UNIVERSITY OF SOUTHAMPTON

Faculty of Business and Law

Southampton Management School

Essays in Statistical Arbitrage

by

Hamad Alsayed

THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

February 2014

For Shams and Mansoor

Blank Page

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF BUSINESS AND LAW
Financial Economics

Doctor of Philosophy

ESSAYS IN STATISTICAL ARBITRAGE

By Hamad Alsayed

This three-paper thesis explores the important relationship between arbitrage and price efficiency. Chapter 3 investigates the risk-bearing capacity of arbitrageurs under varying degrees and types of risk. A novel stochastic process is introduced to the literature that is capable of jointly capturing fundamental risk factors which are absent from extant specifications. Using stochastic optimal control theory, the degree to which arbitrageurs' investment behaviour is affected by aversion to these risks is analytically characterized, as well as conditions under which arbitrageurs cut losses, effectively exacerbating pricing disequilibria. Chapter 4 explores the role of arbitrage in enforcing price parity between cross-listed securities. This work employs an overlooked mechanism by which arbitrage can maintain parity, namely pairs-trading, which is cheaper to implement than the mechanism most commonly employed in the literature on cross-listed securities. This work shows that arbitrage is successful at enforcing parity between cross-listed securities, and also documents the main limits to arbitrage in this market setting. Chapter 5 examines the extent to which arbitrage contributes to the flow of information across markets. It is shown that microscopic lead/lag relationships of the order of a few hundred milliseconds exist across three major international index futures. Importantly, these delays last long enough, and induce pricing anomalies large enough, to compensate arbitrageurs for appropriating pricing disequilibria. These results accord with the view that temporary disequilibria incentivise arbitrageurs to correct pricing anomalies.

Blank Page

Table of Contents

Blank Page

List of Figures

Chapter 3:

Chapter 4:

Chapter 5:

Blank Page

List of Tables

Chapter 3:

(no Tables)

Chapter 4:

Chapter 5:

Declaration of Authorship

I, Hamad Alsayed, declare that the thesis entitled "Essays in Statistical Arbitrage" and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- 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;
- where I have consulted the published work of others, this is always clearly attributed;
- 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;
- I have acknowledged all main sources of help;
- 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;
- Papers 1 and 2 have been published, respectively, as:

Alsayed, Hamad and Frank McGroarty (2013). Optimal Portfolio Selection in Nonlinear Arbitrage Spreads. *European Journal of Finance* 19(3), pp. 206-227.

Alsayed, Hamad and Frank McGroarty (2012). Arbitrage and the Law of One Price in the Market for American Depository Receipts. *Journal of International Financial Markets, Institutions and Money* 22(5), pp. 1258-1276.

Signed:

Date:

Blank Page

Acknowledgements

Thank you to Dr. Frank McGroarty for his invaluable guidance and support throughout my time at the University of Southampton.

Chapter 1

Introduction

1.1 Arbitrage and Statistical Arbitrage

Arbitrage is defined by Shleifer and Vishny (1997) as "the simultaneous purchase and sale of the same, or essentially similar, security in two different markets for advantageously different prices".

The basic concept of arbitrage is simple: In perfectly efficient markets, if two assets command identical cash flows, then the two should be identically priced. Specifically, any discrepancy between the two prices which exceeds the transaction costs associated with executing the trades will immediately be *arbitraged* away. Arbitrage has the effect of enforcing price parity and of bringing prices back to fundamental values. It is therefore extremely important to the field of Financial Economics to understand the role played by arbitrageurs (those who engage in arbitrage activities) in financial markets.

The idea of statistical arbitrage is a generalization of its traditional "riskless" interpretation. Unlike the traditional interpretation, statistical arbitrage identifies mispricings based on deviations from common stochastic trends, and relies on a predictive modelling framework that attempts to exploit consistent regularities in the movements of asset prices. Because the assets involved in a statistical arbitrage trade are often not perfect substitutes, arbitrage of this type will necessarily entail a degree of risk.

In a frequently-cited paper, Gatev et al. (2006) examine the risk and return characteristics of a popular statistical arbitrage strategy with at least a 30-year history on Wall Street, namely pairs trading. The authors define pairs trading as finding "two stocks whose prices have moved together historically. When the spread between them widens, short the winner and buy the loser. If history repeats itself, prices will converge and the arbitrageur will profit". Other examples of statistical arbitrage strategies are volatility arbitrage, convertible arbitrage, and trading based on lead-lag effects.

In this thesis, the wider concept of statistical arbitrage as it is defined above is referred to as either *arbitrage* or *statistical arbitrage*, whereas the traditional riskless interpretation is referred to as *pure arbitrage* or *direct arbitrage*.

1.2 Scope

This thesis answers three main questions surrounding the role of arbitrageurs in enforcing pricing efficiency: First, to what extent does aversion to fundamental risk diminish the extent to which arbitrageurs are willing to correct pricing disequilibria? Second, in the absence of fundamental risk, how effective is arbitrage at exploiting pricing anomalies? Third, to what extent does arbitrage appropriate geographical informational delays between exchanges?

This thesis contributes to the state of the art by focussing on the incentives and disincentives faced by arbitrageurs in a diverse range of market settings. The work presented here adopts a bottom-up perspective from the point of view of arbitrageurs, principally by way of trading exercises in a range of different markets. It is therefore quantitative in nature.

Chapter 3 provides a theoretical contribution to the literature by introducing a novel stochastic model of mispricing time-series aimed at exploring an important aspect of arbitrage trading, namely fundamental risk. Fundamental risk denotes the risk of a breakdown in the pricing relationship between two or more assets which are the subject of a statistical arbitrage trade. This risk is an important consideration for arbitrageurs, but has only been represented qualitatively in the extant literature. The work presented here offers the first quantitative incorporation of this risk into an explicit trading strategy, and in so doing yields new insights into the behaviour of arbitrageurs: First, arbitrageurs who are averse to fundamental risk approach increasingly large mispricings with increased scepticism - a fact which explains why sometimes a number of “obvious” arbitrage trades go underexploited. Second, an arbitrageur who is averse to fundamental risk will unwind losing positions far sooner than an arbitrageur who is only averse to the uncertainty of the lifetime of the trade but who has complete confidence in the fundamental pricing relationship between the assets being traded.

Chapter 4 provides an empirical contribution to the literature by way of a pairs-trading exercise aimed at exploiting deviations from the Law of One Price in the market for American Depository Receipts. By focussing on this market, the analysis presented in chapter 4 explores the capacity of arbitrage to enforce price efficiency in a setting with virtually no fundamental risk - a setting which intuitively should constitute fertile ground for arbitrageurs. Interestingly, it is found that even when arbitrageurs face no fundamental risk, they are averse to the uncertainty surrounding the duration of arbitrage trades. Furthermore, in settings where the main impediments to arbitrage are the costs associated with establishing the arbitrage position, profitable trading opportunities are limited to market participants who can execute trades quickly and cheaply, such as large hedge funds and proprietary trading desks at investment banks. On the whole, the work presented in this chapter finds, in contrast to the prior literature, that arbitrage is highly successful at appropriating pricing disequilibria, and hence is an important component of the price-correction mechanism related to deviations from the Law of One Price.

Chapter 5 contributes to the literature by applying an existing statistical methodology of measuring lead-lag relationships to a new empirical setting namely the futures markets. The methodology applied has its roots in Hayashi and Yoshida (2005), but the precise implementation is based on a recent advance by Hoffman et al (2010). The analysis draws inspiration from a body of literature documenting an increase in geographical correlation in global markets. However, the question of whether informational delays across international markets give rise to arbitrage opportunities is relatively under-explored. This chapter addresses this gap by way of a trading exercise across international index futures. It is found that such temporal disequilibria occur mainly in the sub-second space, and moreover, arbitrage is successful at appropriating these disequilibria across markets. Hence, arbitrage is an important mechanism in the transmission of information across global markets.

This thesis borrows ideas from a variety of academic disciplines, ranging from Optimal Control Theory, Statistics, and Empirical Finance; and presents a synthesis of these fields to facilitate a reflection on the research questions which it tackles.

1.3 Motivation

This thesis is driven by the need for a deeper understanding of the role that arbitrageurs play in financial markets, and how that role relates to pricing efficiency.

The emphasis on arbitrage trading exercises from a profitability point of view is underrepresented in the literature, and yet is crucial to achieving deeper understanding of the role of arbitrage for four main reasons: First, understanding the conditions under which arbitrageurs are incentivised to exploit pricing inefficiencies from a quantitative perspective is fundamental to understanding the extent to which pricing efficiency can be restored by their actions. Second, while it is widely reported in a deep body of literature documenting the “limits to arbitrage" that factors such as risks and costs can deter arbitrageurs from exploiting inefficiencies, the philosophy espoused in this thesis is that the precise nature of how these factors act as limits to arbitrage is best understood by way of a quantitative treatment of the subject. Third, placing emphasis on arbitrageurs trading behaviour necessitates a careful treatment of the trading costs faced in undertaking a real-world exploitation of pricing inefficiencies. Such a careful treatment of trading costs is absent from much of the relevant literature, and the work presented here illustrates the importance of incorporating realistic trading costs to drawing accurate conclusions which link arbitrage with market efficiency. Fourth, undertaking a trading exercise necessitates the utilization of high-quality market data in drawing robust conclusions. The view adopted in this thesis is that advances in computational technology have enhanced the ability of arbitrageurs to explore and exploit regularities in prices. Yet it is relatively uncommon for works written before the early 2000's to employ intraday data at all, or intraday data that is granular enough to illustrate the competitive and short-lived nature of arbitrage opportunities.

1.4 Thesis Overview

The principal aim of this thesis is to provide a deeper understanding of the role that arbitrageurs play in enforcing price efficiency. To this end, each chapter examines the role of arbitrage within a different market setting.

Chapter 3 analytically solves the portfolio optimization problem of an investor faced with a risky arbitrage opportunity (e.g. relative mispricing in equity pairs). This chapter introduces a nonlinear generalization of the commonly-employed Ornstein-Uhlenbeck (OU) process as a way to explicitly incorporate fundamental risk into the mispricing time-series a feature absent from the extant literature. Fundamental risk is the risk that mispricings can exhibit long and persistent departures from economic fundamentals. Incorporating fundamental risk yields several novel insights regarding the investment behaviour of arbitrageurs: First, arbitrageurs who are averse to fundamental risk exhibit a diminishing propensity to exploit increasingly large mispricings. Second, investment behaviour in light of these additional risk factors precipitates the gradual unwinding of losing trades far sooner than is suggested by current models. In practice, the nonlinear generalization yields superior risk-management capabilities relative to OU, when applied empirically to FTSE 100 stock data.

Chapter 4 investigates the mechanism by which the Law of One Price is enforced between UK shares and their American Depository Receipt (ADR) counterparts: the two are virtually the same asset. This work is motivated by observations in the empirical literature that markets for cross-listed securities are less than fully integrated. In theory, environments such as this should constitute fertile ground for arbitrageurs, who are compensated for exploiting these disequilibria. However, the existing literature on stock-ADR arbitrage overwhelmingly concludes in favour of market efficiency, precluding the possibility of arbitrage. The reason for this is that the literature examining stock-ADR arbitrage typically relies on a trading strategy whereby stocks are physically converted into ADRs, and vice versa. This chapter explores the efficacy of pairs trading, a strategy that is under-represented in the literature within the context of ADRs, in enforcing price parity. It is found that pairs trading yields statistically significant returns after adjusting for taxes and costs. This chapter concludes that arbitrage is an important component of restoring stock-ADR parity, contrasting previous assertions that trading in either the US or domestic market alone results in efficient ADR pricing.

Chapter 5 employs a recent advance in the statistical estimation of correlation and uncovers the existence of sub-second lead/lag relationships across international index futures. The methodology is based on work by Hayashi and Yoshida

(2005) and Hoffman et al (2010), and is chosen for its ability to handle the nonsynchronous nature of price observations which is typical of sub-second data. It is found that lead-lag effects across futures contracts last long enough, and induce pricing anomalies large enough, to compensate arbitrageurs for exploiting pricing disequilibria. This chapter concludes that arbitrage is an important mechanism in the transmission of information across international markets.

1.5 Relevance of this Thesis

This thesis focusses on the explicit incentives and deterrents to arbitrage directly from the point of view of arbitrageurs. In so doing, it offers a fuller understanding of the relation between arbitrage and price efficiency. This work may be useful to researchers who wish to quantitatively explore arbitrage within other contexts - for example the role of arbitrage trading in market contagion (an example of which is the collapse of hedge fund Long Term Capital Management in 1998) or agent-based models of arbitrageurs with heterogeneous characteristics. It would be important in the latter context to incorporate one of the central insights of this thesis, which is that arbitrageurs' behaviour is substantially affected not only by the extent, but also by the *types* of risk to which they are averse.

Further, this work highlights the important role arbitrage plays within market settings in which it was previously concluded that arbitrage played little part in the parity enforcement mechanism. Arbitrage is successful at eliminating cross-border deviations from the Law of One Price and at appropriating informational delays across geographies. These observations are relevant to the regulatory debate surrounding the role of arbitrageurs in financial markets.

Finally, because the concepts presented in this thesis are illustrated via the development of explicit trading rules, these parts of the thesis may be useful to market practitioners who engage in arbitrage trading. In particular, the ideas presented could be useful to the development of novel quantitative trading strategies and to the choice of markets on which they are deployed.

Chapter 2

A Review of the Relevant Literature

2.1 - Introduction

The overall aim of this chapter is to summarize our current understanding of the important relationship between arbitrage and pricing efficiency, with the intention of exposing gaps in this understanding. This chapter is split into four subsequent sections, which combine together to achieve the overall aim. Section 2.2 presents a history of the debate between arbitrage and market efficiency, then presents examples from empirical works which document the existence of deviations from price parity in financial markets. Section 2.3 complements this discussion by summarizing the body of literature that documents various impediments that act to preclude arbitrageurs from enforcing price efficiency the so-called "limits to arbitrage". Section 2.4 reviews the theoretical literature on the quantification of arbitrage risks and the establishment of arbitrageurs optimal trading rules from a quantitative perspective. Section 2.5 summarizes the emergent gaps in the literature that this thesis addresses. In all, this chapter provides the framework upon which the contributions of this thesis are based.

2.2 - On the Existence of Pricing Inefficiencies

Friedman (1953) was the first to argue against the plausible existence of prolonged mispricings away from fundamental values. Such deviations from price parity amount to free money left on the table - a reward to those who subsequently bring prices back into line. Though the origins of the relationship between price and fundamental value are evident in writings as early as Hume (1741), it was

Friedman (1953) who characterized the basic mechanism by which arbitrage acts to correct prices within the context of capital asset markets. The author splits the universe of investors into two groups: irrational traders who trade based on measures other than fundamental values (e.g. mistaken beliefs about future utility, chartists); and rational traders who act to exploit differences between prices and fundamental values. In this setting, irrational traders cannot have a substantial effect on prices for two main reasons: First, any mispricing caused by the actions of irrational traders will be immediately exploited by rational arbitrageurs who profit from appropriating this mispricing. Second, the failure of rational agents to correct mispricings is nevertheless tantamount to market participants "buying high and selling low", an environment in which irrational traders would eventually die out.

Closely related to the above concept is the assertion that past prices cannot contain information about future prices, an idea termed "The Random-Walk Hypothesis", which was first modelled by French mathematician Louis Bachelier (Bachelier, 1900) in his PhD Thesis "The Theory of Speculation". In a seminal paper that has become the cornerstone for proponents of this theory, Fama (1965) carries out the first examination of the validity of this model, and finds strong empirical support for the Random-Walk Hypothesis. This result is the basis of what is known as the Efficient Markets Hypothesis; and asserts that one cannot consistently achieve returns in excess of average market returns on a risk-adjusted basis (except through luck), given the information available at the time the investment is made. In other words, the Efficient Markets Hypothesis states that the market is "informationally efficient", such that prices properly reflect all available information and are equal to fundamental values. Further work that supports the implications of the Efficient Markets Hypothesis is Fama (1991), Fama and French (1992), Fama (1995).

The Efficient Markets Hypothesis may appear to be inconsistent with the existence of arbitrage opportunities asserted by Friedman (1953), but the two stipulations are reconcilable in an environment where arbitrageurs react instantaneously to any mispricing caused by irrational investors.

Despite the intuitive nature of these early market models, there is a paradox underlying their assertions. Grossman and Stiglitz (1976, 1980) take exception to validity of the Efficient Markets Hypothesis, and base their counter-arguments on the notion that arbitrage is costly since resources must be expended on the collection of (private) information. If prices fully reflect all available information, then all arbitrage profits are necessarily eliminated. Therefore, there is no incentive for arbitrageurs (who the authors equate to "rational investors") to remain in the market. This creates a paradox whereby on the one hand, arbitrageurs profit from instantaneously exploiting disequilibria caused by irrational traders, but on the other hand, arbitrageurs make no profits. Indeed, if the payoff to the costly activity of collecting information is less than the value of the costs, then there is no incentive to trade. Grossman and Stiglitz (1976, 1980) conclude that the only solution to this paradox is for financial markets to exhibit "an equilibrium level of disequilibrium" that incentivises arbitrageurs to restore price parity.

There is a wide body of empirical literature which supports the view that prices can significantly diverge from fundamental values. The remainder of this section presents several real-world examples of this phenomenon in a range of market settings relevant to this thesis. In the subsequent section, a discussion is presented as to why such mispricings exist, and how they might reconcile with the Efficient Markets Hypothesis.

2.2.1 Deviations from the Law of One Price

Fundamental to the notion of efficient markets is the concept of the Law of One Price, which conveys the intuitive idea that identical goods (in separate markets) should command identical prices. Arbitrage is a key driver of the Law of One Price (De Long et al (1990), Gromb and Vayanos (2010)).

In a hypothetical example wherein (identical) coffee beans are sold on two different continents, any price differential between the two markets which exceeds the cost of physically transporting the coffee represents price inefficiency. In this situation, arbitrageurs are financially incentivised to purchase the coffee in the cheaper location, transport it to, and sell onto the dearer market. This arbitrage trade has the effect of creating excess demand in the cheaper market, and of creating excess supply in the dearer market, which in turn forces prices in each of the two markets to converge. But despite the logical intuitiveness of the above theoretical example, there are instances in real financial markets where the Law of One Price is violated for extended periods of time.

Lamont and Thaler (2003b) present a case-study of the Palm-3Com equity carve-out (spin off). In March 2000, 3Com sold 5% of its stake in Palm through an IPO, and further announced that its remaining stake would be spun off to 3Com shareholders by the end of the year 2000. 3Com shareholders would receive 1.525 Palm shares for each share of 3Com they owned.

From the perspective of the Law of One Price, the price of 3Com shares should have had a minimum bound of 1.525 times the price of Palm. This is because each 3Com share was equivalent to 1.525 Palm shares plus a claim on 3Com's non-Palm future cash flows. Because of limited liability, this cash flow claim is non-negative. However, on the date of the IPO, Palm closed at $95.06 per share, implying a minimum price of $144.97 per 3Com share. However, 3Com closed at $81.81, having dropped from $104.13 on the previous day. In all, 3Com went from a 28% discount to a 41% discount in a clear violation of the Law of One Price.

Rosenthal and Young (1990), Froot and Dabora (1999), and De Jong et al. (2009) show that significant mispricings can exist in dual-listed companies (DLCs). DLCs are companies which issue shares on two or more exchanges in separate countries. These shares have either identical cashflows, or arrangements that set out clearly the cashflow terms that investors in each market are entitled to. Either way, a clear pricing relationship exists between the different classes of shares for these companies. Examples of such companies are Royal Dutch Shell, Unilever, and BHP Billiton.

