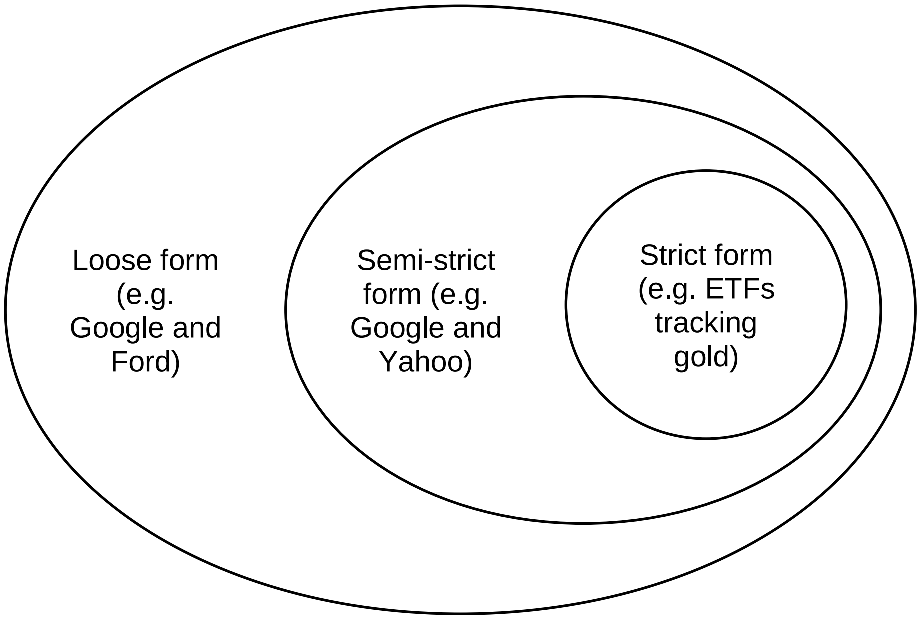

Table 3.1. Classification of the literature on pairs trading. Some authors study more than one type of pairs trading.

Type	Author	Market	Period	Data Type	Data Frequency	Findings
Loose form	Gatev et al. (2006)	US	1962 – 2002	Transaction price	Daily	It is possible to make a return of 11% p.a. from self-financing pairs trades.
	Do and Faff (2010)	US	1962 – 2009	Transaction price	Daily	Pairs trading has become less profitable since the 1990s.
	Do and Faff (2012)	US	1963 – 2009	Transaction price	Daily	Pairs trading is slightly profitable after considering trading costs.
	Broussard and Vaihekoski (2012)	Finland	1987 – 2008	Transaction price	Daily	The annualised return from pairs trading can be up to 12.5% and the profits do not relate to systematic risks.
	Jacobs and Weber (2015)	34 countries	2000 – 2013	Transaction price	Daily	Pairs trading is consistently profitable, generating more than 12% p.a..
Semi-strict form	Gatev et al. (2006)	US industry groups	1962 – 2002	Transaction price	Daily	Pairs trading is most profitable for utility stocks.
	Do and Faff (2010)	US industry groups	1962 – 2009	Transaction price	Daily	Utility and financial stocks are the most profitable. Finer industry classification can increase profitability.
	Mori and Ziobrowski (2011)	US REIT (real estate investment trusts)	1987 – 2008	Transaction price	Daily	The REIT market is more profitable for pairs trading than the general market during the 1993 – 2000 period.
	Do and Faff (2012)	US industry groups	1963 – 2009	Transaction price	Daily	The best pairs earn 3.4% p.a. on average.
Strict form	Schultz and Shive (2010)	US dual-class shares	1993 – 2006	Transaction price and bid-ask quote	2-minute	Pairs trading between share classes of the same firm generates abnormal profits after transaction costs.
	Alsayed and McGroarty (2012)	UK stocks and ADRs	2011	Bid-ask quote	Tick	Pairs trading activities help maintain the stock-ADR price parity.
	Broussard and Vaihekoski (2012)	Finnish common and preferred stocks	1987 – 2008	Transaction price	Daily	Pairs trading between common and preferred stocks of the same company can create significant profits.
	Marshall et al. (2013)	US equity ETFs	2001 – 2010	Bid-ask quote	Tick	The average profit of pairs trading is 6.57% p.a. after considering the bid-ask spread.

3.3. Data

I focus on the most liquid pair of US gold ETFs, namely SPDR Gold Shares (ticker symbol GLD, provided by State Street Global Advisors) and iShares Gold Trust (ticker symbol IAU, provided by BlackRock). I choose the most liquid instruments because of two reasons. Firstly, liquid securities are traded more often, which enables the timely capture of short-lived arbitrage opportunities. Secondly, liquidity is an important determinant of the profit potential since actively traded assets allow more money to be made. Introduced on 18th November 2004 and 21st January 2005 respectively, both GLD and IAU aim to track the price performance of gold bullion. I collect quote data (time-stamped to milliseconds) from February 2005 to May 2010 from Thomson Reuters Tick History. The dataset ends in 2010 because of data unavailability. I have access to a certain amount of tick data and decide to focus on the early period since the introduction of the more recent ETF (IAU), which may be more interesting than the period when it has become more mature, and 2010 is as far as I can get. Similar to Marshall et al. (2013), I consider only the core trading session (i.e. 9:30am – 4pm) to maximise liquidity.

To address potential errors in the data, I follow the data cleaning process used by Schultz and Shive (2010) and Marshall et al. (2013). Letting the bid and ask subscripts denote the bid price and ask price respectively, an observation at time t is removed if at least one of the following conditions is met:

$$GLD_{bid,t} \geq GLD_{ask,t} \quad \text{or} \quad IAU_{bid,t} \geq IAU_{ask,t}$$

$$GLD_{bid,t} \leq 0.25 \times GLD_{ask,t} \quad \text{or} \quad IAU_{bid,t} \leq 0.25 \times IAU_{ask,t}$$

$$\left| ln\left(\frac{GLD_{bid,t}}{GLD_{bid,t-1}}\right) \right| > 0.25 \quad \text{or} \quad \left| ln\left(\frac{GLD_{ask,t}}{GLD_{ask,t-1}}\right) \right| > 0.25 \quad \text{or} \qquad (1)$$

$$\left| ln\left(\frac{IAU_{bid,t}}{IAU_{bid,t-1}}\right) \right| > 0.25 \quad \text{or} \quad \left| ln\left(\frac{IAU_{ask,t}}{IAU_{ask,t-1}}\right) \right| > 0.25$$

$$\frac{GLD_{bid,t}}{IAU_{ask,t}} > 1.5 \quad \text{or} \quad \frac{IAU_{bid,t}}{GLD_{ask,t}} > 1.5$$

Following Marshall et al. (2013), I also remove quotes posted during the first and last five minutes of trading. However, unlike them, I do not exclude the Flash Crash (6th

May 2010) because unexpected situations are an important part of trading. Initially, there are 144,836,859 observations. After cleaning the data, there remain 144,099,869 valid observations. Table 3.2 presents the descriptive statistics of the clean data. Ranging from -0.55% to 0.49%, the mid-quote log returns are negatively skew, leptokurtic and non-normal as shown by the significant Jarque-Bera statistic.

Table 3.2. Descriptive statistics of mid-quote returns. The returns are in percentage. *** superscript denotes significance at 1%.

	GLD	IAU
Mean	1.36E-05	3.91E-05
Median	0	0
Maximum	0.48	0.49
Minimum	-0.52	-0.55
Standard deviation	0.04	0.05
Skewness	-0.18	-0.29
Kurtosis	16.7	19.53
Jarque-Bera normality	4663840 ***	6126751 ***

3.4. Methodology

3.4.1. Pairs trading based on full convergence

I apply a trading strategy similar to that of Marshall et al. (2013) with some adjustments as follows.

1. At time t_0, I $$\begin{cases} \text{sell GLD and buy IAU} \quad \text{if} \quad \dfrac{GLD_{bid,t_0}}{IAU_{ask,t_0}} \geq 1.002 \\ \text{sell IAU and buy GLD} \quad \text{if} \quad \dfrac{IAU_{bid,t_0}}{GLD_{ask,t_0}} \geq 1.002 \\ \text{do nothing otherwise} \end{cases} \tag{2}$$
2. I use contingent marketable limit orders so that each ETF trade is executed only if the other ETF trade can be executed at a pre-determined price. Such execution ensures that the exact mispricing observed is captured. The trigger value of 1.002 helps exclude a large number of small mispricings.
3. Regarding execution speed, Hasbrouck and Saar (2013) find that high-frequency traders in the US stock market can execute their orders in two milliseconds and their sample period starts in October 2007. Moreover, the US Securities and Exchange Commission documents in 2010 the ability of traders to operate in the microsecond environment (Goldstein et al., 2014),

which suggests that execution speed tends to increase over time. Therefore, I use the speed of two milliseconds in my sub-sample starting in October 2007, which is conservative towards the end of the sample. Regarding the period before October 2007, the introduction of the Hybrid Market trading system in October 2006 allows US equity traders to reduce their execution time from 10 seconds to under one second (Hendershott and Moulton, 2011). As a result, I use the execution speed of 10 seconds for the sub-sample before October 2006 and one second for the sub-sample from October 2006 to September 2007. My actual pairs trade is opened at the first quote set available after the entry signal plus the time required for execution. If the ETF prices have moved against me during execution, this trade will not be opened.

4. At time t_1, I $$\begin{cases} \text{close 'short GLD} - \text{long IAU' trades} & \text{if } \frac{IAU_{bid,t_1}}{GLD_{ask,t_1}} \geq 1 \\ \text{close 'short IAU} - \text{long GLD' trades} & \text{if } \frac{GLD_{bid,t_1}}{IAU_{ask,t_1}} \geq 1 \end{cases} \quad (3)$$
5. The trade is actually closed at the first quote set available after the exit signal plus the time required for execution. Unlike the trade entry, even if there has been adverse price movement during execution, my trade will still be closed because it is important to be able to exit the trade, even at the expense of lower profits.
6. Following Marshall et al. (2013), the trading process above employs only fresh quote sets in which quotes of both ETFs have changed in the last five minutes.
7. To ensure that my trades take place within the core trading session, from 3:50pm in any trading day (i.e. ten minutes before the core trading session ends), any existing pairs trade will be closed at the first available quote set. Because of this mandatory position liquidation, I do not enter new trades after 3:30pm to allocate sufficient time for each trade. Strict form pairs trades are generally short (e.g. Alsayed and McGroarty, 2012, Marshall et al., 2013) so 20 minutes should suffice.
8. The one-way commission per share traded at \$1.00 or more (GLD and IAU have always been traded above \$1.00) is 0.1 – 0.3 cent, depending largely

on trading activities of the trader and the order type used [4]. Given the activities of high frequency traders and the order type used in my analysis, I use the commission of 0.2 cent.

9. Regarding the short-selling cost, Engelberg et al. (2008) find that it is not a major friction. Moreover, Stratmann and Welborn (2012) state that the average short-selling cost of more than a thousand ETFs including many illiquid ones is only 0.0035% per day. The cost of short-selling the highly liquid gold ETFs should be much lower and hence negligible.
10. Regarding short-selling restriction, a security cannot be short-sold at or below the current best bid once its price has declined by 10% or more from the close price of the previous day [5]. This restriction can only be triggered in the core trading session (not in the sessions before or after core trading) and will last until the end of the next trading day. I exclude all arbitrage opportunities which require short sales of GLD (IAU) while GLD (IAU) is under restriction.
11. Finally, each position requires a 50% margin (earning zero return) (Mitchell et al., 2002).

Figure 3.3 illustrates an example of my pairs trading. This trade takes place on 18th September 2007. At 14:18:51, when the two ETFs have shown sufficient divergence, the trade is opened by selling the overpriced GLD at bid price and buying the underpriced IAU at ask price. At 14:20:12, when they have converged (i.e. ask GLD has crossed bid IAU), the trade is closed at a profit by buying back GLD at ask price and selling IAU at bid price.

[4] The fee schedule is available at https://www.nyse.com/publicdocs/nyse/markets/nyse-arca/NYSE_Arca_Marketplace_Fees.pdf
[5] The short-sales information is available at http://www1.nyse.com/pdfs/8764_NYSEArca_FAQ_110225.pdf

Figure 3.3. Pairs trading example. The vertical axis shows the price in US dollar.

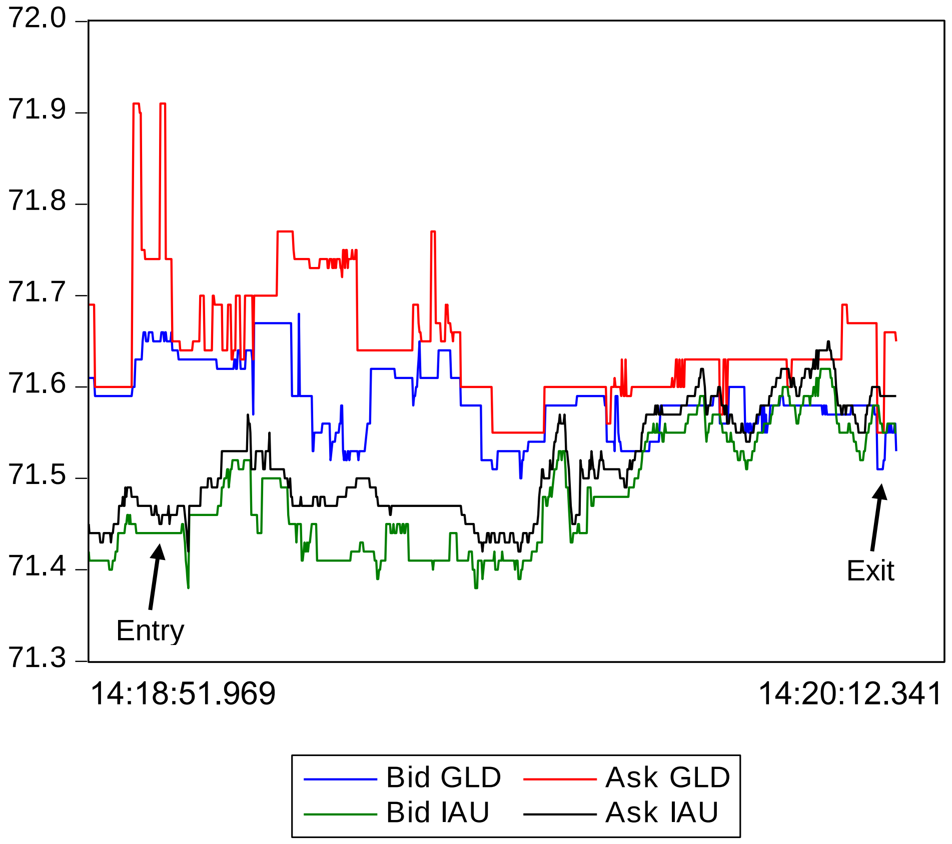

3.4.2. Pairs trading based on partial convergence

The standard pairs trading rule in the literature (e.g. Gatev et al., 2006, Do and Faff, 2010, Jacobs and Weber, 2015) requires complete elimination of the relative mispricing (i.e. trades will be closed only if the two stocks converge completely). I examine an alternative exit rule, namely partial convergence, which only requires partial elimination of mispricing. Specifically, trades are closed after a certain profit target has been reached during convergence.

3.4.2.1. Justification for partial convergence

There are ex-ante reasons why the partial-convergence strategy may be better than the full-convergence strategy. Firstly, I am motivated by findings in the literature regarding the speed of convergence during a pairs trade. Kondor (2009) finds that convergence during a pairs trade is increasingly slow because during convergence, the gap becomes smaller and less appealing to new arbitrageurs. Therefore, fewer traders will attempt to exploit the tightening gap and the gap takes longer to disappear due to the lack of trading pressure. Similarly, Alsayed and McGroarty (2012) show that arbitrageurs face uncertainty about the duration of trades and this uncertainty depends on the convergence target. Full convergence is not worth waiting for if the increased duration uncertainty outweighs the extra profit, which may explain why some mispricings are not exploited. As a result, the full-convergence strategy may not be optimal. Secondly, the partial-convergence strategy has inherent advantages compared to the full-convergence strategy, namely (i) I can capture more opportunities because a partial target allows trades to end quickly and thus allows the funds to become available quickly for the next trade and (ii) I can eliminate the possibility of the current winning trade turning into a losing trade while waiting for full convergence. Thirdly, the full-convergence exit rule employed in the literature may be unnecessarily strict, as illustrated in Figure 3.4.

Figure 3.4. A simplified pairs trade. The vertical axis shows the price and the horizontal axis shows the time which includes three periods. I and II are two ETFs tracking the same index or asset. There are two prices associated with each ETF, namely the bid and the ask.

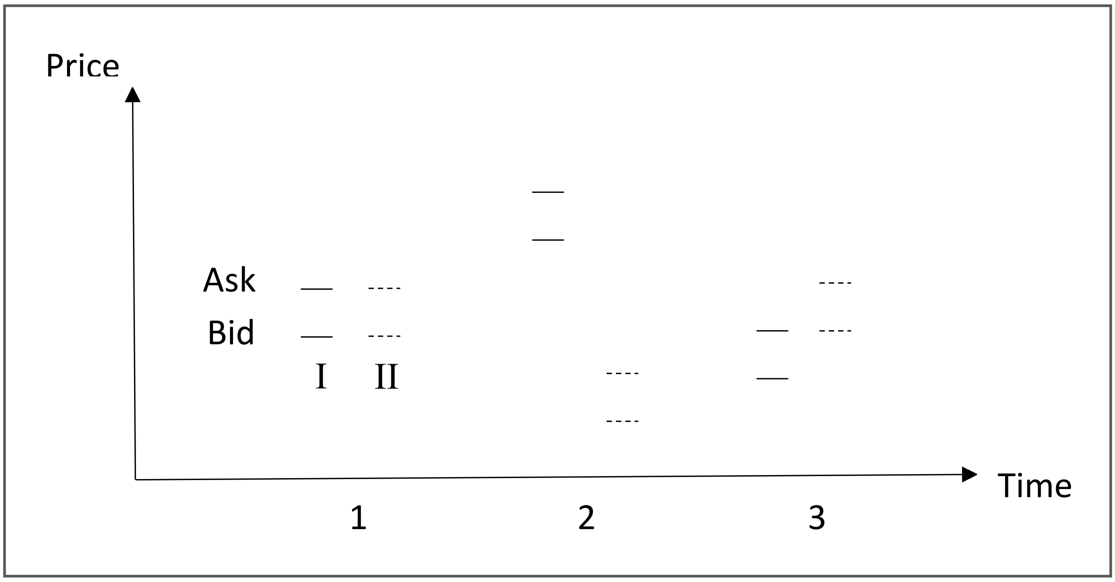

In period 1, the two ETFs are in equilibrium and since they track the same index or asset, their prices should be similar. In period 2, they diverge temporarily, which triggers a pairs trade (i.e. selling I at the bid and buying II at the ask). In period 3, I close the trade only when they converge completely (i.e. the ask of I is equal to the bid of II) so that I can close both positions at the same price. However, the situation in period 3 is not the equilibrium in period 1. In fact, after returning to equilibrium, they must move further before I can exit the trade. Although the ETFs should return to equilibrium, there is no reason to expect them to move further after that. Therefore, the requirement of full convergence described above may be too strict and it is natural to relax it by using the partial-convergence exit rule.

3.4.2.2. The partial-convergence exit rule

The previous trading strategy in section 3.4.1, which is based on full convergence (evidenced by the exit condition in step 4), will be repeated with my partial-convergence exit condition as follows. Let us start with the 'short GLD – long IAU' trade. Letting P denote the profit ($) from convergence (excluding commission) of a given trade, α denote the profit target defined as a percentage of the profit from full convergence ($0 < \alpha < 1$), α and 100% subscripts denote the case of partial and full convergence, t_0 and t_1 denote the time of trade entry and exit; I have the following system.

$$\begin{aligned} P_{\alpha} &= GLD_{bid,t_0} - IAU_{ask,t_0} + IAU_{bid,t_1,\alpha} - GLD_{ask,t_1,\alpha} \\ P_{100\%} &= GLD_{bid,t_0} - IAU_{ask,t_0} + IAU_{bid,t_1,100\%} - GLD_{ask,t_1,100\%} \\ P_{\alpha} &= \alpha \times P_{100\%} \end{aligned} \tag{4}$$

It follows that

$$IAU_{bid,t_1,\alpha} - GLD_{ask,t_1,\alpha} = (\alpha - 1)\left(GLD_{bid,t_0} - IAU_{ask,t_0}\right) + \alpha\left(IAU_{bid,t_1,100\%} - GLD_{ask,t_1,100\%}\right) \tag{5}$$

By definition of full convergence, for 'short GLD – long IAU' trades, I have

$$IAU_{bid,t_1,100\%} \geq GLD_{ask,t_1,100\%} \tag{6}$$

It follows that

$$IAU_{bid,t_1,\alpha} - GLD_{ask,t_1,\alpha} \geq (\alpha - 1)\left(GLD_{bid,t_0} - IAU_{ask,t_0}\right) \tag{7}$$

This condition is my partial-convergence exit rule for 'short GLD – long IAU' trades. Similarly, at time t_1, I close 'short IAU – long GLD' trades if the following condition is met.

$$GLD_{bid,t_1,\alpha} - IAU_{ask,t_1,\alpha} \geq (\alpha - 1)\left(IAU_{bid,t_0} - GLD_{ask,t_0}\right) \tag{8}$$

While examining partial convergence, I mitigate the data mining issue by not testing too many values of α. Nevertheless, too few profit targets may provide less reliable results so I choose three evenly spaced targets (i.e. 25%, 50% and 75%). I also apply the Bonferroni adjustment to the tests. Moreover, data mining is less of an issue because I focus less on the optimal target but more on the general performance of partial convergence as a whole compared to full convergence.

3.4.3. Evaluation of pairs trading performance

Letting t_0 and t_1 denote the time of trade entry and exit, the profit (%) of a given pairs trade (whether based on full or partial convergence) is as follows.

$$\begin{aligned} &\text{short GLD} - \text{long IAU:} \ \frac{GLD_{bid,t_0} - IAU_{ask,t_0} + IAU_{bid,t_1} - GLD_{ask,t_1} - 4 \times 0.002}{0.5 \times \left(GLD_{bid,t_0} + IAU_{ask,t_0}\right)} \times 100 \\ &\text{short IAU} - \text{long GLD:} \ \frac{IAU_{bid,t_0} - GLD_{ask,t_0} + GLD_{bid,t_1} - IAU_{ask,t_1} - 4 \times 0.002}{0.5 \times \left(IAU_{bid,t_0} + GLD_{ask,t_0}\right)} \times 100 \end{aligned} \tag{9}$$

In each case, the last term in the numerator is the total commission which is four times one-way commission because a pairs trade requires four orders (two for entry and two for exit). The denominator is the margin requirement.

In terms of market efficiency, Jensen (1978) states that the most general interpretation of the Efficient Market Hypothesis is that if it is impossible to make economic profits (i.e. risk-adjusted returns after costs) using a given information set, the market is efficient regarding that information set. To arrive at implications for market efficiency, I calculate excess returns of pairs trading (over the risk-free rate of US dollar deposit) and adjust them for risks using the Sharpe (1994) and Sortino (2010) ratio. While the Sharpe ratio adjusts returns for general volatility, the Sortino ratio adjusts them for only downside volatility calculated from returns below the desired target return (DTR) which is set to zero. Because the Sortino ratio reflects the nature of risks more precisely than the Sharpe ratio, it might be the better risk adjustment. After considering risks using these ratios, if pairs trading outperforms the buy-and-hold strategy of these ETFs, the gold ETF market is inefficient. Letting $\overline{ER}$ denote the mean excess return, σ_{ER} denote the standard deviation of excess returns and T denote the number of observations; the Sharpe and Sortino ratios are as follows.

$$Sharpe\ ratio = \frac{\overline{ER}}{\sigma_{ER}} = \frac{\frac{1}{T}\sum_{t=1}^{T} ER_t}{\sqrt{\frac{\sum_{t=1}^{T}\left(ER_t - \overline{ER}\right)^2}{T-1}}} \tag{10}$$

$$Sortino\ ratio = \frac{\frac{\sum_{ER_t > DTR}(ER_t - DTR)}{T}}{\sqrt{\frac{\sum_{ER_t \le DTR}(DTR - ER_t)^2}{T}}} = \frac{\sum_{ER_t > DTR}(ER_t - DTR)}{\sqrt{T . \sum_{ER_t \le DTR}(DTR - ER_t)^2}} \quad (11)$$

To provide further insights into how the pairs trading returns relate to risks, I also apply another risk adjustment method to adjust the returns for exposure to risk factors. Following Engelberg et al. (2008), I regress the returns of pairs trading on (i) market-wide risks including Fama-French three-factor model (i.e. market return, size and value) plus two style factors (i.e. momentum and reversal) and (ii) liquidity risk whose proxy is the Treasury bill – Eurodollar (TED) spread (i.e. the difference between the US T-bill rate and the USD LIBOR rate). A high TED spread means that borrowing is difficult and thus liquidity is low and vice versa. If the market is efficient, the abnormal returns (i.e. intercept of the regression) should not be positive and statistically significant. Table 3.3 shows the summary statistics of these risk factors. The market return, size and momentum factors are negatively skew while the others are positively skew. All of the factors are leptokurtic and non-normal, as shown by the Jarque-Bera statistic.

Table 3.3. Descriptive statistics of the daily risk factors (in percentage). *** superscript denotes significance at 1%.

	Market	Size	Value	Momentum	Reversal	TED spread
Mean	0.006	0.01	0.011	-0.01	0.039	0.719
Median	0.08	0.01	0.00	0.07	-0.01	0.48
Maximum	11.35	3.85	4.80	7.05	11.22	4.58
Minimum	-8.95	-3.78	-4.22	-8.22	-7.16	0.09
Standard deviation	1.488	0.632	0.841	1.268	1.142	0.64
Skewness	-0.019	-0.06	0.623	-0.781	1.63	2.066
Kurtosis	12.176	8.00	11.248	10.629	23.127	8.921
Jarque-Bera normality	4704 ***	1398 ***	3888 ***	3388 ***	23228 ***	2913 ***

Engelberg et al. (2008) also mention divergence risk and default risk. However, in my ETF setting, divergence risk (i.e. the two stocks do not converge or they take a long time to do so) is low since (i) investors buy ETFs to track the underlying index or asset so the ETFs' management must try to minimise tracking errors to attract investors and (ii) ETF shares can be exchanged for the underlying asset or component stocks to benefit from potential mispricing, which should keep prices in equilibrium (Engle and

Sarkar, 2006). On the other hand, default risk is also not an issue because ETFs are not company stocks and gold ETFs are backed by physical gold which can be redeemed by investors.

3.5. Results

Table 3.4, 3.5 and 3.6 summarise the pairs trading performance in the three periods defined by trading speed, considering all trades and two subsets of winning and losing trades.

Table 3.4. Pairs trading performance with different profit targets from February 2005 to September 2006 (i.e. when the trade execution time is 10 seconds). My targets are defined as a percentage of full convergence (e.g. 50% means that trades are closed when the pair has converged by half of its initial divergence). Panel A shows the trading profit and the break-even transaction cost, panel B shows the risk adjustment and panel C shows the trade duration. The break-even transaction cost is the one-way commission per share which reduces the trading profit to zero. The excess return is equal to the trading profit plus the interest earned from the capital when not trading minus the risk-free rate. The Sharpe and Sortino ratio are compared between partial convergence and full convergence (vs. 100%) and between both strategies and the buy-and-hold benchmark (vs. b&h). I test for statistical significance of the Sharpe ratio following Pui-Lam and Wing-Keung (2008) and of the Sortino ratio using the bootstrapping technique in Levich and Thomas (1993). The speed of convergence is measured by the return per trading hour because the trading return relates directly to the distance of convergence. I report the break-even transaction cost, risk adjustment and speed of convergence based on all trades, which is more meaningful than based on only winners or losers. For comparison to the other two periods (with different length), I annualise the number of trades, total profit, excess return and total duration. The *, ** and *** superscript denote statistically significant difference at 10%, 5% and 1% level respectively.