Both Rosenthal and Young (1990) and Froot and Dabora (1999) find that significant mispricings in three DLCs (Royal Dutch Shell, Unilever, and Smithkline Beecham) have existed over a long period of time. Both studies conclude that factors such as currency risk, legal structures, and taxation) do not adequately explain the magnitude of the price deviations. Froot and Dabora (1999) show that the relative prices of the twin stocks are correlated with the stock indices of the markets on which each of the twins has its main listing. For example, if the UK market rises relative to Amsterdam, then the stock price of Reed International PLC would tend to rise relative to Elsevier NV (the two names refer to the same company). De Jong et al. (2009) study a sample of 12 dual listed companies. The authors find that, after adjusting for costs, margin requirements, and taxes, returns to a simple arbitrage strategy yields 10% per annum.

Shleifer (1986) and Chen (2004) investigate the effect on stock prices on the stocks inclusion or exclusion from an index. Chen (2004) studies the period 1989-2000, and finds that on average, a stocks price would rise by 5.5% on the day of the announcement that it would be added to the S&P index. Furthermore, the price of that same stock would rise a further 3.5% from the date of the announcement until the day it is actually added. On the other hand, a stocks price would fall by an average of 14% from the day it is announced that it would be removed from an index, to the actual date of removal. It is clear that these effects are economically significant.

A similar concept to a stock index is a closed-end fund. One type of closed-end fund is a country fund. Country funds are exchange-traded closed-end investment companies that hold portfolios of equities of particular foreign countries. They do not attempt to time the market or achieve returns superior to the market for their investors. Rather, they allow investors the opportunity to gain exposure to the countries that the funds represent, while offering investors the convenience of trading in their local market. Most country funds are denominated in dollars, and trade on the New York Stock Exchange (NYSE). Klibanoff et al. (1998) highlight the growing popularity of this kind of investment vehicle. The authors state that in 1984, only four US-listed country funds existed, whereas by 1995, this number had grown to 54, each specializing in one of 30 countries.

Like most closed-end funds that trade on an exchange, country funds do exhibit premiums or discounts relative to the equities that constitute the fund portfolio. Managers of closed-end funds maintain pricing parity with the underlying basket by constantly tuning the supply of "units" of the fund in the market against the demand shocks generated in its constituent equities. For example, if the fund is trading at a premium, the fund manager issues and sells additional shares in the fund, and vice versa. Nevertheless, discounts and premia on closed-end funds contradict the notion of market efficiency. Klibanoff et al. (1998) study a sample of 39 single-country funds and find evidence of significant mispricings. For example, in 1990, the premium on the German country fund traded in New York reached in excess of 100 percent. The authors make a further finding: during periods where unexpected news permeates through markets, the prices of country funds react much more quickly and fully in incorporating the news informational content; whereas during periods of relative calm, country funds exhibit greater (temporary) mispricings. There are further examples of works which document premia and discounts in closed-end funds. Hardouvelis et al. (1994) also find evidence of mispricings in country funds, while Malkiel (1977) and Lee et al. (1991) document evidence of mispricings in the general class of closed-end funds beyond just country funds.

Futures contracts are derivative instruments which speculate on the future price of a particular stock index or commodity. Figlewski (1984), Stoll and Whaley (1990), and Brennan and Schwartz (1990) are the earliest works which study the pricing efficiency of futures contracts relative to their underlying index. Brennan and Schwartz (1990) study a sample of 16 index futures contracts over a period of four years during 1983-1987, and document evidence of violations from the spot-futures relationship. The spot futures relationship is a pricing relationship between a futures contract, and the index or commodity on which it is based. In Cummings and Frino (2011), the authors examine mispricings between the Australian Stock Index (ASX) and its associated futures contract. The authors find, in contrast to the earlier studies, that futures contracts are efficiently priced relative to the ASX. Brooks et al. (2001) study the pricing relationship between the futures contract and index in the UK market, and find evidence that the futures contract leads the stock market index. However, their results accord with Cummings and Frino (2011), in the sense that this lead/lag relationship is not exploitable, and does not present a violation of efficient pricing.

2.2.2 Depository Receipts

Introduced by the investment bank JP Morgan in 1927 as a way for US investors to diversify their holdings internationally, American Depository Receipts (ADRs) are a dollar-denominated representation of ownership in a non-US company. They trade as conventional shares on US exchanges, and provide identical cashflows as their corresponding underlying shares.

From the point of view of arbitrage, ADRs represent a very interesting market. Because ADRs and their underlying stocks are virtually the same instrument, one would expect that arbitrageurs would be very active in this market. Furthermore, ADRs are two-way convertible, or "fungible", into their underlying stocks. This eliminates a significant portion of risk from an arbitrage trade.

Several works exist which document the returns to arbitrage in this setting. Suarez (2005) analyses 10 French stocks that also trade as American Depository Receipts (ADRs) in New York. The author employs a direct arbitrage technique, wherein upon observing a mispricing (let us suppose the stocks are underpriced), the arbitrageur will buy the stocks, which are converted to ADRs through the depository custodian bank. The ADRs are then sold onto the US market to realize the profit stemming from the initial price differential. The largest banks that offer DR conversion services are JP Morgan and Bank of New York Mellon. Suarez (2005) finds that in a period of 11 months, an arbitrageur can extract USD 70,000 from the given sample of stock-ADR pairs. The author compares this profit to the opportunity cost of hiring a financial expert to monitor this market.

The above result shows that although arbitrage opportunities do exist in this market, they yield a relatively modest level of profit. Moreover, Suarez (2005) is a rare example in the literature that documents any profitability from stock-ADR arbitrage. The seminal work of Maldonado and Saunders (1983) and Rosenthal (1983) find that costs are too prohibitive for arbitrage to be a worthwhile pursuit in this market. Miller and Morey (1996) analyse 3 months of high-frequency intraday data for Glaxo-Wellcome, a UK company which is now known as GlaxoSmithKline. The authors find that the mispricing time-series between the Glaxo-Wellcome stock and its ADR lies within the transaction cost bound of executing a successful arbitrage trade at all times throughout their sample. The authors conclude that the market for cross-listed stocks is efficient.

The above results are surprising: If this market in theory provides such fertile ground for arbitrageurs, why is it that arbitrage is unprofitable? One plausible explanation is that both domestic and US markets on which the ADRs trade are fully integrated. However, evidence from the works of Werner and Kleidon (1996) and Chen et al. (2009) suggest otherwise. By observing distinct U-shaped volatility patterns in the prices of UK stocks and their ADRs, these works conclude that investors in each of the two markets are segmented, which implied that investors do not base their trading decisions solely on the cheaper location to trade.

The works of Hong and Susmel (2003) and Broumandi and Reuber (2012) apply a pairs-trading strategy to the market for ADRs. Hong and Susmel (2003) find significant profitability in the market for Asian ADRs. However, the authors do not consider the transaction costs associated with executing the trades. Furthermore, the authors use daily data, whereas the Asian and US stock markets do not share any overlap. Gagnon and Karloyi (2010) point out that illusory mispricings in this setting are inevitable, due to the inevitable variation in price between the close of each market.

2.2.3 Lead/Lag Relationships

Lead/Lag relationships are induced when the prices of highly correlated securities do not move in perfect contemporaneity. The seminal research of Zeckhauser and Niederhoffer (1983) was the first to document the possibility of lead/lag relationships between stock indices and futures contracts. Kawaller (1987) provides evidence that US futures prices tend to lead stock indices with a lag time of around 45 mins. The work of Herbst et al. (1987) and Stoll and Whaley (1990) documents that futures prices lead stock indices by around 8 mins.

The notion that futures prices lead stock indices is intuitive: Futures contracts tend to be highly liquid, much more so than any of the constituent stocks that comprise the index on which the futures contract is based. Information arriving into the market can reasonably be expected to affect the futures price first. As Brooks et al. (2001) point out, stock indices can only reflect information when every single constituent of the index observes a price change. Interestingly, the authors in Brooks et al. (2001) show a sharp decrease in the duration of the lag in the spot/futures relationship. This is likely due to the advent of algorithmic trading and technological infrastructure since the seminal works were written.

Although one would expect that the lead/lag relationship between futures and indices would induce arbitrageurs to be active in this setting, there is little evidence in the extant literature of any profitable disequilibria. Brooks et al. (2001) achieve a 65% forecasting accuracy of futures contracts over stock indices in the UK market, but find no evidence of exploitable temporal mispricings once transaction costs are taken into account. Wahab and Lashgari (1993), Abhyankar (1998), Brooks and Garrett (2002) analyse the UK market; Tse (1995) studies the Japanese spot-futures lead/lag relation; Andreou and Pierides (2008) and Kenourgios (2004) analyse the Greek market. Further, works such as Eun and Shim (1989), Hamao et al. (1990), Antoniou et al. (2003), and Innocenti et al. (2011) extend the lead/lag study to indices across different geographies. Each of these works concludes that the temporal disequilibria that almost universally exists in these studies admit no arbitrage opportunities.

2.3 - The Limits of Arbitrage

Section 2.2 presented examples from several market settings of prices diverging from fundamental values, as well as examples of works which have documented these phenomena. This section complements these empirical observations by documenting the various impediments that act to preclude arbitrageurs from enforcing pricing efficiency, the so-called "limits of arbitrage".

The concept of limited arbitrage is relatively new to the literature on pricing efficiency, but is a powerful reconciliatory tool between the resolute belief in efficient prices on the one hand, and the view that irrational investment creates exploitable mispricing opportunities on the other. Black (1986) and De Long et al. (1990) were the first to consider the implications on the relationship between prices and fundamental values of arbitrageurs *competing* against irrational investors - each inflicting negative externalities on the other through the effect of their trading behaviour on prices. Even in a simple setting where the fundamental value of an asset is known *a priori*, disagreement between arbitrageurs and irrational investors means that arbitrageurs are unable to bring prices back to fundamental values.

With regards to the work presented in this thesis, the factors which act as limits to arbitrage can be split into three broad categories: horizon risk, divergence risk, and fundamental risk. This section presents a description of each of these types of risk.

2.3.1 Horizon Risk

Horizon risk is the risk that a mispricing can take so long to correct itself, that it is not worth exploiting from an arbitrageurs point of view. The risk to the arbitrageur in this situation stems from considerations toward reporting frequencies. Firms such as hedge funds must report their results to investors at least annually, and if the arbitrageur has not yet realised the return from an arbitrage position, investors might view this as indicative of poor skill. Shleifer and Vishny (1997) and Jurek and Yang (2007) point out that fund flows chase performance. Jurek and Yang (2007) provide the first analytical characterization of arbitrageurs behaviour when arbitrageurs are mindful of this fact. In their work, the authors show that arbitrageurs restrict their investment activity in mispriced assets which they to contain horizon risk.

Black (1986) is the first work to assert that irrational investors (Friedman (1956)) contribute to arbitrageurs horizon risk. Simply stated, irrational investors can keep the fundamental relationship between mispricing securities from restoring itself to parity, due to their erroneous beliefs about the individual securities future utility. Abreu and Brunnermeier (2002) suggest that delayed synchronization in the actions of arbitrageurs exploiting the same mispricing opportunity may cause this mispricing to persist.

Gatev et al. (2006) study the returns to arbitrage of exploiting temporary mispricings across correlated pairs of stocks, while De Jong et al. (2009) perform the same exercise for DLCs. In each of these studies, horizon risk is an important impediment to arbitrage. In De Jong et al. (2009), some trading positions are left open for 9 years before observing a price convergence sufficient to cover the cost of executing the trades.

2.3.2 Divergence Risk

Divergence risk, or “noise trader risk” is the risk that after an arbitrage position is established, the actions of irrational investors can exacerbate the mispricing between the securities which are the subject of the arbitrage trade. De Long et al. (1990) and Shleifer and Vishny (1997) were the first to emphasize that this risk acts as an important impediment to arbitrage. Brunnermeier and Pedersen (2005) highlight the fact that besides irrational investors, predatory trading can contribute to noise trader risk in financial markets. To see this, consider a situation in which an arbitrageur reveals their trading position. In this case, predatory trading involves trading to deliberately drive the mispricing further, creating a capital loss for the arbitrageur. The arbitrageur is then forced to liquidate the position at a highly unfavourable point in time (for the arbitrageur), while benefitting the predatory traders.

Edwards (1999) provides an excellent description of the fate of Long Term Capital Management, a hedge fund which in 1998, which suffered from predatory trading as described by Brunnermeier and Pedersen (2005). After suffering significant losses in DLCs and fixed-income arbitrage, John Meriweather, the Fund Manager, revealed his trading positions in an open letter to their investors, hoping that other arbitrageurs would see the financial opportunities of these mispricings.

Lee et al. (1991) find that noise trader risk, as explored in De Long et al. (1990) accords with stylized facts about closed-end fund premia/discounts. Gatev et al. (2006) find that noise trader risk is a significant source of volatility in arbitrage returns across correlated pairs of stocks. However, the same is not true for the study in De Jong et al. (2009), Hong and Susmel (2003), and Suarez (2005). Noise trader risk in these settings is not prevalent, since the securities being traded are fundamentally the same. In the case of ADRs, the two-way convertibility between stocks and ADRs acts as a natural risk-management tool against divergence risk.

2.3.3 Fundamental Risk

Fundamental risk denotes the risk of a deterioration in the fundamental relationship between two or more securities which are the subject of a statistical arbitrage trade. For example, consider two financial institutions X and Y that operate the same markets and are domiciled in the same country. A price deviation from historical norms between X and Y would induce arbitrageurs to short the expensive stock while buying the cheap one. Suppose now that only one of these institutions is found to be conducting fraudulent activities. This news event would have the effect of breaking down the fundamental pricing relationship which has up to then existed between X and Y.

The prevalence of fundamental risk is well-documented in the literature. Shleifer and Vishny (1997), and Klibanoff et al. (1998) characterize this risk as unique to noise trader risk, as the latter implies only a temporary divergence from fundamental value.

While the way in which arbitrageurs trading behaviour is affected by aversion to divergence and horizon risk is clearly documented, what remains unclear is the precise nature of, and to what extent, aversion to fundamental risk dissuades arbitrageurs from enforcing price parity. If arbitrageurs were to exploit a mispricing opportunity based on some predictive model that spots a deviation between two securities from common stochastic trends, then the existing theoretical framework necessarily entails complete confidence in the specification of the model. In other words, the arbitrageur will, with probability 1, realize the profit of the investment in the long term. Lamont and Thaler (2003a) characterize the effect of fundamental risk by saying "An arbitrageur who shorts technology companies and buys oil companies runs the risk that peace breaks out in the Middle East, causing the price of oil to plummet". Note that in this case, it is irrelevant to the arbitrageur's profitability whether the original judgement that oil appeared cheap was correct. The authors continue, "in contrast, if A and B have identical cash flows but different prices, the arbitrageur eliminates fundamental risk." To that end, the market studied in Gatev et al. (2006) carries inherent fundamental risk, whereas the DLC and ADR markets studied in De Jong et al. (2009) and Suarez (2005), respectively, do not.

The intuitive effect of fundamental risk would be to dissuade arbitrageurs from investing fully in exploitable opportunities, a point that runs parallel to considerations of leverage and margin constraints (Liu and Longstaff (2004), Kondor (2009)). In the absence of fundamental risk, or at least the aversion to it, arbitrageurs are most aggressive when prices are furthest away from fundamentals. This point relates to Friedman's (1953) famous assertion that "to say that arbitrage is destabilizing is equivalent to saying that arbitrageurs lose money on average," which is implausible.

2.4 - Theoretical Models of Arbitrage

This section presents a summary of the theoretical models of arbitrage employed in the extant literature. For each model, a mathematical formulation is presented, as well as a discussion about how the model captures the various impediments to arbitrage discussed in section 2.3.

2.4.1 The Ornstein-Uhlenbeck Process

The Ornstein-Uhlenbeck (OU) process has widespread applications in academic Finance other than the specific context of arbitrage. It has been applied to mean-reverting returns in stock portfolios (Merton (1971), Wachter (2002)) and interest rate modelling (Vasicek (1977)).

In Xiong (2001), Kargin (2003), Boguslavsky and Boguslavskaya (2004) and Jurek and Yang (2007), the authors make use of the OU process for modelling a time-series of mean-reverting mispricings, akin to the spread between equity pairs (Gatev et al. (2006)).

The OU process is simple in nature, and can be characterized as follows: Given a finite horizon of time $t \in [0, T]$, a mispricing process X_t follows the process:

$$dX_t = -\mu X_t dt + \sigma dW_t, \tag{1}$$

where μ is a parameter indicating the strength of mean-reversion, σ indicates the volatility, and dW_t indicates innovations of a Geometric Brownian Motion (GBM).

The appeal of using this specification, as the authors show, is in that it admits closed-form solutions to the arbitrageur's portfolio optimization problem. Furthermore, because the dynamics in equation (1) do not pre-specify a period at which the process X_t would necessarily return to zero, this model captures the notion of horizon risk. Divergence risk is also captures because the magnitude of X_t is unbounded. However, because the reversion term is linear in X_t, this model does not capture the notion of fundamental risk. Specifically, the OU model exhibits a constant half-life in the mispricing, which in turn implies that arbitrageurs who are faithful to this model have complete confidence in an eventual price convergence.

2.4.2 The Brownian Bridge Process

The Brownian Bridge (BB) process is utilized by Liu and Longstaff (2004) to model deviations from put-call option parity relationships. Using the same parametrization as section 2.4.1, the BB process is specified as:

$$dX_t = \frac{-\mu X_t}{T-t} + \sigma dW_t. \tag{2}$$

Like OU, this model captures the notion of divergence risk; but nlike OU, this model specifies a time in the future wherein the process X_t is guaranteed to return to its mean. Therefore, the BB process does not capture horizon risk.

2.5 The Gaps in the Extant Literature

This section looks back on the rest of this chapter and outlines the gaps present in the extant literature that this thesis aims to address.

2.5.1 Gaps in the Theoretical Literature

The work of Xiong (2001, Kargin (2003), Liu and Longstaff (2004), Boguslavsky and Boguslavskaya (2004), and Jurek and Yang (2007) presents a gap in the theoretical literature which this thesis aims to address. Specifically, while the models presented in these papers adequately capture the notions of divergence risk and horizon risk, they do not capture the important notion of fundamental risk. Fundamental risk is an important component of the propensity of arbitrage to correct pricing disequilibria, and denotes the risk of deterioration in the fundamental relationship between prices of stocks which are the subject of an arbitrage trade. While there are many empirical works documenting the prevalence of this risk as a deterrent to arbitrage (Shleifer and Vishny (1997), Brennan and Schwartz (1990) among many others presented in this chapter), this risk has not yet been quantitatively captured into an arbitrage trading model. Doing so can potentially yield significant insights into the important relationship between arbitrage and price efficiency, by examining the extent to which arbitrageurs behaviour is affected by aversion to fundamental risk. Clearly, the intuitive effect of this aversion would be for arbitrageurs to trade more sceptically, but the exact mechanics of this are not fully understood.

2.5.2 Gaps in the Empirical Literature

In the works of Rosenthal (1983), Maldonado and Saunders (1983), Kato et al. (1991), Wahab and Lashgari (1992), Park and Tavakkol (1994), Miller and Morey (1996), Eun and Sabherwal (2003), Suarez (2005), and Gagnon and Karloyi (2010), it is found that the pricing relationship in the market for cross-listed securities (stocks versus their ADRs) exhibits deviations from parity. This is particularly striking, given that stocks and ADRs are virtually the same security. In theory, environments such as this should constitute fertile ground for arbitrageurs, who are compensated for appropriating these pricing disequilibria. However, the same body of literature suggests that arbitrage is unprofitable in this market setting, primarily due to the costs associated with executing the arbitrage trades. Furthermore, Werner and Kleidon (1996) and Chen et al. (2009) show that investors in both the domestic and US markets are segmented, which implies that investors in each of the two separate markets do not base their investment decisions solely on the cheaper location to trade. This in turn implies that an efficient pricing relationship between stocks and ADRs cannot be maintained by investors in each of the two markets alone. The gap in the literature stems from the fact that the extant literature has not considered that pairs trading is an alternative mechanism by which arbitrage can enforce price parity in this setting. Though the works of Hong and Susmel (2003) and Broumandi and Reuber (2012) have applied a pairs trading strategy in the context of ADRs, these works use daily data, which Suarez (2005) argues cannot yield significant insight into the prevalence of arbitrage in maintaining the stock-ADR pricing relationship, since opportunities in this market are fast and fleeting. A study with realistic transaction costs and employing high-frequency intraday data has not yet been undertaken.

While the seminal research of Zeckerhauser and Niederhoffer (1983), along with the works of Kawaller et al. (1987), Stoll and Whaley (1990), Wang and Wang (2001) find that futures prices can lead stock prices, the same works, along with Brooks et al. (2001) find that this relationship cannot be exploited by arbitrageurs, primarily due to the costs associated with trading all of the constituents that comprise the portfolio upon which the futures contract is based. A parallel strand of research in Eun and Shin (1989), Abhyankar et al. (1997), Antoniou et al. (2003), and Innocenti et al. (2011) document the existence of lead/lag relationships across international markets. This is inspired by works such as Ayuso and Bianco (2001) and Kearney and Lucey (2004) which document evidence of increasing global market correlations in recent years. A natural question to ask is whether the increase in global market correlations give rise to stable lead/lag patterns across market indices in different countries, and whether these lead/lag patterns can be exploited by arbitrageurs. An investigation into the profitability of arbitrage using futures contracts has not yet been done. Examining this concept with futures contracts circumvents the costs associated with trading all the constituents of a stock index.

Blank Page

Chapter 3

Optimal Portfolio Selection in Nonlinear Arbitrage Spreads

Abstract

This chapter analytically solves the portfolio optimization problem of an investor faced with a risky arbitrage opportunity (e.g. relative mispricing in equity pairs). Unlike the extant literature, which typically models mispricings through the Ornstein-Uhlenbeck (OU) process, this work introduces a nonlinear generalization of OU which jointly captures several important risk factors inherent in arbitrage trading. While these factors are absent from the standard OU, it is shown that considering them yields several new insights into the behaviour of rational arbitrageurs: Firstly, arbitrageurs recognizing these risk factors exhibit a diminishing propensity to exploit large mispricings. Secondly, optimal investment behaviour in light of these risk factors precipitates the gradual unwinding of losing trades far sooner than is entailed in existing approaches including OU. Finally, an empirical application to daily FTSE100 pairs data shows that incorporating these risks renders the presented model's risk-management capabilities superior to both OU and a simple threshold strategy popular in the literature. These observations are useful in understanding the role of arbitrageurs in enforcing price efficiency.

Keywords: Pairs trading, Hamilton-Jacobi-Bellman equation, Statistical Arbitrage, Stochastic optimal control, Stability bounds.

JEL classification: G11, G12.

3.1 - Introduction

In the pricing of related securities (for example shares with either identical or similar characteristics), market efficiency is enforced by the presence of rational arbitrageurs (Shleifer and Vishny (1997), and Mitchell et al. (2002)). Friedman (1953) argues that if the actions of noise traders cause such a price relation to be violated, arbitrageurs seek to profit from this by betting on the elimination of this violation. By demanding the cheap security and supplying the expensive security, the actions of arbitrageurs exert forces on the mispricing to revert back to its mean or *natural level*,[1] which in turn generates a profit for the arbitrageurs and enforces price efficiency. However, unlike "textbook arbitrage" opportunities (Bjork (2004)) which lock in a riskless profit and require no capital commitments, these types of investments present risks to those who seek to engage in them. To see this, note that if arbitrageurs position themselves to benefit from the convergence of a mispricing, it is possible for this mispricing to diverge further before converging, or not converge at all, resulting in substantial losses for arbitrageurs. In 1998, the near-collapse of the hedge fund Long Term Capital Management is frequently cited as an example of this phenomenon (Kondor (2009), Edwards (1999), and MacKenzie (2003)).

Several works have sought to examine the role of arbitrageurs in markets by modelling arbitrage opportunities and deriving corresponding portfolio strategies in light of these inherent risks. Theoretical works favor the Ornstein-Uhlenbeck[2] process (henceforth "OU") or the Brownian Bridge process as candidates to

[1] If a mispricing is between two identical shares quoted on different markets, then economic arguments suggest a natural level of zero (Froot and Dabora (1999)). A mispricing between any two cointegrated assets (Engle and Granger (1987)), can have its natural level take any *constant* value.

[2] It is well-known that the Ornstein-Uhlenbeck process has widespread applications in finance, e.g. Wachter's (2002) and Merton's (1971) application to mean-reverting returns in stock portfolios, applications to stochastic interest rates (Korn and Kraft (2002), Vasicek (1977)), and applications to arbitrage opportunities in Xiong (2001), Boguslavsky and Boguslavskaya (2004), Jurek and Yang (2007), and Lv and Meister (2009).

model the evolution of the mispricing (e.g. Jurek and Yang (2007), Xiong (2001), and Liu and Longstaff (2004)), while empirical works examine profitability in statistical arbitrage trading based on simple threshold trading rules (e.g. Gatev et al. (2006)). A major commonality in the existing works lies in the assumptions regarding the correction of mispricings: the larger the mispricings are, the more dramatically they tend to be corrected. This assumption has clear implications regarding the behaviour of arbitrageurs toward mispricings: the greater the magnitude, the greater the investible opportunity. But this notion seems inconsistent with economic arguments that mispricings can exhibit large and long-lasting departures from fundamentals (e.g. Abreu and Brunnermeier (2003), Brunnermeier and Pedersen (2005), De Long et al. (1990)), or that the fundamentals themselves can change (Lamont and Thaler (2003b)). While these risk factors are largely ignored by existing models, considering them is essential to fully examining the role of arbitrageurs in enforcing price efficiency. To what extent do these additional risk factors influence the behaviour of arbitrageurs towards mispricings? How do they affect what arbitrageurs interpret as attractive investible opportunities?

This chapter develops a theoretical model to provide answers to these questions. Its main contribution lies in the introduction of a novel stochastic process, whose dynamics are aimed specifically at capturing the fundamental risk factors ignored by existing models of arbitrage. This process is a nonlinear generalization of OU, and builds on Wong's (1964) physics-based *repulsive Wong process*. By considering the problem under a partial equilibrium setting, this chapter is able to fully characterize a representative arbitrageur's optimal portfolio policy. The closed-form availability of the optimal portfolio policy considering these additional risk factors allows us to make a number of interesting and novel observations. Firstly, arbitrageurs recognizing the possibility of persistently large mispricings exhibit a diminishing incremental propensity to exploit them, as they develop an aversion to fundamental risk factors and lose confidence in the model's ability to predict mean-reversion. Interestingly, this behaviour is counter-intuitive from the point of view of OU, which ignores fundamental risk factors and translates large mispricings into attractive opportunities and hence portfolio allocations linearly. Secondly, since widening mispricings imply both an improvement in the investment set and a capital loss on a current position, this work explicitly shows that beyond certain bounds in the magnitude of mispricing, a further widening of the mispricing prompts a *reduction* in the capital allocated to exploiting it. This chapter refers to these bounds as stability bounds. This artifact of the presented model complements the corresponding OU-specific result in Boguslavsky and Boguslavskaya (2004) and Jurek and Yang (2007). It is shown that the stability bounds in the nonlinear model are always tighter than those implied by OU. More importantly, this result suggests that when considering fundamental risk factors ignored by OU and other existing models, arbitrageurs seek to unwind losing trades sooner. Each of these observations is useful for researchers examining the capacity and role of arbitrageurs in enforcing price efficiency.

By way of a pairs-trading exercise, this chapter examines the effect of incorporating these fundamental risk factors on the nonlinear model's empirical performance. This model is tested against OU and a "2-sigma" strategy common in the literature (Gatev et al. (2006)), on daily FTSE100 pairs data obtained from DataStream. The results are in favor of the nonlinear model, whereas OU and the "2-sigma" strategy which frequently go bankrupt. These results are interpreted as evidence that fundamental risk factors are significant to the capacity of arbitrageurs to enforce price efficiency. Interestingly, since model parameters are typically estimated from historical data, it is also found that the nonlinear model is more robust to the adverse effects of parameter misspecification than OU.

The rest of this chapter is organized as follows: Section 3.2 provides a review of the related literature, further motivating approach taken in this chapter. Section 3.3 derives the optimal portfolio strategy and stability bounds in the model presented in this work, then compares these with the corresponding results assuming OU. This section also presents a test to determine the effect of parameter misspecification on the optimal portfolio strategy. Section 3.4 presents a FTSE100 pairs-trading exercise. Section 3.5 concludes.

3.2 - Related Literature

There is considerable evidence that prices can diverge from fundamental values and similarly, that relative prices can diverge from their natural levels. Shleifer and Summers (1990) and Black (1986) highlight the effect of noise traders' irrational trading behaviour on the formation of prices. Noise traders react disproportionately to news in the belief that they have insider information regarding the future direction of prices, creating a mispricing. On the one hand, Friedman (1953) argues that such mispricings cannot exist for long, as rational arbitrageurs will trade against noise traders, hence pushing prices toward fundamental values. On the other hand, Lamont and Thaler (2003a) find examples of persistent deviations from *the law of one price*, Froot and Dabora (1999) find mispricings between *Siamese twin* shares, Malkiel (1977) finds mispricings in the valuation of closed-end funds, and Lamont and Thaler (2003b) find mispricings in tech stock carve-outs. Consistent with the latter view, the nonlinear model presented in this chapter is based on the notion that mispricings can persist for significant periods of time. But why in general does capital not materialize to eliminate such mispricings?

The persistence of mispricings can be attributed to many factors. Abreu and Brunnermeier (2002, 2003) show that when individual arbitrageurs attempt to "time the market", informational asymmetries can cause coordination problems that consequently allow mispricings to persist. Brunnermeier and Pedersen (2005) highlight the role of predatory trading behaviour in the persistence of mispricings. Shleifer and Vishny (1997), Xiong (2001), and De Long et al. (1990) show that impediments to arbitrage can arise endogenously. When arbitrageurs allocate capital aimed at exploiting a mispricing, further divergences may cause them to unwind positions at a loss to preserve capital, a phenomenon described by Friedman (1953) as "buying high and selling low". This creates a price feedback mechanism which exacerbates the mispricing. Lamont and Thaler (2003b), Liu and Longstaff (2004), Gromb and Vayanos (2002), and Basak and Croitoru (2000) suggest that portfolio constraints can allow mispricings to persist. Moreover, Kondor (2009) shows that arbitrageurs do not necessarily act to eliminate a mispricing even when portfolio constraints do not bind, because "opportunities might get better tomorrow".

The partial equilibrium model is aimed at capturing these intuitions into an exogenously defined mispricing process. This is motivated in part by the need for more robust quantitative risk management models in the wake of the recent financial crisis (see Shaw and Schofield (2012)). The need for more robust risk-management ties with arbitrageurs' profit-seeking behaviour. In the empirical section of this chapter, it is shown that commonly-employed arbitrage strategies which are not robust to fundamental risk frequently go bankrupt. After specifying the model, this chapter develops an optimal portfolio policy for a representative arbitrageur to adopt. This relates the model presented here to Boguslavsky and Boguslavskaya (2004), Jurek and Yang (2007), Liu and Longstaff (2004), and more generally to the portfolio optimization methodology pioneered by Merton (1969, 1971).[3]

[3]Similar works exploring optimal portfolio selection in arbitrage include Mudchanatongsuk et al. (2008), Kim et al. (2008), Lv and Meister (2009), and Elliott et al. (2005).

Liu and Longstaff (2004) model the arbitrage opportunity using a Brownian Bridge process, which has two fixed values *a priori*, typically at the start and end of the investment horizon. Under this setup, arbitrageurs have perfect *ex-ante* knowledge regarding the date of the elimination of a mispricing. Jurek and Yang (2007) observe that, apart from finitely-lived opportunities such as deviations from option put-call parity, arbitrageurs typically would not have perfect knowledge regarding the elimination of a mispricing. They suggest that general mispricings including relative mispricings between shares are better modelled using OU, which implies uncertainty regarding the mispricing at all future dates. Jurek and Yang (2007) also introduce the concept of horizon risk and divergence risk which imply, respectively, uncertainty whether a mispricing will converge prior to the reporting period, and a worsening in the mispricing after an arbitrageur has opened a position. They argue that the Brownian Bridge process considers only divergence risk, whereas OU considers both. Nonetheless, a commonality between these approaches is that arbitrageurs have perfect knowledge regarding the natural level of a mispricing. But this assumption ignores Lamont and Thaler's (2003b) *fundamental risk*, which highlights changes in economic fundamentals that can affect the natural mispricing level itself. The presented model considers this risk. Even though the natural level is assumed fixed, the fact that an arbitrageur faithful to this model exhibits a diminishing incremental exploitation of the mispricing can be interpreted in terms of his decreased confidence that the mispricing will revert to its natural level. On the other hand, the fact that models based on OU or Brownian Bridge call for aggressively increasing the capital allocation as divergences get large (linearly in the case of OU) implies that arbitrageurs faithful to these models have complete confidence in the estimated natural level, which completely ignores fundamental risk.

The dynamics of the model presented in this chapter are simple enough to yield a closed-form solution to the arbitrageur's allocation problem and to include OU as a special case. Thus, these results can be readily compared to Jurek and Yang (2007) or Boguslavsky and Boguslavskaya (2004). The work here also relates to Xiong (2001), Basak and Croitoru (2000), and Liu and Longstaff (2004) in the sense that log-utility over terminal wealth is assumed.

More generally, considering the arbitrageur's portfolio allocation problem in partial equilibrium gives us two main advantages over the general equilibrium models discussed in this section. Firstly, this is able to construct an analytically tractable and empirically implementable optimal portfolio policy for a representative arbitrageur, which yields novel insights into the behaviour of rational arbitrageurs. This is particularly important in illustrating one of the central novelties in the model presented in this chapter, namely the scepticism toward exploiting large mispricings. Secondly, the precise conditions are ascertained under which arbitrageurs cease exploiting a mispricing and unwind losing positions, complementing works on the impediments to arbitrage.

3.3 - The Model

This section considers the behaviour of a rational arbitrageur who has access to an arbitrage opportunity and a riskless asset. The riskless asset yields a continuously compounded return. This work models the arbitrage opportunity as a mean-reverting asset S_t. Intuitively, this asset can be thought of as a long/short position in an equity pair. In this context, holding long one unit of S_t is equivalent to holding one unit long in some undervalued security and one unit short in an overvalued security at time $t \in [0, T]$.

3.3.1 The Investment Opportunity:

Consistent with Merton (1969, 1971), Wachter (2002), and Kim and Omberg (1996), it is assumed that the investible universe consists of two assets. The first is a risky asset S_t which describes the evolution of the arbitrage opportunity (i.e. mispricing). The second is the risk-free asset B_t. The risky asset is shaped by the hyperbolic tangent function, and this is precisely what this chapter relies upon to model the notion of fundamental risk. This assertion is justified in the next section and figure 3.1. The dynamics of these two assets, respectively, are given by the following equations:

$$
\begin{aligned}
(1) \qquad\qquad dS_t &= -\frac{k}{c}\tanh\big(c(S_t - \overline{S})\big)\,dt + \sigma dZ_t, \\
(2) \qquad\qquad dB_t &= rB_t dt.
\end{aligned}
$$

where $\overline{S}$ is the natural level of the mispricing, r is the risk-free rate, Z_t is a Brownian Motion with respect to the real-world probability measure, $k > 0$ is the parameter of mean reversion, and $\sigma > 0$ is the volatility parameter. The investment horizon is continuous and finite from $0 \leq t \leq T < \infty$. The parameter $c > 0$ measures the nonlinearity in the mean-reversion component of (1). This parameter is explored further below.

3.3.2 Contrasting the Nonlinear Dynamics with OU:

The parameter $c > 0$ is a novel feature in this model. This parameter specifies the nonlinearity of mean-reversion in the stochastic process (1). To make clear the connection between the model presented in this chapter and OU, it is observed that in the special case as $c \to 0$, the dynamics of the mispricing (1) reduce to:

$$\text{(3)} \qquad \lim_{c\to 0}\left[-\frac{k}{c}\tanh\left(c(S_t-\overline{S})\right)dt\right] = -k(S_t-\overline{S})dt.$$

The right hand side of (3) is precisely the mean-reversion component of an OU process with parameters k and σ. Indeed, this would reduce (1) to:

$$\text{(4)} \qquad dS_t \to -k(S_t-\overline{S})dt + \sigma dZ_t \text{ as } c \to 0,$$

which implies that this model is now reduced to precisely the specification in Jurek and Yang (2007). Further setting $r = 0$ and $\overline{S} = 0$ reduces the model to precisely the specification in Boguslavsky and Boguslavskaya (2004).

As long as c remains strictly positive, the stochastic process in (1) is nonlinear in the spread. Specifically, the rate at which the strength of reversion increases in S_t weakens, implying that the reversion term attains a limit as the mispricing diverges significantly from its natural level, i.e. $|S_t| \gg \overline{S}$. To see this, let us force S_t infinitely far above its natural level, and obtain:

$$\text{(5)} \qquad \lim_{S_t\to\infty}\left[-\frac{k}{c}\tanh\left(c(S_t-\overline{S})\right)dt\right] \to -\frac{k}{c}dt.$$

The right hand side of (5) implies that the strength of mean-reversion pulling S_t to its natural level has attained its limit. This reduces the dynamics of the mispricing (1) to:

$$\text{(6)} \qquad dS_t \to -\frac{k}{c}dt + \sigma dZ_t \text{ as } S \to \infty.$$

The dynamics now are a diffusion process with a constant drift. But what is the economic significance of properties (3) and (5)?

To adequately capture the realistic nature of a mispricing process, it is desirable for a model to command dynamics simple enough to be practically implementable, yet rich enough to capture realistic market characteristics. It has been shown that the model presented in this chapter addresses an important phenomenon absent from the dynamics of OU and the Brownian Bridge, namely fundamental risk. Consistent with Shaw and Schofield (2012), the "hyperbolic OU" process (1) is approximately OU for a small mispricing (i.e. when S_t is close to $\overline{S}$), but the reversion term tails off as the mispricing becomes arbitrarily large, and tends to a constant. This artifact captures the intuition that investors may have diminished confidence in that the mispricing will mean-revert. Conversely, OU necessarily implies a mean-reversion that strengthens linearly with the mispricing.

3.3.3 Modelling Risk-Preferences:

A central part of establishing an optimal portfolio strategy is to obtain an idea of the risk preferences of arbitrageurs. The representative arbitrageur is assumed to maximize expected log utility over terminal wealth. Translating this to the investment actions of a typical hedge fund, for example, terminal wealth can be defined as the fund's performance at the end of a reporting period. Specifically, the arbitrageur seeks to maximize the *value function*:

$$V(S_t, W_t, t) = \max_{N_t} \left[\mathbb{E}_t \left(\ln W_T \right) \right], \tag{7}$$

Here N_t is the number of units held in the mispricing S_t, and W_t denotes wealth at time t. Equation (7) is an economic statement that, through optimally exploiting the mispricing, an arbitrageur maximizes log utility over terminal wealth. In the next section an explicit expression for N_t is derived.