	25%			50%			75%			100%		
	All trades	Winners	Losers	All trades	Winners	Losers	All trades	Winners	Losers	All trades	Winners	Losers
Number of trades	37.2	34.2	3	26.4	20.4	6	19.2	12.6	6.6	16.2	7.8	8.4
Panel A: Trading Profit (%)												
Total	1.564	1.988	-0.424	1.105	2.167	-1.062	0.572	1.633	-1.061	-0.228	0.903	-1.131
Mean	0.042	0.058	-0.141	0.042	0.106	-0.177	0.030	0.130	-0.161	-0.014	0.116	-0.135
Median	0.049	0.050	-0.088	0.100	0.103	-0.088	0.046	0.149	-0.087	-0.013	0.057	-0.055
Standard deviation	0.064	0.019	0.104	0.158	0.025	0.220	0.198	0.085	0.215	0.204	0.115	0.196
Break-even TC (cents)	0.936	-	-	0.966	-	-	0.820	-	-	0.139	-	-
Panel B: Risk Adjustment												
Excess return (%)	1.378	-	-	0.919	-	-	0.386	-	-	-0.414	-	-
Sharpe ratio (vs. 100%)	0.112 ***	-	-	0.055 ***	-	-	0.027 ***	-	-	-0.036	-	-
Sortino ratio (vs. 100%)	0.358 ***	-	-	0.150 ***	-	-	0.090 ***	-	-	0.034	-	-
Sharpe ratio (vs. b&h)	0.112	-	-	0.055	-	-	0.027	-	-	-0.036	-	-
Sortino ratio (vs. b&h)	0.358 ***	-	-	0.150 ***	-	-	0.090	-	-	0.034 ***	-	-
Panel C: Duration (hours)												
Total	17.107	7.500	9.606	35.718	17.768	17.950	53.829	26.922	26.907	67.096	29.442	37.655
Mean	0.460	0.219	3.202	1.353	0.871	2.992	2.804	2.137	4.077	4.142	3.775	4.483
Median	0.071	0.061	3.622	0.327	0.078	3.534	2.363	1.070	4.433	5.434	5.383	5.488
Standard deviation	1.033	0.460	1.755	1.810	1.437	2.056	2.356	2.318	1.935	2.171	2.467	1.884
Speed of convergence	0.091	-	-	0.031	-	-	0.011	-	-	-0.003	-	-

Table 3.5. Pairs trading performance with different profit targets from October 2006 to September 2007 (i.e. when the trade execution time is 1 second). My targets are defined as a percentage of full convergence (e.g. 50% means that trades are closed when the pair has converged by half of its initial divergence). Panel A shows the trading profit and the break-even transaction cost, panel B shows the risk adjustment and panel C shows the trade duration. The break-even transaction cost is the one-way commission per share which reduces the trading profit to zero. The excess return is equal to the trading profit plus the interest earned from the capital when not trading minus the risk-free rate. The Sharpe and Sortino ratio are compared between partial convergence and full convergence (vs. 100%) and between both strategies and the buy-and-hold benchmark (vs. b&h). I test for statistical significance of the Sharpe ratio following Pui-Lam and Wing-Keung (2008) and of the Sortino ratio using the bootstrapping technique in Levich and Thomas (1993). The speed of convergence is measured by the return per trading hour because the trading return relates directly to the distance of convergence. I report the break-even transaction cost, risk adjustment and speed of convergence based on all trades, which is more meaningful than based on only winners or losers. For comparison to the other two periods (with different length), I annualise the number of trades, total profit, excess return and total duration. The *, ** and *** superscript denote statistically significant difference at 10%, 5% and 1% level respectively.

	25%			50%			75%			100%		
	All trades	Winners	Losers	All trades	Winners	Losers	All trades	Winners	Losers	All trades	Winners	Losers
Number of trades	21	20	1	18	18	0	18	17	1	14	13	1
Panel A: Trading Profit (%)												
Total	4.861	5.901	-1.040	6.656	6.656	0.000	7.327	7.354	-0.027	4.519	4.546	-0.027
Mean	0.231	0.295	-1.040	0.370	0.370	0.000	0.407	0.433	-0.027	0.323	0.350	-0.027
Median	0.064	0.078	-1.040	0.147	0.147	0.000	0.208	0.218	-0.027	0.252	0.253	-0.027
Standard deviation	0.726	0.682	0.000	0.701	0.701	0.000	0.692	0.704	0.000	0.251	0.239	0.000
Break-even TC (cents)	4.024	-	-	6.403	-	-	7.028	-	-	5.571	-	-
Panel B: Risk Adjustment												
Excess return (%)	4.717	-	-	6.511	-	-	7.183	-	-	4.374	-	-
Sharpe ratio (vs. 100%)	0.078 **	-	-	0.109 ***	-	-	0.116 ***	-	-	0.146	-	-
Sortino ratio (vs. 100%)	0.507 ***	-	-	1.511 ***	-	-	1.727 ***	-	-	0.680	-	-
Sharpe ratio (vs. b&h)	0.078	-	-	0.109	-	-	0.116	-	-	0.146	-	-
Sortino ratio (vs. b&h)	0.507 ***	-	-	1.511 ***	-	-	1.727 ***	-	-	0.680 ***	-	-
Panel C: Duration (hours)												
Total	0.210	0.209	0.000	0.332	0.332	0.000	3.384	1.790	1.594	6.620	5.026	1.594
Mean	0.010	0.010	0.000	0.018	0.018	0.000	0.188	0.105	1.594	0.473	0.387	1.594
Median	0.001	0.001	0.000	0.001	0.001	0.000	0.018	0.016	1.594	0.162	0.115	1.594
Standard deviation	0.020	0.020	0.000	0.040	0.040	0.000	0.394	0.185	0.000	0.701	0.648	0.000
Speed of convergence	23.148	-	-	20.048	-	-	2.165	-	-	0.683	-	-

Table 3.6. Pairs trading performance with different profit targets from October 2007 to May 2010 (i.e. when the trade execution time is 2 milliseconds). My targets are defined as a percentage of full convergence (e.g. 50% means that trades are closed when the pair has converged by half of its initial divergence). Panel A shows the trading profit and the break-even transaction cost, panel B shows the risk adjustment and panel C shows the trade duration. The break-even transaction cost is the one-way commission per share which reduces the trading profit to zero. The excess return is equal to the trading profit plus the interest earned from the capital when not trading minus the risk-free rate. The Sharpe and Sortino ratio are compared between partial convergence and full convergence (vs. 100%) and between both strategies and the buy-and-hold benchmark (vs. b&h). I test for statistical significance of the Sharpe ratio following Pui-Lam and Wing-Keung (2008) and of the Sortino ratio using the bootstrapping technique in Levich and Thomas (1993). The speed of convergence is measured by the return per trading hour because the trading return relates directly to the distance of convergence. I report the break-even transaction cost, risk adjustment and speed of convergence based on all trades, which is more meaningful than based on only winners or losers. For comparison to the other two periods (with different length), I annualise the number of trades, total profit, excess return and total duration. The *, ** and *** superscript denote statistically significant difference at 10%, 5% and 1% level respectively.

	25%			50%			75%			100%		
	All trades	Winners	Losers	All trades	Winners	Losers	All trades	Winners	Losers	All trades	Winners	Losers
Number of trades	53.25	48	5.25	32.25	30.75	1.5	25.875	23.625	2.25	22.125	19.875	2.25
Panel A: Trading Profit (%)												
Total	2.735	3.129	-0.395	3.201	3.444	-0.243	3.092	3.318	-0.227	2.933	3.127	-0.194
Mean	0.051	0.065	-0.075	0.099	0.112	-0.162	0.119	0.140	-0.101	0.133	0.157	-0.086
Median	0.047	0.051	-0.042	0.101	0.102	-0.119	0.146	0.150	-0.032	0.128	0.155	-0.026
Standard deviation	0.091	0.079	0.098	0.112	0.094	0.154	0.135	0.114	0.146	0.154	0.135	0.149
Break-even TC (cents)	1.363	-	-	2.398	-	-	2.844	-	-	3.140	-	-
Panel B: Risk Adjustment												
Excess return (%)	2.622	-	-	3.088	-	-	2.979	-	-	2.820	-	-
Sharpe ratio (vs. 100%)	0.157	-	-	0.159	-	-	0.153	-	-	0.147	-	-
Sortino ratio (vs. 100%)	4.482 ***	-	-	2.808 ***	-	-	2.290	-	-	2.208	-	-
Sharpe ratio (vs. b&h)	0.157 **	-	-	0.159 **	-	-	0.153 **	-	-	0.147 **	-	-
Sortino ratio (vs. b&h)	4.482 ***	-	-	2.808 ***	-	-	2.290 ***	-	-	2.208 ***	-	-
Panel C: Duration (hours)												
Total	7.836	5.093	2.743	18.636	14.872	3.764	36.766	28.291	8.475	50.552	43.125	7.427
Mean	0.147	0.106	0.523	0.578	0.484	2.509	1.421	1.198	3.767	2.285	2.170	3.301
Median	0.009	0.009	0.009	0.013	0.012	2.676	0.075	0.059	3.707	1.413	1.209	2.404
Standard deviation	0.459	0.298	1.119	1.272	1.152	2.187	2.098	1.941	2.437	2.329	2.336	2.191
Speed of convergence	0.349	-	-	0.172	-	-	0.084	-	-	0.058	-	-

Regarding full convergence, high frequency traders, able to execute their orders with minimal delay, make a positive profit on 79% of their trades and generate an excess return of 11.2% in total from all three sub-periods. This documented profit is consistent with previous studies on pairs trading which generally find that this strategy is profitable. However, intuitively, there should be differences in terms of profitability among the three types of pairs trading (i.e. the loose form, the semi-strict form and the strict form) due to their differences in the level of fundamental relationship between the two securities in a pair. More specifically, if the two assets are less closely related, they are more likely to show frequent instances of significant mispricing, which will lead to higher returns from pairs trading. In line with this intuition, loose form studies tend to report greater profits than semi-strict form and strict form studies. Similarly, the strict form research in this paper finds profits of less than 5% p.a., which are lower than those in loose form studies (Gatev et al., 2006, Broussard and Vaihekoski, 2012, Jacobs and Weber, 2015), which are often more than 10% p.a..

On the other hand, Table 3.4 – 3.6 show that when the profit target decreases from full convergence to partial convergence, the average trade duration decreases (because normally it takes less time to reach a closer target) so the number of trades increases (because if a trade ends quickly, the capital becomes available quickly and I can enter the next trade soon). Moreover, full convergence generally takes longer than partial convergence not only because of the distance of the target but also because convergence becomes more and more difficult as the two ETFs approach full convergence. The reason is that (i) on one hand, full convergence requires the bid price of one ETF to converge to the ask price of the other but (ii) on the other hand, since these ETFs track the same asset, they should have similar bid and ask price so the bid price of one should remain lower and not converge to the ask price of the other. This explanation is supported not only by findings in the literature that convergence during a pairs trade is increasingly slow (e.g. Kondor, 2009, Alsayed and McGroarty, 2012) but also by my measure of convergence speed. The speed of convergence is measured by the return per trading hour (i.e. total return divided by total duration of trades) since the total return (i.e. the numerator) relates directly to the distance of convergence and the total duration of trades (i.e. the denominator) is the time of convergence. I find that the partial-convergence strategy has higher speed of convergence than the full-convergence strategy. The difficulty in achieving full convergence may also explain why the winning rates of the partial convergence targets are higher than that of full convergence because a currently winning trade may turn into a losing trade while waiting for full convergence.

Regarding profitability, the total profit shows that partial convergence outperforms full convergence as a trade exit criterion, which means that it may be a good idea to give up some profits per trade in exchange for faster trades and the ability to enter more trades. The most profitable criterion is 25%, 75% and 50% for the first, second and third sub-sample respectively. Moreover, the relationship between the profit target and total profit is not often monotonic. In the second and third sub-period, initial reduction of the target increases the profit but excessive reduction decreases it. This means when reducing the profit target, the higher trading frequency can compensate for the lower per-trade profit, but only to some extent. However, to mitigate the data mining issue, I do not focus on finding the optimal partial target but instead I am interested in the outperformance of partial convergence in general compared to full convergence. For the partial-convergence strategy, the break-even transaction cost ranges from 0.820 to 7.028 cents and is relatively high compared to the applicable cost on the stock exchange (i.e. less than 0.5 cent), which suggests that the pairs trading strategy is robust to transaction cost. With regards to trade duration, all of its statistics (i.e. total, mean, median and standard deviation) increase monotonically with the profit target. Finally, the speed of convergence, measured by the return per trading hour, shows how much return is generated per unit of time in a trade and thus shows the effectiveness of in-trade capital utilisation. The higher convergence speed of the partial-convergence strategy allows it to use the trading capital during trades more effectively than the full-convergence strategy.

As for the risk-adjusted performance, the Sharpe ratios show mixed results. Compared to full convergence, partial convergence outperforms in the first sub-period, underperforms in the second and performs similarly in the third. However, the Sortino ratios show that partial convergence outperforms full convergence at 1% level in most cases across the sub-samples. Moreover, because the Sortino ratio considers only downside volatility instead of general volatility like the Sharpe ratio, it might be a more accurate reflection of risks and thus a more appropriate adjustment for risks. Also importantly, my pairs trading outperforms the buy-and-hold strategy (at 5% level in the last sub-sample as shown by the Sharpe ratio and at 1% level in all sub-samples as shown by the Sortino ratio). The low significance of the Sharpe ratio may be because this ratio also penalises the upside volatility (i.e. favourable returns). Because pairs trading outperforms the buy-and-hold strategy on a risk-adjusted basis, the gold ETF market may be inefficient.

Regarding the influence of speed on trading performance, there is evidence that a higher speed can improve performance. The highest speed in the third sub-sample helps capture more opportunities, resulting in the highest number of trades. Meanwhile, the lowest speed in the first sub-sample leads to not only the lowest profitability but also the longest duration of trades. More importantly, both the Sharpe and Sortino ratio suggest that the risk-adjusted performance is enhanced when trading becomes faster.

Table 3.7. Exposure of pairs trading returns to daily risk factors from February 2005 to September 2006 (i.e. when the trade execution time is 10 seconds). This table shows the regression results of the trading returns from different convergence targets (i.e. 25%, 50%, 75% and 100%) on the risk factors, namely (i) the market-wide risks including Fama-French three-factor model (i.e. market return, size and value) plus two style factors (i.e. momentum and reversal) and (ii) the liquidity risk whose proxy is the Treasury bill – Eurodollar (TED) spread. The numbers in brackets show the standard error of the estimated coefficients. The *, ** and *** superscript denote statistical significance at 10%, 5% and 1% level respectively.

	25%	50%	75%	100%
Intercept	-3.3E-05	-1E-04	-0.0001	-8.8E-05
	(7.05E-05)	(9.55E-05)	(8.32E-05)	(6.58E-05)
Market	7.88E-05 *	2.2E-05	-4.6E-05	-3.9E-05
	(4.45E-05)	(6.04E-05)	(5.26E-05)	(4.16E-05)
Size	-0.00012 *	-0.00015 *	-0.00012	-7.5E-05
	(6.3E-05)	(8.54E-05)	(7.44E-05)	(5.88E-05)
Value	9.16E-05	-0.00011	-0.00011	-0.00013
	(0.000113)	(0.000154)	(0.000134)	(0.000106)
Momentum	-0.00015 ***	-0.00015 *	-6.7E-05	2.01E-05
	(5.68E-05)	(7.7E-05)	(6.7E-05)	(5.3E-05)
Reversal	-1.6E-05	-4.3E-05	-3.4E-05	-1.1E-05
	(5.04E-05)	(6.83E-05)	(5.95E-05)	(4.7E-05)
TED spread	0.000201	0.000337	0.000296	0.000186
	(0.000159)	(0.000216)	(0.000188)	(0.000149)
Adjusted R^2	0.030934	0.030674	0.026941	0.002147

Table 3.8. Exposure of pairs trading returns to daily risk factors from October 2006 to September 2007 (i.e. when the trade execution time is 1 second). This table shows the regression results of the trading returns from different convergence targets (i.e. 25%, 50%, 75% and 100%) on the risk factors, namely (i) the market-wide risks including Fama-French three-factor model (i.e. market return, size and value) plus two style factors (i.e. momentum and reversal) and (ii) the liquidity risk whose proxy is the Treasury bill – Eurodollar (TED) spread. The numbers in brackets show the standard error of the estimated coefficients. The *, ** and *** superscript denote statistical significance at 10%, 5% and 1% level respectively.

	25%	50%	75%	100%
Intercept	0.000442 *	0.000289	0.000279	0.000116
	(0.00026)	(0.000256)	(0.000266)	(0.000129)
Market	0.000175	0.000194	0.000196	-0.00011
	(0.000189)	(0.000186)	(0.000193)	(9.36E-05)
Size	-0.00055	-0.00035	-0.00031	3.02E-06
	(0.000383)	(0.000376)	(0.00039)	(0.000189)
Value	-0.00088	-0.00084	-0.00082	-0.00038
	(0.000639)	(0.000629)	(0.000652)	(0.000316)
Momentum	-8.9E-05	-0.00035	-0.00041	-0.00024
	(0.000386)	(0.00038)	(0.000394)	(0.000191)
Reversal	0.000219	0.000352	0.000376	0.000148
	(0.000329)	(0.000323)	(0.000335)	(0.000163)
TED spread	-0.0005	-0.00011	-3.9E-05	8.93E-05
	(0.000358)	(0.000353)	(0.000366)	(0.000177)
Adjusted R^2	-0.00136	-0.00888	-0.00973	-0.00339

Table 3.9. Exposure of pairs trading returns to daily risk factors from October 2007 to May 2010 (i.e. when the trade execution time is 2 milliseconds). This table shows the regression results of the trading returns from different convergence targets (i.e. 25%, 50%, 75% and 100%) on the risk factors, namely (i) the market-wide risks including Fama-French three-factor model (i.e. market return, size and value) plus two style factors (i.e. momentum and reversal) and (ii) the liquidity risk whose proxy is the Treasury bill – Eurodollar (TED) spread. The numbers in brackets show the standard error of the estimated coefficients. The *, ** and *** superscript denote statistical significance at 10%, 5% and 1% level respectively.

	25%	50%	75%	100%
Intercept	-7.8E-05 *	-8.8E-05 *	-0.00011 **	-7.2E-05
	(3.97E-05)	(4.58E-05)	(4.55E-05)	(4.55E-05)
Market	-1E-05	2.58E-05	1.69E-05	2.24E-05
	(1.74E-05)	(2.01E-05)	(2E-05)	(2E-05)
Size	2.86E-05	2.41E-05	6.95E-06	4.63E-05
	(3.31E-05)	(3.82E-05)	(3.8E-05)	(3.79E-05)
Value	-5.6E-05 *	-5.2E-05	-8.6E-05 **	-9.7E-05 ***
	(3.05E-05)	(3.52E-05)	(3.5E-05)	(3.49E-05)
Momentum	-4.8E-05 **	-4.4E-05 *	-6.5E-05 ***	-6.9E-05 ***
	(2.14E-05)	(2.47E-05)	(2.45E-05)	(2.45E-05)
Reversal	2.82E-05	4.43E-05 **	4.67E-05 **	4.03E-05 **
	(1.77E-05)	(2.04E-05)	(2.03E-05)	(2.03E-05)
TED spread	0.000188 ***	0.000219 ***	0.000234 ***	0.000189 ***
	(3.18E-05)	(3.68E-05)	(3.65E-05)	(3.65E-05)
Adjusted R^2	0.059964	0.070874	0.084385	0.067161

Although some risk factors are statistically significant, they do not explain the variation in trading returns well, as shown by the low adjusted R^2 in all sub-samples. Compared to the full-convergence strategy (i.e. 100% target), the partial-convergence strategy (i.e. 25%, 50% and 75% target) is more exposed to the risk factors in the first sub-period (i.e. some more statistically significant coefficients of the factors) and less exposed in the last sub-period (i.e. less significant coefficients). In the middle sub-sample, no coefficient is significant for both strategies.

Figure 3.5 shows the cumulative wealth of high-frequency traders using different profit targets to exit trades, starting with $100. The lines exhibit similar upward movements and low volatility. The first and second half of the sample period roughly corresponds to the period before and after the global financial crisis respectively. The cumulative wealth increases gradually in both the pre- and post-crisis period, which suggests that the pairs trading strategy performs similarly in both periods.

Figure 3.5. Cumulative wealth from different profit targets. The starting point is $100. The vertical axis shows the wealth in US dollar.

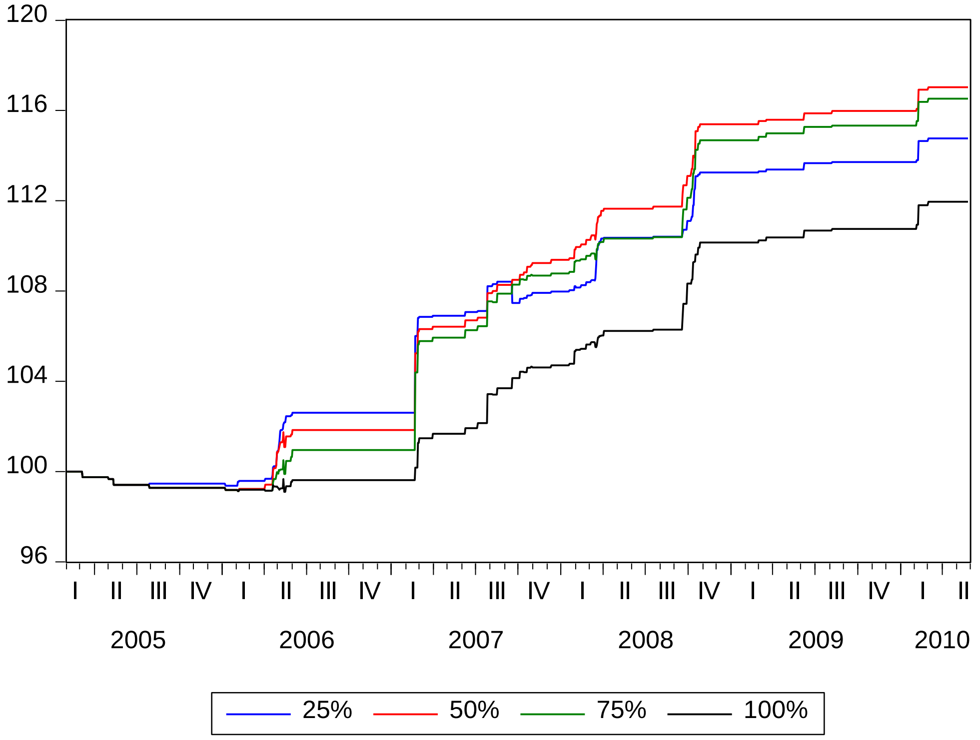

Table 3.10 examines the variation of returns among the three sub-samples defined by speed. Compared to the first sub-period, trading returns from both full and partial convergence are higher in the other two (especially in the last one where the outperformance is statistically significant even at 1% level), which suggests that a higher trading speed may improve performance.

Table 3.10. Variation of returns over time. This table shows the test results for statistically significant differences in returns among the three sub-samples for both the partial-convergence and full-convergence strategy. The numbers are the test statistic. The *, ** and *** superscript denote statistical significance at 10%, 5% and 1% level respectively.

	First vs. Second	Second vs. Third	Third vs. First
25%	-0.861	0.557	1.314
50%	-1.459	0.915	1.895 *
75%	-1.724 *	1.082	2.453 **
100%	-2.409 **	0.786	3.453 ***

Regarding the profit potential (i.e. how much money is available to obtain in the market), it depends on the liquidity of the two ETFs. During the sample period, the average daily trading volume and trading value of GLD are 3,339,335 shares and $290,291,641 respectively and those of IAU are 131,953 shares and $12,594,036 respectively. Both ETFs are highly liquid but because pairs trading requires trading both of them at the same time, the profit potential is determined by the less liquid instrument which is IAU. Even if only 1% of IAU's trading value (and an equivalent amount of GLD's trading value) is involved in pairs trading, the excess return of 2% p.a. generated by the baseline strategy can translate to an annual profit of more than $1.2 million. Moreover, this estimate is conservative because the trading volume is lower than the quoted volume which shows the total availability of these instruments. Table 3.11 below shows the average trading volume of both ETFs on days with and without opportunities for pairs trading. The trading volume on days with opportunities is significantly different from that on days without, at 1% level for GLD and 10% level for IAU. Given the active trading of both ETFs, liquidity should not be an issue.

Table 3.11. Trading volume conditional on trading opportunities. This table shows the average trading volume (i.e. number of shares traded) of both ETFs on days with and without opportunities for pairs trading. The *, **, *** superscripts denote statistically significant difference (if any) of days with opportunities from days without at the 10%, 5% and 1% level respectively.

	GLD volume	IAU volume
Days with opportunities	5,985,254***	166,830*
Days without opportunities	3,173,166	129,763

However to be even more conservative about the potential profitability, one also needs to examine liquidity at the trade level because it is an important limit to arbitrage. Table 3.12 shows the summary statistics of available volume at trade level. Compared to the first two sub-periods, the last one has a lower mean volume but a higher maximum volume. The reason is that additional trading opportunities captured by the higher speed in the third sub-sample have a smaller volume on average but sometimes have a really large volume. The minimum volume is 100 shares which may be discouraging but one should consider the mean value and the number of trades for a better idea about the overall liquidity. Fortunately, it has been reported that the high speed in the last sub-sample indeed helps exploit more arbitrage signals. Finally, the trading volume exhibits positive skewness, leptokurtosis and non-normality (i.e. significant Jarque-Bera statistic) throughout the sample period.

Table 3.12. Summary statistics of available volume at trade level for each sub-sample based on different convergence targets. The available volume at trade level refers to the number of shares available for the less liquid stock in a given pairs trade. The reason is that a pairs trade requires trading the same number of shares in both stocks so the available volume is determined by the less liquid stock (i.e. the stock with fewer shares available for that particular trade). For instance, if a trade requires (i) buying (i.e. trading at the ask price) GLD whose ask volume is 700 shares and (ii) selling (i.e. trading at the bid price) IAU whose bid volume is 400 shares, then the available volume for that trade is 400 shares. The *, **, *** superscript denote statistical significance at 10%, 5% and 1% level respectively.

	February 2005 – September 2006				October 2006 – September 2007				October 2007 – May 2010			
	25%	50%	75%	100%	25%	50%	75%	100%	25%	50%	75%	100%
Mean	440	470	447	474	486	478	478	450	246	298	338	339
Median	200	200	200	200	200	200	200	200	100	100	200	200
Maximum	2000	2000	2000	2000	1800	1800	1800	1500	3000	3000	3000	3000
Minimum	100	100	100	100	100	100	100	100	100	100	100	100
Standard deviation	454	500	493	527	505	521	521	442	393	489	539	568
Skewness	1.428	1.311	1.524	1.349	1.301	1.364	1.364	1.110	5.263	4.201	3.718	3.691
Kurtosis	4.309	3.728	4.629	3.908	3.607	3.748	3.748	3.158	34.280	21.812	17.426	16.554
Jarque-Bera statistic	25.483***	13.568***	15.921***	9.119**	6.249**	6.001**	6.001**	2.888	6444.6***	1521.0***	757.3***	585.5***

Another factor affecting profitability is the initial investment in high frequency trading facilities. It can be argued that high frequency trading firms are unlikely to invest heavily in their infrastructure only to trade a single strategy so pairs trading can take advantage of the existing infrastructure used for other strategies and its upfront cost is reduced significantly. Furthermore, when the pairs trading strategy is established, more pairs can be added without incurring much extra cost.

Despite the potential issue of data mining, it is interesting to see whether there is an optimal convergence target. Figure 3.6 shows the performance of multiple targets from 5% increasing at 5% increments to 100%.

Figure 3.6. Performance of multiple targets in each sub-sample. This figure shows the performance of multiple targets including the number of trades, the winning rate, the total profit and the total duration of trades. The blue, orange and green line represents the first, second and third sub-sample respectively. For comparison purposes, I annualise the number of trades, total profit and total duration. The horizontal axis shows the convergence targets (%).