Log-utility occurs as a special case of the general CRRA family of power utility (See e.g. Wachter (2002)). Here, log-utility is assumed for two main reasons. Firstly, Breiman (1961), Kyle and Xiong (2001), and Lv and Meister (2009) show that investors who assume log-utility will asymptotically outperform any other strategy with certainty. Secondly, assuming log-utility admits a mathematically tractable solution to the resulting stochastic optimization problem (see e.g. Kim and Omberg (1996), Boguslavsky and Boguslavskaya (2004), and Kargin (2003)). This is particularly relevant to the optimization problem, given the nonlinearity present in (1). Though assuming general CRRA would allow us to explore inter-temporal hedging demands as in Kim and Omberg (1996) and Liu (2007), it is noteworthy to point out that since OU occurs as a special case of the process (1), and CRRA results for OU are available in Jurek and Yang (2007), conclusions regarding inter-temporal hedging demands are likely to be maintained in a more complex model such as the one presented in this chapter.

3.3.4 The Optimal Portfolio Strategy:

Here the optimal portfolio policy for the representative arbitrageur is derived. First, the self-financing condition is imposed. This requires that changes in wealth are directly attributable to investment in the risky and risk-free asset, and no external source. If N_t is the number of units held in the mispricing, and M_t is the number of units held in the riskless asset, then the self-financing condition implies that the arbitrageur's budget constraint is given by:

$$W_t = N_t S_t + M_t B_t,\ t \in [0, T]\,. \tag{8}$$

Differentiating the budget constraint with respect to time t yields:

$$
\begin{aligned}
dW_t &= N_t dS_t + M_t dB_t \\
&= N_t dS_t + \frac{W_t - N_t S_t}{B_t} dB_t,
\end{aligned}
\tag{9}
$$

where the second equality in (9) imposes that whatever remains after investment in the mispricing is invested in the riskless asset.[4] To facilitate the analysis, a new time variable $\tau = T - t$ is defined as the time left for investing, so (7) is re-written as:

$$
V(S_t, W_t, \tau) = \max_{N_t} \left[\mathbb{E}_t \left(\ln W_T\right)\right]. \tag{10}
$$

Maximizing (10) subject to (9) yields the optimal portfolio strategy. The optimal number of units held in the mispricing depends on the current mispricing level, current wealth, and time left for trading. Therefore, it is denoted by $N_t(S_t, W_t, \tau)$. It is derived as follows:

Theorem 1. (OPTIMAL PORTFOLIO STRATEGY) *The optimal portfolio strategy for a log-utility maximizing arbitrageur facing investment opportunities (1) and (2) is given by*:

$$
N_t(S_t, W_t, \tau) = \left(\frac{-\frac{k}{c} \tanh\left[c(S_t - \overline{S})\right] - rS_t}{\sigma^2} \right) W_t. \tag{11}
$$

Proof. The proof is given in the Technical Appendix.

□

[4]Similarly, an arbitrageur willing to make an investment in the mispricing larger than current wealth can do so by shorting the riskless asset.

Remark 1. It is interesting to contrast the optimal portfolio strategy with the equivalent OU strategy. Firstly, taking the limit $c \to 0$ in (11) yields the corresponding OU optimal strategy, which is labelled N_{OU}:

$$
\begin{aligned}
N_{OU} &= \lim_{c \to 0} \left(\frac{-\frac{k}{c} \tanh \left[c(S_t - \overline{S}) \right] - rS_t}{\sigma^2} \right) W_t \\
&= \left(\frac{-k(S_t - \overline{S}) - rS_t}{\sigma^2} \right) W_t.
\end{aligned} \tag{12}
$$

N_{OU} is the specification in Jurek and Yang (2007). Setting $r = 0$ yields the specification in Boguslavsky and Boguslavskaya (2004).

While c is positive, the optimal holding is nonlinear in the mispricing, unlike OU. This point is illustrated by comparing strategy (11), labeled N_{TANH}, to the equivalent OU strategy N_{OU} using common parameter values. This is illustrated in Figure 3.1.

Figure 3.1 illustrates an interesting property: Unlike OU, this strategy implies that, as the mispricing diverges away from its natural level, an arbitrageur becomes less interested in further exploiting it. This idea represents an important difference between the model presented here and OU. The model presented in this chapter interprets large divergences as potential regime breaks, consistent with the notion of fundamental risk (Lamont and Thaler (2003b)). On the other hand, OU interprets large divergences as attractive investment opportunities.

3.3.5 Do Arbitrageurs Always Exploit a Mispricing?

It seems intuitive to think that an arbitrageur will always exploit a mispricing in pursuit of a profit. Indeed, Figure 3.1 illustrates this notion for a given level of wealth. However, once an arbitrageur allocates capital aimed at exploiting a mispricing, further divergence of the mispricing causes him to lose capital, and hence his risk-bearing capacity. Since the optimal strategy calls for investment

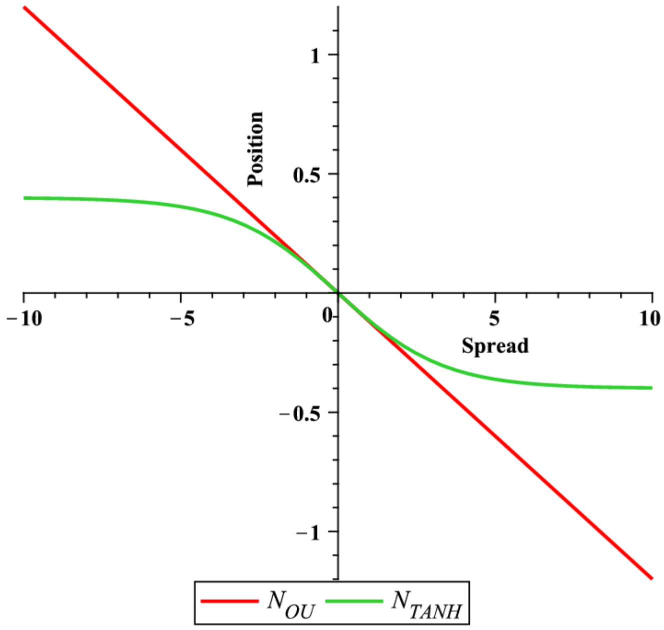

Figure 3.1 - Illustration of position size in the mispricing (spread). The OU (red) and TANH (green) optimal portfolio strategies with parameters for nonlinearity $c = 0.3$, mean-reversion $k = 3$, volatility $\sigma = 5$, wealth $W = 1$, risk-free rate $r = 0$, and natural mispricing level $\overline{S} = 0$. The mispricing range is $S \in [-10, 10]$, and the position in the mispricing taken by both strategies is on the vertical axis.

immediately following a mispricing's departure from its natural level, it is shown here that there is a critical level in the mispricing beyond which arbitrageurs unwind positions at a loss to preserve capital. These are called *stability bounds*, so-called because an arbitrageur only seeks to enforce price efficiency (trade against the mispricing) as long as the mispricing lies within these bounds. Consequently, a scrutiny of the optimal portfolio policy (11) is done in order to determine the direction in which an arbitrageur trades as a response to the evolution of the mispricing process. It turns out that as long as the magnitude of the mispricing relative to its natural level lies within the stability bounds, further divergence causes the arbitrageur to continue investing against the direction of the spread asset (i.e. exploiting the mispricing). However, once this stability bound is breached, further divergence causes the arbitrageur to unwind a losing position, effectively trading *with* the mispricing. The stability bounds are analytically characterized in the following theorem. Without loss of generality, the riskfree rate is set at $r = 0$.

Theorem 2. (STABILITY BOUNDS) *An arbitrageur will only trade against a mispricing if the magnitude of the mispricing beyond its natural level,* $\left|S_t - \overline{S}\right|$, *lies within a fixed bound. Beyond this bound, an arbitrageur will unwind a losing position. The bound is as follows*:

$$\left|(S_t - \overline{S})\right| < \frac{1}{c} \operatorname{arcsinh}\left(\frac{\sigma c}{\sqrt{k}}\right). \tag{13}$$

Proof. The proof is given in the Technical Appendix.

□

Remark 2. Taking the limit $c \to 0$ shows that the stability bounds (13) tend to those implied by OU (Jurek and Yang (2007), Boguslavsky and Boguslavskaya (2004)):

$$\lim_{c\to 0}\left[\frac{1}{c} \operatorname{arcsinh}\left(\frac{\sigma c}{\sqrt{k}}\right)\right] = \frac{\sigma}{\sqrt{k}}. \tag{14}$$

Direct calculation in (14) shows that for any values of reversion $k > 0$, volatility $\sigma > 0$, and coefficient of nonlinearity $c > 0$, the stability bounds in the model presented here are tighter than those implied by OU. This implies that when arbitrageurs are averse to fundamental risk factors, small losses are cut sooner rather than potentially large losses later (Shleifer and Vishny (1997), De Long et al. (1990)).

3.3.6 Robustness to Parameter Uncertainty:

It has been implicit up to this point that the investor has perfect *ex-ante* knowledge regarding the values of the parameters k, c, and σ. This assumption is maintained in Brieman's (1961) and Lv and Meister's (2009) theorems regarding non-bankruptcy of the trading strategy. In reality, parameters are often estimated from historical data, and these theories seem predicated on the assumption that the future perfectly mimics the past, and that the model used describes the time series perfectly. Of course one cannot affirm that the future will mimic the past perfectly, nor that any diffusion model can describe market evolution perfectly. Thus the risk of bankruptcy persists in reality, which is in part due to model risk or Knightian uncertainty (Knight (1921)). This realization leads us to an interesting empirical question, namely what is the effect of a parameter misspecification on the risk-management capabilities of the optimal strategy?

To answer this question, the optimal strategy (11) is compared against the OU strategy on synthetically generated data by forcing trading in both models based on deliberately wrong reversion and volatility parameters. The methodology is as follows:

Here, 20,000 paths are simulated from both a discretized OU process and the generalized Wong process (1) over one year with a time increment of $\frac{1}{252}$ (corresponding to 252 trading days). The synthetic data is generated by a Euler-Maruyama discretization scheme (see Kloeden and Platen (1999)). For both processes, let us set $c = 0.3$, $k = 4$, $\sigma = 4$, $r = 0$, and $\overline{S} = 0$. The initial value for each simulated path is set at zero, corresponding to no immediate arbitrage opportunity. The optimal strategy (11) is then applied (now referred to as N_{TANH}), against the optimal strategy implied by OU, which is denoted by N_{OU}. For clarity, these are given explicitly as follows:

$$N_{TANH} = \left(\frac{-\frac{k}{c}\tanh\left(c(S-\overline{S})\right) - rS}{\sigma^2} \right) W. \tag{15}$$

$$N_{OU} = \left(\frac{-k(S-\overline{S}) - rS}{\sigma^2} \right) W. \tag{16}$$

To test the effect of parameter misspecification in (15) and (16), this section trades using deliberately wrong parameters k and σ. Specifically, the parameters $k = k(1 + K)$ and $\sigma = \sigma(1 + Z)$ are set, with $K \in [0, 3]$ and $Z \in [-0.5, 0]$. This means that the two strategies each trade with k ranging from its true value 4 up to quadruple its true value 16, and σ ranging from its true value 4 down to half its true value 2. Let us define a parameter grid such that it contains 16 values for each parameter, implying 256 parameter combinations in total.

For each parameter combination and each strategy, two quantities are measured: the proportion of trials which result in bankruptcy, and the mean terminal wealth over all trials. The results for OU-generated data are presented in Figure 3.2, and the results for Wong-generated data are presented in Figure 3.3.

The top-left panel in Figure 3.2 shows the proportion of trials of the OU strategy that result in bankruptcy. Trading using the correct parameters universally causes none of the trials to result in bankruptcy. However, note that overestimating the reversion parameter leads to an increase in the proportion of trials that result in bankruptcy. Underestimating the volatility parameter has the same effect. This is intuitive, as the portfolio strategy seeks to exploit a mispricing which it believes will revert faster than it actually would, or exhibit less volatility than it actually would, respectively. The findings here are consistent with Lv and Meister (2009) in the sense that quadrupling the estimated reversion parameter has the identical detrimental effect to halving the volatility estimate. For N_{OU} in the top-left panel, 58% of trials ended in bankruptcy as a result of trading wrongly assuming half the actual volatility or quadruple the actual reversion rate. The corresponding statistic for N_{TANH} in the center-left panel is 41%. This shows that, given a common set of model parameters, the optimal TANH strategy (15) is more robust to parameter misspecification than the corresponding OU strategy (16) in terms of bankruptcy risk. This point is illustrated in the bottom-left panel, which subtracts the proportion of trials resulting in bankruptcy of the

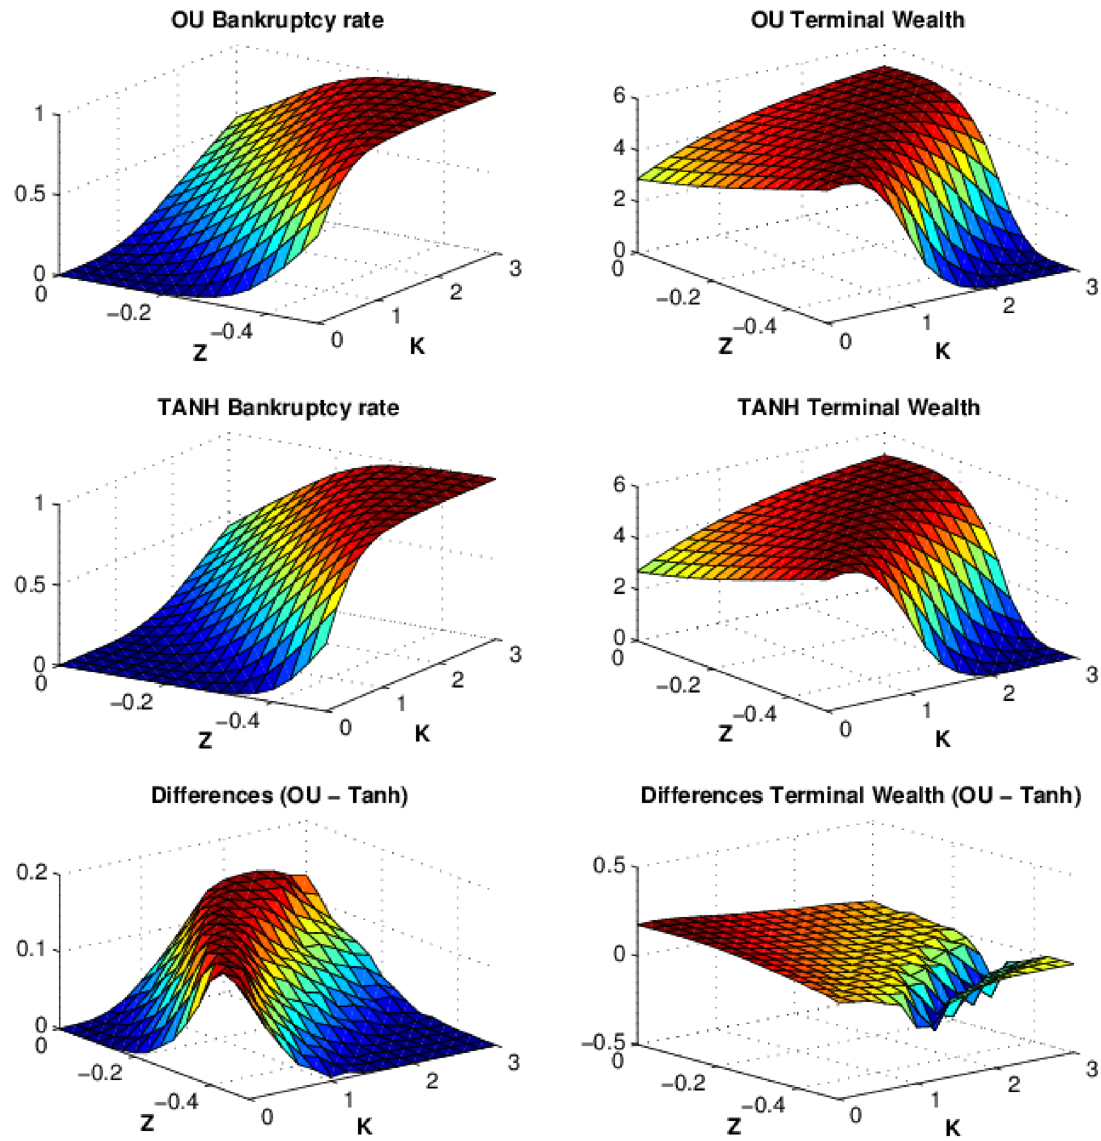

Figure 3.2 - Effect of parameter misspecification on the optimal strategy assuming OU mispricing dynamics. Simulation of 20,000 OU paths over 252 trading days with parameters $k = 4$, $\sigma = 4$, $r = 0$, $c = 0.3$, and $\overline{S} = 0$. Both OU and TANH optimal strategies (11) and (12) trade assuming deliberately wrong parameter values. $K \in [0, 3]$ implies reversion $k \in [4, 16]$, while $Z \in [-0.5, 0]$ implies volatility $\sigma \in [4, 2]$. Initial wealth is set to 1. The left panels show the effect of parameter misspecification on the probability of bankruptcy, while the right panels show this effect on mean terminal wealth.

TANH strategy from that of the OU strategy for each combination of k and σ. That the bottom-left panel is positive reflects the notion that the TANH strategy is less susceptible to parameter misspecification.

The top-right and center-right panels of Figure 3.2 show the effect of parameter misspecification on terminal wealth in both strategies. These panels show that OU results in higher mean terminal wealth. The reason for this is that the diminishing marginal exploitation of the spread in the N_{TANH} strategy for a common set of parameters implies that the TANH strategy always invests relatively less in the spread than N_{OU} for a given wealth. Therefore, should the larger position assumed by N_{OU} not result in bankruptcy, mean terminal wealth would be greater, reflecting the greater risk OU takes. This fact is illustrated in the bottom-right panel, which subtracts the average terminal wealth of the TANH strategy from the OU strategy for each combination of k and σ. That the bottom-right panel is largely positive implies that, conditional on survival, the OU strategy outperforms the TANH strategy in terms of maximizing terminal wealth. However, the OU strategy goes bankrupt quicker than the TANH strategy on average, which explains the "dip" below zero in the bottom right panel.

Next, the results of running the same exercise as above are presented but on synthetic data generated from the generalized Wong process (1). These are shown in Figure 3.3.

The layout of Figure 3.3 is identical to that of Figure 3.2. As one might expect, overestimating the actual reversion and underestimating the actual volatility each have detrimental effects on trading performance. When assuming Wong mispricing dynamics, the OU strategy N_{OU} goes bankrupt 59% of the time when either the reversion parameter is misspecified to be quadruple its actual value, or the volatility parameter half its actual value. The corresponding bankruptcy rate for N_{TANH} is 46%. These statistics are almost identical to those obtained assuming OU mispricing dynamics.

Although very similar, comparing Figures 3.2 and 3.3 reveals an interesting finding. In analyzing the effect of parameter misspecification on terminal wealth, there is one qualitative difference between assuming OU mispricing dynamics

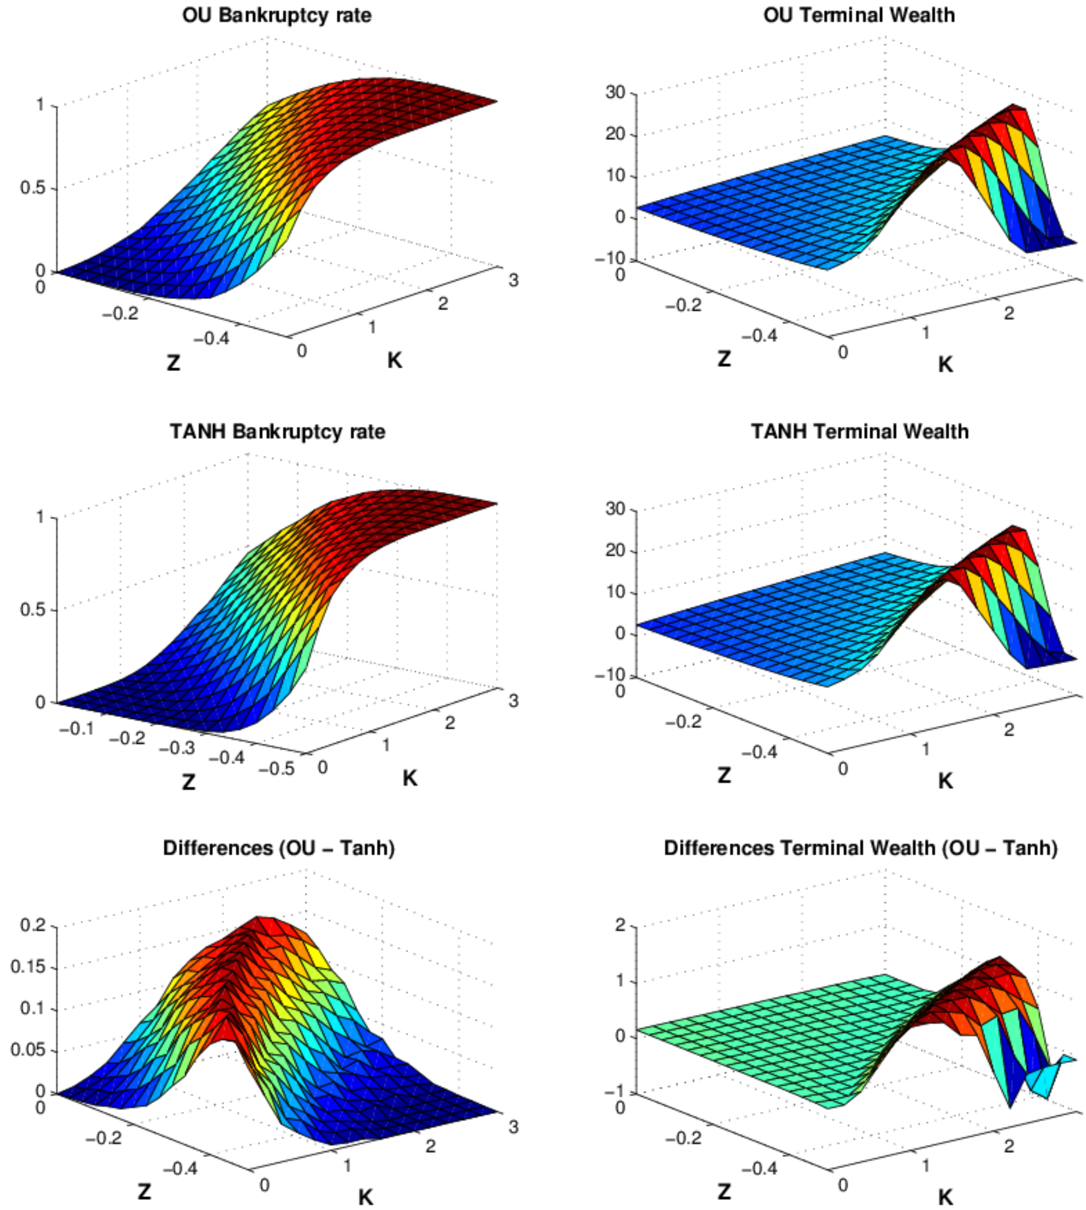

Figure 3.3 Effect of parameter misspecification on the optimal strategy assuming Wong mispricing dynamics. Simulation of 20,000 paths generated from (1) over 252 trading days with parameters $k = 4$, $\sigma = 4$, $r = 0$, $c = 0.3$, and $\overline{S} = 0$. Both OU and TANH optimal strategies (11) and (12) trade assuming deliberately wrong parameter values. $K \in [0, 3]$ implies reversion $k \in [4, 16]$, while $Z \in [-0.5, 0]$ implies volatility $\sigma \in [4, 2]$. Initial wealth is set to 1. The left panels show the effect of parameter misspecification on the probability of bankruptcy, while the right panels show this effect on mean terminal wealth.

and Wong mispricing dynamics. The bottom-right panel of Figure 3.3 shows by how much N_{OU} outperforms N_{TANH}. It seems surprising at first that this number should be positive, especially when the trading parameters are highly misspecified.

Indeed by Theorem 1, N_{TANH} is demonstrably the optimal trading strategy suited towards the generalized Wong process (1). How then can N_{OU} outperform N_{TANH} when trading with wrong parameters? Consider the following facts for a fixed set of parameters: First, N_{OU} is naturally more aggressive than N_{TANH}. Second, the generalized Wong process (1) is naturally less mean-reverting than the OU process (see section 3.3.2). Taken together, the N_{OU} returns massive profits in a very small number of "lucky" cases, even though it goes bankrupt in the vast majority of others. The handful of high-return cases have a disproportionately large effect on mean terminal wealth. Nevertheless, because N_{OU} goes bankrupt quicker than N_{TANH}, the latter outperforms in both terminal wealth and bankruptcy aversion when the error in parameter estimation becomes sufficiently large.

Overall in the simulation, bankrupted trials were not excluded from the calculation of average terminal wealth. Excluding bankrupted trials results in a large positive bias in terminal wealth. This is intuitive, as the surviving trials would have taken extremely large positions in the mispricing, which "luckily" paid off, and exaggerate average terminal wealth.

3.4 - Empirical Application

Now that this chapter has characterized the optimal portfolio strategy (14) and explored its properties, let us implement these results to historical data. This is done by testing the empirical efficacy of the N_{TANH} strategy (15) against two alternative strategies. The first is the N_{OU} optimal strategy (16). The second is a popular "rule of thumb" strategy based on Gatev et al. (2006). For consistency this chapter calls this strategy N_{GGR}, and it works as follows:

> N_{GGR}: Open a position in the mispricing when the mispricing deviates more than 2 standard deviations away from its historical mean. Hold the position constant until the mispricing reverts fully to its natural level.

The size of the position opened and held is the equal to what N_{OU} would suggest (i.e. the position of an equivalent OU arbitrageur).

Remark 3. The nature of N_{GGR} is such that a position is not opened in a mispricing immediately as a mispricing diverges from its natural level. Because of this, some tests result in N_{GGR} not trading at all. Also, because there is no intrinsic stop-loss strategy to the 2-sigma approach (Gatev et al. (2006), De Jong et al. (2009)), the analysis presented here does not impose one.

3.4.1 Data, Historical Period, and Methodology:

The data is obtained from DataStream, and chosen for four pairs of FTSE100 companies using closing daily prices over the period Jan 1, 2001 - Jun 30, 2008 (the *training* period), and Jul 1, 2008 - Jun 30, 2009 (the *trading* period). When these days are non-trading (holidays, new years), the previous day's closing value is chosen as a proxy.

The four pairs are chosen based on the highest correlation of daily returns over the training period. Beyond the statistical relationship, they are commercially intuitive choices, and are as follows:

- HSBC and Barclays Plc.
- Legal & General and Prudential.
- Royal Dutch Shell-B and BP.
- Royal Bank of Scotland and Lloyds Plc.

To construct the mispricing S_t, the linear difference in daily closing prices for each pair is taken. Denoting these prices by $P_{i,t}$ for $i = \{1, 2\}$ implies:

$$S_t = P_{1,t} - P_{2,t}.$$

Here, an OU process is calibrated for its reversion, volatility, and natural level parameters $\{\hat{k}, \hat{\sigma}, \hat{\overline{S}}\}$ respectively, over the training period Jan 1, 2001 - Jun 30, 2008 (1956 trading days) using Maximum Likelihood Estimation. Once estimates for $\{\hat{k}, \hat{\sigma}, \hat{\overline{S}}\}$ are obtained, they are substituted into both N_{OU} and N_{TANH}[5] and set the coefficient of nonlinearity $c = 0.02$ into N_{TANH}. For N_{GGR}, the historical mean is set equal to the natural level implied by the OU calibration, i.e. $\mu_{\hat{GGR}} = \hat{\overline{S}}$, and calculate the historical standard deviation $\sigma_{\hat{GGR}}$ based on the 1956 daily data points over the training period.

With this in hand, let us proceed by backtesting the three strategies (15), (16), and N_{GGR} over the trading period Jul 1, 2008 - Jun 30, 2009, a total of 261 trading days. Initial wealth is set arbitrarily at 1.

3.4.2 Application to FTSE100 data:

The results of the empirical backtest are shown in Figures 3.4 through 3.7. Throughout the backtest, the parameter $r = 0$ is set to clearly illustrate the differences between the three strategies without the complication of an additional asset. Further, this section assumes no transaction costs, full use of the short proceeds, and no other institutional frictions. The performance of each strategy in this analysis is measured by the total profit at the end of the trading period, which can easily be seen in figures 3.4 - 3.7. Furthermore, the quality of mean-reversion is assessed qualitatively based on the number of zero-crossings of the mispricing process. This can also be easily seen from the figures.

[Figures 3.4-3.7 presented in the Technical Appendix of this chapter]

[5] Implementing N_{TANH} as a stand-alone model requires its parameters to be estimated independently through numerical methods (Hurn et al. (2003), Picchini (2007)). The application here makes use of the parameter c as a proxy to Knightian uncertainty, and hence use a common set of parameters for N_{TANH} and N_{OU}. Consequently, an interpretation of this methodology is a replacement of the risk aversion parameter inherent in general CRRA utility, in favor of a parameter of model-risk aversion: The higher the value of c set by the arbitrageur, the less confidence he has in the ability of OU to capture the dynamics of the mispricing.

The results show that, when mean reversion is poor, or there are apparent signs of a regime-break, both the N_{OU} and N_{GGR} strategies lead to very large losses, or bankruptcy. Conversely, assuming a coefficient of curvature $c = 0.02$, the N_{TANH} strategy never resulted in bankruptcy, despite poor mean reversion and regime breaks in the data. Figure 3.7 shows that although N_{OU} returns are volatile, they are very high when reversion is strong about the natural level $\overline{S}$. Given that the nature of OU is such that an arbitrageur has complete confidence in reversion to the natural level, this result is expected.

The backtesting period is chosen specifically to test the models' ability to handle the tumultuous market conditions characteristic of that period. Nevertheless, the four examples in Figures 3.4 through 3.7 represent a wide range of market characteristics regarding the evolution of the spread. Overall the results are consistent with the intuition from the works mentioned in sections 3.1 and 3.2 in this chapter.

3.5 - Conclusion

A partial equilibrium, nonlinear model of arbitrage is developed; aimed specifically at capturing the fundamental market- and sentiment-based risks inherent in arbitrage trading. In so doing, it is demonstrated that the incorporation of these risks yields novel insights into the optimal behaviour of rational arbitrageurs. Most importantly, this work derives an optimal portfolio strategy in closed form, and hence the simplicity of this framework provides potential for its use in further investigating the role of arbitrageurs in enforcing price efficiency.

The approach taken here draws insight directly from the empirical and behavioural perspectives regarding the conduct of rational arbitrageurs facing irrational noise traders. As a result, the specification presented in this chapter captures far broader market risk factors related to arbitrage than otherwise assuming OU or a Brownian Bridge. Moreover, a proxy to the notion of Knightian uncertainty is provided. It is shown that the incorporation of these factors results in a portfolio strategy exhibiting a diminishing appetite for exploiting widening mispricings: a feature absent in existing specifications. Consistent with empirical observations, the model presented here incorporates fundamental risks which manifest as regime-breaks or persistent mispricings, whereas OU and the Brownian Bridge do not. When tested empirically, the model presented in this chapter delivers superior overall performance compared to OU or "2-sigma" strategies. This result is attributed to the notion that this model better captures realistic market phenomena.

Though the role of arbitrageurs is typically aimed at exploiting mispricings, widening mispricings clearly have adverse effects on arbitrageurs' wealth. The model presented in this chapter shows that under certain conditions, arbitrageurs will choose to unwind losing positions by effectively trading *with* (not against) the mispricing in order to preserve capital, even when investible mispricing opportunities are at their greatest. The closed-form availability of these conditions complements the existing literature on arbitrage in both partial and general equilibrium. Specifically, the presented model calls for small losses to be cut sooner, in contrast to OU which calls for large losses to be cut later, or indeed the "2-sigma" strategy which has no intrinsic risk-management functionality whatsoever. In the empirical application, these artifacts of OU and the "2-sigma" strategies frequently lead them to bankruptcy, unlike the model presented here, which remains solvent. Based on this, one can conjecture that the specification presented here results in lower price inefficiency and volatility following the liquidation of funds.

The empirical period over which the model presented in this chapter is tested contains a 4-month period during which short-selling was banned in financial stocks in the UK (September 19, 2008 to January 16, 2009). Nevertheless, the bulk of the empirical period studied in this chapter had short-selling permitted. It could be argued that while incorporating the short-selling ban would obviously impact the numerical result of arbitrage profits, the theoretical and qualitative insights derived in this chapter would remain unaffected.

Although the four pairs studied in this chapter were chosen to represent a wide range of market behaviour, it could be argued that these were chosen "after the fact". However, the choice nevertheless clearly illustrates how the performance of the model presented in this chapter contrasts with those of the extant literature - and that is the thrust of this chapter's contributions.

Future research in this area would benefit from incorporating frictions such as transaction costs and portfolio constraints into the presented model, which would allow an investigation into market efficiency based on the kind of portfolio algorithm which is the subject of this chapter. The results from such a study would complement existing studies, which are predominantly based on threshold-type trading strategies. Further, the framework developed here could be generalized to consider the effects of arbitrage activity on market liquidity, price-feedback, and contagion. Since the model shows that recognizing fundamental risk factors inherent in arbitrage trading reformulates arbitrageurs' behaviour toward the exploitation of perceived mispricings as compared to existing models, the results of such a generalized study can yield new insight into the effects of arbitrage activity on price stability in the wider market.

3.6 - Technical Appendix

In this section, Theorems 1 and 2 are presented, along with accompanying proofs, as well as Figures 3.4 through 3.7.

Theorem 1. (OPTIMAL PORTFOLIO STRATEGY) *The optimal portfolio strategy for a log-utility maximizing arbitrageur facing investment opportunities (1) and (2) is given by*:

$$N_t(S_t, W_t, \tau) = \left(\frac{-\frac{k}{c} \tanh\left[c(S_t - \overline{S})\right] - rS_t}{\sigma^2} \right) W_t. \tag{17}$$

Proof. First, let us substitute (1) and (2) into budget constraint (9) to obtain:

$$\begin{aligned} dW_t &= N_t dS_t + \frac{W_t - N_t S_t}{B_t} dB_t \\ &= \left(-\frac{k}{c} N_t \tanh\left(c(S_t - \overline{S})\right) + r(W_t - N_t S_t) \right) dt + \sigma N_t dZ_t, \end{aligned} \tag{18}$$

Now let us employ stochastic optimal control. First, the value function (10) is expanded using Ito's Lemma and Ito multiplication. Suppressing the time subscripts:

$$dV = \frac{\partial V}{\partial S} dS + \frac{\partial V}{\partial W} dW + \frac{\partial^2 V}{\partial S^2}(dS)^2 + \frac{\partial^2 V}{\partial W^2}(dW)^2 + \frac{\partial^2 V}{\partial S \partial W}(dS \cdot dW) - \frac{\partial V}{\partial \tau} dt.$$

Substituting from (1) and (18):

$$
\begin{aligned}
dV = \; & \frac{\partial V}{\partial W}\left(\left(-\frac{k}{c}N\tanh\left(c(S-\overline{S})\right)+r(W-NS)\right)dt+\sigma N dZ\right) \\
& +\frac{\partial V}{\partial S}\left(-\frac{k}{c}\tanh\left(c(S-\overline{S})\right)dt+\sigma dZ\right)+\frac{1}{2}\sigma^2\frac{\partial^2 V}{\partial S^2}dt+\frac{1}{2}\sigma^2N^2\frac{\partial^2 V}{\partial W^2}dt \\
& +\frac{1}{2}\sigma^2N\frac{\partial^2 V}{\partial S\partial W}dt-\frac{\partial V}{\partial \tau}dt.
\end{aligned}
$$

Next, let us employ Bellman's principle of optimality $\mathbb{E}_t(dV) = 0$ to derive the Hamilton-Jacobi-Bellman equation (Bellman (1957)) for this optimization problem:

$$
\begin{aligned}
0 = & \frac{\partial V}{\partial W}\left(-\frac{k}{c}N\tanh\left(c(S-\overline{S})\right)+r(W-NS)\right)+\frac{1}{2}\sigma^2N\frac{\partial^2 V}{\partial S\partial W}-\frac{\partial V}{\partial \tau} \\
& +\frac{\partial V}{\partial S}\left(-\frac{k}{c}\tanh\left(c(S-\overline{S})\right)\right)+\frac{1}{2}\sigma^2\frac{\partial^2 V}{\partial S^2}+\frac{1}{2}\sigma^2N^2\frac{\partial^2 V}{\partial W^2},
\end{aligned}
\tag{19}
$$

along with the terminal condition:

$$
V(S, W, 0) = 0. \tag{20}
$$

Condition (20) implies that at the end of the investment horizon $\tau = 0$, any arbitrage opportunity is worthless to the arbitrageur, since the trading period is over. Deriving (19) has used the fact that Z_t is a martingale under the real-world measure. Equation (19) is the Hamilton-Jacobi-Bellman (henceforth "HJB") equation for the optimization problem. In order to derive the optimal policy, let us maximize the right hand side of (19) with respect to N. To achieve this, let us use the first-order optimality condition with respect to N by differentiating (19) with respect to N then setting this derivative equal to zero and solving for N. Differentiating:

$$0 = \frac{\partial V}{\partial W}\left(-\frac{k}{c}\tanh\left(c(S-\overline{S})\right) - rS\right) + \sigma^2 N \frac{\partial^2 V}{\partial W^2} + \frac{1}{2}\sigma^2 \frac{\partial^2 V}{\partial S \partial W},$$

and finally solving for N while using subscripts to denote partial derivatives:

(21) $$N = -\left(\frac{V_W}{V_{WW}}\right)\left(\frac{-\frac{k}{c}\tanh\left(c(S-\overline{S})\right) - rS}{\sigma^2}\right) - \left(\frac{V_{SW}}{V_{WW}}\right).$$

Equation (21) gives us the optimal portfolio policy N. However, in its current form, it is of little use, since it depends explicitly on partial derivatives of the value function $V(S, W, \tau)$. Without eliminating the dependence of the optimal portfolio N on the derivatives of the value function, equation (21) remains of little practical use. In order to proceed in finding an explicit form for N, this step employs the method of separation of variables. Let us postulate a trial solution, in terms of a functional form for $V(S, W, \tau)$, then verify whether it solves the HJB equation (19). It is asserted that it is likely $V(S, W, \tau)$ inherits some structural properties from its components, namely the log-utility function (Bjork (2004)). To this end, let us postulate the following solution:

(22) $$V^*(S, W, \tau) = \ln(W) + f(S, \tau).$$

Here f is a function of S and τ only. Condition (20) translates to $f(S, 0) = 0$. The next step now checks whether this solution is successful. The derivatives of (22) are:

- $V^*_W = \frac{1}{W}$
- $V^*_{WW} = -\frac{1}{W^2}$
- $V^*_S = f_S$
- $V^*_{SS} = f_{SS}$
- $V^*_\tau = f_\tau$

- $V_{SW}^{*} = 0$.