Figure 3.6.1. Number of trades. The vertical axis shows the number of trades.

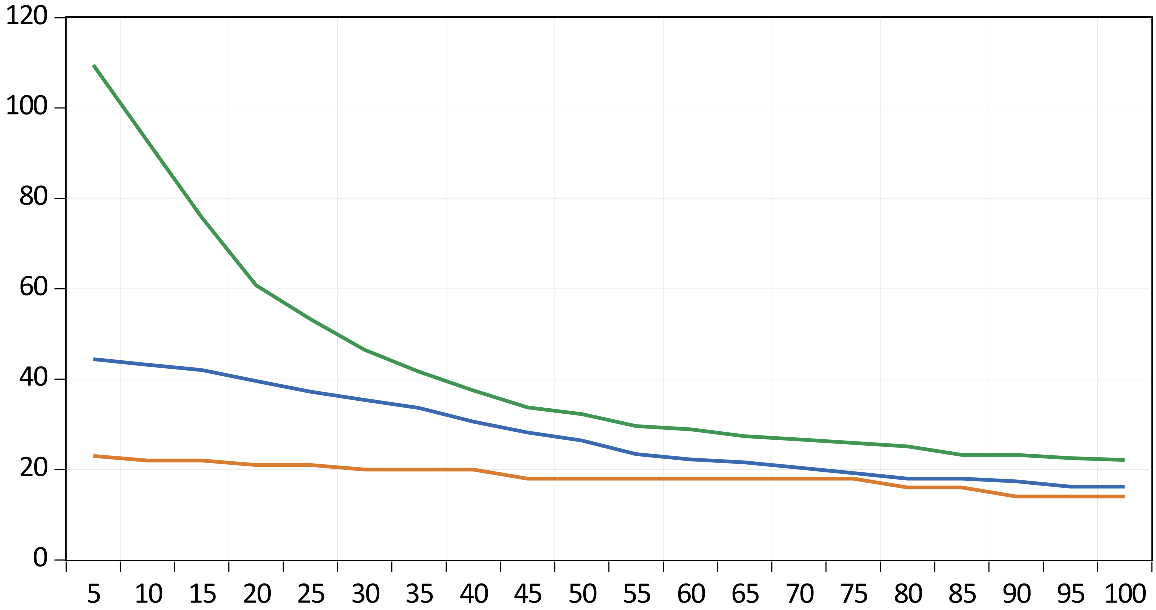

Figure 3.6.2. Winning rate. The vertical axis shows the winning rate (%). The winning rate is the number of profitable trades divided by the number of trades.

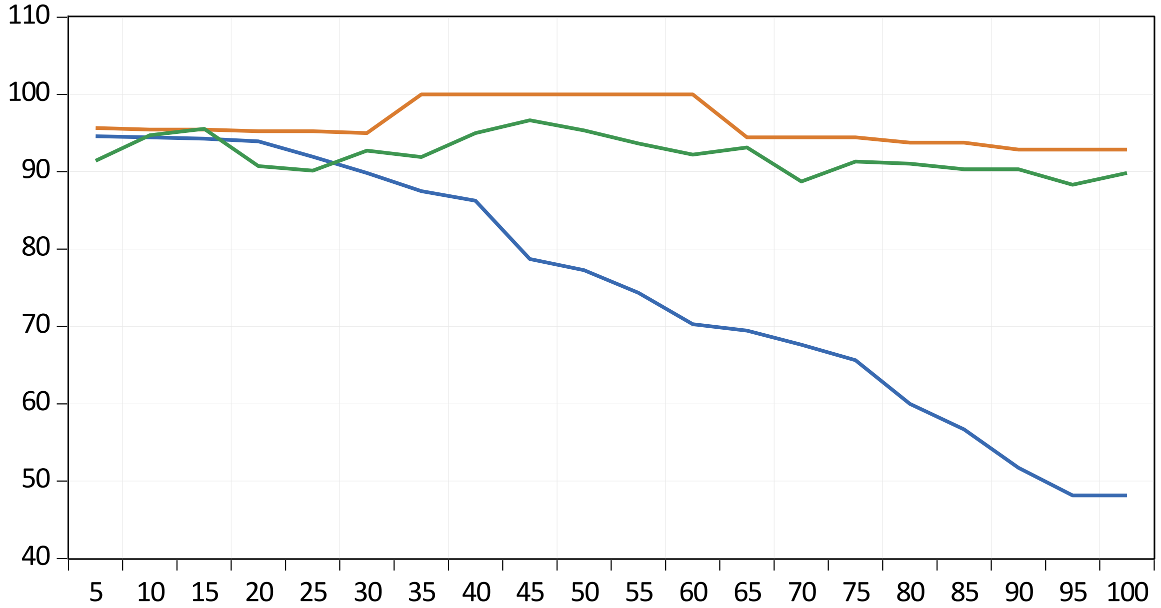

Figure 3.6.3. Total profit. The vertical axis shows the total profit (%).

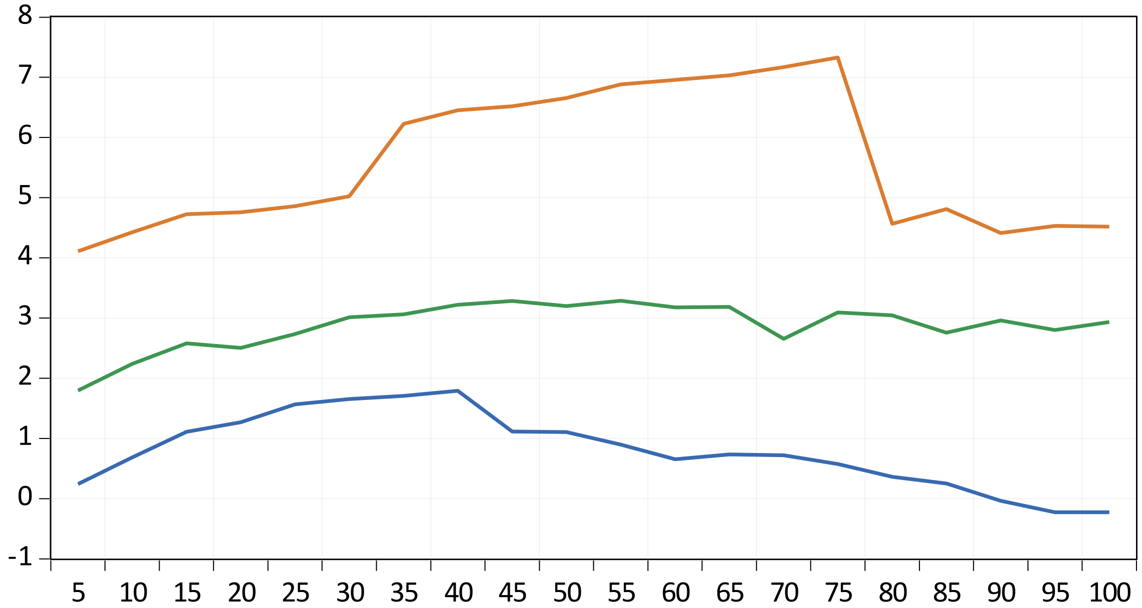

Figure 3.6.4. Total duration of trades. The vertical shows the total duration of trades (hours).

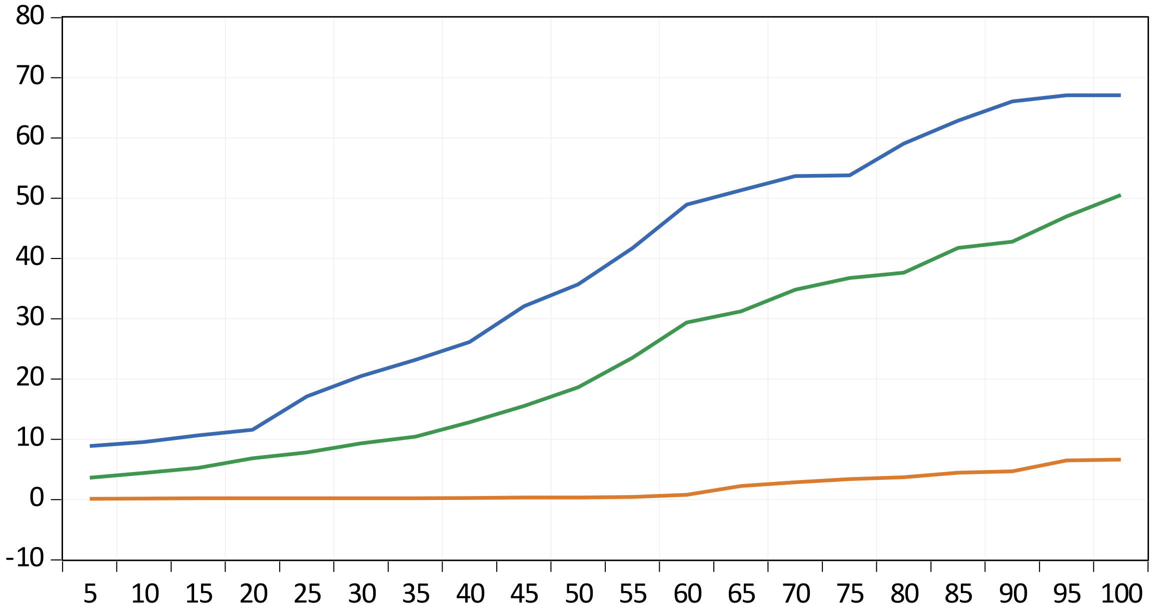

In each chart, the same pattern is observed for all three sub-samples (to different extents), which suggests that they are stable and robust. The number of trades tends to decrease when increasing the target, which shows that the partial targets can capture more opportunities than the full 100% target. Moreover, most partial targets lead to higher profits than the full target. The highest profit in the three sub-periods is generated by the 40%, 75% and 55% target respectively as a combined result of a decent profit per trade and a sufficiently large number of trades. On the other hand, the duration of trades increases monotonically when increasing the target, showing that convergence is increasingly slow when using higher targets so partial convergence may be preferable to full convergence. Finally, there is evidence that speed is positively correlated with trading results since (i) the highest speed can utilise the most opportunities, generating the largest number of trades while (ii) the lowest speed is associated with not only the lowest winning rate and smallest profit but also the longest duration and exposure of trades.

3.6. Conclusion

Using a large dataset of gold ETFs, I find that high frequency traders, who can execute their orders with minimal delay, can profit from gold ETFs with pairs trading. To my knowledge, this paper is the first to analyse high frequency pairs trading of gold ETFs. The fact that very fast order execution can capture more arbitrage opportunities and enhance profitability suggests that these opportunities are short-lived, which is consistent with other strict form pairs trading studies (e.g. Alsayed and McGroarty, 2012, Marshall et al., 2013). There is also evidence that a higher speed can enhance trading performance. According to Grossman and Stiglitz (1976) and Grossman and Stiglitz (1980), the profitability of my pairs trading may be compensation for the risks and costs involved in arbitrage and thus may motivate arbitrage activities. However, my excess return of 11.2% over the sample period or 2.1% p.a. is lower than that of equity ETFs in Marshall et al. (2013).

More importantly, I explain why the trade exit rule of full convergence used in previous studies may not be optimal and propose a rule based on partial convergence which outperforms the standard full-convergence rule, both with and without risk adjustment. In addition, the partial-convergence strategy has higher speed of convergence than the full-convergence strategy, which enables more effective use of the trading capital during trades. The outperformance of my rule is consistent in the sub-samples, showing that pairs trading can exploit market inefficiency better when it requires only partial elimination of the relative mispricing. On the other hand, pairs trading outperforms the buy-and-hold strategy after adjusting for risks, which suggests that the gold ETF market may be inefficient.

Chapter 4

High frequency lead-lag relationship and information arrival

Abstract

To my knowledge, this paper is the first study on the effect of information arrival on the lead-lag relationship among related spot instruments. Based on a large dataset of ultra-high frequency transaction prices time-stamped to the millisecond of the S&P 500 index and its two most liquid tracking ETFs, I find that their lead-lag relationship is affected by the rate of information arrival whose proxy is the unexpected trading volume of these instruments. Specifically, when information arrives, the leadership of the leading instrument may weaken when the lagging instrument responds to that information, as shown by the unexpected volume. In addition to the strength of leadership, a change in the unexpected volume in response to information arrival may also have an impact on the lead-lag correlation coefficient whether that volume change belongs to the leader or the lagger.

4.1. Introduction

The efficient market hypothesis (EMH) suggests that return predictability and arbitrage opportunities should not exist in financial markets. Accordingly, the returns of related instruments (e.g. an equity index and its futures) should show contemporaneous correlations in efficient and frictionless markets (Stoll and Whaley, 1990, Brooks et al., 1999). However the lead-lag effect, a phenomenon where a security follows the movements of another with a time delay (Huth and Abergel, 2014), is often found in the literature. Yet according to the EMH, even if returns of security A are correlated with past returns of security B, it should still be impossible to use price changes of B to forecast and make abnormal profits from price changes of A due to practical constraints such as transaction costs.

Motivated by the literature on lead-lag effect in returns (e.g. Kawaller et al., 1987, Fleming et al., 1996, Chen and Gau, 2009, Alsayed and McGroarty, 2014, Curme et al., 2015), my paper is the first to investigate the effect of information arrival on the lead-lag relationship between related stocks. To examine related stocks, I focus on ETFs because the purpose of ETFs is to track the performance of some index or asset and ETFs tracking the same index or asset can be considered very much related. Moreover, the ETF market has been growing rapidly (Shin and Soydemir, 2010, Kearney et al., 2014) since their introduction with many investors choosing ETFs as their investment vehicle. ETFs track different asset classes (e.g. stock, bond, commodity) and I choose equity ETFs for my analysis because they are easily accessible to investors, very liquid (Marshall et al., 2013) and should be representative of the economy. Equity ETFs, which track stock indices or economic sectors, are more liquid than single stocks (Ruan and Ma, 2012).

My study on the lead-lag effect in relation to information is also motivated by the fact that the lead-lag relationship exists because some financial instruments reflect information faster than others. In general, information plays an important role in financial markets. Hanousek and Podpiera (2003) state that informed trading affects the bid-ask spread because market makers set the spread to compensate for the risk of adverse selection which they face when trading with informed traders. Similarly, Gregoriou et al. (2005) argue that market makers mitigate their informational disadvantage compared to informed traders by increasing the bid-ask spread. On the other hand, the flow of information to the market also helps explain the ARCH and GARCH effects in daily stock returns (Lamoureux and Lastrapes, 1990, Sharma et al., 1996, Aragó and Nieto, 2005). In addition to prices and volatility, Frank and Kenneth (2005) suggest that the relative trade size (i.e. trade size scaled by market depth) is also affected by information, showing that informed traders prefer to trade larger volume.

I find that the lead-lag relationship is indeed influenced by the rate of information arrival. When information arrives, the leadership of the leading instrument may weaken when the lagging instrument responds to that information, as shown by the unexpected volume. In addition to the strength of leadership, a change in the unexpected volume in response to information arrival may also have an impact on the lead-lag correlation coefficient whether that volume increase belongs to the leader or the lagger.

My research is conducted in the high frequency context which has become more and more important in recent times. In current financial markets, speed is considered such an essential competitive edge that many market participants are willing to make significant technological investments to increase their speed of analysis and execution (even by only a small amount), in both absolute and relative terms (i.e. trying to be faster than their competitors). This competition for speed has pushed the boundary to extreme levels; specifically, Hasbrouck and Saar (2013) find that high frequency traders can operate with a latency of only a few milliseconds while an eye blink takes a few hundred milliseconds. Even more extreme, Goldstein et al. (2014) report that it is possible to trade in the microsecond environment. In any case, high frequency traders are among the most active and important market participants and thus I use the high frequency setting to examine the lead-lag relationship. This setting is appropriate for my analysis because the lead-lag relationship is a common phenomenon in high frequency data (Huth and Abergel, 2014).

Moreover, using high frequency data to analyse the lead-lag relationship is suitable since the increasing electronification of financial markets and high frequency trading activities have reduced the lead-lag time dramatically, to the point where data sampled at regular intervals can no longer capture this time delay (Huth and Abergel, 2014). In other words, it is not suitable to use regularly sampled data to measure high-frequency correlation (Zhang, 2011), especially when one security is traded more often than the other (De Long et al., 1990). However, in order to measure correlation, using high frequency data requires an approach different from that used for regularly sampled data. Following Alsayed and McGroarty (2014), I apply the model of Hayashi and Yoshida (2005) to calculate the contemporaneous correlation, and its extension by Hoffmann et al. (2013) to include leads and lags. This model uses the original tick data and does not require any modification such as interpolation or resampling at regular intervals (Huth and Abergel, 2014).

I contribute to the literature on lead-lag effects in returns by analysing the role of information in the lead-lag relationship among related spot instruments (i.e. equity index and ETFs) which, to my knowledge, has not been studied before. I examine the most liquid ETFs that track the S&P 500 index and thus are representative of the performance of the US economy. I show that information arrival has a part of play in the lead-lag effect among these instruments. Moreover, I conduct my analysis in the high frequency context, which is increasingly important in financial markets, using a large dataset and a novel approach proposed by Hayashi and Yoshida (2005) and Hoffmann et al. (2013).

The structure of this chapter is as follows. Section 4.2 reviews the literature on lead-lag effects in returns. Section 4.3 presents my dataset. Section 4.4 describes my estimation of the lead-lag relationship and my analysis of the effect of information arrival on this relationship. Section 4.5 and 4.6 provide the results and conclusions respectively.

4.2. Literature review

The literature on lead-lag effects is plentiful and I classify the literature in Figure 4.1.

Figure 4.1. Classification of the literature on lead-lag effects in returns. Related securities refer to securities which have the same underlying asset (i.e. stocks and their derivatives). Unrelated securities refer to stocks in general.

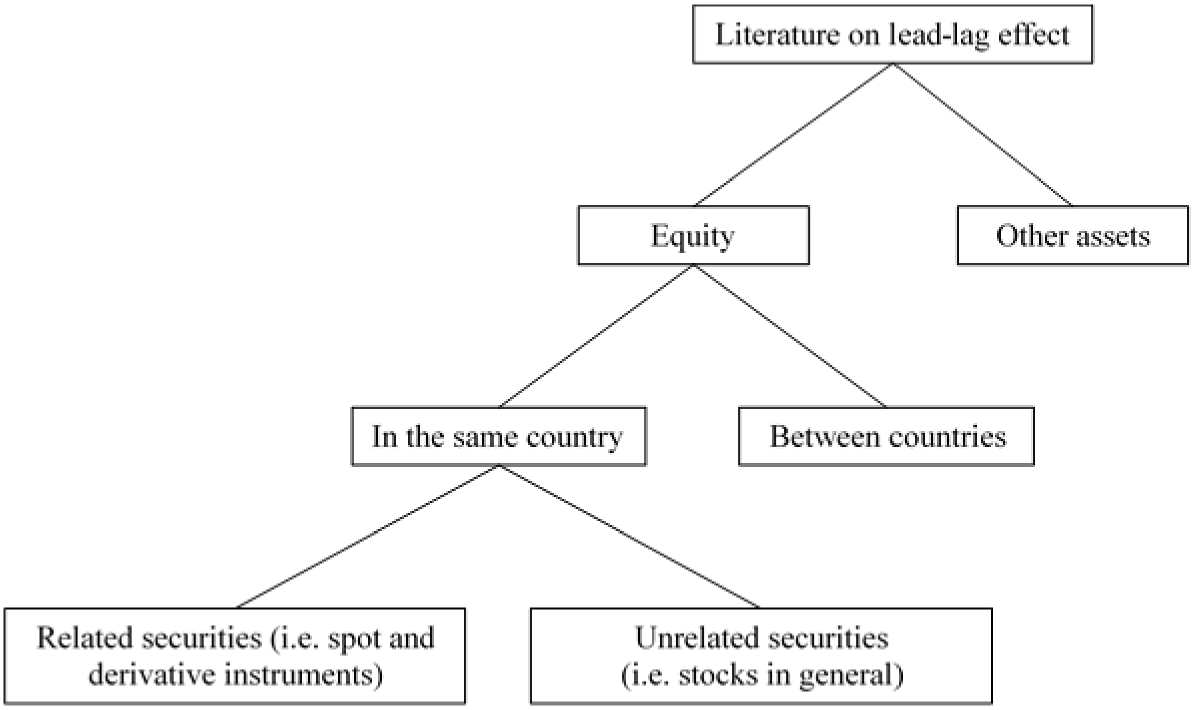

Regarding Figure 4.1, a large part of the literature on lead-lag effects focuses on equity markets so I divide the literature into studies on equity markets and studies on other assets. Studies on equity markets can be divided further into studies on instruments in the same country and studies on instruments between countries. Finally, studies on instruments in the same country include studies on related securities (i.e. spot and derivative instruments) and studies on unrelated securities (i.e. stocks in general). Because my paper investigates related securities (i.e. ETFs tracking the same index), the following literature review will focus on the lead-lag effect between related securities. I mainly look at the spot – futures and spot – options relationship since there are few studies on other relationships.

The spot – futures relationship is the most extensively studied relationship in the literature on lead-lag effects of related securities. Many papers find that the lead-lag effect is bi-directional (i.e. futures lead the index and vice versa), although the futures' lead is stronger and longer than the index's lead (e.g. Chiang and Fong, 2001, Nam et al., 2006, Ergün, 2009). While futures' lead can be up to 45 minutes (Kawaller et al., 1987), the index's lead does not exceed 15 minutes (Chan, 1992). Regarding US indices, many papers find a bi-directional spot – futures lead-lag relationship (e.g. Chan, 1992, Pizzi et al., 1998, Ergün, 2009) and only a few find a uni-directional effect where futures lead the index (e.g. Fleming et al., 1996). Regarding non-US indices, it is common to find both a bi-directional effect (e.g. Brooks et al., 1999, Frino and West, 1999) and a uni-directional effect (e.g. Najand and Min, 1999, Frino et al., 2000). However, Brooks et al. (2001) find that although the lead-lag effect can be used to produce accurate forecasts, trading these forecasts does not outperform the benchmark after considering transaction costs.

Some researchers attribute the lead-lag relationship to infrequent and nonsynchronous trading (Shyy and Vijayraghavan, 1996, Brooks et al., 1999) while others still find the lead-lag effect after considering infrequent and nonsynchronous trading (Stoll and Whaley, 1990, Grünbichler et al., 1994, Martikainen and Perttunen, 1995, Fleming et al., 1996). Another explanation for the lead-lag effect, especially the futures' lead, is the trading cost hypothesis (Nam et al., 2008). Specifically, because trading the index is cheaper in the derivative markets than in the spot market, new information should be updated in the derivative markets before the spot market (Martikainen and Perttunen, 1995, Fleming et al., 1996). Consistent with the trading cost hypothesis, Chen and Gau (2009) find that when the bid-ask spread in the spot market decreases (due to a decrease in the minimum tick size), the spot index's contribution to price discovery becomes more significant. In addition to nonsynchronous trading and trading costs, the trading mechanism also affects the lead-lag relationship. When the futures are screen-traded and the index is floor-traded, the futures' lead is longer than when both are floor-traded, since screen trading increases the price discovery speed (Grünbichler et al., 1994).

In terms of time variation, the lead-lag effect is regime-switching and non-linear (Chung et al., 2011). The futures' lead has weakened and the spot – futures integration has strengthened over time (Frino and West, 1999, Brooks et al., 1999). Lien et al. (2003) use daily data and even find that in more recent time, information flows from the spot market to the futures market, which means the index leads the futures.

However, Nam et al. (2008) warn that using low-frequency data may lead to information loss and incorrect results, which is why I use tick data in this paper.

In addition to the spot – futures relationship, the spot – options relationship has often been studied and the findings are mixed. Some authors find no lead-lag effect (Panton, 1976, Chan et al., 1993) while others find a uni-directional effect (Stephan and Whaley, 1990, Fleming et al., 1996) or a bi-directional effect (Chiang and Fong, 2001, Nam et al., 2006). However, the spot market tends to have a longer lead than the options. The spot market can lead the options by up to 20 minutes whereas the options lead the spot market by up to only 10 minutes (Chiang and Fong, 2001). Interestingly, bi-directional effects are usually found in non-US markets (e.g. Chiang and Fong, 2001, Nam et al., 2006).

Although option prices contain information not reflected in stock prices (Manaster and Rendleman, 1982), this information is not lucrative enough to cover transaction costs and search costs (Bhattacharya, 1987). On the other hand, the spot market's lead over options might be due to the infrequent trading and illiquidity of options (Chan et al., 1993, Fleming et al., 1996, Chiang and Fong, 2001). Lead-lag effects can also be explained by trading costs. Generally, information will be updated faster where it is cheaper to trade. Fleming et al. (1996) find that stocks lead options because trading stocks in the spot market is cheaper than in the option market and that futures lead options since trading costs in the future market are lower than in the option market. Ryu (2015) also finds that in the futures – options relationship, futures play a more significant role in price discovery than options. However, regarding the options – options relationship, there is no lead-lag effect between call and put options because of their similar trading costs (Fleming et al., 1996).

In summary, my literature review has focused on the lead-lag effect in returns of related securities (i.e. stocks, futures and options). The findings range from no lead-lag effect to a uni-directional effect to a bi-directional effect, with liquid and cheap instruments often leading illiquid and expensive ones. The lead-lag relationship may be affected by the trading mechanism of securities (i.e. screen-traded or floor-traded). However, this relationship might not be exploited profitably after considering transaction costs. Finally, there is also evidence of a weakening lead-lag effect and strengthening integration between markets over time.

4.3. Data

My dataset includes the S&P 500 index and the two most liquid equity ETFs in the US market which are constructed to reflect the performance of the S&P 500 index. They are SPDR S&P 500 ETF Trust (ticker symbol SPY) and iShares Core S&P 500 ETF (ticker symbol IVV), provided by State Street Global Advisors and BlackRock respectively. Using Thomson Reuters Tick History database, I collect index value and ETF transaction price data (time-stamped to milliseconds) between August 2014 and July 2015. The end of the sample period is July 2015, which is the most recent data available at the start of this study. I decide to use data from the most recent period so that it is more likely for the research results and findings to reflect accurately and be applicable to the current environment of financial markets.

In addition, due to constraints in terms of computation, the length of the sample period is only one year, which can be considered to be fairly short. To be more specific, there are two reasons for these constraints as follows. Firstly, the algorithm used to carry out the numerical analysis of the data is highly demanding with regards to computing power. Secondly, the financial instruments examined in this work (i.e. the S&P 500 index and its tracking ETFs) are among the most liquid securities not only in the US market in particular but also in international markets in general, which means that they are able to generate a lot of trading activities on a regular basis. As a result of their great popularity, the volume of their data is huge and requires a high degree of computing power to process and analyse properly. It may also be worth mentioning that in many cases, computational limitations are the main reason why the sample period of high frequency studies tends to be relatively short. Finally, the period used in my research is a period of relative tranquillity in the market (Jin, 2016) so it may be representative of typical behaviours of security prices under normal market conditions.

Following Marshall et al. (2013), in every trading day, only the main trading session from 9:30am to 4pm is considered to ensure maximum liquidity. Following the data cleaning procedure employed by Marshall et al. (2013) to remove potential data errors, I exclude observations where (i) the logarithmic return of price is higher than 25% or lower than -25% [6], or (ii) the time-stamp is within the first or last five minutes of the trading session. Table 4.1 shows the summary statistics of the data after the cleaning process. The index has the highest number of observations because it consists of a large number of stocks and moves whenever one or more component stocks move. The mean returns are small due to the large number of observations while the median returns are zero since there are many instances where consecutive transactions occur at the same price, resulting in zero returns. The returns range from -1.5% to 1.5%, with SPY and IVV showing the lowest and highest standard deviation respectively. Because of the positive skewness and leptokurtosis, the non-normality of the data is confirmed by the Jarque-Bera statistic, which is statistically significant at 1%.