A sufficient condition for optimality of the portfolio policy is the concavity of the value function with respect to the state variable W. From above, one can see that the sign of V_{WW} and the fact that wealth is non-negative ensure that this ansatz, if successful in the sense below, would yield the optimal portfolio strategy. Substituting the relevant partial derivatives into the optimal portfolio policy in (21):

$$N(S, W, \tau) = \left(\frac{-\frac{k}{c} \tanh\left(c(S - \overline{S})\right) - rS}{\sigma^2} \right) W, \tag{23}$$

and substituting (23) together with the relevant partial derivatives into the HJB equation (19) yields a simplification of the HJB equation, namely:

$$\begin{aligned} 0 = &\left(-\frac{k}{c} \tanh\left(c(S - \overline{S})\right)\right) f_S - f_\tau + \frac{1}{2}\sigma^2 f_{SS} - \frac{\left(-\frac{k}{c}\tanh(c(S - \overline{S})) - rS\right)^2}{2\sigma^2} \\ &+ r\left(1 - \frac{\left(-\frac{k}{c}\tanh(c(S - \overline{S})) - rS\right) S}{\sigma^2}\right) \\ &- \frac{k \tanh\left(c\left(S - \overline{S}\right)\right)\left(-\frac{k}{c}\tanh\left(c\left(S - \overline{S}\right)\right) - rS\right)}{\sigma^2 c}. \end{aligned} \tag{24}$$

This step has successfully eliminated W from (24), meaning the HJB equation is reduced in dimensionality, and now only depends on S and τ. Moreover, the optimal portfolio policy in (23) is fully characterized and stripped of its dependence on the value function or any of its derivatives. Since this theorem is only interested in the optimal portfolio strategy, the optimization procedure is complete. One needs only to make the technical assumption that a smooth function $f(S, \tau)$ twice differentiable in S, once in τ, exists and satisfies (24).[6]

[6]The explicit solution $f(S, \tau)$ would allow us to characterize an optimal portfolio policy for a general CRRA arbitrageur. See e.g. Jurek and Yang (2007) and Boguslavsky and Boguslavskaya (2004).

This assumption is used in Xiong (2001). This completes the proof of Theorem 1.

□

Theorem 2. (STABILITY BOUNDS) *An arbitrageur will only trade against a mispricing if the magnitude of the mispricing beyond its natural level,* $\left|S_t - \overline{S}\right|$, *lies within a fixed bound. The bound is as follows*:

$$\left|(S_t - \overline{S})\right| < \frac{1}{c} \operatorname{arcsinh}\left(\frac{\sigma c}{\sqrt{k}}\right). \tag{25}$$

Proof. Let us begin by expanding the optimal strategy (11) using Ito's lemma:

(26)
$$dN = \frac{\partial N}{\partial S} dS + \frac{\partial N}{\partial W} dW + \frac{\partial^2 N}{\partial S^2}(dS)^2 + \frac{\partial^2 N}{\partial W^2}(dW)^2 + \frac{\partial^2 N}{\partial S \partial W}(dS \cdot dW) - \frac{\partial N}{\partial \tau} dt.$$

Next, substituting from (1), (9), and (11), along with the partial derivatives of N, and using $(dS)^2 = \sigma^2 dt$, $(dW)^2 = \sigma^2 N^2 dt$, and $(dS \cdot dW) = \sigma^2 N dt$, let us rewrite (26) as:

$$\begin{aligned} dN = &\left[\frac{-wk \operatorname{sech}^2(c(S - \overline{S}))}{\sigma^2} + \frac{wk^2 \tanh^2(c(S - \overline{S}))}{\sigma^4 c^2}\right] dS \\ &+ \left[\frac{wk(2c^2\sigma^2 + k) \operatorname{sech}^2(c(S - \overline{S})) \tanh(c(S - \overline{S}))}{c\sigma^2}\right] dt. \end{aligned} \tag{27}$$

Equation (27) separates the instantaneous changes in the optimal portfolio allocation as a response to changes in the mispricing dS and the time variable dt. Therefore, the direction of the arbitrageur's portfolio allocation in response to changes in the mispricing is governed by the sign of the first set of square brackets in (27), namely:

$$\left[\frac{-wk\ \mathrm{sech}^2(c(S-\overline{S}))}{\sigma^2}+\frac{wk^2\tanh^2(c(S-\overline{S}))}{\sigma^4c^2}\right],$$

and this is negative when:

$$\frac{wk^2\tanh^2(c(S-\overline{S}))}{\sigma^4c^2}-\frac{wk\ \mathrm{sech}^2(c(S-\overline{S}))}{\sigma^2}<0.$$

Simplifying this expression, the following result is obtained:

$$\left|(S-\overline{S})\right|<\frac{1}{c}\ \mathrm{arcsinh}\left(\frac{\sigma c}{\sqrt{k}}\right). \tag{28}$$

$\square$

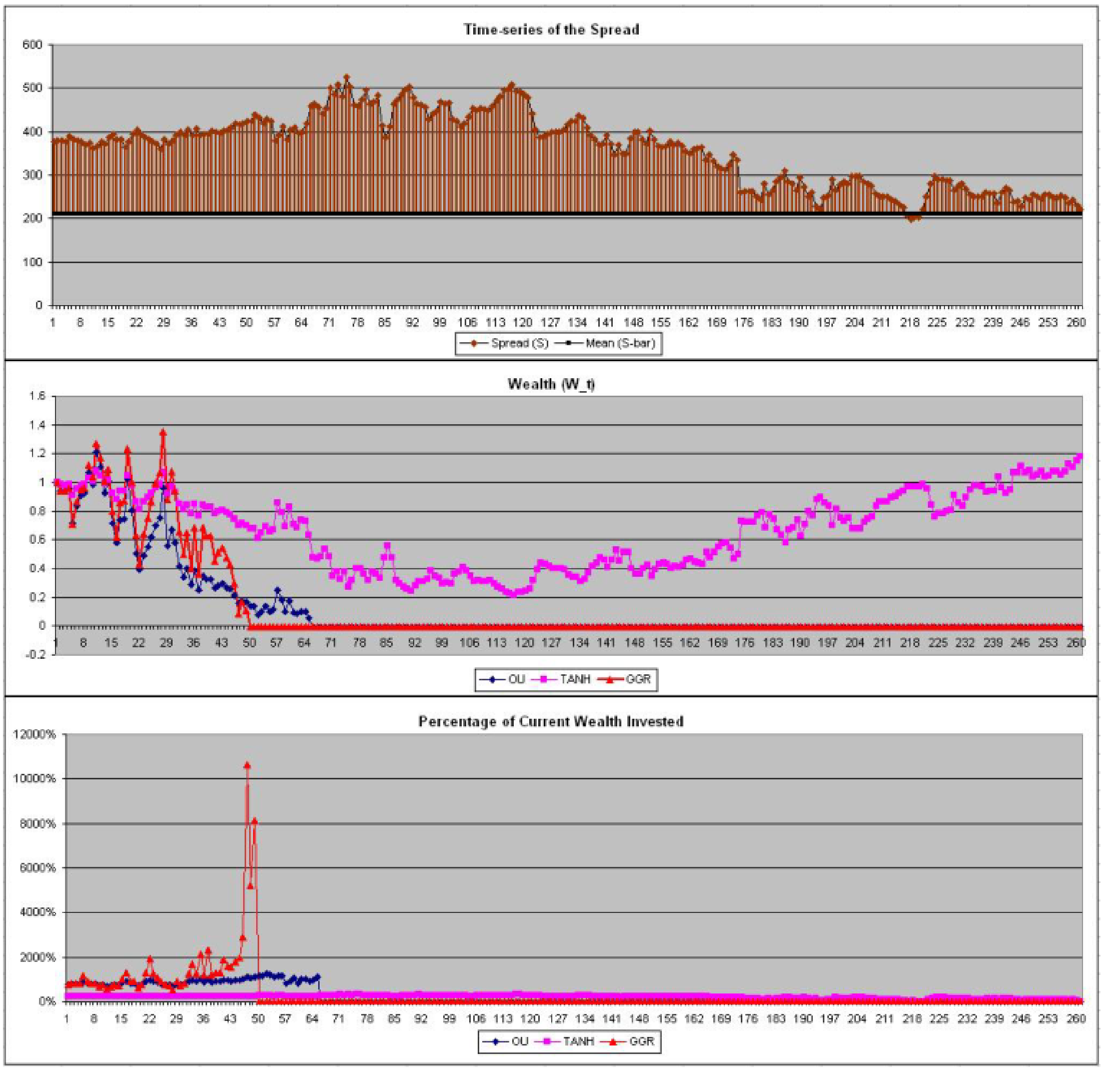

Figure 3.4 - HSBC - Barclays: Poor mean reversion. The top panel shows the evolution of the daily spread during Jul 1, 2008 - Jun 30, 2009 (Brown) and the natural spread level (Black). The middle panel shows the corresponding evolution of wealth resulting from trading using N_{TANH} (Pink), N_{OU} (Blue), and N_{GGR} (Red). The bottom panel shows the percentage of wealth invested in the spread.

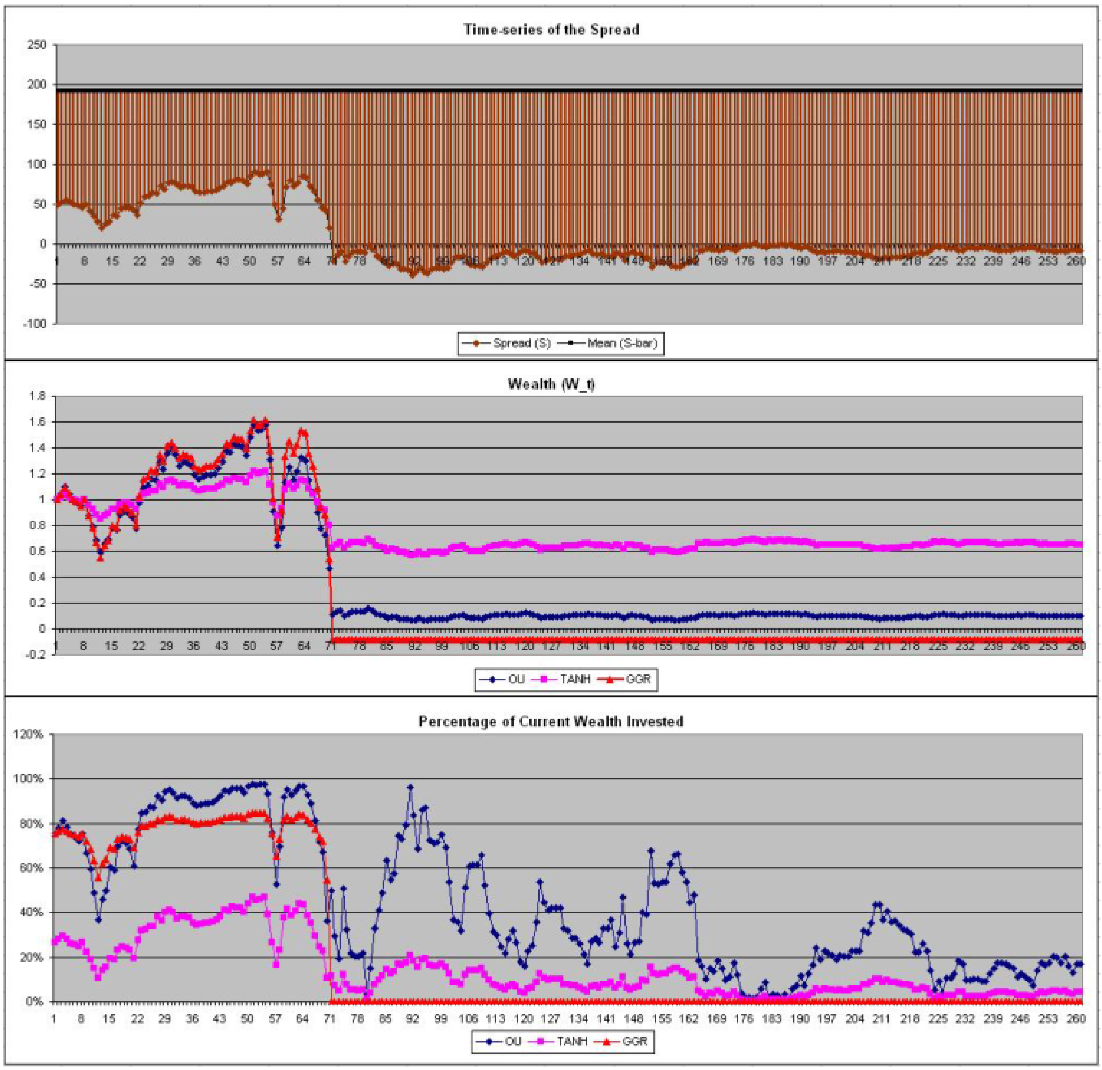

Figure 3.5 - RBS - Lloyds: Regime break. The top panel shows the evolution of the daily spread during Jul 1, 2008 - Jun 30, 2009 (Brown) and the natural spread level (Black). The middle panel shows the corresponding evolution of wealth resulting from trading using N_{TANH} (Pink), N_{OU} (Blue), and N_{GGR} (Red). The bottom panel shows the percentage of wealth invested in the spread.

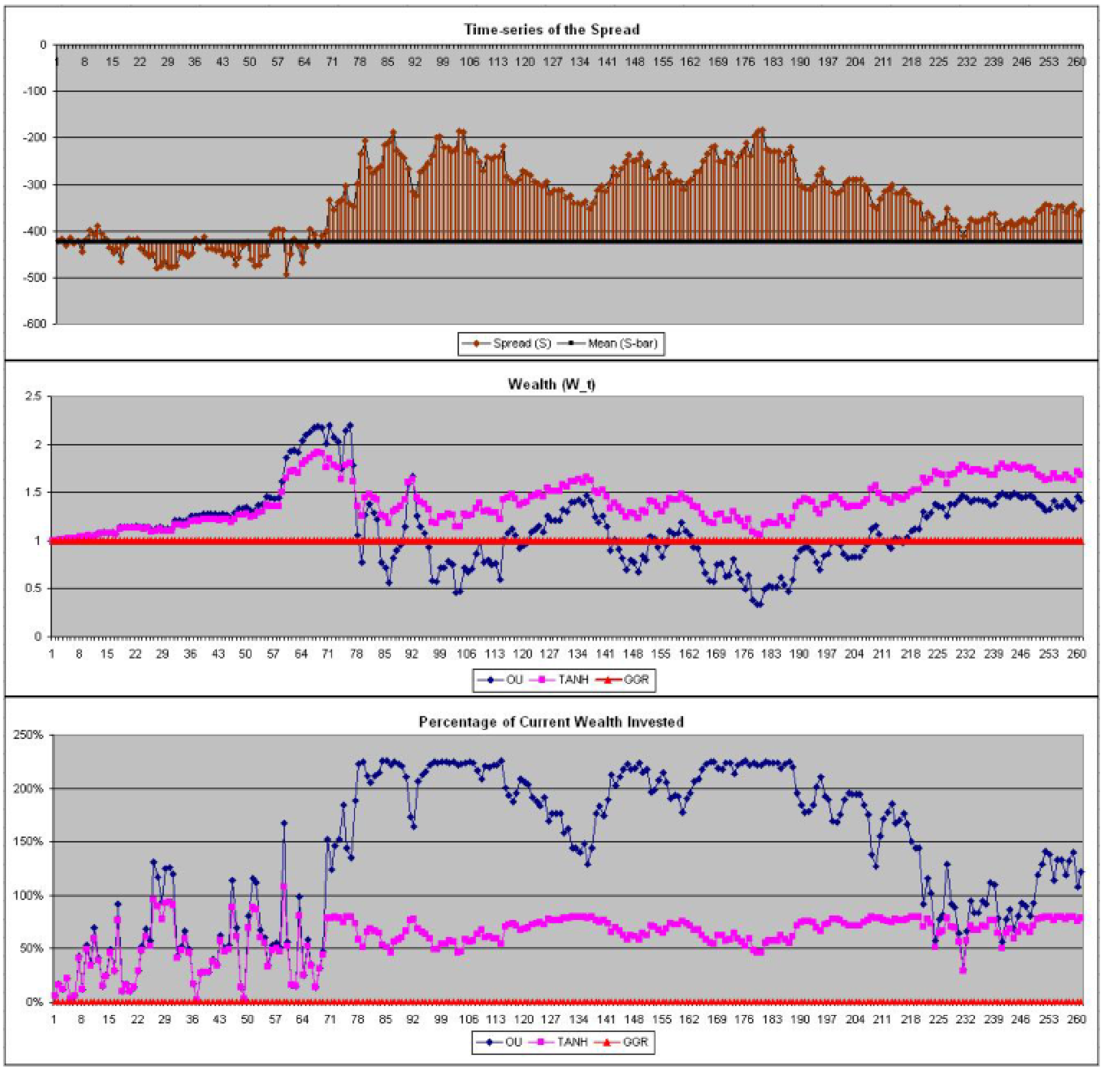

Figure 3.6 - Legal&General - Prudential: Poor mean reversion. The top panel shows the evolution of the daily spread during Jul 1, 2008 - Jun 30, 2009 (Brown) and the natural spread level (Black). The middle panel shows the corresponding evolution of wealth resulting from trading using N_{TANH} (Pink), N_{OU} (Blue), and N_{GGR} (Red). The bottom panel shows the percentage of wealth invested in the spread.

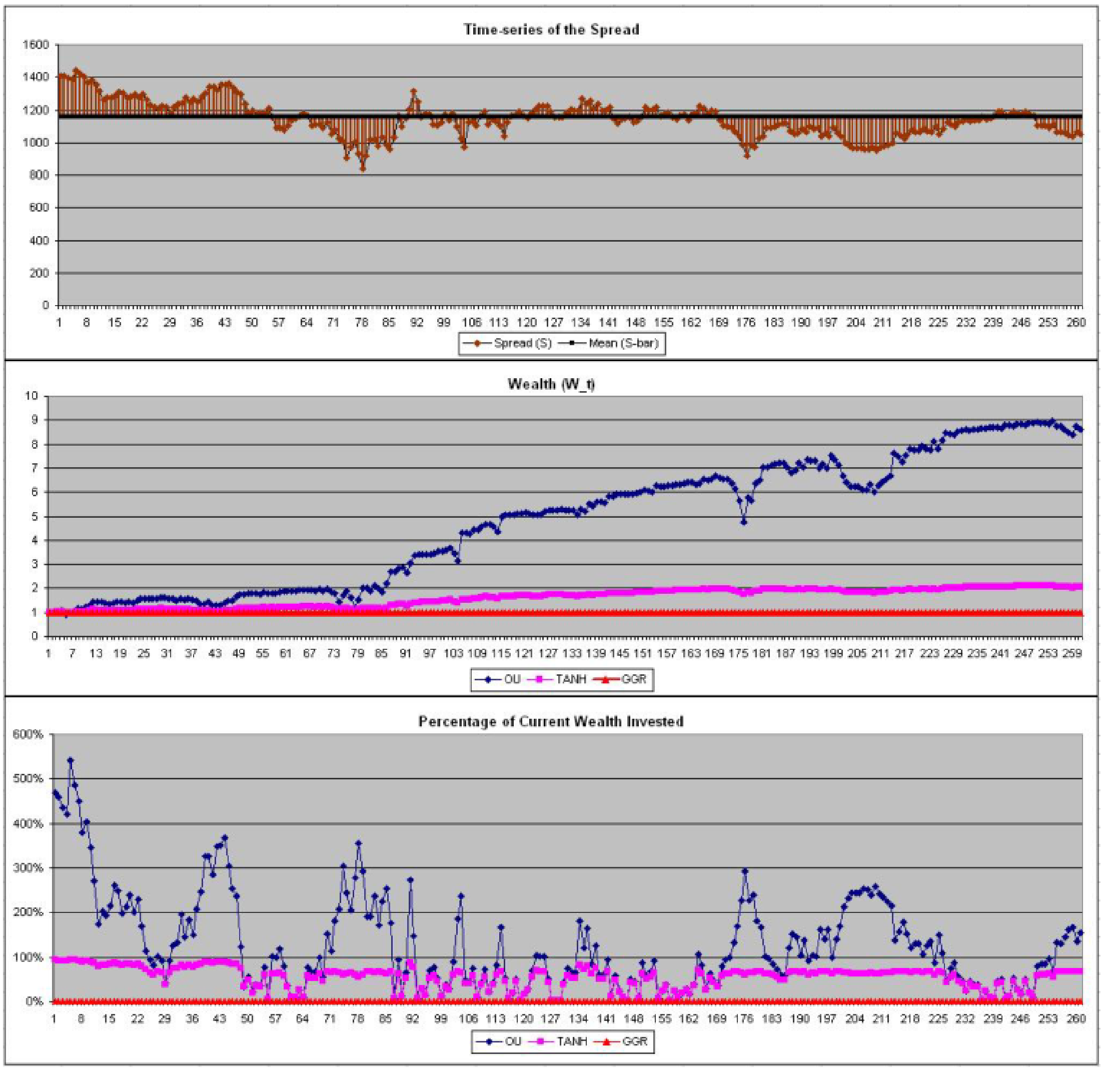

Figure 3.7 - Royal Dutch Shell B - BP: Good mean reversion. The top panel shows the evolution of the daily spread during Jul 1, 2008 - Jun 30, 2009 (Brown) and the natural spread level (Black). The middle panel shows the corresponding evolution of wealth resulting from trading using N_{TANH} (Pink), N_{OU} (Blue), and N_{GGR} (Red). The bottom panel shows the percentage of wealth invested in the spread.