Table 4.1. Summary statistics of logarithmic returns. The returns are in percentage terms. *** superscript denotes statistical significance at 1% level.

	S&P 500	SPY	IVV
Mean	4.93E-07	2.41E-06	3.1E-05
Median	0	0	0
Maximum	1.48	1.31	1.30
Minimum	-1.36	-1.44	-1.47
Standard deviation	4.75E-03	2.60E-03	0.01
Skewness	30.04	10.03	2.65
Kurtosis	30349.33	41274.72	2271.93
Jarque-Bera normality	1.35E+14 ***	1.25E+15 ***	2.10E+11 ***
Number of observations	17745864	4117085	281734

4.4. Methodology

4.4.1. Estimation of the lead-lag relationship

I analyse the lead-lag relationship between the S&P500 index and each ETF as well as between the two ETFs. To examine the lead-lag relationship between two series of non-synchronous tick data, I use the method of Hayashi and Yoshida (2005) and Hoffmann et al. (2013). My purpose is to calculate the correlation coefficients between one series and timestamp-adjusted versions of the other to find the time adjustment which maximises their correlation. I describe the specific steps below using the SPY

[6] In addition to the 25% threshold, I have used alternative cut-off points (i.e. 5%, 15%, 35% and 45%) and still got the same results.

– IVV pair as an example. These steps are equally applicable to the S&P500 – SPY and S&P500 – IVV pairs.

The first step is to estimate the contemporaneous correlation coefficient between SPY and IVV (i.e. correlation where the timestamps of both series are kept unchanged). Because of the non-synchronicity of the data (illustrated in Figure 4.2), I use the method of Hayashi and Yoshida (2005). Their method does not require data synchronisation (e.g. through interpolation) and thus can avoid potential biases.

Figure 4.2. Non-synchronicity of the data. Each dot is a data point; t is the arrival time of observations; I is the interval between two consecutive observations.

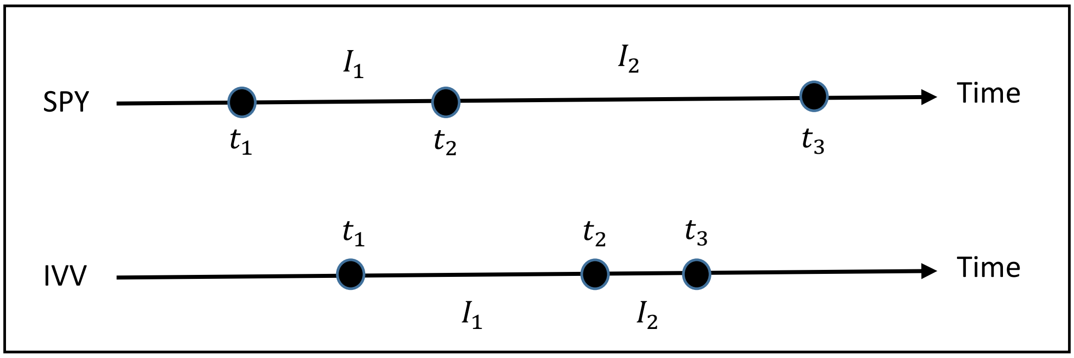

Letting R denote the return and I denote the interval between two consecutive observations, the covariance C between SPY and IVV is given by

$$C = \sum_{i,j} R_{SPY}^{I_i} R_{IVV}^{I_j} \mathbb{I} \tag{1}$$

$$\text{where} \quad \mathbb{I} = \begin{cases} 1 & if \;\; I_i \cap I_j \neq \emptyset \\ 0 & if \;\; I_i \cap I_j = \emptyset \end{cases}$$

Equation (1) shows that the covariance is calculated by summing the products of every SPY interval return and its overlapping IVV interval returns. For example, in Figure 4.2, the covariance is calculated by summing the products of the following pairs of returns: $(R_{SPY}^{I_1}, R_{IVV}^{I_1})$, $(R_{SPY}^{I_2}, R_{IVV}^{I_1})$ and $(R_{SPY}^{I_2}, R_{IVV}^{I_2})$.

The standard deviation σ of SPY and IVV is given by

$$\sigma_{SPY} = \sqrt{\sum_i (R_{SPY}^{I_i})^2} \quad (2)$$

$$\sigma_{IVV} = \sqrt{\sum_j (R_{IVV}^{I_j})^2} \quad (3)$$

The correlation coefficient ρ of SPY and IVV is given by

$$\rho = \frac{C}{\sigma_{SPY}\ \sigma_{IVV}} \quad (4)$$

After estimating the contemporaneous correlation by equation (4), the next step is to adjust the timestamp of either SPY or IVV to allow for leads and lags, and re-estimate their correlation as suggested by Hoffmann et al. (2013). I choose to fix the timestamp of SPY and adjust that of IVV. Regarding the S&P500 – SPY and S&P500 – IVV pairs, I choose to fix the timestamp of the ETFs and adjust that of the index. Figure 4.3 illustrates this process.

Figure 4.3. Example of time adjustment and correlation re-estimation. Each dot is a data point; t is the arrival time of observations; I is the interval between two consecutive observations. IVV' is created by moving every IVV observation backward in time by the same amount Δt. Then the correlation is re-estimated between SPY and IVV'.

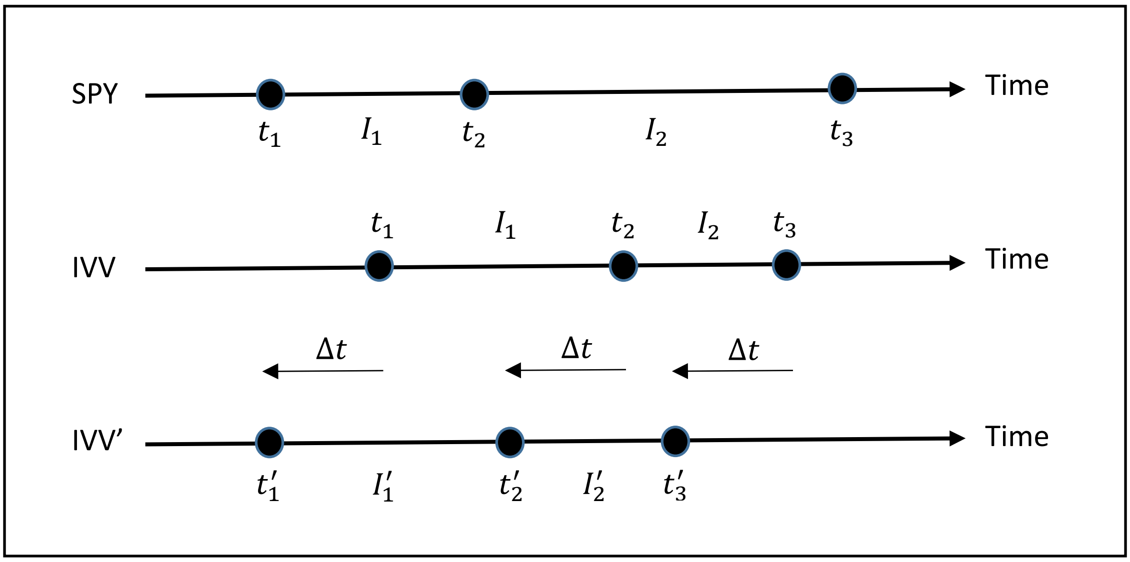

Letting IVV' denote the timestamp-adjusted IVV series, I' denote the interval between two consecutive IVV' observations and C' denote the covariance between SPY and IVV'; the correlation coefficient ρ' between SPY and IVV' is given by

$$\rho' = \frac{C'}{\sigma_{SPY}\,\sigma_{IVV'}} = \frac{\sum_{i,j} R_{SPY}^{I_i} R_{IVV'}^{I'_j}\,\mathbb{I}'}{\sqrt{\sum_i (R_{SPY}^{I_i})^2 \sum_j (R_{IVV'}^{I'_j})^2}} \tag{5}$$

$$\text{where } \mathbb{I}' = \begin{cases} 1 & if \;\; I_i \cap I'_j \neq \emptyset \\ 0 & if \;\; I_i \cap I'_j = \emptyset \end{cases}$$

Calculating ρ' with different time adjustments Δt of IVV produces the correlation curve which shows the correlation coefficient between SPY and IVV at different leads and lags of IVV. To capture the ultra-high frequency lead-lag relationship, I consider Δt from -100 milliseconds (i.e. moving IVV backward) to 100 milliseconds (i.e. moving IVV forward) with 10-millisecond increments.

After producing the correlation curve, the final step is to find the Δt which maximises the correlation. This Δt shows the temporal relationship between SPY and IVV. If it is zero, there is no lead-lag relationship; if it is negative, IVV lags SPY by Δt; if it is positive, IVV leads SPY by Δt. Also, for ease of reference, I refer to the maximum correlation on the correlation curve as the lead-lag correlation coefficient hereafter. In addition, following Huth and Abergel (2014), I calculate the lead-lag ratio (LLR) as follows.

$$LLR = \frac{\sum_i (\rho_{(\Delta t)_i})^2}{\sum_i (\rho_{-(\Delta t)_i})^2} \quad (\Delta t > 0) \tag{6}$$

The numerator of LLR is the sum of squared correlation coefficients at all leads of IVV while the denominator is the sum of squared correlation coefficients at all lags of IVV. LLR measures the relative strength of leadership (i.e. if LLR is higher than one, IVV tends to lead SPY more than lag and vice versa). Figure 4.4 shows an example of the correlation curve.

Figure 4.4. Example of the correlation curve. This curve is obtained by calculating the correlation between two instruments while fixing the timestamp of one and adjusting that of the other. The horizontal axis shows the time adjustment Δt. The vertical axis shows the correlation coefficient ρ. The peak of the curve is the lead-lag correlation and its corresponding Δt is the lead-lag time.

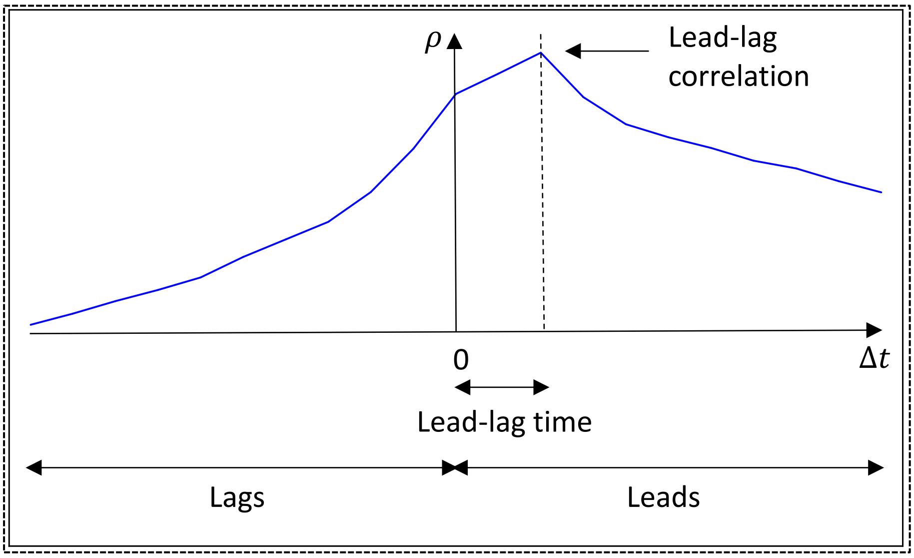

4.4.2. The effect of information arrival on the lead-lag relationship

4.4.2.1. Dependent variables

To study the effect of information arrival on the lead-lag relationship, I use regression analysis. The dependent variables are variables which represent the lead-lag relationship, namely the lead-lag correlation coefficient, the lead-lag time and the strength of leadership (measured by the lead-lag ratio). Although previous studies focus on the lead-lag correlation and the lead-lag time (e.g. Fleming et al., 1996, Nam et al., 2006, Ergün, 2009), the lead-lag ratio (calculated by Equation 6) is necessary to provide a more comprehensive analysis of the lead-lag relationship. Let us consider the following example.

Figure 4.5. Another example of a correlation curve. This curve is obtained by calculating the correlation between two instruments while fixing the timestamp of one and adjusting that of the other. The horizontal axis shows the time adjustment Δt. The vertical axis shows the correlation coefficient ρ. The peak of the curve is the lead-lag correlation and its corresponding Δt is the lead-lag time.

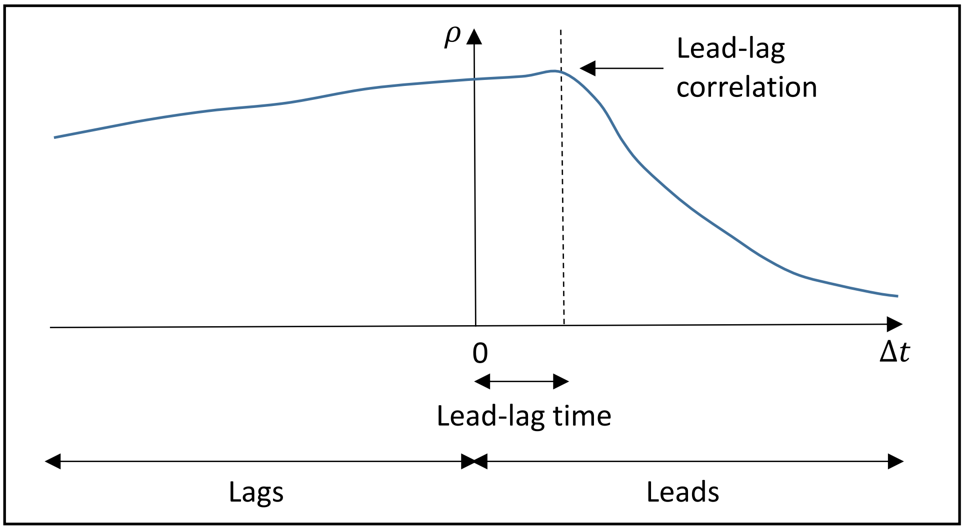

If I focus on the lead-lag correlation and the lead-lag time, I may conclude that the time-adjusted instrument leads the time-fixed one because the peak of the curve is on the 'leads' side. However, the lead-lag ratio, which examines not only the peak but also a range of leads and lags, results in a different conclusion. Since the correlation is generally higher on the 'lags' side than on the 'leads' side, the lead-lag ratio suggests that the time-adjusted instrument tends to lag the time-fixed one. To cover this type of situation, it is important to consider the lead-lag ratio in addition to the lead-lag correlation and the lead-lag time. For each pairwise combination of the S&P 500 index and the two ETFs, I follow the steps in section 4.4.1 to obtain these lead-lag quantities for every trading day in the sample period. As a result, each pair of instruments has three daily series corresponding to the three lead-lag variables which are used as the dependent variables in my regression analysis.

4.4.2.2. Independent variables

Previous research has found that the lead-lag relationship is affected by factors such as the trading cost and the trading mechanism of the instruments. Regarding the trading cost, information is generally updated faster where it is cheaper to trade. For example, because trading the index is cheaper in the derivative markets than in the spot market, new information should be updated in the derivative markets before the spot market (Martikainen and Perttunen, 1995, Fleming et al., 1996, Nam et al., 2008). However, as for the derivative markets, Fleming et al. (1996) show that there is no lead-lag effect between call and put options because of their similar trading costs. Regarding the trading mechanism, Grünbichler et al. (1994) find that when the leading instrument changes from being floor-traded to being screen-traded, its leadership strengthens since screen trading increases the price discovery speed. In my study, these factors should not contribute to the lead-lag effect because I only examine electronically traded spot instruments and no derivative.

Unlike the trading cost and trading mechanism, the information flow to the market may influence the lead-lag effect in my study. Motivated by (i) the fact that the lead-lag relationship exists because some instruments reflect information faster than others and (ii) the importance of information in financial markets (e.g. Hanousek and Podpiera, 2003, Gregoriou et al., 2005, Frank and Kenneth, 2005), I hypothesise that changes in the information flow have an impact on the lead-lag relationship. Therefore, my independent variables are variables which represent information arrival. A common proxy for the rate of information arrival is trading volume (e.g. Lamoureux and Lastrapes, 1990, Sharma et al., 1996, Aragó and Nieto, 2005) so my independent variables are daily trading volume of the S&P 500 index and the two ETFs. Table 4.2 shows the summary statistics of the trading volume. The volume is highest for the index and lowest for the IVV ETF. The index volume is platykurtic while the ETFs' volume is leptokurtic. The volume of all instruments is positively skew and non-normal, as shown by the Jarque-Bera statistic.

Table 4.2. Summary statistics of the daily trading volume. The volume is in million shares. *** superscript denotes statistical significance at 1% level.

	S&P 500	SPY	IVV
Mean	525.074	29.944	0.770
Median	509.568	26.997	0.684
Maximum	927.539	100.688	2.673
Minimum	224.945	10.861	0.175
Standard deviation	101.001	13.250	0.432
Skewness	0.742	1.719	1.715
Kurtosis	2.286	4.274	3.763
Jarque-Bera normality	27.887 ***	138.382 ***	127.118 ***

Before doing the regression analysis, it is important to ensure that the variables are stationary so I test for unit roots in the time series using the Augmented Dickey – Fuller test. If unit roots are present, the affected variables need differencing before further analysis. Table 4.3 confirms the absence of unit roots in both the lead-lag quantities and the trading volumes because the null hypothesis of a unit root is rejected at 1% level.

Table 4.3. Unit root test. This table shows the test statistic of the Augmented Dickey – Fuller test. Using daily data, I set the maximum number of lags to 5, which translates to a trading week. The reported statistic is based on the number of lags that gives the best value of Akaike Information Criterion. The null hypothesis is that a unit root is present. *** denotes statistical significance at 1% level.

	Index – IVV	Index – SPY	IVV – SPY	Index	IVV	SPY
Lead-lag correlation	-13.274 ***	-12.571 ***	-14.398 ***	-	-	-
Lead-lag time	-16.716 ***	-8.961 ***	-9.043 ***	-	-	-
Lead-lag ratio	-15.676 ***	-15.588 ***	-12.526 ***	-	-	-
Trading volume	-	-	-	-8.317 ***	-3.847 ***	-4.607 ***

To reflect the information flow more accurately, I divide trading volume into the expected and unexpected component, as suggested by Bessembinder and Seguin (1993) and Aragó and Nieto (2005). Aragó and Nieto (2005) point out that the expected and unexpected volume capture the normal level of market activity and the arrival of new information respectively. Following Bessembinder and Seguin (1993), I use an autoregressive model for each of the three volume series (i.e. S&P500 index, IVV and SPY) to obtain the corresponding expected volume. Using daily data, I set the number of AR lags to 5, which translates to a trading week. In addition to AR(5), I also account for potential weekday effects in trading volume by regressing the volume on dummy variables for different weekdays.

$$V_t = \alpha + \sum_{i=1}^{5} \beta_i V_{t-i} + \sum_{j=1}^{4} \gamma_j D_j + \varepsilon_t \tag{7}$$

where V_t is a 3x1 vector of volume because there are three instruments in my study.

α is a 3x1 vector of intercepts.

D_j is a 3x1 dummy variable which takes the value of one for a given weekday and zero otherwise. There are four variables for Monday to Thursday. When an observation is on Friday, all of them are zero.

ε_t is a 3x1 vector of residuals.

The expected volume is the sum of all terms on the right hand side of regression (7) except the error term. The unexpected volume is the error term.

4.4.2.3. Regression analysis

My hypothesis is that the information flow (i.e. unexpected volume) affects the lead-lag relationship (i.e. lead-lag correlation coefficient, lead-lag time and lead-lag ratio). For each index – ETF pair, the independent variables are the trading volume of each instrument in that pair. However, for the ETF – ETF pair, the independent variables include not only the volume of each ETF but also the index volume, because both ETFs track the index so the index volume may be relevant to the ETFs. For each pair, I estimate a separate regression for each of the three lead-lag variables (i.e. lead-lag correlation coefficient, lead-lag time or lead-lag ratio). More specifically, I estimate regression (8) for each index – ETF pair and regression (9) for the ETF – ETF pair.

$$Y_t = \alpha + \beta_1 EV_{t,SP500} + \beta_2 UV_{t,SP500} + \gamma_1 EV_{t,ETF} + \gamma_2 UV_{t,ETF} + \varepsilon_t \qquad (8)$$

$$Y_t = \alpha + \beta_1 EV_{t,SP500} + \beta_2 UV_{t,SP500} + \gamma_1 EV_{t,IVV} + \gamma_2 UV_{t,IVV} + \delta_1 EV_{t,SPY} + \delta_2 UV_{t,SPY} + \varepsilon_t \qquad (9)$$

where Y_t is a lead-lag variable.

α is the intercept.

EV_t is the expected volume.

UV_t is the unexpected volume.

ε_t is the residuals.

At this point in the regression analysis, a potential question that one may ask is whether it is advisable to adopt the Vector Autoregression (VAR) framework to capture the temporal relationship among the dependent variables (i.e. how each of them is affected by its own lagged values as well as the lagged values of the other dependent variables). Although the idea may seem tempting, I decide not to use the VAR approach in this particular case since it has been argued that there are a number of important drawbacks associated with it (Brooks, 2014). First of all, by definition the use of VAR is often just a numerical exercise that is not sufficiently motivated and supported by theoretical considerations so researchers run the risk of finding and reporting a spurious relationship which does not really exist simply due to the process of data mining (also called data snooping). Secondly, the lack of a good theoretical basis may also make it more difficult and less intuitive to correctly interpret the estimated coefficients of variables in a VAR system in a meaningful way.

Thirdly, if the VAR model is to be employed, there may be too many parameters to estimate. Specifically, the model will consist of the following list of components – (i) 9 regressions (i.e. 3 pairs of instruments with 3 dependent variables for each pair), (ii) 6 additional independent variables (i.e. 3 instruments with 2 volume variables for each), (iii) a number of lags (e.g. for these daily variables, the number of lags can be set to 5, which translates to one trading week and is still quite conservative) and (iv) 9 intercepts (i.e. one for each of the 9 regressions). Therefore, the total number of coefficient estimates is very large and thus a huge challenge to calculate, report and make sense of. Finally, in the increasingly fast-paced environment of financial markets nowadays, especially high frequency markets, one should generally expect new information to be more and more quickly incorporated into asset prices and as a result, information in the current period is likely to make a larger contribution to current behaviours of the market than information in previous periods. Because of the various reasons explained above, it may not be worth it to use the VAR method so separate regressions are used instead.

4.5. Results

4.5.1. Lead-lag relationship

Table 4.4 shows the lead-lag relationships among the S&P 500 index and its tracking ETFs, namely SPY and IVV. These results are based on the whole sample, as opposed to the daily estimation used for the regression analysis. The lead-lag relationship of each pairwise combination of the three instruments is represented by three lead-lag quantities. After obtaining the correlation curve using Equation 5, the lead-lag correlation and the lead-lag time are measured at the peak of the curve while the lead-lag ratio is calculated by Equation 6.

Table 4.4. Lead-lag relationship among the S&P 500 index and its tracking ETFs. The first name in each pair is the leader. The lead-lag time is in milliseconds. For example, that the lead-lag time of the S&P500 – IVV pair is 10 means that the S&P 500 index leads the IVV ETF by 10 milliseconds.

	Lead-lag correlation	Lead-lag time (ms)	Lead-lag ratio
S&P500 – IVV	0.1943	10	1.0058
S&P500 – SPY	0.1886	30	1.0230
IVV – SPY	0.2370	20	1.0284

Although the S&P 500 index and its tracking ETFs are highly correlated using daily data (i.e. their daily correlation coefficient is 99.9%), Table 4.4 shows that they are only moderately correlated (i.e. around 20% correlation) using tick data due to the Epps effect (i.e. an effect documented by Epps (1979) where financial instruments become less correlated at higher sampling frequencies). The ETF – ETF pair is more correlated than the two index – ETF pairs. The lead-lag time is relatively short, ranging from 10 to 30 milliseconds. The index leads both ETFs, so price discovery starts from the index and the ETFs follow, which makes sense because these ETFs are designed to track the index. However, it should be noted that in practice, sometimes it is still possible for the ETFs to lead the index if the former reflect accurate forecasts about the latter. More specifically, if market participants are confident in their prediction about the index, they can act on it by trading the ETFs, which moves them. Then if the prediction turns out to be correct, the index will follow suit. Easley et al. (2016) argue that a number of high frequency traders are able to access public information a little earlier than other traders so for a short period, this public information has become private information which will help its owner make better forecasts about the market. Regarding the lead-lag ratio, it is slightly higher than one and highest for the ETF – ETF pair. Figure 4.6 shows the correlation curves of the three pairs of instruments.

Figure 4.6. Correlation curves. This figure shows the correlation curves of the three pairs of instruments. The vertical axis shows the correlation coefficient. The horizontal axis shows the time adjustment in milliseconds (I adjust the timestamp of the first instrument in each pair). Because the correlation range of each pair is small, if I plot all the three curves on the same chart, they look like straight lines. Therefore, I plot each curve on a separate chart.

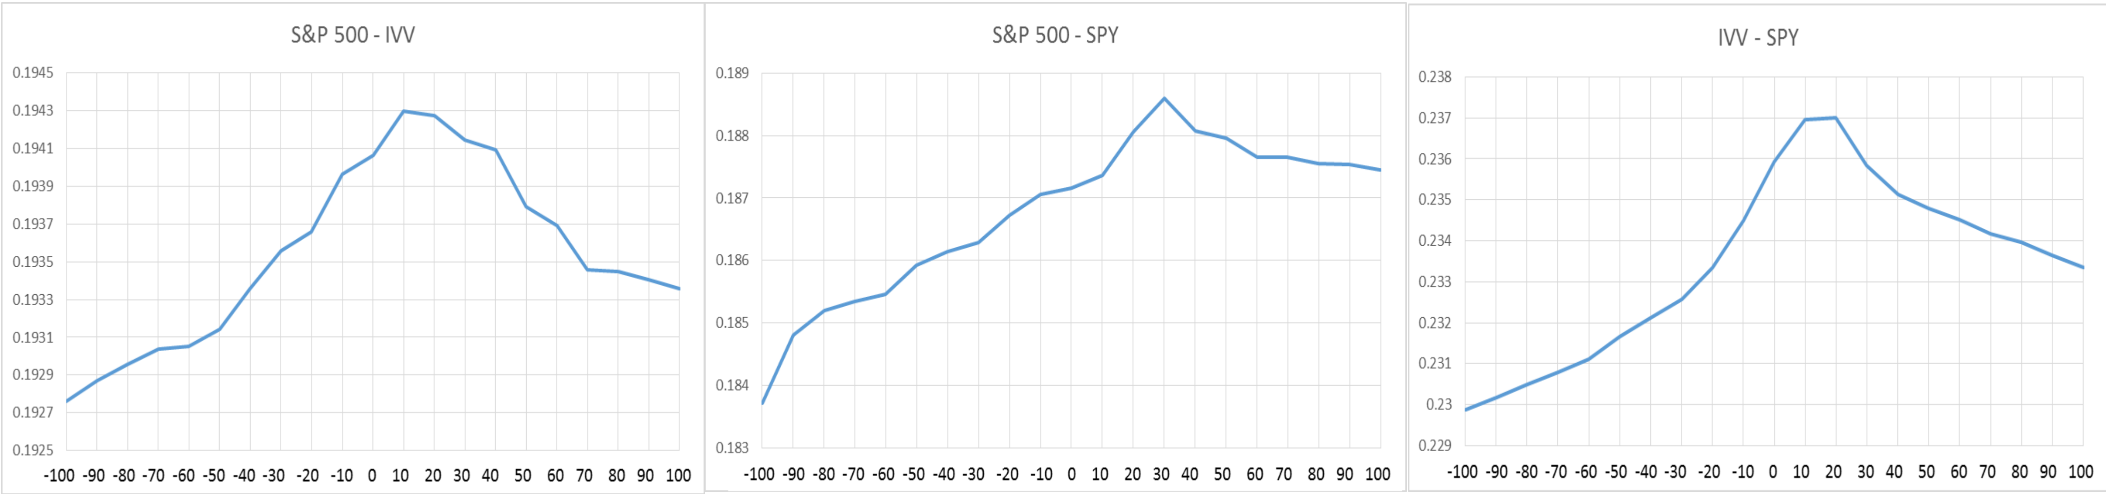

According to Figure 4.6, the peaks of all three curves correspond to a positive time adjustment on the horizontal axis, which means that the instrument whose timestamp is adjusted (i.e. the first instrument in each pair) leads the instrument whose timestamp is fixed (i.e. the second instrument in each pair). When I move away from the peak on both sides, the correlation coefficient decreases gradually. The ETF – ETF curve is smoother than the two index – ETF curves and the two IVV curves are more symmetrical on both sides of the peak than the S&P500 – SPY curve.

4.5.2. The effect of information arrival on the lead-lag relationship

Table 4.5. Autoregressive and weekday effects in trading volume. This table shows the estimated coefficients for regression (7). The volume is in million shares. The numbers in brackets are standard errors of the coefficients. *, ** and *** superscript denote statistical significance at 10%, 5% and 1% level respectively.

	Index volume		IVV volume		SPY volume	
Intercept	639.968 ***	(23.935)	0.763 ***	(0.115)	32.081 ***	(3.207)
Monday	-137.376 ***	(40.872)	-0.103	(0.062)	-4.556 ***	(1.268)
Tuesday	-120.982 ***	(36.874)	0.030	(0.065)	-3.568 ***	(1.322)
Wednesday	-118.345 ***	(33.858)	0.039	(0.068)	-0.902	(1.346)
Thursday	-124.125 ***	(40.707)	0.034	(0.068)	-1.640	(1.275)
Lag 1	0.221 ***	(0.061)	0.442 ***	(0.044)	0.639 ***	(0.069)
Lag 2	0.173 ***	(0.063)	0.160 **	(0.069)	0.153 **	(0.073)
Lag 3	0.024	(0.061)	-0.038	(0.077)	0.038	(0.079)
Lag 4	0.049	(0.058)	-0.009	(0.067)	-0.002	(0.082)
Lag 5	-0.039	(0.063)	0.181 ***	(0.049)	-0.026	(0.070)
Adjusted R^2	0.169	-	0.352	-	0.598	-

Table 4.5 reports the results of regressing trading volume on AR terms and weekday variables to obtain the expected and unexpected volume. Regarding weekday effects, the dummy variables are all zero for Friday observations so their coefficients show the difference between Friday and other weekdays. For example, the coefficient of Monday variable shows that on average the index volume is lower on Monday than on Friday and the difference is 137 million shares. Friday is higher than other days for the index and SPY volume, but only higher than Monday for the IVV volume. However, the difference is significant at 1% level for the index and SPY, but not significant for IVV. Regarding autoregressive effects, all three instruments exhibit a positive relationship with their first and second lag, which is significant at 5% and 1% level. Moreover, IVV is also positively related to its fifth lag, which is significant at

1%. In general, the adjusted R^2 indicates that autoregressive and weekday effects can help explain trading volume, especially for SPY.

Table 4.6. Effect of information arrival on the lead-lag relationship. This table shows the estimated coefficients for regression (8) and (9). The first instrument in each pair is the leader. The lead-lag time is in milliseconds. The numbers in brackets are standard errors of the coefficients. The ***, ** and * superscripts denote statistical significance at 1%, 5% and 10% level respectively.

	Index – IVV		Index – SPY		IVV – SPY	
Panel A: lead-lag correlation						
Intercept	0.1444 ***	(0.0344)	0.2094 ***	(0.0231)	0.2400 ***	(0.0273)
Expected S&P500 volume	0.0000	(0.0001)	0.0000	(0.0000)	0.0000	(0.0001)
Unexpected S&P500 volume	0.0000	(0.0000)	0.0000	(0.0000)	-0.0001 **	(0.0000)
Expected IVV volume	0.0262 *	(0.0138)	-	-	0.0178	(0.0126)
Unexpected IVV volume	0.0422 **	(0.0188)	-	-	0.0590 ***	(0.0213)
Expected SPY volume	-	-	-0.0005	(0.0004)	-0.0008	(0.0005)
Unexpected SPY volume	-	-	-0.0007 *	(0.0003)	-0.0016 ***	(0.0006)
Adjusted R^2	0.2052	-	0.1597	-	0.2631	-
Panel B: lead-lag time (ms)						
Intercept	4.8147	(3.9086)	24.7412 **	(9.5964)	20.9045 **	(9.2803)
Expected S&P500 volume	0.0089	(0.0077)	-0.0014	(0.0203)	-0.0119	(0.0194)
Unexpected S&P500 volume	-0.0053	(0.0039)	-0.0047	(0.0099)	0.0024	(0.0095)
Expected IVV volume	0.3590	(2.1331)	-	-	-2.7801	(7.2658)
Unexpected IVV volume	1.9665	(1.5652)	-	-	-3.3305	(4.2887)
Expected SPY volume	-	-	0.0593	(0.1443)	0.1191	(0.2028)
Unexpected SPY volume	-	-	0.0601	(0.1600)	0.0088	(0.1798)
Adjusted R^2	0.1302	-	0.0393	-	0.0700	-
Panel C: lead-lag ratio						
Intercept	1.0336 ***	(0.0333)	1.0540 ***	(0.0164)	0.9860 ***	(0.0446)
Expected S&P500 volume	0.0000	(0.0001)	0.0000	(0.0000)	0.0000	(0.0001)
Unexpected S&P500 volume	0.0000	(0.0000)	0.0000	(0.0000)	0.0000	(0.0000)
Expected IVV volume	-0.0122	(0.0182)	-	-	0.0129	(0.0349)
Unexpected IVV volume	-0.0088	(0.0133)	-	-	-0.0290	(0.0206)
Expected SPY volume	-	-	-0.0001	(0.0002)	0.0006	(0.0010)
Unexpected SPY volume	-	-	-0.0001	(0.0003)	-0.0020 **	(0.0009)
Adjusted R^2	0.0842	-	0.0867	-	0.1826	-

Table 4.6 reports the results of regression (8) and (9), which estimate the effect of information arrival on the daily lead-lag relationship among the S&P500 index and ETFs. The daily trading volume of each instrument (i.e. the independent variable, in million shares) is divided into the expected and unexpected component which are used as a proxy for the normal market activity and information arrival respectively. The lead-lag relationship (i.e. the dependent variable) is represented by three lead-lag quantities. After obtaining the correlation curve using Equation 5, the lead-lag

correlation (panel A) and the lead-lag time (panel B) are measured at the peak of the curve while the lead-lag ratio (panel C) is calculated by Equation 6.

According to Table 4.6, there is evidence that the lead-lag relationship among the S&P 500 index and its tracking ETFs is influenced by the rate of information arrival which is captured by the unexpected trading volume of these instruments. In particular, the impact of information is most significant on the lead-lag correlation coefficient. For all pairs, increased trading of the leader leads to a higher correlation coefficient while for most pairs, increased trading of the lagger leads to a lower correlation coefficient. Compared to the ETFs, the unexpected index volume has a less pronounced effect on the lead-lag correlation (statistically significant in only the IVV – SPY pair). Regarding the strength of leadership (measured by the lead-lag ratio), although an increase in information intensity may lead to an increase in the unexpected trading volume of both the leader and the lagger, their changes have opposite effects on the strength of leadership. If the unexpected volume of the leader increases, its leadership often strengthens while if the unexpected volume of the lagger increases, this leadership weakens. This effect is significant at 5% level for the IVV – SPY pair.

Looking at the results of the regression analysis reported in Table 4.6 above, one might notice that there is a difference between the results of the lead-lag time (i.e. panel B) and the results of the lead-lag correlation coefficients and the lead-lag ratio (i.e. panel A and C) since the results in panel B are not statistically significant with the exception of the intercepts. A direct implication of this difference is that there are more than one way to analyse the lead-lag effect (e.g. based on the lead-lag time or the lead-lag ratio) and they may or may not provide similar results to one another. Hence, researchers should examine this effect from multiple points of view and also be aware of potential discrepancies among them in order to reach a more comprehensive conclusion.

On the other hand, a possible explanation for why the results in panel B are different from those in panel A and C could be the granularity of the lead-lag analysis. As described earlier in the methodology section, the lead-lag relationship between two instruments is established by (i) keeping the timestamp of one of them unchanged while at the same time (ii) adjusting the timestamp of the other (both forward and backward) one step at a time by a fixed amount (which is set to 10 milliseconds), then

(iii) estimating the correlation between the time-fixed security and the new time-adjusted version of the other security, and finally (iv) finding the adjustment of time that leads to the highest correlation (i.e. the peak of the correlation curve). However, the shape (and the peak) of the curve may change if the incremental adjustment of time is set to a value other than the current value of 10 milliseconds. In particular, if the new value is lower than the current one, it is possible that a new peak will appear which cannot be observed using the current value because the granularity (or resolution) is not high enough.

For the purpose of illustration, let us consider two scenarios where the incremental time adjustment is set to 1 millisecond and 10 milliseconds respectively. In the case of 1-millisecond increments, it is assumed that the peak of the correlation curve is found with an adjustment of 4 milliseconds. In the case of 10-millisecond increments, it is assumed that the peak is found with zero adjustment (i.e. using the original series). As a result, there is a lead-lag effect in the former case but such an effect cannot be captured in the latter case. One can say that even if no lead-lag relationship is confirmed based on incremental adjustments of a certain value, it does not necessarily mean that such a relationship does not actually exist at all, instead it may only mean that the currently used increments are not small enough to show the effect hidden in the data. As for the reason why the value used for the adjustment of time is not set to be lower than 10 milliseconds in this study, it is in fact due to computational constraints which have been explained in more details in an earlier section.

Table 4.7. Magnitude of statistically significant results. This table shows the changes in the lead-lag correlation and the lead-lag ratio (in absolute and relative terms) caused by an increase of one standard deviation in the trading volume of the leader and the lagger. The first instrument in each pair is the leader.

	Index - IVV		Index - SPY		IVV - SPY	
	Absolute	Relative	Absolute	Relative	Absolute	Relative
Lead-lag correlation						
Leader	0.000	0.00%	0.000	0.00%	0.025	10.43%
Lagger	0.021	10.93%	-0.011	-5.80%	-0.021	-9.05%
Lead-lag ratio						
Leader	0.000	0.00%	0.000	0.00%	-0.006	-0.63%
Lagger	-0.006	-0.62%	-0.002	-0.18%	-0.010	-1.00%

Table 4.7 summarises the magnitude of statistically significant results, namely the lead-lag correlation and the lead-lag ratio. It shows the changes in these lead-lag variables caused by a given increase in the trading volume of the leader and the lagger. Because each instrument has a different level of trading activities, I report the changes caused by a volume increase of one standard deviation to make the results more comparable. The changes in the lead-lag variables are reported in absolute and relative terms (i.e. as a percentage of the overall levels in Table 4.4). According to the results, trading volume affects the lead-lag correlation more than the lead-lag ratio for all pairs in both absolute and relative terms. Except for IVV – SPY correlation, the leader has a less pronounced influence than the lagger. The IVV volume has the largest impact on the lead-lag correlation in absolute terms (0.025 for the IVV – SPY pair) and relative terms (10.93% for the index – IVV pair). On the other hand, the SPY volume has the largest effect on the lead-lag ratio also in both absolute and relative terms (-0.01 and -1% respectively for the IVV – SPY pair).

4.5.3. Robustness check

To make the results more reliable, I carry out a few robustness checks in this section. The objective of these tests is to address a number of potential issues and concerns with regards to the previous regression analysis such as (i) whether or not there is serial correlation in the dependent variables, (ii) whether or not the error terms of the regressions are contemporaneously correlated, (iii) whether the independent variables are endogenous or exogenous, and (iv) whether or not there is multicollinearity among the independent variables.

4.5.3.1. Endogeneity issue

In essence, endogeneity is an issue and needs to be considered because endogenous variables have a high correlation with the error terms (Brooks, 2014). The process of investigating this issue includes two steps as follows – (i) in the first step, the matrix of correlation coefficients between the independent variables and the corresponding residuals in regression (8) and (9) is presented with the purpose of giving readers an initial idea about the overall level of correlation between them, and then (ii) in the second step, with that knowledge I will proceed to conduct a test of endogeneity with the use of instrumental variables and two-stage least squares regressions.

Table 4.8. Correlation between independent variables and residuals. This table shows the contemporaneous correlation between the independent variables and the corresponding errors in all of the regressions. In total, there are 9 regressions which are numbered from 1 to 9 in the first row and 6 independent variables which are numbered from 1 to 6 in the first column. Letting each regression be identified by its lead-lag dependent variable, the regressions are reported in the following order: index – IVV correlation, index – IVV time, index – IVV ratio, index – SPY correlation, index – SPY time, index – SPY ratio, IVV – SPY correlation, IVV – SPY time and IVV – SPY ratio. On the other hand, the independent variables are reported in the following order: expected index volume, unexpected index volume, expected IVV volume, unexpected IVV volume, expected SPY volume and unexpected SPY volume.

	1	2	3	4	5	6	7	8	9
1	-1.66E-14	-3.53E-15	-3.97E-14	-3.12E-14	-3.81E-15	-5.79E-14	-1.10E-14	-3.11E-15	7.43E-15
2	-1.02E-16	-6.91E-17	3.69E-16	1.67E-16	8.40E-17	1.49E-16	9.86E-17	5.70E-18	7.65E-17
3	-6.90E-15	8.23E-16	-3.00E-14	-	-	-	-4.22E-16	2.35E-17	-2.53E-16
4	-7.48E-17	-7.37E-17	-3.59E-17	-	-	-	4.85E-17	2.28E-18	1.91E-16
5	-	-	-	-1.06E-14	-1.36E-15	-2.83E-14	8.60E-16	-1.51E-15	-2.52E-15
6	-	-	-	-9.27E-17	-7.86E-18	-1.06E-15	2.80E-17	-2.41E-17	-5.32E-16

As shown in Table 4.8 above, the correlation coefficients of the independent variables with the error terms are very small and they can be considered uncorrelated with each other so at first glance, it is not likely that these variables are actually endogenous. However, to confirm endogeneity or exogeneity of the variables, it is necessary to utilise instrumental variables which, to put it simply, may be treated as viable substitutes for the independent variables in the original regressions. To be able to take this role, there are generally two important conditions which a variable needs to meet in order to be chosen as an instrumental variable. The first condition is that it must not be correlated with the residuals and the second condition is that it has to be correlated with the independent variables. Taking these criteria into account, I choose the first lag of the independent variables as instrumental variables.

Table 4.9. Correlation between lagged independent variables and current residuals. This table shows the correlation between the first lag of the independent variables and the corresponding current errors in all of the regressions. In total, there are 9 regressions which are numbered from 1 to 9 in the first row and 6 independent variables which are numbered from 1 to 6 in the first column. Letting each regression be identified by its lead-lag dependent variable, the regressions are reported in the following order: index – IVV correlation, index – IVV time, index – IVV ratio, index – SPY correlation, index – SPY time, index – SPY ratio, IVV – SPY correlation, IVV – SPY time and IVV – SPY ratio. On the other hand, the independent variables are reported in the following order: expected index volume, unexpected index volume, expected IVV volume, unexpected IVV volume, expected SPY volume and unexpected SPY volume.

	1	2	3	4	5	6	7	8	9
1	0.1604	-0.0125	-0.0479	-0.0294	-0.0492	-0.0895	0.1658	-0.0521	-0.1402
2	0.0183	0.0242	0.0214	-0.0490	0.0583	0.0014	-0.0609	0.0501	-0.0256
3	0.0718	0.0334	0.0517	-	-	-	0.0577	-0.0293	-0.0149
4	-0.0396	-0.0561	-0.0617	-	-	-	-0.0519	0.0344	-0.0227
5	-	-	-	-0.0096	-0.0139	-0.0563	0.0694	-0.0047	-0.0178
6	-	-	-	0.0121	0.0163	0.0531	-0.0389	-0.0053	-0.0413

Table 4.9 shows that generally speaking, the first lag of the independent variables has a low correlation with the current errors. To be more specific, 93% of the correlation coefficients (39 out of 42) are within -10% and 10%, and all of them are within -20% and 20%. Furthermore, it is worth pointing out that by definition, the lagged values of a variable was already available and pre-determined for any given period so they are obviously not influenced by current values of other variables and not correlated with the residuals in the current period. As a result, the first lag can satisfy the first condition of instrumental variables.

Table 4.10. Relationship between the independent variables and their first lag. This table shows the correlation coefficients between the independent variables and their first lag as well as the test statistic of the Breusch – Godfrey test, which is used to confirm the presence (or absence) of serial correlation in time series. The null hypothesis is that there is no autocorrelation for the first lag of the variables. The $^{*, **, ***}$ superscripts denote statistical significance of the results at the 10%, 5% and 1% level respectively (if any).

	Correlation	Test statistic
Expected S&P500 volume	0.1624	6.734***
Unexpected S&P500 volume	0.0007	0.000
Expected IVV volume	0.7474	313.586***
Unexpected IVV volume	-0.0112	0.031
Expected SPY volume	0.8304	551.992***
Unexpected SPY volume	0.0017	0.001

As shown by the results in Table 4.10, the expected volumes are more correlated with their first lag compared to the unexpected volumes. For the expected volumes, the null hypothesis of no autocorrelation is rejected at 1% level and the average correlation is 58%. Overall, the average coefficient is 29% and the highest correlation is 83% (for expected SPY volume). The results provide some evidence of a positive relationship between the independent variables and their first lag and hence the first lag can reasonably satisfy the second condition of instrumental variables. After confirming the validity and reliability of the first lag to be used as instrumental variables, the next step is to estimate the regressions again based on these instrumental variables and the method of two-stage least squares. As the name suggests, there are two different phases in a two-stage least squares regression as follows – (i) in the first phase, the independent variables are regressed on the instrumental variables, and then (ii) in the second phase, the fitted values from the first phase will be used as a replacement for the independent variables in order to estimate the original regression. Based on the output of the two-stage least squares regressions, the last step is to implement the Durbin – Wu – Hausman test to finally confirm whether the independent variables are endogenous or exogenous. The results of this test are presented in Table 4.11.

Table 4.11. Variable endogeneity test. This table shows the test statistic of the Durbin – Wu – Hausman test, which is used to confirm endogeneity or exogeneity of variables in the lead-lag regressions. The null hypothesis is that the variables are exogenous. The *, **, *** superscripts denote statistical significance of the results at the 10%, 5% and 1% level respectively (if any). In total, there are 9 regressions which are numbered from 1 to 9 in the first row and 6 independent variables which are numbered from 1 to 6 in the first column. Letting each regression be identified by its lead-lag dependent variable, the regressions are reported in the following order: index – IVV correlation, index – IVV time, index – IVV ratio, index – SPY correlation, index – SPY time, index – SPY ratio, IVV – SPY correlation, IVV – SPY time and IVV – SPY ratio. On the other hand, the independent variables are reported in the following order: expected index volume, unexpected index volume, expected IVV volume, unexpected IVV volume, expected SPY volume and unexpected SPY volume.

	1	2	3	4	5	6	7	8	9
1	0.4920	0.0133	0.0561	0.0994	0.0565	1.5561	0.4771	0.0004	0.1264
2	2.8628*	0.0005	0.0219	0.0338	0.5258	0.7701	2.6267	0.2440	0.1503
3	3.6095*	0.0153	0.0165	-	-	-	0.0288	0.0044	0.0029
4	4.4148**	0.0285	0.1780	-	-	-	0.3917	0.1662	5.4602**
5	-	-	-	0.0544	0.1040	3.0251*	2.9555*	0.0727	0.6031
6	-	-	-	0.0378	1.0011	1.9686	6.2635**	0.7842	0.7521

The above table shows that for 83% of the cases (35 out of 42), the value of the test statistic is small and the null hypothesis of exogeneity cannot be rejected. For the remaining cases, the null hypothesis is rejected at the 10% or 5% level but not at the 1% level. Therefore, one may say that the independent variables are actually exogenous. In summary, this section has addressed the issue of endogeneity and the next section will address some other issues related to the regression analysis.

4.5.3.2. Other issues

Table 4.12. Test of autocorrelation for the dependent variables in regression (8) and (9). This table shows the test statistic of the Breusch – Godfrey test, which is used to confirm the presence (or absence) of serial correlation in time series. The null hypothesis is that there is no autocorrelation. Since the variables are observed on a daily basis, the number of lags is set to 5 which is equivalent to a trading week. The *, **, *** superscripts denote statistical significance of the results at the 10%, 5% and 1% level respectively (if any).

	Test statistic
Panel A: index – IVV pair	
Lead-lag correlation	1.668
Lead-lag time	0.524
Lead-lag ratio	0.942
Panel B: index – SPY pair	
Lead-lag correlation	1.698
Lead-lag time	2.404 **
Lead-lag ratio	0.886
Panel C: IVV – SPY pair	
Lead-lag correlation	0.125
Lead-lag time	2.094 *
Lead-lag ratio	4.234 ***

It has already been discussed earlier why it may not be necessary to include lagged variables as additional independent variables in the regression analysis, and the reported results in Table 4.12 can further reinforce this argument. According to these results, it has been confirmed that serial correlation is not present in most cases. Therefore, it is safe to say that in general the dependent variables in regression (8) and (9) are not significantly correlated with the lagged versions of themselves so their own lagged values are not likely to be able to offer much help in explaining changes in the dependent variables and those lagged variables may be excluded from the regressions.

Table 4.13. Correlation of residuals. This table shows the matrix of contemporaneous correlation among the error terms of the lead-lag regressions. There are 9 regressions in total which are numbered from 1 to 9. Letting each regression be identified by its lead-lag dependent variable, the regressions are reported in the following order: index – IVV correlation, index – IVV time, index – IVV ratio, index – SPY correlation, index – SPY time, index – SPY ratio, IVV – SPY correlation, IVV – SPY time and IVV – SPY ratio. The correlation coefficients are shown in percentage terms. By definition of a correlation matrix, all of the diagonal values are 100% which is the correlation between a time series and itself. Also by default, the lower half of the table is a mirror image of its upper half and thus the lower half is not shown for the sake of simplicity.

	1	2	3	4	5	6	7	8	9
1	100.00	9.68	-0.72	10.37	-19.90	-6.00	4.64	-25.31	-33.41
2	-	100.00	-4.62	-8.23	32.22	1.82	2.91	-9.20	-7.39
3	-	-	100.00	-1.36	-3.63	-5.80	-5.22	-2.06	0.69
4	-	-	-	100.00	-8.96	12.71	1.97	-5.77	-17.86
5	-	-	-	-	100.00	11.90	-28.66	9.12	34.08
6	-	-	-	-	-	100.00	-11.58	11.88	6.47
7	-	-	-	-	-	-	100.00	-31.43	-30.65
8	-	-	-	-	-	-	-	100.00	39.08
9	-	-	-	-	-	-	-	-	100.00

Regarding the issue of correlated errors among the regressions, based on the results in Table 4.13, one can say that on average the series of residuals from the regressions are not highly correlated with one another. More specifically, 58% of the correlation coefficients (i.e. 21 out of 36) are within -10% and 10%, 78% of them (i.e. 28 out of 36) are within -20% and 20%, and all of them are within -40% and 40%. Importantly, the coefficients are distributed pretty evenly across the board in terms of magnitude (i.e. the high values are spread out and not concentrated in just a couple of series).

Table 4.14. Correlation among independent variables. This table shows the matrix of contemporaneous correlation among the independent variables of the lead-lag regressions. In total, there are 6 variables reported in the following order: expected index volume, unexpected index volume, expected IVV volume, unexpected IVV volume, expected SPY volume and unexpected SPY volume. The correlation coefficients are shown in percentage terms. By definition of a correlation matrix, all of the diagonal values are 100% which is the correlation between a time series and itself. Also by default, the lower half of the table is a mirror image of its upper half and thus the lower half is not shown for the sake of simplicity.

	1	2	3	4	5	6
1	100.00	-0.53	35.94	-2.98	49.22	-7.51
2	-	100.00	23.18	28.96	26.82	33.92
3	-	-	100.00	0.27	77.23	8.02
4	-	-	-	100.00	10.56	54.03
5	-	-	-	-	100.00	0.45
6	-	-	-	-	-	100.00

Table 4.15. Variance inflation factors. This table shows the variance inflation factors for the independent variables in the lead-lag regressions, which are regarded as an indicator of the level of multicollinearity. In total, there are 9 regressions which are numbered from 1 to 9 in the first row and 6 independent variables which are numbered from 1 to 6 in the first column. Letting each regression be identified by its lead-lag dependent variable, the regressions are reported in the following order: index – IVV correlation, index – IVV time, index – IVV ratio, index – SPY correlation, index – SPY time, index – SPY ratio, IVV – SPY correlation, IVV – SPY time and IVV – SPY ratio. On the other hand, the independent variables are reported in the following order: expected index volume, unexpected index volume, expected IVV volume, unexpected IVV volume, expected SPY volume and unexpected SPY volume.

	1	2	3	4	5	6	7	8	9
1	1.159	1.159	1.159	1.358	1.358	1.358	1.362	1.362	1.362
2	1.169	1.169	1.169	1.254	1.254	1.254	1.267	1.267	1.267
3	1.230	1.230	1.230	-	-	-	2.658	2.658	2.658
4	1.097	1.097	1.097	-	-	-	1.534	1.534	1.534
5	-	-	-	1.467	1.467	1.467	3.184	3.184	3.184
6	-	-	-	1.142	1.142	1.142	1.586	1.586	1.586

Table 4.14 and 4.15 are related to the issue of multicollinearity. Regarding correlation among the independent variables (Table 4.14), 67% of the coefficients (10 out of 15) are lower than 30%, 80% of them (12 out of 15) are lower than 40% and 93% of them (14 out of 15) are lower than 55%. Meanwhile, Table 4.15 reports the variance inflation factor. The inflation factor for each of the variables is the same across the first three regressions because the same set of independent variables is used in each of these regressions, and the same explanation also applies to the next three as well as

the last three regressions. If the inflation factor is high, then the degree of multicollinearity among the independent variables is also high, and vice versa. According to the results, as much as 93% of the inflation factors are lower than 3 and the factor is slightly higher than 3 in only one case (i.e. expected SPY volume for the IVV – SPY pair), which suggests that the extent of relationship among the independent variables is still acceptable.

4.6. Conclusion

To my knowledge, this paper is the first study on the effect of information arrival on the lead-lag relationship among related spot instruments. Based on a large dataset of ultra-high-frequency transaction prices time-stamped to the millisecond and a novel approach proposed by Hayashi and Yoshida (2005) and Hoffmann et al. (2013), I find lead-lag effects among the S&P500 index and its two most liquid tracking ETFs. The lead-lag correlation coefficients are relatively low (i.e. around 20%) due to the Epps effect (Epps, 1979) where financial instruments become less correlated at higher sampling frequencies. The index leads both ETFs by 10 and 30 milliseconds respectively so price discovery starts from the index. Using daily unexpected trading volume of the index and ETFs to proxy for the arrival rate of information, I find that information intensity affects the lead-lag relationship, especially the lead-lag correlation coefficient and the strength of leadership. Regarding the strength of leadership, the leadership of the leader may weaken when the lagger responds to the arrival of information, as shown by the unexpected volume. Regarding the lead-lag correlation coefficient, it is influenced by changes in the unexpected volume of both the leader and the lagger.

All in all, this study has identified an important factor that has an impact on the lead-lag effect, namely the rate of information arrival. The lead-lag effect is an essential topic to study in the existing literature because it is closely related to the assimilation of new information and the process of price discovery in financial markets, both of which are key subjects to further enrich my understanding of how markets work. The findings of this work suggest that when doing research on the lead-lag relationship in the future, especially in high frequency data, researchers should pay careful attention to the influence of information arrival on this relationship and take this variable into consideration in their analysis where applicable.