Blank Page

Chapter 4

Arbitrage and the Law of One Price in the Market for American Depository Receipts

Abstract

This is the first work to highlight pairs trading as the main price-correcting mechanism by which arbitrage can maintain stock-ADR parity in the high-frequency intraday domain. It is shown that arbitraging stock-ADR pairs extracts small per-trade profits which accumulate to a substantial aggregate return. The observed strong tendency of pricing disequilibria to mean-revert, along with the two-way convertibility between stocks and ADRs, mean that arbitrageurs face minimal risks toward price divergence. They do, however, face uncertainty about the duration of individual trades. The magnitude of this uncertainty relates directly to the profit target arbitrageurs set after a long/short position is established. This fact can explain why some disequilibria go unexploited. Overall, the work presented here provides evidence against automatically efficient prices (the notion that price disequilibria are eliminated in absence of arbitrageurs), and supports the view that mispricings incentivise arbitrageurs to enforce market efficiency.

Keywords: American Depository Receipts, Arbitrage, Law of One Price, Pairs Trading.

JEL classification: F36, G14, G15.

4.1 - Introduction

The principal aim of this chapter is to highlight, in a high-frequency setting, a previously overlooked market mechanism by which stock-ADR parity is enforced, namely, pairs trading. This is contrasted with an alternative parity enforcement mechanism which is the one typically assumed in the ADR literature, and which is known as "direct arbitrage". The latter entails actually converting the stock into the ADR or vice-versa.

This chapter is motivated by an observation made by Chen et al. (2009) and also by Werner and Kleidon (1996), that markets for cross-listed securities are less than fully integrated. That fact suggests that investment decisions in each of the two markets are not based solely on which is the cheaper location to trade. Hence, aggregate "local" trading is not enough to enforce stock-ADR parity. It is therefore intuitive that these markets should constitute fertile ground for arbitrageurs, who are compensated for correcting price disparities across international markets (e.g. Kondor (2009)). Several works explore whether these disparities are profitable from an arbitrageur's point of view. On the one hand, Kato et al. (1991) and Miller and Morey (1996) find that arbitrage is unprofitable due to costs and regulatory impediments. On the other hand, Suarez (2005) finds evidence of profitable arbitrage in French stock-ADR pairs. However, all of these studies use the direct arbitrage method as the means of arbitrage that they study. In other words, they assume that the only way to realise a profit from a stock-ADR mispricing is to convert the stock into the ADR, or vice versa. However, as is explain in detail later, the costs of conversion are very high. For this reason, it is found that in the UK sample, direct arbitrage is not a profitable arbitrage strategy. With this in mind, what other market mechanisms can restore stock-ADR parity? What are the characteristics, costs, and impediments facing these mechanisms? What inferences can be drawn about the auto-efficiency of the stock-ADR market?

Using 131 million quote price observations, this chapter sets out to provide answers to these questions. The quote price data covers an exhaustive sample of 25 firms domiciled in the UK, which also trade as ADRs on US exchanges. This work contributes to the existing literature in two ways. Firstly, this chapter introduce to the ADR literature an overlooked market mechanism by which the Law of One Price is enforced, namely pairs trading[1] - a trading technique explored in the literature by Gatev et al. (2006), Jurek and Yang (2007), and De Jong et al. (2009). Pairs trading is an arbitrage strategy based on establishing a long/short position in a pair of relatively mispriced securities, then liquidating the position upon price convergence. Compared to direct arbitrage, pairs trading enjoys fewer cross-border regulatory constraints and lower trading costs. These features allow the pairs trader to exploit pricing anomalies which, comparatively, would be prohibitively expensive or even infeasible for the direct arbitrageur.[2] Clearly, pairs trading relies on the establishment of a long-short position in the stock and ADR. While this strategy is susceptible to short-selling bans, there are no short-selling bans in effect during the historical period studied here.

The second contribution of this chapter involves an examination of the limits to stock-ADR arbitrage. Why do seemingly profitable disequilibria appear and persist in such a highly developed and liquid market? What are the most prevalent risks facing arbitrageurs in this market? If an exploitable opportunity

[1]This technique in the context of ADRs bears superficial similarity to Gagnon and Karolyi's (2010) "perfect foresight arbitrage". Its major difference is in the fact that the methodology presented here has no ex-ante knowledge of price convergence. Hence one is able to fully characterize trading risks. The technique is also similar to that employed in Hong and Susmel (2003) and Broumandi and Reuber (2012), except that this chapter employs high-frequency intraday data.

[2]A typical regulatory constraint limits conversion of the home stock into ADRs. Puthenpurackal (2006) describes how the Infosys ADR traded at over double the price of its home-market stock. Unlike pairs trading, direct arbitraging cannot extract profits in this environment.

is either missed or ignored in such fertile ground for arbitrageurs, then it is important to understand why.

The results show definitively that small but frequent profitable disparities appear in stock-ADR price pairs. When a simple threshold pairs trading strategy is applied, some 640 arbitrage trades yield in total 1.45% in excess of the risk-free rate. Arbitrage returns are uncorrelated with broad market returns. Furthermore, since stocks and ADRs are two-way convertible, stock-ADR pairs trading entails no fundamental risk. In contrast, a direct arbitrage strategy applied to the sample yields only a single profitable trading opportunity. Hence from the point of view of a direct arbitrageur, this market gives the illusion of being auto-efficient, without any need for arbitrage as a parity enforcement mechanism.

The results also show that the main disincentive facing arbitrageurs in this market is an uncertainty towards the duration of open arbitrage positions. Half of all arbitrage positions last under nine minutes. However, for each additional US-cent demanded by arbitrageurs in per-share profit, the expected trade duration more than doubles, and the standard deviation of trade durations increases by some 20%. Arbitrageurs who are averse to this uncertainty, and widen their trade entry bounds, observe lower duration uncertainty but also fewer exploitable opportunities. These results are consistent with works documenting arbitrageur heterogeneity, synchronization problems between arbitrageurs, and the foregoing of present opportunities in the hopes of "even better ones tomorrow" (Abreu and Brunnermeier (2002, 2003) and Kondor (2009)). Finally, this work examine the mean-reverting characteristics of stock-ADR mispricings by way of an Ornstein-Uhlenbeck calibration exercise. Across the sample, stock-ADR mispricing evolutions exhibit remarkably strong mean-reversion, with half-lives lasting no more than a few minutes. Hence, the risk posed by a worsening of the mispricing after a long/short position is established, is minimal.

Overall, this work provide evidence against the auto-efficiency of prices in markets for cross-listed securities, and suggest that arbitrage is an important mechanism in maintaining stock-ADR parity. It is concluded that pricing anomalies incentivise arbitrageurs to restore price parity, consistent with Grossman and Stiglitz (1976, 1980).

In the rest of this chapter, section 4.2 reviews the literature, section 4.3 describes the dataset and methodology, section 4.4 presents the main empirical results and robustness tests, and section 4.5 concludes.

4.2 - Related Literature

Introduced by JP Morgan in 1927 as a way for US investors to diversify their holdings internationally, ADRs are a dollar-denominated representation of ownership in a non-US company. They trade as conventional shares on US exchanges, and provide identical cashflows as their corresponding domestic stocks. In this regard, the two are arguably the same security (see Pulatkonak and Sofianos (1999)). Moreover, since ADRs are two-way convertible into their domestic stocks, the Law of One Price suggests that the two securities should trade at parity. However, several works document factors such as foreign capital controls, trading costs, high T-bill rates which reduce the incentive to arbitrage, and heterogeneity in investor preferences, as factors keeping the stock-ADR relation significantly away from parity (Grossmann et al. (2007), Gagnon and Karolyi (2010) and Rabinovitch et al. (2003)).

Since the seminal work on ADRs of Rosenthal (1983) and Maldonado and Saunders (1983), research surrounding ADRs can broadly be split into two categories: The first involves investigations into the contribution of cross-listings to international price discovery, and the determinants of premia and discounts of ADRs relative to their home-market (underlying) shares. Eun and Sabherwal

(2003) examine US stocks which also trade on the Toronto Stock Exchange, and find that price adjustments to disequilibria occur on both markets. They suggest that the speed of the adjustment process relates positively to competition for order flow, and that prices in the thinner market adjust quicker. Grammig et al. (2005) examine 3 German stock-ADR pairs and conclude that the majority of price discovery occurs in the home market, while price adjustment following an exchange rate shock occurs on the ADR side. Grossmann et al. (2007) examine 74 stocks from nine countries against their ADRs, and show that transaction costs and high T-bill rates act to augment stock-ADR mispricings. Rabinovitch et al. (2003) and Gagnon and Karolyi (2010) suggest that the prevalence of ADR premia/discounts relate negatively to the degree of economic home-market development: The less economically developed the home-market, the greater the magnitude of ADR premia/discounts relative to the home-market stock. Furthermore, foreign-ownership restrictions act to augment ADR mispricings and a frequently-cited case study of India's Infosys is discussed in Puthenpurackal (2006). Overall, this strand of research highlights factors which can keep the stock-ADR relation away from price parity.

The second strand of research recognizes explicitly the role of arbitrage in enforcing the Law of One Price. This group examines to what extent mispricings between ADRs and their home-market shares constitute arbitrage opportunities. It is this group to which this chapter contributes. Are ADR premia/discounts exploitable; or attributable to costs and other regulatory and institutional frictions?

Frequently using daily closing prices, early research aimed at answering this question largely concludes in favour of price efficiency, and finds little if any exploitable price discrepancies between ADRs and their underlying shares. Examples of such works include Kato et al. (1991), Wahab and Lashgari (1992), and Park and Tavakkol (1994). However, as Suarez (2005) points out, if stock-ADR markets are indeed price-efficient, then one would expect exploitable mispricings to be short-lived and hence invisible to the daily observer. Furthermore, Gagnon and Karolyi (2010) point out that since international markets trade at different hours, conclusions based on closing prices are unlikely to be definitive due to an inevitable relative price change between the close of each market.

To address the problems surrounding daily sampling, recent works have used higher-frequency datasets to examine whether arbitrage opportunities exist in the stock-ADR market. Suarez (2005) uses quote-level data for 12 French stocks and their ADRs. He concludes that arbitrage opportunities can exist but that trading costs massively impede profitability, and the profits that typically remain are comparable to the opportunity costs of hiring a financial expert to monitor the market. Miller and Morey (1996) use a 3-month high-frequency dataset of a single British stock and its ADR, namely Glaxo-Wellcome, and conclude firmly that no arbitrage opportunities exist. Hence the conclusions of high-frequency works are unclear, and the question remains: Are there arbitrage opportunities in ADR markets?

This chapter sets out to definitively answer this question. Inspired by Chen et al. (2009) and Werner and Kleidon's (1996) conclusion of heterogeneity between UK and US investor preferences, it is puzzling that the majority of works conclude that very few arbitrage opportunities exist. Indeed, if the two markets were fully integrated, then competition for order flow alone should be enough to enforce price efficiency, as investors would be indifferent toward holding stocks or ADRs, and base their decisions solely on the cheaper trading venue. Conversely, in segmented markets, arbitrage opportunities should be visible and profitable enough to attract capital and enforce price efficiency. The notion that arbitrage is unprofitable and yet stock-ADR markets are segmented is precisely what motivates us to pursue alternative mechanisms by which the Law of One Price can be enforced.

This work employs a pairs trading approach, which by nature is contrarian, since it relies on exploiting a price divergence which the arbitrageur hopes will converge. This ties this work directly to the substantial body of literature examining short-run mean-reversion in securities prices. Inspired by Jegadesh and Titman's (1993) observation that previous "loser" stocks tend to outperform previous "winner" stocks, contrarian strategies such as pairs trading have become a very prevalent tool in enforcing price efficiency in many contexts. Gatev et al. (2006) employ this methodology in the context of statistically cointegrated pairs of stocks, and De Jong et al. (2009) employ the same methodology in the context of dual-listed shares such as Royal Dutch Shell or Unilever. As Pontiff (1996, 2006) points out, this methodology is not without risk. Pairs trading involves going long one underpriced security while simultaneously shorting the expensive security in the hope that both prices "converge". If they diverge further, the pairs trader suffers a capital loss. The appeal of applying this methodology to ADRs is in the fact that ADRs are arguably more economically similar to their underliers than a pair of dual-listed stocks, which cannot be "converted" into one another.

4.3 - Data and Methodology

This section presents a description of the dataset and data refinements, together with details of the pairs trading strategy employed.

4.3.1 UK Stock and ADR Data

Contemporaneous intraday quotes from Bloomberg are obtained for a comprehensive list of 25 UK ADR-stock pairs in existence through the period

April 4, 2011 through August 5, 2011. To make the price comparison between UK stocks and ADRs, the contemporaneous GBP/USD exchange rates are also obtained from Bloomberg. obtained. Details of the stock-ADR pairs can be found in Table 1. In total there are 130.7 million quotes.

Company Name	Exchange	ADR Ratio	Industry	Number of Stock Quotes	Number of ADR Quotes	Average Stock Liquidity / sec	Average ADR Liquidity / sec
ARM Holdings	NASDAQ	1:3	Technology	1,468,988	4,742,430	$5,268	$5,056
Astrazeneca	NYSE	1:1	Pharma. & Biotech	1,325,833	378,192	$2,626	$1,178
Aviva	NYSE	1:2	Nonlife Insurance	1,947,902	2,371,604	$1,399	$112
Barclays Bank	NYSE	1:4	Banks	2,867,024	4,041,689	$7,257	$12,498
BP	NYSE	1:6	Oil & Gas	3,434,790	8,416,350	$15,682	$51,347
Brit American Tobacco	AMEX	1:2	Tobacco	1,172,230	516,329	$4,447	$2,021
BT Group	NYSE	1:10	Fixed Line Telecom	1,704,348	2,387,739	$18,020	$1,048
Carnival	NYSE	1:1	Travel & Leisure	863,555	1,045,753	$659	$310
Diageo	NYSE	1:4	Beverages	1,712,417	8,283,247	$24,663	$1,958
GlaxoSmithKline	NYSE	1:2	Pharma. & Biotech	2,177,768	3,916,843	$19,190	$5,360
HSBC	NYSE	1:5	Banks	2,788,279	5,570,529	$17,793	$19,138
InterContinental Hotels	NYSE	1:1	Travel & Leisure	738,072	718,493	$871	$152
Lloyds Banking Group	NYSE	1:4	Banks	2,168,910	1,882,436	$2,674	$7,983
National Grid	NYSE	1:5	Multi-Utility	1,061,762	1,124,402	$24,199	$1,387
Pearson	NYSE	1:1	Media	597,922	530,397	$1,481	$114
Reed Elsevier PLC	NYSE	1:4	Media	922,594	398,341	$9,758	$406
Rio Tinto	NYSE	1:1	Mining	4,551,116	7,206,543	$4,094	$2,566
Royal Bank of Scotland	NYSE	1:20	Banks	1,554,139	1,065,509	$10,220	$4,265
Royal Dutch Shell - A	NYSE	1:2	Oil & Gas	3,342,762	8,546,975	$14,759	$10,212
Royal Dutch Shell - B	NYSE	1:2	Oil & Gas	3,881,505	5,929,302	$14,166	$6,309
Shire	NASDAQ	1:3	Pharma. & Biotech	851,216	999,404	$4,399	$1,069
Smith & Nephew	NYSE	1:5	Healthcare	603,953	885,867	$3,761	$893
Unilever	NYSE	1:1	Food and Staples	1,741,127	3,643,162	$7,137	$3,562
Vodafone Group	NASDAQ	1:10	Mobile Telecom	4,018,636	6,406,635	$149,268	$165,110
WPP	NASDAQ	1:5	Media	1,065,606	1,268,050	$12,246	$1,982
			Aggregates	**48,562,454**	**82,276,221**	**$15,041**	**$12,241**

Table 4.1. List of UK ADRs. The sample over the period April 4 through August 5, 2011. For each company, the ADR ratio shows how many UK shares trade as a single ADR. The 5th and 6th columns show the total number of quotes for each firm. The 7th and 8th columns express the average per-second liquidity (in dollars).

Table 1 describes the list of companies used in this sample. These are highly liquid large capitalization companies. The average per-second liquidity provided to the market is in excess of $15,000. Liquidity in this context refers to the volume of orders submitted to the market.

Each daily trading window was chosen to match the period where the London and New York markets were both open. Through the sample period, this window length is equal to two hours, specifically 9.30am to 11.30am New York time which corresponds to 2.30pm to 4.30pm London time. Estimates of broker commissions are obtained from research published by Investment Technology Group,[3] a Canada-based trade cost analysis firm, while estimates of exchange commissions are obtained from the London Stock Exchange and New York Stock Exchange websites.

To arrive at a final list of companies, a list of all UK stocks is compiled, which trade as ordinary shares and have ADRs which trade on the NYSE, NASDAQ, or AMEX. This process achieves two crucial refinements: Firstly, it filters out ADRs which are only issued over-the-counter at the depository bank, and are hence not suitable for pairs trading because they are not exchange-traded. Secondly, to ensure maximum liquidity, this work considers only ordinary shares and not preferred shares. Overall this process ensures a deliberately conservative approach in terms of the likelihood of finding exploitable mispricings.

This section applies a number of quote-filtering algorithms to filter out erroneous quotes. The process is aimed at ensuring that the reported quotes are executable in sufficient quantities to make the arbitrage operation economically viable. First, inspired by Schultz and Shive (2010) and Marshall et al. (2010), this section filters quotes where the ask price exceeds the bid price by 25% or more, in addition to bids being higher than asks. This section also filters bid-bid or ask-ask returns exceeding 25%.

The second filtering process is aimed at eliminating execution risk (see Kozhan and Tham (2012)). Quotes arrive irregularly and often quickly. An arbitrageur issuing a market order against a quote is not guaranteed to be executed at the quoted price, if another order was submitted a short time before. Execution risk is particularly important as one would expect the stock-ADR market to be heavily monitored, and heavily arbitraged. To address this risk, the quotes are processed

[3]Available at www.itg.com/news_events/papers

into one-second bins based on their volume-weighted average price (VWAP). To illustrate this, consider the following four quotes arriving in the 9:30:00 am to 9:30:01 am interval:

- Ask 100 shares at £85.30.
- Bid 180 shares at £85.20.
- Ask 150 shares at £85.35.
- Bid 120 shares at £85.25.

This one-second time interval carries a VWAP bid price of £85.22 with a depth of 300 shares, and a VWAP ask price of £85.33 with a depth of 250 shares. There are two advantages in this process. First this eliminates execution risk. Second, it provides the arbitrageur with the full depth of the limit order book, rebuilt every second. Overall, this process is conservative towards finding exploitable mispricings, since the arbitrageur is never trading at the best available prices unless all quotes arriving in any given 1-second window are equal.