Chapter 5

The Brexit vote and high frequency currency markets

Abstract

This paper studies the effect of the Brexit vote on the intraday correlation and volatility transmission among major currencies. I find that the vote causes an increase in the correlation among the safe-haven currencies of the Swiss franc and Japanese yen as well as gold, and also find a decrease in their correlation with the directly involved currencies of British sterling and the Euro. These changes are due to the appreciation of the former group and the depreciation of the latter group which represents a flight to quality of investors. I also observe a substantial decrease in volatility transmission between British sterling and the Euro following the Brexit vote due to lower levels of market integration. However the volatility transmission among the currencies has increased in general and their net spillover is positively correlated with their level of volatility and trading activities. Therefore I document the impact of the politically important Brexit vote on the high frequency correlation and volatility spillover in the foreign exchange market, which is significant at 1% level.

5.1. Introduction

The decision of the UK on June 24, 2016 to leave the EU is often referred to as the Brexit vote. This unique event is significant because this is the first time a member of the EU has voted to leave. In response to the vote, both the FTSE and other major international stock indices suffered negative returns. To be more specific, the FTSE 100, the Dow Jones Industrial Average and the DAX decreased by 3.2%, 3.4% and 6.8% respectively. In particular, banking stocks depreciated more substantially than they did when Lehman Brothers collapsed (Schiereck et al., 2016). Moreover, the vote has led to increased political and economic uncertainty since it remains to be seen what exit deal the UK will be able to negotiate with the EU (Ramiah et al., 2016). Motivated by the significance of the Brexit vote and the uncertainty after the vote, I study the effect of the Brexit vote on the intraday correlation and volatility transmission of major currencies to provide insights into investor behaviour and the resultant effect on the foreign exchange (FX) market.

The Brexit vote is an international event which will affect the global financial system, especially the FX market. This market is the largest in the world and the currencies directly involved in the vote, namely sterling and Euro, are two of the major currencies. According to the survey of the Bank of International Settlement (BIS) in 2016, the global foreign exchange market generates an average turnover of $5.1 trillion on a daily basis.[7] To truly grasp the incredible magnitude of this market, one can refer to the fact that the largest stock market in the world, namely the New York Stock Exchange (NYSE), is able to generate an average turnover of only $65.8 billion per day, which is equal to just 1.3% of the typical level of trading activities in the currency market.[8] The survey of the BIS also shows that the Euro and sterling are the second and fourth most traded currency respectively, accounting for 31.3% and 12.8% of the FX market turnover. As well as sterling and the Euro, I analyse other major currencies which are often considered safe-haven assets, namely the Swiss franc and Japanese yen (Angelo and Paul, 2010, De Bock and Filho, 2015, Grisse and Nitschka, 2015, Fatum and Yamamoto, 2016) as well as gold which has long been considered a safe-haven asset (Baur and Lucey, 2010, Baur and McDermott, 2010, Bredin et al., 2015).

[7] The survey is available at http://www.bis.org/publ/rpfx16fx.pdf

[8] The NYSE trading statistics can be found at https://www.nyse.com/data/transactions-statistics-data-library

Motivated by the importance of the Brexit vote and the literature on the effect of political and unique events on financial markets (e.g. Korkeamäki, 2011, Boutchkova et al., 2012, Pastor and Veronesi, 2012), I am the first to investigate the correlation and volatility transmission in the currency market after the vote. My study is conducted in the high-frequency context which has received less attention than the daily and lower-frequency context. Ederington and Lee (1993) and Tanner (1997) show that information is reflected quickly in exchange rates so intraday data can capture market behaviours which are invisible in daily data. Moreover, in the same sample period, using a high sampling frequency increases the number of observations and leads to more reliable results. In spite of these advantages, high frequency data may pose a problem due to the high level of noise involved. In order to help alleviate this issue, I use the 5-minute sampling frequency as per the suggestion of Andersen (2000), who argues that it is probably the highest frequency possible that is not associated with an excessive amount of noise.

I contribute to the literature by being the first to examine the effect of the politically significant Brexit vote on the sterling, Euro and major safe-haven currencies in terms of their correlation and volatility transmission. The relationship between financial instruments have become more and more important due to the continuous expansion of current financial markets and the increasingly diverse choices of assets for investment readily available to market participants at any time. Investors now have to take into consideration a number of asset classes (e.g. equity, bond, currency, to name just a few) and each of these classes consists of a large and growing number of individual assets as well as numerous derivative products (e.g. futures, options and more) based on those assets. Moreover, since the correlation structure and volatility spillover among financial assets are not static and tend to vary over time, it is even more important to consider these aspects. They play an essential role in evaluating the risks involved in investments and thus can be highly useful in a wide range of practical applications such as risk management, portfolio selection, diversification, hedging and trading. For example, based on the level of correlation of financial securities with one another, one can (i) find potentially profitable opportunities for pairs trading (also known as statistical arbitrage) which is a relatively popular trading strategy that aims to take advantage of a temporary deviation from the usual co-movement of closely related instruments or (ii) identify the optimal hedge for an asset during its entire holding period using one or more other assets, which is often necessary in multiple aspects of financial planning.

Following Antonakakis (2012), I apply the Dynamic Conditional Correlation (DCC) model of Engle (2002) for the correlation analysis and the approach of Diebold and Yilmaz (2012) to study the volatility transmission. Regarding the DCC model, I also apply the model of Aielli (2013), which ensures the consistency of the estimator. Meanwhile, Diebold and Yilmaz (2012) utilise the generalised VAR framework to produce improved variance decomposition which does not depend on the ordering of variables and hence is able to describe the relationship among markets more accurately (Dekker et al., 2001, Yang et al., 2003).

I find that since the directly involved currencies depreciate and the safe-haven currncies appreciate as a result of the vote, the correlation between the two groups decreases while the correlation within the latter group increases. Trading activities of the safe-haven currencies also increase, which is a sign of flight to quality. On the other hand, in line with the findings of Baele (2005) and Christiansen (2007) that volatility spillover has a positive correlation with the level of market integration, I find a large decrease of 64% in the volatility transmission between GBP and EUR after the Brexit vote. However, there is still an overall increase in transmission among the currencies although the main source of forecast error variance of each currency is its own volatility persistence. In line with Lee and Rui (2002) and Baele (2005), there is a positive correlation between the net transmission of the currencies and their level of volatility and trading activities. Overall, I observe that the Brexit vote, which is a recent political event of great importance, has had an impact on key aspects of the high frequency relationship among major instruments in the largest market of all, namely the foreign exchange market, and this effect is significant at 1% level.

This chapter proceeds as follows. Section 5.2 reviews the literature on the effect of political and unique events on financial markets. Section 5.3 and 5.4 describe the data and methodology used to study the correlation and volatility transmission among markets. Section 5.5 and 5.6 provide the results and conclusion respectively.

5.2. Literature review

In this section I review the literature on how financial markets are affected by political and unique events in general as well as by the Brexit vote in particular.

5.2.1. Effect of political events

Elections are among the most widely studied political events. Pantzalis et al. (2000) examine 33 countries from 1974 to 1995 and find that during the two-week period before an election, the market observes positive abnormal returns which depend on the country's level of political freedom and the election timing. In another multi-country study, Foerster and Schmitz (1997) point out that the returns also show some predictability after an election. They investigate 18 countries in four decades (1957 – 1996) and find that similar to the US, other stock markets show lower returns in the second year after a US presidential election. Therefore, US elections are an important systematic factor which affects international stock returns. However, in stark contrast to these findings, Döpke and Pierdzioch (2006) find that elections in Germany have no effect on the German equity market using data from 1960 to 2003. On the other hand, Santa-Clara and Valkanov (2003), based on a long dataset of more than 70 years (1927 – 1998), show that the excess return of the US market may change depending on which party wins the election. Interestingly, this difference exists not only around the election date and it cannot be explained by the business cycle or a risk premium. Yet again, Döpke and Pierdzioch (2006) find that this phenomenon does not hold in Germany and conclude that politics does not have much impact on the German market.

Elections affect not only market returns but also the return variance. Białkowski et al. (2008), who study 27 OECD countries during the 1980 – 2004 period, document that the country-specific component of variance may double in the election week when the results surprise the market. The magnitude of this change depends on factors such as the margin of victory, changes in the political direction of the government and the maturity of the market. In addition to the equity market, the foreign exchange market is also affected by elections. In a study on Latin American countries from 1980 to 1996, Sibley (2001) find that presidential elections are related to currency depreciation which is most significant around inauguration and argue that the political cycle is an important determinant of exchange rates. Foerster and Schmitz (1997) observe a similar pattern in the longer term where the US dollar often depreciates more in the second year after the election.

Elections are not the only important political factor. Pastor and Veronesi (2012) note that when a policy change is announced, there is often a decrease in stock prices which is positively correlated with the uncertainty about the policy and also depends on whether that policy change follows an economic recession. Moreover, policy changes are associated with an increase in stock volatility, correlation and risk premium. Another source of political uncertainty, namely an uprising, is examined by Ahmed (2014), who confirms a uni-directional volatility transmission from the stock market to the foreign exchange market in Egypt before and during the uprising in 2011, showing the influence of equity on currency. Importantly, Boutchkova et al. (2012) argue that political uncertainty has more impact on some industries than on others. Industries which rely more on trade and labour become more volatile when political risks increase locally or in the countries of trading partners. Volatility decomposition indicates that the systematic and industry-specific volatility are affected by local and global political risks respectively.

5.2.2. Effect of unique events

In addition to political events, a number of unique events have been studied in the literature such as the Euro introduction, the collapse of Lehman Brothers and terrorist attacks. Dijk et al. (2011) investigate European currencies from 1994 to 2003 and discover that after the Euro introduction in 1999, there are significant structural breaks in the correlation among exchange rates and to a lesser extent, their volatility. McGroarty et al. (2006) document the important role of various microstructure effects in determining the bid-ask spread and volatility of currencies after the European Monetary Union (EMU) came into existence. In another paper, Legrand (2014) study seven countries adopting the Euro and three EU countries maintaining their own currencies and document that since 2000, the Euro-adopting countries have experienced a substantial increase in the influence of monetary shocks. In a study of 25 currencies during the 1999 – 2013 period, Eun et al. (2015) observe a decreasing level of exchange rate risks for Euro-based agents as well as increasing influence of the Euro on other currencies since its introduction.

Regarding the effects of the Euro on equity markets, Haselmann and Herwartz (2010) show a decrease in the home bias of German investors since the Euro was introduced which leads to a decrease in national investments and an increase in investments in other EMU countries and the US. Korkeamäki (2011) points out that for most Western European countries, the correlation between stock returns and interest rate changes is no longer negative since 1999. Using daily data of the 10 most developed equity markets in Europe between 1988 and 2012, Urquhart (2014) confirms that the Euro introduction is not an essential determinant of market efficiency because after this event, some markets have become more efficient while others have become less efficient.

Another unique event, namely the failure of Lehman Brothers, has also been studied. Following the event, Baba and Packer (2009) look into the FX swap market between the US dollar and three major European currencies and find a negative correlation between deviations from the covered interest parity and the creditworthiness of US and European financial institutions. On the other hand, Chesney et al. (2011) analyse terrorist attacks in 25 countries in 11 years and conclude that terrorism has a significant negative impact on stock markets, especially on the airline and insurance industry. Similar to financial crashes and unlike natural disasters, terrorist attacks generally result in extreme returns on the event day but their influence weakens after the event.

5.2.3. Effect of the Brexit vote

Research on the Brexit vote, my event of interest, is limited because of its recency. Pain and Young (2004) argue that the UK will be affected negatively by its withdrawal from the EU due to a lower level of incoming foreign direct investment in the future. Moreover, Schiereck et al. (2016) find that the short-term negative stock return caused by the announcement of the vote is more significant than that caused by the bankruptcy of Lehman Brothers, especially for EU banks. Ramiah et al. (2016), who examine the British economy from June to July 2016, add that the effect of the Brexit vote is sector-specific and that the banking and travel sector in particular suffer from negative effects. In a related study, Tielmann and Schiereck (2017) focus on the logistics industry using a large cross-section of 107 firms in the EU and UK and show that UK companies suffer more from the vote and perform worse than their EU counterparts.

In summary, I have reviewed the literature on the effect on financial markets of political and unique events as well as of the Brexit vote itself. These events have been found to have significant impacts on markets (i) of different types, especially equity and currency, (ii) in different aspects including returns, volatility, correlation, transmission, risks and efficiency, (iii) at different horizons (i.e. during and after the event) and (iv) with differences from sector to sector. I contribute to the literature by analysing the effect of the Brexit vote on the high-frequency correlation and volatility transmission among the sterling, Euro and other major currencies.

5.3. Data

My dataset includes the currencies directly involved in the vote, namely the sterling (GBP) and Euro (EUR), and major safe-haven currencies, namely the Swiss franc (CHF), Japanese yen (JPY) and gold. I use the Thomson Reuters Tick History database to collect their 5-minute exchange rates against USD from June 2015 to June 2017. I choose the 5-minute frequency because Andersen (2000) argues that frequencies higher than 5-minute are seriously distorted by noise. The sample period is chosen such that the period before and after the event are equally long as suggested by Antonakakis (2012).

Following Marshall et al. (2013), I remove potential data errors by excluding unrealistic returns given the sampling frequency used (i.e. higher than 25% or lower than -25%). Table 5.1 shows the summary statistics of the clean data. The mean returns of GBP, CHF and gold are negative while those of EUR and JPY are positive. The range of returns is largest for GBP (i.e. from -4.6% to 1.3%) and smallest for CHF (i.e. from -0.9% to 0.9%). On the other hand, the standard deviation is highest for gold (0.06%) and lowest for EUR (0.04%). Whereas the skewness differs among the currencies (i.e. negative for GBP, CHF and gold but positive for EUR and JPY), all of them are leptokurtic. As a result, the significant Jarque-Bera statistic confirms that these variables have a non-normal distribution.

Table 5.1. Summary statistics of 5-minute returns. The currencies are priced in USD. The returns are in percentage terms. *** superscript denotes statistical significance at 1% level.

	GBP	EUR	CHF	JPY	Gold
Mean	-6.36E-05	4.06E-05	-4.29E-05	8.88E-05	-2.78E-05
Median	0.000	0.000	0.000	0.000	0.000
Maximum	1.348	0.999	0.897	2.893	1.381
Minimum	-4.594	-0.982	-0.852	-1.666	-2.131
Standard deviation	0.044	0.037	0.042	0.042	0.056
Skewness	-10.509	0.312	-0.107	4.256	-0.168
Kurtosis	999.777	44.556	17.130	293.766	41.195
Jarque-Bera normality	6.10E+09 ***	1.06E+07 ***	1.22E+06 ***	5.19E+08 ***	8.95E+06 ***

Figure 5.1 below plots the time series throughout the entire sample period. The most striking feature with regards to the trend is that prior to the Brexit vote, only GBP and JPY show a clear trend but in opposite directions (i.e. downward for GBP and upward for JPY). However, after the vote, only GBP shows a relatively clear trend, which is not surprising because GBP should have been affected more strongly by the vote than the other currencies. In addition, the post-vote trend of GBP is also downward and hence it is the only currency that has maintained a clear trend throughout the sample period. At the same time, GBP is also the only currency that has lost value during the whole period (i.e. a depreciation of 14.3%) while all of the remaining currencies have gained value (i.e. EUR by 4.6%, CHF by 1.3%, JPY by 11% and gold by 4.4%). GBP is different from the rest not only in terms of the direction of the overall change in value but also in terms of the magnitude of this change (i.e. 14.3%).

Figure 5.1. Exchange rates over time. This figure shows changes in the exchange rates during the whole sample period. The vertical axis shows the price of each currency in terms of the US dollar (i.e. how many dollars are equivalent to one unit of the currency in question). The sample period is from June 2015 to June 2017 and the Brexit vote in June 2016 is in the middle area of the graphs.

5.4. Methodology

5.4.1. Correlation

Assuming that the currency returns follow an AR(p) process, they can be written as

$$r_t \;=\; \mu + \sum_{i=1}^{p} \alpha_i \, r_{t-i} \;+\; \varepsilon_t \tag{1}$$

where r_t is a 5x1 vector of returns because there are five currencies in my study.

μ is the 5x1 vector of intercepts.

ε_t is the 5x1 vector of residuals.

The residuals can be written as

$$\varepsilon_t \;=\; D_t \, u_t \tag{2}$$

D_t is the 5x5 diagonal matrix of volatility (i.e. square root of variance). The variances follow a univariate GARCH process. I use GARCH(1,1) since Hansen and Lunde (2005) show that for the currency market, the performance of GARCH(1,1) is comparable to that of more sophisticated models. u_t is the 5x1 vector of standardised residuals and $u_t \sim N(0, I)$ where I is the 5x5 identity matrix.

I employ the DCC model of Engle (2002), which measures the correlation among the currencies over time using the covariance of standardised residuals from the AR – GARCH process.

$$Q_t \;=\; (1 - \alpha - \beta)\bar{Q} \;+\; \alpha u_{t-1} u'_{t-1} \;+\; \beta Q_{t-1} \tag{3}$$

where Q_t is the 5x5 covariance matrix of u_t.

$\bar{Q}$ is the 5x5 unconditional covariance matrix of u_t.

u_t is the 5x1 vector of standardised residuals.

α and β satisfy $\alpha \geq 0, \beta \geq 0$ and $\alpha + \beta < 1$.

Letting Q_t^* be the 5x5 diagonal variance matrix of u_t, Aielli (2013) suggests a slightly different specification of Q_t to make the model more tractable.

$$Q_t \;=\; (1 - \alpha - \beta)\bar{Q} \;+\; \alpha[Q_{t-1}^{*1/2} u_{t-1} u'_{t-1} Q_{t-1}^{*1/2}] \;+\; \beta Q_{t-1} \tag{4}$$

The 5x5 correlation matrix R_t is given by

$$R_t \;=\; Q_t^{*-1/2} \, Q_t \, Q_t^{*-1/2} \tag{5}$$

I test different model parameters (i) with different AR lags, (ii) with or without the asymmetric GARCH term used by Cappiello et al. (2006) and (iii) with or without the improvement by Aielli (2013). Based on the Akaike information criterion, the best model is AR(0) – GARCH(1,1) with the improvement by Aielli (2013).

5.4.2. Volatility transmission

Adopting the approach of Diebold and Yilmaz (2012) to analyse the volatility transmission among currencies, I employ the generalised VAR framework to produce improved variance decomposition which does not depend on the ordering of variables. For each variable, variance decomposition shows how much of its forecast error variance is due to shocks to itself and the other variables. Let us start with the following VAR system.

$$y_t = \sum_{i=1}^{p} \Theta_i \, y_{t-i} + \varepsilon_t \tag{6}$$

where y_t is the 5x1 vector of the currency variables.

Θ_i is the 5x5 coefficient matrix.

ε_t is the 5x1 vector of residuals and $\varepsilon_t \sim (0, \sum)$.

The VAR system can be transformed to the following moving average model.

$$y_t = \sum_{i=0}^{\infty} A_i \, \varepsilon_{t-i} \tag{7}$$

where A_i is the 5x5 coefficient matrix which satisfies $A_i = \sum_{j=1}^{p} \Theta_j A_{i-j}$ and A_0 is the identity matrix. According to (7), the value of y at time t + s is given by

$$y_{t+s} = \sum_{i=0}^{\infty} A_i \, \varepsilon_{t+s-i} \tag{8}$$

On the other hand, the forecast of y_{t+s} based on the information available at time t is

$$\hat{y}_{t+s|t} = \sum_{i=s}^{\infty} A_i \, \varepsilon_{t+s-i} \tag{9}$$

Therefore, the forecast error is

$$y_{t+s} - \hat{y}_{t+s|t} = \sum_{i=0}^{s-1} A_i \, \varepsilon_{t+s-i} \tag{10}$$

Diebold and Yilmaz (2012) show that the contribution of variable m's shocks to the s-step-ahead forecast error variance of variable n is given by

$$\theta_{m\to n}^{(s)} = \frac{\sigma_{mm}^{-1} \sum_{i=0}^{s-1} (e_n' \, A_i \, \sum e_m)^2}{\sum_{i=0}^{s-1} (e_n' \, A_i \, \sum A_i' \, e_n)} \tag{11}$$

where σ_{mm} is the standard deviation of shocks to variable m.

$\sum$ is the 5x5 covariance matrix of the shocks. As a result, $A_i \sum A_i'$ is the 5x5 covariance matrix of the effects of shocks.

e_n and e_m are 5x1 selection vectors whose elements take the value of one when corresponding to variable n and m, and zero otherwise. $n = m$ refers to the contribution of a variable's own shocks while $n \neq m$ refers to the contribution of shocks from the other variables.

Because the shocks to the variables may be correlated, the contributions of all variables to the forecast error variance of each of them may not add to one. Hence, the contribution of each variable is normalised so that their total contribution to each of them is equal to one.

$$\tilde{\theta}_{m \to n}^{(s)} = \frac{\theta_{m \to n}^{(s)}}{\sum_{m=1}^{5} \theta_{m \to n}^{(s)}} \tag{12}$$

Based on the variance decomposition, I can calculate several quantities of interest including the total volatility spillover index, net spillovers and net pairwise spillovers.

1. The total volatility spillover index measures how much of the total forecast error variance of all variables is due to shocks to all the other variables (i.e. not due to shocks to themselves).

$$V^{(s)} = \frac{\sum_{m,n=1;m \neq n}^{5} \tilde{\theta}_{m \to n}^{(s)}}{\sum_{m,n=1}^{5} \tilde{\theta}_{m \to n}^{(s)}} * 100 \tag{13}$$

2. The net spillover of each variable measures the difference between the volatility it gives to and the volatility it receives from all the other variables.

$$V_{n - \sum m}^{(s)} = \sum_{m=1, m \neq n}^{5} \left(\tilde{\theta}_{n \to m}^{(s)} - \tilde{\theta}_{m \to n}^{(s)} \right) * 100 \tag{14}$$

3. The net pairwise spillover measures the difference between the volatility one currency in the pair gives to and the volatility it receives from the other currency.

$$V_{n-m}^{(s)} = \left(\tilde{\theta}_{n \to m}^{(s)} - \tilde{\theta}_{m \to n}^{(s)} \right) * 100 \tag{15}$$

5.5. Results

5.5.1. Correlation

Figure 5.2. Correlation among currencies. This figure shows the correlation coefficients over time among the five currencies before and after the Brexit vote. The correlation is based on the Dynamic Conditional Correlation model of Engle (2002) and its improvement by Aielli (2013). The vote corresponds to the spikes in the middle area of the graphs.

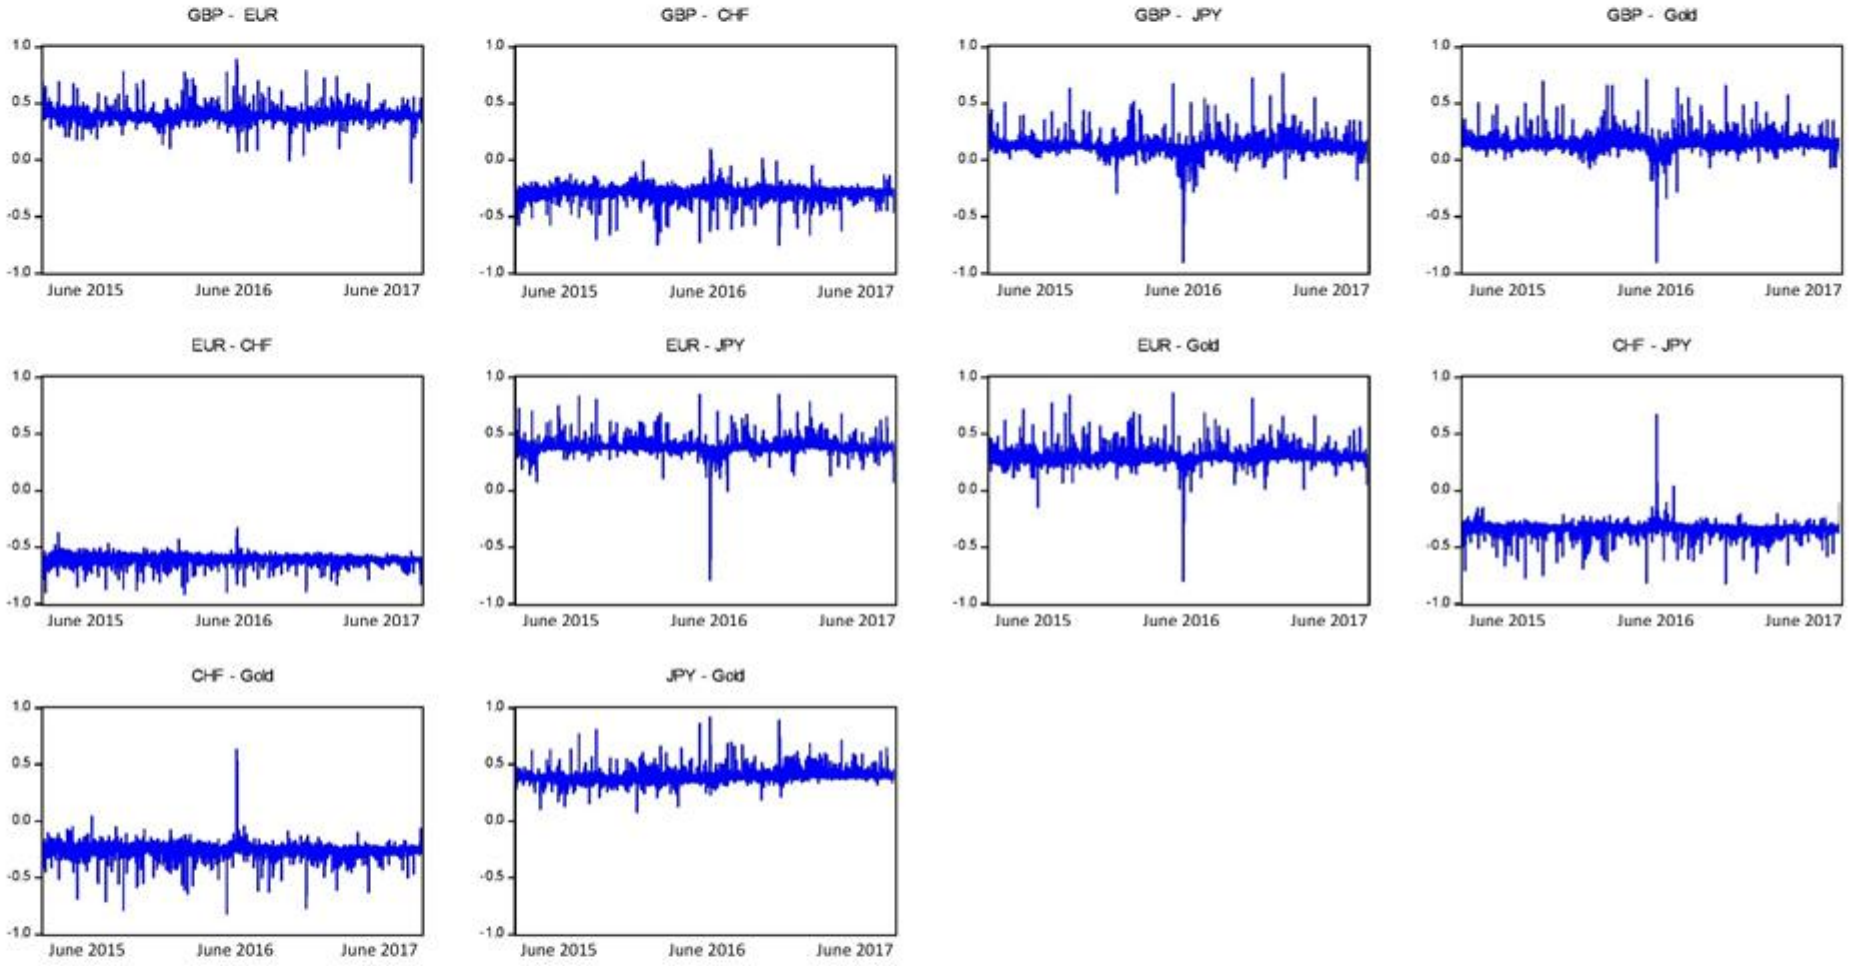

Figure 5.2 shows that among all 10 pairs, the pair of EUR and CHF is the least volatile since their correlation coefficient tends to remain in a relatively tight range compared to the other pairs, ranging only from around -50% to around -80%. More importantly, the Brexit vote triggers significant reaction from many currency pairs, evidenced by large spikes around the event (i.e. in the middle area of the graphs) which stand out from the remaining period. The most substantial reaction is observed in the pairs between the three European currencies and JPY or gold (i.e. GBP – JPY, GBP – gold, EUR – JPY, EUR – gold, CHF – JPY and CHF – gold). Among these six pairs, the correlation increases for the CHF – JPY and CHF – gold pair while it decreases for the other pairs. Interestingly, the reaction of these six pairs to the vote has changed the sign of their usual correlation coefficients. To be more specific, (i) for the GBP – JPY pair and the GBP – gold pair, their usual positive correlation of around 15% decreases to as low as -90%, (ii) for the EUR – JPY pair and the EUR – gold pair, their usual positive correlation of around 35% decreases to around -75%, and (iii) for the CHF – JPY pair and the CHF – gold pair, their usual negative correlation of around -35% increases to as high as 60%. In spite of the great magnitude of those changes, this effect is temporary because afterwards the correlation coefficients resume their pre-vote behaviours.

Table 5.2. Statistical significance of changes in correlation. This table shows the test results for statistically significant differences in correlation between the event day and other days. The correlation is between instruments in the first column and those in the first row. The numbers are the test statistic; positive values mean higher correlation in the event day than in other days and vice versa. The *, ** and *** superscript denote statistical significance at 10%, 5% and 1% level respectively.

	GBP	EUR	CHF	JPY	Gold
GBP	-	1.828 *	11.586 ***	-21.505 ***	-24.574 ***
EUR	-	-	27.909 ***	-24.280 ***	-20.499 ***
CHF	-	-	-	27.036 ***	28.865 ***
JPY	-	-	-	-	31.884 ***
Gold	-	-	-	-	-

Although Figure 5.2 shows substantial changes in correlation in response to the vote, we need to test whether these changes are statistically significant. Using the test for mean equality, I test whether correlation among the five instruments is different on the event day compared to other days. Table 5.2 confirms statistical significance for all pairs including those whose reaction to the event is not clearly visible in Figure 5.2. All results are significant at 1% level except the GBP – EUR pair which is significant at 10%. The sign of test statistics is consistent with Figure 5.2 because positive values correspond to upward movements of correlation on the event day and vice versa.

Table 5.3. Returns and changes in trading activities. This table shows the returns and changes in the number of trades and in the trading volume of the currencies on the event day. Since the spot FX volume is not available to me, I use the FX futures volume of the Chicago Mercantile Exchange as its proxy, following Bessembinder (1994), who argues that they are highly correlated. The change in trading activities on the event day is calculated relative to the normal level of activities which is the average daily number of trades and trading volume of the period before the announcement of the Brexit vote in February 2016.

	GBP	EUR	CHF	JPY	Gold
Return	-8.05%	-2.35%	1.46%	3.85%	4.69%
Change in number of trades	58.16%	103.71%	51.08%	87.46%	40.95%
Change in trading volume	502.11%	72.93%	87.58%	113.50%	319.08%

To better understand the reaction, I examine the returns and changes in trading activities of the currencies on the event day in Table 5.3. The changes in number of trades as well as in trading volume show that trading activities increase for all currencies. Regarding the number of trades, gold has the smallest increase (40.95%) while EUR has the largest increase (103.71%). Regarding the trading volume, EUR and GBP have the smallest and largest increase respectively (72.93% and 502.11%). On the other hand, GBP has the lowest return (-8.05%) whereas gold has the highest return (4.69%). In terms of the relationship between returns and trading activities, Evans and Lyons (2002) find that currency returns depend on the signed volume. Specifically, positive returns result from net purchases and negative returns result from net sales. Therefore, the depreciation of the currencies directly involved in the vote (i.e. -8.05% for GBP and -2.35% for EUR) indicates their net sales and the appreciation of the safe-haven currencies (i.e. 1.46% for CHF, 3.85% for JPY and 4.69% for gold) indicates their net purchases, both of which represent a flight to quality of investors.

5.5.2. Volatility transmission

In addition to the correlation among the currencies, I also investigate another aspect of their relationship in Table 5.4, namely the volatility transmission which is measured by their contribution to the forecast error variance of one another.

Table 5.4. Volatility transmission. Panel A and B show the variance decomposition of the currencies before and after the Brexit vote respectively. Each row shows the contribution (%) of the five instruments to the s-step-ahead forecast error variance of each of them (e.g. the first row shows how much of the GBP variance comes from each of the five instruments). Using 5-minute data, I set s to 12 which translates to one hour. The 'to others' and 'from others' values are the sum of the corresponding columns and rows excluding the diagonal values and they show how much volatility each instrument gives to and receives from the others. The 'net' values are 'to others' minus 'from others'. The spillover index (%) is the sum of all 'to others' values divided by the total volatility of all instruments and it shows how much of their volatility is due to volatility transmission.

	GBP	EUR	CHF	JPY	Gold	From others
Panel A: pre-vote						
GBP	77.45	13.26	6.79	0.67	1.84	22.55
EUR	9.35	54.38	21.76	9.70	4.82	45.62
CHF	5.15	23.32	60.35	7.69	3.50	39.65
JPY	0.62	12.62	9.34	70.83	6.59	29.17
Gold	1.89	7.13	4.80	7.45	78.73	21.27
To others	17.01	56.32	42.68	25.50	16.74	Spillover index
Net	-5.55	10.71	3.03	-3.66	-4.53	31.65
Panel B: post-vote						
GBP	78.40	13.38	7.06	0.42	0.74	21.60
EUR	9.58	56.23	23.53	5.72	4.93	43.77
CHF	5.22	24.10	59.13	6.58	4.96	40.87
JPY	0.27	6.91	7.73	67.44	17.65	32.56
Gold	0.68	6.08	5.99	18.07	69.18	30.82
To others	15.75	50.47	44.32	30.79	28.28	Spillover index
Net	-5.85	6.70	3.45	-1.77	-2.54	33.92

According to Table 5.4, the largest contribution to the forecast error variance of each variable comes from its own volatility (i.e. the diagonal values) which ranges from around 55 (i.e. EUR both before and after the vote) to around 80 (i.e. GBP both before and after the vote as well as gold before the vote). However, the spillover index is 31.65 prior to the Brexit vote and 33.92 following the vote, which shows that the overall volatility transmission has increased by 7.2% after this event. In particular, EUR gives and receives the most volatility in both the pre-vote period (i.e. 56.32 and

45.62 respectively) and the post-vote period (i.e. 50.47 and 43.77 respectively). Meanwhile, gold gives and receives the least pre-vote volatility (i.e. 16.74 and 21.27 respectively) and GBP gives and receives the least post-vote volatility (i.e. 15.75 and 21.6 respectively). After the vote, the net volatility transmission of GBP and EUR decreases (i.e. from -5.55 to -5.85 for GBP and from 10.71 to 6.7 for EUR) whereas that of the safe-haven currencies increases (i.e. from 3.03 to 3.45 for CHF, from -3.66 to -1.77 for JPY and finally from -4.53 to -2.54 for gold).

Although Table 5.4 shows that volatility transmission among the five assets has changed after the Brexit vote, we are not yet able to test statistical significance of the results. The reason is that every value in Table 5.4 is derived from the whole period before or after the vote, so the pre-vote as well as post-vote sample is condensed into and represented by only a single number. Meanwhile to test statistical significance, we need a time series to have an idea about the variation of the quantity in question. Therefore, I generate time series by re-estimating volatility transmission with the approach in section on a rolling basis. Using intraday data, I set the rolling window to one day, which means volatility transmission is estimated for the first day and then that one-day sub-sample will move forward one observation at a time until the end of the sample period. Doing this rolling estimation for both the pre-vote and post-vote period will produce two time series for us to apply the test for mean equality. Table 5.5 presents the test results for statistically significant difference in means before and after the event for key quantities representative of volatility spillover. All test statistics are significant at 1% level. Except for the volatility GBP receives from others and the net spillover of CHF, the signs of test statistics are consistent with Table 5.4 (i.e. a positive value indicates that the quantity of interest increases following the vote and vice versa).

Table 5.5. Statistical significance of changes in volatility transmission. This table shows the test results for statistically significant post-vote changes in various quantities related to volatility transmission. 'To others' and 'from others' refer to the volatility each instrument gives to and receives from the others. Net spillover is the difference between 'to others' and 'from others'. Spillover index is the proportion of all instruments' total volatility that is due to transmission. The numbers are the test statistic; positive values mean an increase in the quantity of interest after the Brexit vote and vice versa. The *, ** and *** superscript denote statistical significance at 10%, 5% and 1% level respectively.

	GBP	EUR	CHF	JPY	Gold	Overall
To others	-15.069 ***	-45.737 ***	34.429 ***	64.273 ***	165.880 ***	-
From others	2.704 ***	-2.814 ***	49.609 ***	52.712 ***	121.993 ***	-
Net spillover	-24.950 ***	-80.667 ***	-13.908 ***	40.401 ***	75.340 ***	-
Spillover index	-	-	-	-	-	53.273 ***

Table 5.6. Changes in volatility and trading activities. This table shows the changes in volatility, number of trades and trading volume of the currencies in the post-vote period. The volatility over time is obtained from the Dynamic Conditional Correlation model. The changes in each quantity are calculated based on the average value in the post-vote period relative to the average value in the pre-vote period.

	GBP	EUR	CHF	JPY	Gold
Change in volatility	22.86%	-16.19%	17.36%	-7.76%	2.05%
Change in number of trades	41.31%	58.17%	6.74%	85.31%	44.81%
Change in trading volume	-0.11%	-12.93%	10.88%	-18.14%	29.69%

Motivated by Lee and Rui (2002) and Baele (2005), I examine the potential relation of changes in volatility spillover among the instruments to changes in their volatility and trading activities in Table 5.6. For EUR, CHF and gold, their volatility changes in the same direction as their net volatility transmission (Table 5.4) with a decrease for EUR (-16.19%) and an increase for CHF (17.36%) and gold (2.05%), which is consistent with Baele (2005), who finds a positive correlation between volatility spillover and the level of volatility. Table 5.6 also shows that for all currencies except GBP, the volatility is positively correlated with the trading volume (i.e. both are negative for EUR and JPY but positive for CHF and gold), which is consistent with Lee and Rui (2002). In turn, the changes in volume in particular and in trading activities in general may be explained by the market's reaction to the Brexit vote. Despite the negligible decrease in the GBP volume (-0.11%), the large increase in its number of trades (41.31%) indicates the market's continued interest in GBP but mainly as a target for selling activities, evidenced by its negative return in the post-vote period (-12.45%). EUR is also affected by the vote but perhaps less than GBP so it is less appealing for selling activities and the market's interest in EUR decreases, as shown by its negative volume change (-12.93%). On the other hand, the market becomes more interested in the safe-haven currencies, which increases their volume and number of trades except the JPY volume (-18.14%). To better understand the pairwise interaction of the currencies, I report their net pairwise volatility transmission in Table 5.7.

Table 5.7. Pairwise volatility transmission. Panel A and B are derived from Table 5.4 and show the net volatility transmission in each pair of instruments before and after the Brexit vote respectively. For example, the first value in the second row is the volatility GBP gives EUR minus the volatility GBP receives from EUR. The diagonal values are zero because the volatility an instrument gives to and receives from itself are equal. Panel C is Panel B minus Panel A and shows the change in net pairwise transmission after the vote.

	GBP	EUR	CHF	JPY	Gold
Panel A: pre-vote					
GBP	0.00	3.91	1.65	0.04	-0.05
EUR	-3.91	0.00	-1.56	-2.92	-2.31
CHF	-1.65	1.56	0.00	-1.65	-1.30
JPY	-0.04	2.92	1.65	0.00	-0.86
Gold	0.05	2.31	1.30	0.86	0.00
Panel B: post-vote					
GBP	0.00	3.80	1.84	0.15	0.06
EUR	-3.80	0.00	-0.57	-1.19	-1.15
CHF	-1.84	0.57	0.00	-1.15	-1.03
JPY	-0.15	1.19	1.15	0.00	-0.43
Gold	-0.06	1.15	1.03	0.43	0.00
Panel C: B – A					
GBP	0.00	-0.11	0.19	0.10	0.12
EUR	0.11	0.00	1.00	1.73	1.16
CHF	-0.19	-1.00	0.00	0.49	0.27
JPY	-0.10	-1.73	-0.49	0.00	0.43
Gold	-0.12	-1.16	-0.27	-0.43	0.00

Table 5.7 shows that the Brexit vote has the most consistent effects on EUR and gold because the changes in net pairwise transmission of EUR (i.e. the EUR column in Panel C) are all negative (i.e. ranging from -0.11 to -1.73) and those of gold (i.e. the gold column in Panel C) are all positive (i.e. ranging from 0.12 to 1.16). On the other

hand, the vote has the most mixed effects on CHF (i.e. the CHF column in Panel C) with two negative values and two positive values which range from -0.49 to 1. More interestingly, these changes of EUR and gold are positively correlated with the changes in their volatility (Table 5.6), which is again consistent with Baele (2005). Among the pairs, we are most interested in the GBP – EUR pair since they are directly involved in the vote. However, in the current five-variable VAR system, their interaction with each other may be obscured by the other variables. Therefore, I redo their variance decomposition in a two-variable VAR system in Table 5.8.

Table 5.8. GBP – EUR volatility transmission. Panel A and B show the variance decomposition of GBP and EUR in a two-variable VAR system before and after the Brexit vote respectively. Each row shows the contribution (%) of both currencies to the forecast error variance of each of them. The ‘to others’ and ‘from others’ values are the sum of the corresponding columns and rows excluding the diagonal values and they show how much volatility each instrument gives to and receives from the other. The ‘net’ values are ‘to others’ minus ‘from others’. The spillover index (%) is the sum of both ‘to others’ values divided by the total volatility of both instruments and it shows how much of their volatility is due to volatility transmission.

	GBP	EUR	From others
Panel A: pre-vote			
GBP	88.88	11.12	11.12
EUR	14.56	85.44	14.56
To others	14.56	11.12	Spillover index
Net	3.44	-3.44	12.84
Panel B: post-vote			
GBP	96.24	3.76	3.76
EUR	5.55	94.45	5.55
To others	5.55	3.76	Spillover index
Net	1.78	-1.78	4.66

Table 5.8 shows that the forecast error variance of GBP and EUR is mostly due to their own volatility (i.e. the diagonal values) which ranges from around 85 (i.e. EUR before the vote) to around 95 (i.e. both GBP and EUR after the vote). Moreover, after the Brexit vote, there was a large decrease in (i) the volatility GBP gives EUR (decreasing 62% from 14.56 to 5.55), (ii) the volatility EUR gives GBP (decreasing 66% from 11.12 to 3.76), (iii) the net volatility transmission of GBP (decreasing 48% from 3.44 to 1.78) and (iv) the overall volatility transmission (decreasing 64% from 12.84 to 4.66). This decrease in the volatility transmission between GBP and EUR

following the vote is consistent with Baele (2005) and Christiansen (2007), who find a positive correlation between volatility spillover and the level of market integration.

Table 5.9. Statistical significance of changes in volatility transmission between GBP and EUR. This table shows the test results for statistically significant post-vote changes in various quantities related to volatility transmission between GBP and EUR. 'To others' and 'from others' refer to the volatility each currency gives to and receives from the other. Net spillover is the difference between 'to others' and 'from others'. Spillover index is the proportion of both currencies' total volatility that is due to transmission. The numbers are the test statistic; positive values mean an increase in the quantity of interest after the Brexit vote and vice versa. The *, ** and *** superscript denote statistical significance at 10%, 5% and 1% level respectively.

	GBP	EUR	Overall
To others	-12.004 ***	4.139 ***	-
From others	4.139 ***	-12.004 ***	-
Net spillover	-35.584 ***	35.584 ***	-
Spillover index	-	-	-3.894 ***

Table 5.8 is conceptually similar to Table 5.4 so statistical significance of the results cannot be tested yet, as explained previously for Table 5.4. Again, I carry out the rolling estimation to overcome the same problem. Table 5.9 reports the test results for mean equality between the pre-vote and post-vote period for key quantities associated with volatility transmission between GBP and EUR. Statistical significance of all the test statistics is confirmed at 1% level. Apart from the volatility GBP receives from EUR, signs of the test statistics are in line with Table 5.8 (i.e. an increase in the quantity of interest after the event is accompanied by a positive test statistic and vice versa).

5.6. Conclusion

Mine is the first study on the effect of the Brexit vote, an international political event, on the foreign exchange market, the largest market currently in existence. Using high-frequency data to accommodate the fast-moving FX market, I examine the correlation and volatility transmission among major currencies before and after the vote. For the correlation analysis, I apply the DCC model of Engle (2002) and Aielli (2013). To analyse the volatility transmission, I utilise the improved variance decomposition technique based on the generalised VAR framework. In general, I show that these important aspects of the relationship among the currencies have experienced changes following the vote which are significant at 1% level.

After the vote, there was a decrease in the correlation between the directly involved currencies (i.e. GBP and EUR) and the safe-haven currencies (i.e. CHF, JPY and gold) while there was an increase in the correlation among the safe-haven currencies. These changes were caused by the depreciation of GBP and EUR and the appreciation of the safe-haven currencies in response to the event. There was also an increase in the trading activities of the safe-haven currencies, which indicates a flight to quality. Regarding the volatility transmission, there was a significant decrease of 64% in the transmission between GBP and EUR following the Brexit vote, which is consistent with Baele (2005) and Christiansen (2007), who find that volatility spillover is positively correlated with the level of market integration. However, volatility transmission among the currencies has increased in general although most of the forecast error variance of each variable still comes from its own volatility persistence. Consistent with Lee and Rui (2002) and Baele (2005), the net transmission of the currencies is positively correlated with their level of volatility and changes in their trading activities as a result of the vote. In summary, I find that the Brexit vote has had significant effects on the foreign exchange market in terms of correlation and volatility transmission.

Chapter 6

Conclusion

This thesis has presented a multifaceted study on a number of different important aspects of the relationship among financial assets in the increasingly significant research context of high frequency markets, in an attempt to extend the current understanding of this area. This chapter is going to conclude the research by discussing the contribution and implication of the three studies in previous chapters while also pointing out their limitation and suggesting potential directions for future research.

6.1. Contribution and implication

Chapter 3 contributes to the literature on pairs trading by identifying an issue with the standard trading mechanism used in earlier studies which is likely to lead to a bias in the estimation of pairs trading profitability and the level of market inefficiency, and then proposing a solution to help correct this bias and improve the estimation. Specifically, I have explained the reason why the standard trade exit rule based on full convergence in the literature may be too strict and may result in an unnecessary increase in the duration of trades and sub-optimal trading performance. To address this issue, I suggest an alternative rule based on partial convergence which is more relaxed than full convergence and should give a more accurate estimate of the profit potential and market inefficiency. I also show the empirical application of this new trading rule in the setting of the ETF market which has become more and more economically significant over time. It has been found that the performance of the new rule is superior to that of the standard rule, whether adjusted for risks or not, and the results are robust to the potential bias of data mining. Importantly, even though the arbitrage opportunities found in this chapter are very short-lived, they can still be exploited profitably by high frequency traders who can have their trade orders executed extremely fast with almost no delay. The returns from these opportunities are more than enough to cover transaction costs and make up for the risks involved in trading, which is evidence against the efficiency of this market at high frequency.

Chapter 3 has implications for both academics and practitioners. Researchers are now aware of the downward bias in estimation that has affected previous works on pairs trading as well as a simple solution to deal with this problem and improve their estimates. Meanwhile, market traders have been shown a simple change that they can immediately apply to the standard specification of this popular arbitrage strategy in order to easily enhance their current trading performance in a reliable way. The empirical analysis in this chapter is conducted in a particular market, namely the ETF market, but the new trading rule should be equally applicable to other markets as well because it is based on consideration of the core strategy of pairs trading which does not depend on any specific context. However, it should be noted that although the observation that partial convergence is a better rule than full convergence may hold in general across various markets and there should be an optimal level of partial convergence, the actual optimal level in a given situation (i.e. market and sample period) is determined by the data and may be different from that in another situation.

Chapter 4 contributes to the literature on lead-lag effects between financial instruments by analysing these effects in a setting where it is inherently challenging to do so and identifying an important factor that has an impact on the lead-lag relationship. Specifically, the chapter has examined this relationship among a major stock index and its tracking ETFs in the US market, which is the largest equity market in the world. Since these ETFs are designed to follow the index as closely as possible, the lead-lag relationship among them should be short-lived and difficult to observe. To overcome this problem, I use ultra-high frequency data, together with a special approach to handle this particular type of data. More importantly, it has been found that the arrival of information in the market is a factor that can affect the lead-lag relationship. It is necessary to be aware of such determinants in order to better understand this relationship which has a close link to key issues of financial markets such as the price discovery process. According to the findings in this chapter, the rate of information arrival is an important variable that should be taken into account by researchers who study the lead-lag effect in the future, particularly in high frequency data.

Chapter 5 contributes to the literature on the effect of political and unique events on financial markets by studying the politically significant Brexit referendum, which is an unprecedented and internationally influential event. To be specific, the chapter has examined the influence of this event on the correlation and spillover of volatility in the foreign exchange market, the most actively traded market around the globe, and has documented substantial changes following the event in these important aspects of the relationship among currencies. These findings can benefit not only academics but also practitioners as well as the government. For academics, understanding the vote may help researchers to better understand the entire Brexit situation when it is over since the vote is an essential step towards the actual Brexit, which is going to take place in the future. The study is useful also because (i) studies on the Brexit vote are somewhat limited due to the recency of the event and (ii) many research papers on this event focus on individual instruments while this study focuses on the relationship and interaction among different instruments.

For practitioners, being aware of changes in correlation and volatility transmission is helpful for their usual activities such as portfolio management and diversification. In general, when there is a decrease in these quantities, financial assets become less closely related to one another so they are better for the purpose of diversification, and vice versa. For the government (specifically that of the UK, which is the country most affected by Brexit), they may expect further impact from the actual Brexit, given the reaction of the market to the vote, and act accordingly. For example, it has been found that there is a considerable decrease in the transmission of volatility between the sterling and the euro after the vote, possibly due to a lower degree of integration between the two markets as a result of the UK's decision to leave the EU. If such an effect is desirable to the UK government, they should try to carry out their future negotiation with the EU in a way that can accelerate the process of Brexit, such as proposing an earlier completion date of the transition phase (i.e. the actual Brexit).

6.2. Limitation

The potential limitation of this thesis is related to the length of the sample period used in Chapter 4 and 5; specifically, one may argue that the sample periods in these chapters are fairly short. Generally speaking, a long sample is often preferable because it can help researchers to be more confident in the results obtained from their studies. However, there are a few reasons why the sample length in Chapter 4 and 5 has been chosen as such. In Chapter 4, the analysis of the lead-lag relationship is computationally intensive so the sample length is chosen in order to overcome computational constraints. On the other hand, in Chapter 5, the sample period is not longer since the Brexit vote was a recent event. Moreover, the short time span of these sample periods is compensated for by the density of the data which is very high due to the use of high frequency data. This type of data can capture a huge amount of information and make the results more reliable, especially when as many as hundreds of millions of observations are used. High frequency data can also help researchers to tackle difficult challenges which cannot be overcome with low frequency data, such as documenting tiny lead-lag effects or fleeting trading opportunities.

6.3. Future research

A possibility for future research is the application of the analyses in this thesis to a variety of different contexts and situations. For example, the pairs trading strategy in Chapter 3 can be applied to other settings such as (i) the futures market using futures contracts with different maturities and / or (ii) arbitrage between the ETF market and the futures market, in addition to trading between ETFs which has already been considered, in order to evaluate the variations in terms of performance of pairs trading in general and the new trading rule in particular across a number of markets. Similarly, the analysis of the lead-lag relationship and information arrival in Chapter 4 can also be applied to the above settings to assess how this relationship and the effect of its influencing factors may vary depending on the research context.

With regards to research contexts, it is not only about spot and derivative instruments but also about other dimensions such as asset classes. There is one asset class which may be particularly interesting and informative to study, namely cryptocurrency. There are at least three reasons why this asset class is promising for research in the future. First of all, although they are called currency, cryptocurrencies may be considered a separate asset class because of some fundamental differences from traditional currencies. For instance, Bitcoin (i.e. the first as well as the biggest cryptocurrency) has a fixed supply as only 21 million units will ever be created. Secondly, compared to other asset classes, cryptocurrency is relatively new and less understood but it has gained a lot of attention in recent times. Last but not least, in addition to Bitcoin, there is a large number of cryptocurrencies available in the market at the moment and more are likely to come, which offers a wide range of opportunities to researchers.

On the other hand, another potential direction for future works is related to the Brexit event examined in Chapter 5. Even though the vote was an important part of Brexit, the situation is still ongoing and not over since the UK has not officially left the EU yet. Therefore, when Brexit is complete, it may be necessary to analyse the response of the market to the actual Brexit which is likely to depend partly on the negotiation between the involved parties regarding the exit deal for the UK. To be more specific, the analysis in Chapter 5 can be extended to evaluate changes in the market following the actual Brexit in terms of the correlation and spillover of volatility among currencies, and possibly other financial assets as well. Such a study will provide a more comprehensive observation of the development over time of the whole Brexit situation and its impact on financial markets in different phases.

List of References

ABREU, D. & BRUNNERMEIER, M. K. 2002. Synchronization risk and delayed arbitrage. *Journal of Financial Economics,* 66**,** 341-360.

ACKERT, L. F. & TIAN, Y. S. 2000. Arbitrage and Valuation in the Market for Standard and Poor's Depositary Receipts. *Financial Management,* 29**,** 71-87.

AHMED, W. M. A. 2014. Dynamic Interactions between Egyptian Equity and Currency Markets Prior to and during Political Unrest. *Applied Financial Economics,* 24**,** 1347-1359.

AIELLI, G. P. 2013. Dynamic Conditional Correlation: On Properties and Estimation. *Journal of Business and Economic Statistics,* 31**,** 282-299.

AITKEN, M., CUMMING, D. & ZHAN, F. 2015. High frequency trading and end-of-day price dislocation. *Journal of Banking and Finance,* 59**,** 330-349.

AKRAM, Q. F., RIME, D. & SARNO, L. 2009. Does the law of one price hold in international financial markets? Evidence from tick data. *Journal of Banking and Finance,* 33**,** 1741-1754.

ALSAYED, H. & MCGROARTY, F. 2012. Arbitrage and the Law of One Price in the market for American depository receipts. *Journal of International Financial Markets, Institutions and Money,* 22**,** 1258-1276.

ALSAYED, H. & MCGROARTY, F. 2014. Ultra-high-Frequency Algorithmic Arbitrage across International Index Futures. *Journal of Forecasting,* 33**,** 391-408.

ANANTH, M. 2012. Exchange-Traded Funds, Market Structure, and the Flash Crash. *Financial Analysts Journal,* 68**,** 20-35.

ANDERSEN, T. G. 2000. Some Reflections on Analysis of High-Frequency Data. *Journal of Business & Economic Statistics,* 18**,** 146-153.

ANDERSEN, T. G. & BOLLERSLEV, T. 1998. Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts. *International Economic Review,* 39**,** 885-905.

ANDERSEN, T. G., BOLLERSLEV, T. & LANGE, S. 1999. Forecasting financial market volatility: Sample frequency vis-à-vis forecast horizon. *Journal of Empirical Finance,* 6**,** 457-477.

ANGELO, R. & PAUL, S. 2010. Safe Haven Currencies. *Review of Finance,* 14**,** 385.

ANTONAKAKIS, N. 2012. Exchange return co-movements and volatility spillovers before and after the introduction of euro. *Journal of International Financial Markets, Institutions & Money,* 22**,** 1091-1109.

ARAGÓ, V. & NIETO, L. 2005. Heteroskedasticity in the returns of the main world stock exchange indices: volume versus GARCH effects. *Journal of International Financial Markets, Institutions & Money,* 15**,** 271-284.

BABA, N. & PACKER, F. 2009. From turmoil to crisis: Dislocations in the FX swap market before and after the failure of Lehman Brothers. *Journal of International Money & Finance,* 28**,** 1350-1374.

BAELE, L. 2005. Volatility Spillover Effects in European Equity Markets. *Journal of Financial & Quantitative Analysis,* 40**,** 373-401.

BAUR, D. G. & LUCEY, B. M. 2010. Is Gold a Hedge or a Safe Haven? An Analysis of Stocks, Bonds and Gold. *The Financial Review,* 45**,** 217-229.

BAUR, D. G. & MCDERMOTT, T. K. 2010. Is gold a safe haven? International evidence. *Journal of Banking and Finance,* 34**,** 1886-1898.

BENOS, E., BRUGLER, J., HJALMARSSON, E. & ZIKES, F. 2017. Interactions among High-Frequency Traders. *Journal of Financial & Quantitative Analysis,* 52**,** 1375-1402.

BENOS, E. & SAGADE, S. 2016. Price discovery and the cross-section of high-frequency trading. *Journal of Financial Markets,* 30**,** 54-77.

BESSEMBINDER, H. 1994. Bid - ask spreads in the interbank foreign exchange markets. *Journal of Financial Economics,* 35**,** 317-348.

BESSEMBINDER, H. & SEGUIN, P. J. 1993. Price Volatility, Trading Volume, and Market Depth: Evidence from Futures Markets. *Journal of Financial and Quantitative Analysis,* 28**,** 21-39.

BHATTACHARYA, M. 1987. Price Changes of Related Securities: The Case of Call Options and Stocks. *Journal of Financial & Quantitative Analysis,* 22**,** 1-15.

BIAIS, B., FOUCAULT, T. & MOINAS, S. 2015. Equilibrium fast trading. *Journal of Financial Economics,* 116**,** 292-313.

BIAŁKOWSKI, J., GOTTSCHALK, K. & WISNIEWSKI, T. P. 2008. Stock market volatility around national elections. *Journal of Banking and Finance,* 32**,** 1941-1953.

BLACK, F. 1986. Noise. *Journal of Finance,* 41**,** 529-543.

BLOSE, L. E. & GONDHALEKAR, V. 2013. Weekend gold returns in bull and bear markets. *Accounting & Finance,* 53**,** 609-622.

BORCH, C. & LANGE, A.-C. 2017. High-frequency trader subjectivity: emotional attachment and discipline in an era of algorithms. *Socio-Economic Review,* 15**,** 283-306.

BOUTCHKOVA, M., HITESH, D., DURNEV, A. & MOLCHANOV, A. 2012. Precarious Politics and Return Volatility. *Review of Financial Studies,* 25**,** 1111-1154.

BREDIN, D., CONLON, T. & POTÌ, V. 2015. Does gold glitter in the long-run? Gold as a hedge and safe haven across time and investment horizon. *International Review of Financial Analysis,* 41**,** 320-328.

BROGAARD, J., HENDERSHOTT, T. & RIORDAN, R. 2014. High-Frequency Trading and Price Discovery. *Review of Financial Studies,* 27**,** 2267-2306.

BROGAARD, J., HENDERSHOTT, T. & RIORDAN, R. 2017. High frequency trading and the 2008 short-sale ban. *Journal of Financial Economics,* 124**,** 22-42.

BROOKS, C. 2014. *Introductory econometrics for finance*, Cambridge : Cambridge University Press, c2014, 3rd ed.

BROOKS, C., GARRETT, I. & HINICH, M. J. 1999. An alternative approach to investigating lead-lag relationships between stock and stock index futures markets. *Applied Financial Economics,* 9**,** 605-613.

BROOKS, C., REW, A. G. & RITSON, S. 2001. A trading strategy based on the lead–lag relationship between the spot index and futures contract for the FTSE 100. *International Journal of Forecasting,* 17**,** 31-44.

BROUSSARD, J. P. & VAIHEKOSKI, M. 2012. Profitability of pairs trading strategy in an illiquid market with multiple share classes. *Journal of International Financial Markets, Institutions & Money,* 22**,** 1188-1201.

BRUNNERMEIER, M. K. & PEDERSEN, L. H. 2005. Predatory trading. *Journal of Finance,* 60**,** 1825-1864.

BUDISH, E., CRAMTON, P. & SHIM, J. 2015. The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *Quarterly Journal of Economics,* 130**,** 1547-1621.

CAGINALP, G., DESANTIS, M. & SAYRAK, A. 2014. The nonlinear price dynamics of U.S. equity ETFs. *Journal of Econometrics,* 183**,** 193-201.

CAPPIELLO, L., ENGLE, R. F. & SHEPPARD, K. 2006. Asymmetric Dynamics in the Correlations of Global Equity and Bond Returns. *Journal of Financial Econometrics,* 4**,** 537-572.

CARRION, A. 2013. Very fast money: High-frequency trading on the NASDAQ. *Journal of Financial Markets,* 16**,** 680-711.

CHABOUD, A. P., VEGA, C., CHIQUOINE, B. & HJALMARSSON, E. 2014. Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market. *Journal of Finance,* 69**,** 2045-2084.

CHAN, K. 1992. A Further Analysis of the Lead--Lag Relationship Between the Cash Market and Stock Index Futures Market. *Review of Financial Studies,* 5**,** 123-152.

CHAN, K., CHUNG, Y. P. & JOHNSON, H. 1993. Why Option Prices Lag Stock Prices: A Trading-based Explanation. *Journal of Finance,* 48**,** 1957-1967.

CHEN, Y. L. & GAU, Y. F. 2009. Tick Sizes and Relative Rates of Price Discovery in Stock, Futures and Options Markets: Evidence from the Taiwan Stock Exchange. *Journal of Futures Markets,* 29**,** 74-93.

CHEN, Z. & KNEZ, P. J. 1995. Measurement of Market Integration and Arbitrage. *Review of Financial Studies,* 8**,** 287-325.

CHESNEY, M., RESHETAR, G. & KARAMAN, M. 2011. The impact of terrorism on financial markets: An empirical study. *Journal of Banking and Finance,* 35**,** 253-267.

CHIANG, R. & FONG, W.-M. 2001. Relative informational efficiency of cash, futures, and options markets: The case of an emerging market. *Journal of Banking and Finance,* 25**,** 355-375.

CHORDIA, T., GOYAL, A., LEHMANN, B. N. & SAAR, G. 2013. High-frequency trading. *Journal of Financial Markets,* 16**,** 637-645.

CHRISTIANSEN, C. 2007. Volatility-Spillover Effects in European Bond Markets. *European Financial Management,* 13**,** 923-948.

CHUNG, H.-L., CHAN, W.-S. & BATTEN, J. A. 2011. Threshold non-linear dynamics between Hang Seng stock index and futures returns. *European Journal of Finance,* 17**,** 471-486.

CONRAD, J. & KAUL, G. 1989. Mean Reversion in Short-Horizon Expected Returns. *Review of Financial Studies,* 2**,** 225-240.

CURME, C., STANLEY, H. E., KENETT, D. Y., TUMMINELLO, M. & MANTEGNA, R. N. 2015. Emergence of statistically validated financial intraday lead-lag relationships. *Quantitative Finance,* 15**,** 1375-1386.

D'AVOLIO, G. 2002. The market for borrowing stock. *Journal of Financial Economics,* 66**,** 271-306.

DAO, T. M., MCGROARTY, F. & URQUHART, A. 2018. Ultra-high-frequency lead–lag relationship and information arrival. *Quantitative Finance,* In press**,** 1-11.

DE BOCK, R. & FILHO, I. D. 2015. The behavior of currencies during risk-off episodes. *Journal of International Money and Finance,* 53**,** 218-234.

DE LONG, J. B., SHLEIFER, A., SUMMERS, L. H. & WALDMANN, R. J. 1990. Noise Trader Risk in Financial Markets. *Journal of Political Economy,* 98**,** 703-738.

DEKKER, A., SEN, K. & YOUNG, M. R. 2001. Equity market linkages in the Asia Pacific region: A comparison of the orthogonalised and generalised VAR approaches. *Global Finance Journal,* 12**,** 1-33.

DEWAN, S. & MIN, C.-K. 1997. The Substitution of Information Technology for Other Factors of Production: A Firm Level Analysis. *Management Science,* 43**,** 1660-1675.

DIEBOLD, F. X. & YILMAZ, K. 2012. Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting,* 28**,** 57-66.

DIJK, D. V., MUNANDAR, H. & HAFNER, C. M. 2011. The euro introduction and noneuro currencies. *Applied Financial Economics,* 21**,** 95-116.

DO, B. & FAFF, R. 2010. Does simple pairs trading still work? *Financial Analysts Journal,* 66**,** 83-95.

DO, B. & FAFF, R. 2012. Are pairs trading profits robust to trading costs? *Journal of Financial Research,* 35**,** 261-287.

DÖPKE, J. & PIERDZIOCH, C. 2006. Politics and the stock market: Evidence from Germany. *European Journal of Political Economy,* 22**,** 925-943.

EASLEY, D., O'HARA, M. & YANG, L. 2016. Differential Access to Price Information in Financial Markets. *Journal of Financial and Quantitative Analysis,* 51**,** 1071-1110.

EASLEY, D., PRADO, L. D., M., M. & O'HARA, M. 2011. The Microstructure of the 'Flash Crash': Flow Toxicity, Liquidity Crashes, and the Probability of Informed Trading. *Journal of Portfolio Management,* 37**,** 118-128.

EDERINGTON, L. H. & LEE, J. H. 1993. How Markets Process Information: News Releases and Volatility. *Journal of Finance,* 48**,** 1161-1191.

ENGELBERG, J., GAO, P. & JAGANNATHAN, R. 2008. An anatomy of pairs trading: the role of idiosyncratic news, common information and liquidity.

ENGLE, R. 2002. Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models. *Journal of Business & Economic Statistics,* 20**,** 339-350.

ENGLE, R. & SARKAR, D. 2006. Premiums-Discounts and Exchange Traded Funds. *Journal of Derivatives,* 13**,** 27-45.

EPPS, T. W. 1979. Comovements in Stock Prices in the Very Short Run. *Journal of the American Statistical Association,* 74**,** 291-298.

ERGÜN, A. T. 2009. NYSE Rule 80A restrictions on index arbitrage and market linkage. *Applied Financial Economics,* 19**,** 1675-1685.

EUN, C. S., KIM, S.-H. & LEE, K. 2015. Currency competition between the dollar and euro: Evidence from exchange rate behaviors. *Finance Research Letters,* 12**,** 100-108.

EVANS, M. D. D. & LYONS, R. K. 2002. Order Flow and Exchange Rate Dynamics. *Journal of Political Economy,* 110**,** 170-180.

FAMA, E. F. & FRENCH, K. R. 1997. Industry costs of equity. *Journal of Financial Economics,* 43**,** 153-193.

FATUM, R. & YAMAMOTO, Y. 2016. Intra-safe haven currency behavior during the global financial crisis. *Journal of International Money and Finance,* 66**,** 49-64.

FLEMING, J., OSTDIEK, B. & WHALEY, R. E. 1996. Trading Costs and the Relative Rates of Price Discovery in Stock, Futures, and Option Markets. *Journal of Futures Markets,* 16**,** 353-387.

FOERSTER, S. R. & SCHMITZ, J. J. 1997. The Transmission of U.S. Election Cycles to International Stock Returns. *Journal of International Business Studies,* 28**,** 1-27.

FRANK, H. & KENNETH, W. S. 2005. Trade Size And Informed Trading: Which Trades Are "Big"? *Journal of Financial Research,* 28**,** 133-163.

FRINO, A., WALTER, T. & WEST, A. 2000. The Lead-Lag Relationship between Equities and Stock Index Futures Markets around Information Releases. *Journal of Futures Markets,* 20**,** 467-487.

FRINO, A. & WEST, A. 1999. The Lead-Lag Relationship Between Stock Indices and Stock Index Futures Contracts: Further Australian Evidence. *Abacus,* 35**,** 333-341.

FROOT, K. A. & DABORA, E. M. 1999. How are stock prices affected by the location of trade? *Journal of Financial Economics,* 53**,** 189-216.

GAGNON, L. & KAROLYI, G. A. 2010. Multi-market Trading and Arbitrage. *Journal of Financial Economics,* 97**,** 53-80.

GATEV, E., GOETZMANN, W. N. & ROUWENHORST, K. G. 2006. Pairs Trading: Performance of a Relative-Value Arbitrage Rule. *Review of Financial Studies,* 19**,** 797-827.

GOLDSTEIN, M. A., KUMAR, P. & GRAVES, F. C. 2014. Computerized and High-Frequency Trading. *Financial Review,* 49**,** 177-202.

GOODHART, C. A. E. & O'HARA, M. 1997. High frequency data in financial markets: Issues and applications. *Journal of Empirical Finance,* 4**,** 73-114.

GREGORIOU, A., IOANNIDIS, C. & SKERRATT, L. 2005. Information Asymmetry and the Bid-Ask Spread: Evidence from the UK. *Journal of Business Finance and Accounting,* 32**,** 1801-1826.

GRISSE, C. & NITSCHKA, T. 2015. On financial risk and the safe haven characteristics of Swiss franc exchange rates. *Journal of Empirical Finance,* 32**,** 153-164.

GROSSMAN, S. J. & STIGLITZ, J. E. 1976. Information and Competitive Price Systems. *American Economic Review,* 66**,** 246-253.

GROSSMAN, S. J. & STIGLITZ, J. E. 1980. On the Impossibility of Informationally Efficient Markets. *American Economic Review,* 70**,** 393-408.

GROSSMANN, A., OZUNA, T. & SIMPSON, M. W. 2007. ADR mispricing: Do costly arbitrage and consumer sentiment explain the price deviation? *Journal of International Financial Markets, Institutions & Money,* 17**,** 361-371.

GRÜNBICHLER, A., LONGSTAFF, F. A. & SCHWARTZ, E. S. 1994. Regular Article: Electronic Screen Trading and the Transmission of Information: An Empirical Examination. *Journal of Financial Intermediation,* 3**,** 166-187.

HAFERKORN, M. 2017. High-frequency trading and its role in fragmented markets. *Journal of Information Technology,* 32**,** 283-296.

HAGSTRÖMER, B. & NORDÉN, L. 2013. The diversity of high-frequency traders. *Journal of Financial Markets,* 16**,** 741-770.

HANOUSEK, J. & PODPIERA, R. 2003. Informed trading and the bid–ask spread: evidence from an emerging market. *Journal of Comparative Economics,* 31**,** 275-296.

HANSEN, P. R. & LUNDE, A. 2005. A forecast comparison of volatility models: Does anything beat a GARCH(1,1)? *Journal of Applied Econometrics,* 20**,** 873-889.

HASBROUCK, J. & SAAR, G. 2013. Low-latency trading. *Journal of Financial Markets,* 16**,** 646-679.

HASELMANN, R. & HERWARTZ, H. 2010. The introduction of the Euro and its effects on portfolio decisions. *Journal of International Money & Finance,* 29**,** 94-110.

HAUPTFLEISCH, M., PUTNIŅŠ, T. J. & LUCEY, B. 2016. Who Sets the Price of Gold? London or New York. *Journal of Futures Markets,* 36**,** 564-586.

HAYASHI, T. & YOSHIDA, N. 2005. On Covariance Estimation of Non-Synchronously Observed Diffusion Processes. *Bernoulli,* 11**,** 359-379.

HENDERSHOTT, T., JONES, C. M. & MENKVELD, A. J. 2011. Does Algorithmic Trading Improve Liquidity? *Journal of Finance,* 66**,** 1-33.

HOFFMANN, M., ROSENBAUM, M. & YOSHIDA, N. 2013. Estimation of the lead-lag parameter from non-synchronous data. *Bernoulli,* 19**,** 426-461.

HOFFMANN, P. 2014. A dynamic limit order market with fast and slow traders. *Journal of Financial Economics,* 113**,** 156-169.

HUDSON, R., MCGROARTY, F. & URQUHART, A. 2017. Sampling frequency and the performance of different types of technical trading rules. *Finance Research Letters,* 22**,** 136-139.

HUTH, N. & ABERGEL, F. 2014. High frequency lead/lag relationships — Empirical facts. *Journal of Empirical Finance,* 26**,** 41-58.

JACOB LEAL, S., NAPOLETANO, M., ROVENTINI, A. & FAGIOLO, G. 2016. Rock around the clock: An agent-based model of low- and high-frequency trading. *Journal of Evolutionary Economics,* 26**,** 49-76.

JACOBS, H. & WEBER, M. 2015. On the determinants of pairs trading profitability. *Journal of Financial Markets,* 23**,** 75-97.

JAIN, P. K. 2005. Financial market design and the equity premium: Electronic versus floor trading. *Journal of Finance,* 60**,** 2955-2985.

JAIN, P. K., JAIN, P. & MCINISH, T. H. 2016. Does high-frequency trading increase systemic risk? *Journal of Financial Markets,* 31**,** 1-24.

JAMES, J. A. & DOUGLAS, M. 2013. Fairness in Financial Markets: The Case of High Frequency Trading. *Journal of Business Ethics,* 112**,** 585-595.

JARNECIC, E. & SNAPE, M. 2014. The Provision of Liquidity by High-Frequency Participants. *Financial Review,* 49**,** 371-394.
JEGADEESH, N. & TITMAN, S. 1995. Overreaction, Delayed Reaction, and Contrarian Profits. *Review of Financial Studies,* 8**,** 973-993.
JENSEN, M. C. 1978. Some anomalous evidence regarding market efficiency. *Journal of Financial Economics,* 6**,** 95-101.
JIN, X. 2016. The impact of 2008 financial crisis on the efficiency and contagion of Asian stock markets: A Hurst exponent approach. *Finance Research Letters,* 17**,** 167-175.
KAWALLER, I. G., KOCH, P. D. & KOCH, T. W. 1987. The Temporal Price Relationship Between S&P 500 Futures and the S&P 500 Index. *Journal of Finance,* 42**,** 1309-1329.
KEARNEY, F., CUMMINS, M. & MURPHY, F. 2014. Outperformance in Exchange-Traded Fund Pricing Deviations: Generalized Control of Data Snooping Bias. *Journal of Financial Markets,* 19**,** 86-109.
KIRILENKO, A. A. & LO, W. A. 2013. Moore's Law versus Murphy's Law: Algorithmic Trading and Its Discontents. *Journal of Economic Perspectives,* 27**,** 51-72.
KONDOR, P. 2009. Risk in Dynamic Arbitrage: The Price Effects of Convergence Trading. *Journal of Finance,* 64**,** 631-655.
KORKEAMÄKI, T. 2011. Interest rate sensitivity of the European stock markets before and after the euro introduction. *Journal of International Financial Markets, Institutions & Money,* 21**,** 811-831.
LAMOUREUX, C. G. & LASTRAPES, W. D. 1990. Heteroskedasticity in Stock Return Data: Volume versus GARCH Effects. *Journal of Finance,* 45**,** 221-229.
LEE, B.-S. & RUI, O. M. 2002. The dynamic relationship between stock returns and trading volume: Domestic and cross-country evidence. *Journal of Banking and Finance,* 26**,** 51-78.
LEGRAND, R. 2014. Euro introduction: Has there been a structural change? Study on 10 European Union countries. *Economic Modelling,* 40**,** 136-151.
LEVICH, R. M. & THOMAS, L. R., III 1993. The Significance of Technical Trading-Rule Profits in the Foreign Exchange Market: A Bootstrap Approach. *Journal of International Money and Finance,* 12**,** 451-474.
LIEN, D., TSE, Y. K. & XIBIN, Z. 2003. Structural change and lead-lag relationship between the Nikkei spot index and futures price: a genetic programming approach. *Quantitative Finance,* 3**,** 136-144.
MADHAVAN, A. & SOFIANOS, G. 1998. An Empirical Analysis of NYSE Specialist Trading. *Journal of Financial Economics,* 48**,** 189-210.
MANAHOV, V., HUDSON, R. & GEBKA, B. 2014. Does high frequency trading affect technical analysis and market efficiency? And if so, how? *Journal of International Financial Markets, Institutions & Money,* 28**,** 131-157.
MANASTER, S. & RENDLEMAN, J. R. J. 1982. Option Prices as Predictors of Equilibrium Stock Prices. *Journal of Finance,* 37**,** 1043-1057.
MARSHALL, B., NGUYEN, N. H. & VISALTANACHOTI, N. 2013. ETF arbitrage: Intraday evidence. *Journal of Banking & Finance,* 37**,** 3486-3498.
MARTIKAINEN, T. & PERTTUNEN, J. 1995. On The Dynamics of Stock Index Futures and Individual Stock Returns. *Journal of Business Finance & Accounting,* 22**,** 87-100.
MCGROARTY, F., AP GWILYM, O. & THOMAS, S. 2006. Microstructure effects, bid–ask spreads and volatility in the spot foreign exchange market pre and post-EMU. *Global Finance Journal,* 17**,** 23-49.
MITCHELL, M., PULVINO, T. & STAFFORD, E. 2002. Limited Arbitrage in Equity Markets. *Journal of Finance,* 57**,** 551-584.
MORI, M. & ZIOBROWSKI, A. J. 2011. Performance of Pairs Trading Strategy in the U.S. REIT Market. *Real Estate Economics,* 39**,** 409-428.

NAJAND, M. & MIN, J. H. 1999. A further investigation of the lead-lag relationship between the spot market and stock index futures: Early evidence from Korea. *Journal of Futures Markets,* 19**,** 217-232.
NAM, S. O., OH, S. & KIM, H. K. 2008. The time difference effect of a measurement unit in the lead–lag relationship analysis of Korean financial market. *International Review of Financial Analysis,* 17**,** 259-273.
NAM, S. O., OH, S., KIM, H. K. & KIM, B. C. 2006. An empirical analysis of the price discovery and the pricing bias in the KOSPI 200 stock index derivatives markets. *International Review of Financial Analysis,* 15**,** 398-414.
NARAYAN, P. K., MISHRA, S., NARAYAN, S. & THURAISAMY, K. 2015. Is Exchange Rate Trading Profitable? *Journal of International Financial Markets, Institutions & Money,* 38**,** 217-229.
O'CONNOR, F. A., LUCEY, B. M., BATTEN, J. A. & BAUR, D. G. 2015. The financial economics of gold — A survey. *International Review of Financial Analysis,* 41**,** 186-205.
O'HARA, M. 2014. High-Frequency Trading and Its Impact on Markets. *Financial Analysts Journal,* 70**,** 18-27.
O'HARA, M., YAO, C. & YE, M. A. O. 2014. What's Not There: Odd Lots and Market Data. *Journal of Finance,* 69**,** 2199-2236.
O'HARA, M. 2015. High frequency market microstructure. *Journal of Financial Economics,* 116**,** 257-270.
PAIN, N. & YOUNG, G. 2004. The macroeconomic impact of UK withdrawal from the EU. *Economic Modelling,* 21**,** 387-408.
PANTON, D. 1976. Chicago Board Call Options As Predictors of Common Stock Price Changes. *Journal of Econometrics,* 4**,** 101-113.
PANTZALIS, C., STANGELAND, D. A. & TURTLE, H. J. 2000. Political Elections and the Resolution of Uncertainty: The International Evidence. *Journal of Banking and Finance,* 24**,** 1575-1604.
PASTOR, L. U. & VERONESI, P. 2012. Uncertainty about Government Policy and Stock Prices. *Journal of Finance,* 67**,** 1219-1264.
PIZZI, M. A., ECONOMOPOULOS, A. J. & O'NEILL, H. M. 1998. An Examination of the Relationship between Stock Index Cash and Futures Markets: A Cointegration Approach. *Journal of Futures Markets,* 18**,** 297-305.
PUI-LAM, L. & WING-KEUNG, W. 2008. On testing the equality of multiple Sharpe ratios, with application on the evaluation of iShares. *Journal of Risk,* 10**,** 15-30.
PULLEN, T., BENSON, K. & FAFF, R. 2014. A Comparative Analysis of the Investment Characteristics of Alternative Gold Assets. *Abacus,* 50**,** 76-92.
RAMIAH, V., PHAM, H. N. A. & MOOSA, I. 2016. The sectoral effects of Brexit on the British economy: early evidence from the reaction of the stock market. *Applied Economics,* 49**,** 2508-2514.
ROBERT, A. J. & PHILIP, P. 2012. A Dysfunctional Role Of High Frequency Trading In Electronic Markets. *International Journal of Theoretical and Applied Finance (IJTAF),* 15**,** 1-15.
RUAN, J. & MA, T. 2012. Ex-Dividend Day Price Behavior of Exchange-Traded Funds. *Journal of Financial Research,* 35**,** 29-53.
RYU, D. 2015. The information content of trades: An analysis of KOSPI 200 index derivatives. *Journal of Futures Markets,* 35**,** 201-221.
SANTA-CLARA, P. & VALKANOV, R. 2003. The Presidential Puzzle: Political Cycles and the Stock Market. *Journal of Finance,* 58**,** 1841-1872.
SCHIERECK, D., KIESEL, F. & KOLARIC, S. 2016. Brexit: (Not) another Lehman moment for banks? *Finance Research Letters,* 19**,** 291-297.
SCHULTZ, P. & SHIVE, S. 2010. Mispricing of dual-class shares: Profit opportunities, arbitrage, and trading. *Journal of Financial Economics,* 98**,** 524-549.

SHARMA, J. L., MOUGOUE, M. & KAMATH, R. 1996. Heteroscedasticity in stock market indicator return data: volume versus GARCH effects. *Applied Financial Economics,* 6**,** 337-342.

SHARPE, W. F. 1994. The Sharpe ratio. *Journal of Portfolio Management,* 21**,** 49-58.

SHIN, S. & SOYDEMIR, G. 2010. Exchange-traded funds, persistence in tracking errors and information dissemination. *Journal of Multinational Financial Management,* 20**,** 214-234.

SHLEIFER, A. & VISHNY, R. W. 1997. The Limits of Arbitrage. *Journal of Finance,* 52**,** 35-55.

SHYY, G. & VIJAYRAGHAVAN, V. 1996. A Further Investigation of the Lead-Lag Relationship between the Cash Market and Stock Index Futures Market with the Use of Bid/Ask Quotes: the Case of France. *Journal of Futures Markets,* 16**,** 405-420.

SIBLEY, M. 2001. The Impact of Presidential Elections on Currency Values in Latin America. *Multinational Business Review,* 9**,** 14-19.

SORTINO, F. 2010. *The Sortino Framework for Constructing Portfolios,* United States of America, Elsevier.

STEPHAN, J. A. & WHALEY, R. E. 1990. Intraday Price Change and Trading Volume Relations in the Stock and Stock Option Markets. *Journal of Finance,* 45**,** 191-220.

STOLL, H. R. & WHALEY, R. E. 1990. The Dynamics of Stock Index and Stock Index Futures Returns. *Journal of Financial & Quantitative Analysis,* 25**,** 441-468.

STRATMANN, T. & WELBORN, J. W. 2012. Exchange-traded funds, fails-to-deliver, and market volatility.

SUN, E., KRUSE, T. & YU, M.-T. 2014. High frequency trading, liquidity, and execution cost. *Annals of Operations Research,* 223**,** 403-432.

TANNER, G. 1997. A Note on Economic News and Intraday Exchange Rates. *Journal of Banking and Finance,* 21**,** 573-585.

TAYLOR, M. P. & ALLEN, H. 1992. The use of technical analysis in the foreign exchange market. *Journal of International Money and Finance,* 11**,** 304-314.

TIELMANN, A. & SCHIERECK, D. 2017. Arising borders and the value of logistic companies: Evidence from the Brexit referendum in Great Britain. *Finance Research Letters,* 20**,** 22-28.

URQUHART, A. 2014. The Euro and European stock market efficiency. *Applied Financial Economics,* 24**,** 1235-1248.

YANG, J., KOLARI, J. W. & MIN, I. 2003. Stock market integration and financial crises: The case of Asia. *Applied Financial Economics,* 13**,** 477-486.

ZHANG, L. 2011. Estimating covariation: Epps effect, microstructure noise. *Journal of Econometrics,* 160**,** 33-47.