4.3.2 Constructing the Mispricing Process

To arrive at a tradable mispricing or "spread" in the stock-ADR pair, the UK stock price is multiplied by the ADR Ratio for each company as shown in Table 1, and then translated to the US dollar. This gives us two dollar-denominated time-series: One ADR time-series, and one UK stock time-series translated by the contemporaneous spot exchange rate. With this in hand, this section extracts for each company the UK stock absolute premium/discount compared to the ADR, and express the result in dollars.

Let us denote the UK stock bid and ask prices as S_{bid} and S_{ask} respectively. Furthermore, the ADR bid and ask prices are denoted as A_{bid} and A_{ask} respectively. Finally the spot GBP/USD exchange rate (amount of dollars per

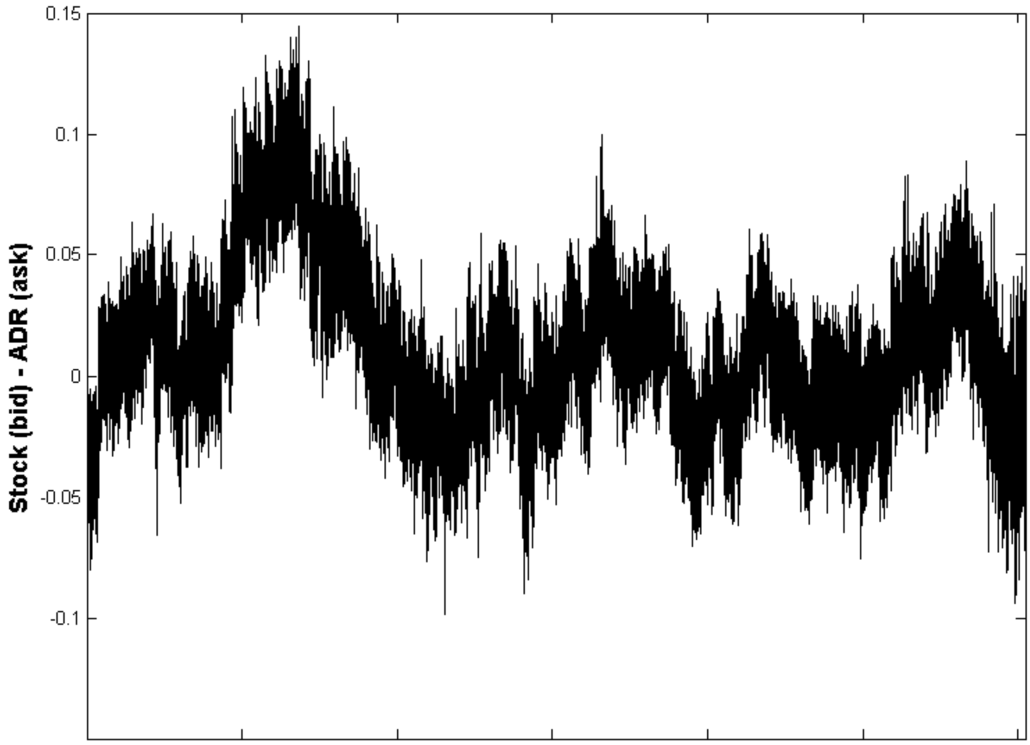

Figure 4.1: The Mispricing Process. This figure shows the stock-ADR mispricing process for Royal Dutch Shell A for each contemporaneous 1-second interval through the period April 4 - August 5, 2011. The vertical axis denotes the magnitude of mispricing in dollars. This chart plots the difference between the stock bid price and the ADR ask price.

British pound) bid and ask prices are denoted as FX_{bid} and FX_{ask} respectively. Let us now define the mispricing schedule M as:

$$M = \begin{cases} S_{\text{bid}} \cdot FX_{\text{ask}} - A_{\text{ask}} & \text{if } A_{\text{ask}} < S_{\text{bid}} \cdot FX_{\text{ask}} \\ A_{\text{bid}} - S_{\text{ask}} \cdot FX_{\text{bid}} & \text{if } A_{\text{bid}} > S_{\text{ask}} \cdot FX_{\text{bid}} \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

The example of the Royal Dutch Shell A mispricing schedule is given in Figure 4.1. The horizontal axis of the figure denotes 604,800 second-by-second observations spanning the entirety of the data period (April 4 - August 5, 2011). The vertical axis indicates the premium/discount of the stock relative to the ADR in dollars. Positive territory suggests the stock is overpriced relative to the ADR.

The first case in equation (1) is a situation where the UK stock is trading at a premium relative to its ADR. To initiate a position in the mispricing, the UK stock is sold short at the bid price, and the ADR is bought long at the ADR ask price. The second case in equation (1) is a situation where the UK stock is trading at a discount relative to its ADR. In this case, the opposite arrangement is executed to initiate a position in the mispricing. The case where $M = 0$ is a situation where there is overlap between the currency-adjusted bid/ask spread of the UK stock and its ADR. In this case, though the mispricing may not be strictly zero, it is nevertheless unprofitable. In other words, stock-ADR pairs are price-efficient whenever $M = 0$. Once a position is established in the mispricing, how an arbitrageur closes that position to realize a profit is precisely what differentiates pairs trading from direct-arbitraging. The pairs trader maintains the individual long-short positions until the mispricing converges, whereas the direct-arbitrageur converts the cheaper security into the dearer security, and sells it immediately onto the corresponding market.

It is important to note that in the context of equation (1), the bid and ask prices account for the entire structure of transaction costs, including exchange and brokers' fees, the bid-ask spread, market impact costs, taxes, etc. Therefore the notion of "bid" and "ask" in equation (1) are best interpreted qualitatively.

4.3.3 Direct Arbitrage vs Pairs-Trading: Costs and Risks

This section illustrates the main differences between the direct-arbitrage strategy employed elsewhere in the literature, and the pairs trading technique which forms the subject of this chapter. In view of equation (1), it is informative to estimate the total profits attainable from following a strategy of direct-arbitrage, i.e. treating the stock and ADR as interchangeable securities. The quantity M can be used to estimate total arbitrage profits. To see this, consider a trading horizon $0 = t_1 \leq t_2 \leq ... \leq t_n = T$ split into n intervals (e.g. n-minute intervals).

If M_{t_i} is the mispricing process observed at time t_i, then one can estimate total profits by the simple sum,

$$\text{Total Profits} = \sum_{i=1}^{i=n} M_{t_i}. \tag{2}$$

Profitable direct arbitrage occurred just once in the entire sample, specifically for Aviva Plc on April 15, 2011 at 14:30:01 GMT. This result is in contrast to Suarez's (2005) data sampled roughly a decade before ours. However, this disappearance of such "riskless" pricing anomalies can be attributed in part to the advent of algorithmic trading and the automatic dissemination of quotes in recent years.[4]

In contrast, pairs trading relies on the converge of a mispricing once a long-short position has been established. Though these results suggest that convergence typically takes place in a matter of minutes, pairs trading nevertheless faces an altogether different set of risk factors relative to the direct-arbitrage approach (see Pontiff (1996)). If a mispricing widens after a position is established, the pairs trader stands to make a capital loss on her position. If a margin call is issued, the pairs trader must liquidate part of her long-short portfolio at a highly unfavorable point in time. As for direct-arbitrage, the ADR conversion facility can be temporarily halted with very little notice. If an arbitrageur is in the process of conversion to take advantage of a price discrepancy, this can leave her with excess inventory: an unwanted open position in one or both securities. This point has received relatively little attention in the literature, including papers which directly perform this type of arbitrage operation such as Suarez (2005).

Another fundamental difference between pairs trading and direct-arbitraging involves the costs associated with the operation. Pairs trading involves four transactions in total: two to take a position in the mispricing, and two to

[4]See e.g. Chordia et al. (2011), Chaboud et al. (2009), and Hendershott et al. (2011).

unwind the position once price parity is restored. This incurs costs of two full bid-ask spreads, plus associated commissions. Furthermore, since The Finance Act (1986), the UK imposes a Stamp Duty Reserve Tax (SDRT) of 0.5% on the purchase of UK stocks.[5] The selling of UK stocks does not incur SDRT. Direct-arbitrage, on the other hand, involves one purchase of the cheaper security and one sale of the expensive security, a total of two half-spreads incurred in cost. The cheaper security needs to be converted to the more expensive security to complete the operation. Custodian banks (major ones include Bank of New York Mellon and JP Morgan) operating ADR schemes charge a fee of $0.05 per conversion. Further, if a UK stock is converted into an ADR, the applicable rate of UK SDRT is levied at 1.50%. This is to allow for the fact that future trades in the newly-created ADR will not incur UK SDRT. With this in mind, an interesting question arises: Which arbitrage method is cheaper to implement?

The cheaper method to arbitrage depends on the full structure of transaction costs faced by the arbitrageur. Intuitively, consider that only the largest and most liquid UK companies can afford maintaining ADR programs. Shares in these companies are likely to exhibit the smallest bid-ask spreads, which compensates for the fact that pairs trading involves four transactions, whereas direct-arbitrage involves two. By obtaining precise estimates of the full spectrum of transaction costs in each case, this work suggests that pairs trading is cheaper to implement. Hence exploitable mispricings are more visible to the pairs trader. The demonstration of this point is the central aim of the following analysis.

4.3.4 Transaction Costs

It is informative to assess the costs associated directly with trading. Table 4.2 shows for each company and each ADR in the sample, the median, 5th and 95th percentiles of all computed bid-ask spreads.

[5]Based on the Oxera (an economics consultancy) report "Stamp Duty: its Impact, and the Benefits of its Abolition" - May 2007. Available at (www.oxera.com).

	Median bid-ask spread (bp)		5th Percentile bid-ask spread (bp)		95th Percentile bid-ask spread (bp)	
Company Name	**UK Stock**	**ADR**	**UK Stock**	**ADR**	**UK Stock**	**ADR**
ARM Holdings	9.19	3.54	7.95	2.43	18.15	6.78
Astrazeneca	2.43	2.02	1.31	1.91	4.25	3.89
Aviva	4.73	17.79	2.19	7.25	9.49	28.49
Barclays Bank	3.91	5.79	1.70	5.03	7.73	7.70
BP	2.47	2.26	1.05	2.05	5.08	3.31
Brit American Tobacco	3.36	4.85	1.54	1.67	5.30	8.53
BT Group	5.19	11.83	4.50	3.97	10.53	19.18
Carnival	8.19	12.37	3.83	5.02	13.02	18.43
Diageo	7.97	3.55	7.29	1.23	15.53	5.90
GlaxoSmithKline	3.88	2.34	3.37	2.16	7.71	3.05
HSBC	2.85	2.00	1.37	1.75	4.64	3.54
InterContinental Hotels	8.57	9.77	7.08	4.72	17.11	18.73
Lloyds Banking Group	3.74	31.80	1.35	24.97	7.99	38.10
National Grid	8.34	4.14	7.53	2.00	16.13	8.15
Pearson	8.86	10.75	7.74	5.29	17.81	20.90
Reed Elsevier PLC	9.21	13.66	8.02	5.70	18.33	19.37
Rio Tinto	3.15	2.67	1.13	1.34	5.56	4.62
Royal Bank of Scotland	5.33	7.65	2.36	7.00	10.37	14.71
Royal Dutch Shell - A	4.45	1.44	2.15	1.25	6.69	2.97
Royal Dutch Shell - B	4.32	2.52	2.09	1.33	6.01	4.19
Shire	5.50	5.55	4.48	2.11	11.04	10.12
Smith & Nephew	7.72	9.54	6.54	3.78	16.05	16.23
Unilever	5.17	3.12	4.60	3.01	10.37	4.30
Vodafone Group	3.10	3.69	2.64	3.41	6.20	3.94
WPP	6.85	11.45	6.04	6.07	14.24	17.94

Table 4.2 Bid-ask spreads. The median, 5th and 95th percentiles of bid-ask spreads for each stock-ADR pair.

What is interesting in table 4.2 is that the median bid-ask spreads seem significantly lower than those assumed elsewhere in the literature (De Jong et al. (2009) for example, calculate a median of 25 basis points (bp)). At first this seems surprising, but this observation can be explained as follows: Firstly, Easley and O'Hara (1987) show that bid-ask spreads are inversely related to trade size. With this in mind, it is reasonable that the calculated median bid-ask spreads here are significantly lower than those assumed elsewhere in the literature, especially since the sample consists of shares in the top two quartiles of the FTSE 100 index. Within the sample, the correlation between UK bid-ask spreads and average per-minute quoted volume is -16%, and the corresponding figure for ADRs is -40%. Secondly, Hendershott and Menkveld (2011) provides evidence that the advent of algorithmic trading, now accounting for some 70% of US equity trades, has improved liquidity particularly in large stocks. One effect of this development is the narrowing of spreads.

Further to the bid-ask spread, other direct trading costs are presented which affect retail investors, namely brokerage commissions, exchange trading fees, and UK taxes. Table 4.3 presents a summary of these costs (in bp).

Broker and Regulatory Costs	Amount (bp)
UK Broker Fee (Stocks)	11.35
US Broker Fee (ADRs)	9.40
London Stock Exchange Fee	1.00
New York Stock Exchange Fee	1.00
UK Stamp Duty (Buying Stocks)	50.00
UK Stamp Duty Reserve (ADR Conversion)	150.00

Table 4.3 Institutional frictions. Direct transaction costs associated with trading in both UK stocks and ADRs.

The exchange trading fees are obtained from published reports on both the LSE and NYSE websites. Both exchanges frequently revise their cost structures, often switching from flat per-trade fees to variable tiered fees. At no point during the sample period do the costs exceed 1bp, hence one can assume this figure conservatively.

An inspection of Tables 2 and 3 reveals that pairs trading is by far the cheaper method of arbitraging than the direct-arbitrage approach. Median spreads are rarely over 10bp, and spreads would have to exceed 100bp for the direct-arbitrage approach to be more cost-efficient. The next section explores a further cost advantage of pairs trading, namely the ability to circumvent UK SDRT, a feature which remains impossible for the direct-arbitrageur.

4.3.5 Different Investors = Different Costs

Contracts-for-difference (CFDs), are SDRT tax-free derivative instruments which give an investor exposure (long or short) to price movements of an underlying security, e.g. a stock. Avoidance of UK SDRT may be attributed to the surge in CFD trading on UK stocks. CFD transactions have increased from accounting for 10% of UK equity transactions in 2001, to 35% in 2007. The largest group of users of CFD contracts are financial institutions and hedge funds, followed by retail investors.[6]

The exercise presented here profiles large institutional investors trading CFDs. This type of investor benefits from the lowest overall transaction costs, since CFD positions are exempt from UK SDRT. A CFD investor may be retail or institutional, and one would reasonably expect institutional CFD traders to benefit from far lower costs than retail CFD traders.

The view espoused in this chapter is that CFD traders at large investment banks or hedge funds are most able to exploit stock-ADR mispricings, because they circumvent UK SDRT and face the smallest broker commissions. The representative arbitrageur in this study is a CFD trader.

4.3.6 The Pairs-Trading Strategy

The pairs trading strategy employed requires an arbitrageur to specify two parameters, namely the entry bound and the exit bound. The entry bound denotes the level of mispricing above or below parity at which a pairs trade is initiated, while the exit bound signals the termination of the arbitrage position, when the mispricing in the stock-ADR pair reaches parity.

[6] Financial Services Authority, Disclosure of Contracts of Difference: Consultation and Draft Handbook Text (November 2007), pp. 11-12.

Different works specify different methods toward selecting these bounds. For example, Gatev et al. (2006) follow a statistical approach where two-standard deviations from a historical mean (the latter interpreted to be the point of price parity) is an entry level. A mispricing level returning to its historical mean is then deemed to be at parity, triggering the unwinding of the pairs trade. On the other hand, De Jong et al. (2009) follow an approach based on economic fundamentals. They specify a range of entry/exit bounds, the exit bound based on the theoretical point of price parity (see Friedman (1953)) - typically a mispricing level of zero. Consistent with the latter, the entry/exit bounds are based on fundamentals. This has the advantage that the entire data period can be considered out-of-sample. Unlike the statistical approach, the approach based on fundamentals does not require a historical "training" period as a point of reference from which to generate trading signals. Akin to De Jong et al., this work considers a wide range of entry/exit bounds in order to assess the sensitivity of arbitrage returns to the profit target set by arbitrageurs.

This work imposes the conservative constraint of assuming that the investor has no access to short-sale proceeds. Further, it is assumed that the investor must be able to cover the short position at all times in cash, a requirement similar to D'Avolio (2002). Hence a short position immobilizes its value in the investor's cash, in the same way a long position does. As a simple example, consider two stocks both trading at \$50. If an investor has a \$100 capital base, the largest long/short portfolio that can be formed is 1 stock short and 1 stock long. This requirement is clearly stricter than the US Regulation T 50% margin requirement as described by De Jong et al. (2009). In reality, this situation may occur if, for example, the investor utilizes different brokers for the stock and ADR positions. In this case, she may not be able to use the stock in which she is long, as consideration for collateral in the short position.

To formalize the pairs trading strategy algebraically, let us start by defining terms: In the trading horizon $0 \leq t \leq T$, the UK bid and ask stock prices are translated (subscripts removed for neatness) S_t into dollars by multiplication with the corresponding bid or ask GBP/USD spot rate FX_t as per equation (1). If the ADR price at time t is denoted by A_t, one can construct a time-series of the absolute stock-to-ADR mispricing, called R_t, as follows:

$$R_t = S_t \cdot FX_t - A_t. \tag{3}$$

Equation (3) is what is used to produce a mispricing process such as that in Figure 4.1. The parameter F_t denotes the level at which the UK stock trades at parity with the ADR. It is proposed that the true price parity in the stock-ADR pair occurs not at $R_t = 0$, but instead when the stock trades at a 25bp discount relative to the ADR. The reasoning behind this is as follows: If one assumes that investors in the UK and US exhibit homogeneous preferences, so that they care only about the cheaper location to trade, and indifferent towards owning a stock or an ADR, it follows that in order to tempt them to buy UK stocks over ADRs, the UK stock will have to trade at 50bp or more below its ADR price. This overcomes the incidence of the SDRT. All R_t for the data are adjusted to incorporate this fact.

If C_t refers to the total costs of completing a round-trip pairs trade, it remains only to define the entry/exit bounds for the pairs trade. This section posits that the arbitrageur enters a pairs trade at a variable arbitrary magnitude a (in dollars) away from parity, and exits when she overcomes trading costs, plus a variable profit margin (in dollars). Let us denote this profit margin by E.

Finally, if W_t is the arbitrageur's wealth at time t in US dollars, then one can denote the arbitrageur's opening position in the stock at time t as N_t^{Stock} as follows:

$$(4) \qquad N_t^{Stock} = \begin{cases} \frac{0.5 \cdot W_t}{S_t \cdot FX_t} & \text{if } R_t < a \\ -\frac{0.5 \cdot W_t}{S_t \cdot FX_t} & \text{if } R_t > a \\ 0, & \text{otherwise.} \end{cases}$$

Similarly, the arbitrageur's open position in the ADR, N_t^{ADR} is as follows:

$$(5) \qquad N_t^{ADR} = \begin{cases} -\frac{0.5 \cdot W_t}{A_t} & \text{if } R_t < a \\ \frac{0.5 \cdot W_t}{A_t} & \text{if } R_t > a \\ 0, & \text{otherwise.} \end{cases}$$

Taken together, equations (4) and (5) suggest that if a mispricing is wide enough to at least compensate the costs of completing a pairs trade, the arbitrageur will allocate half her capital to going long the cheap security, half to short the expensive security. Thus she has a dollar-neutral portfolio. Once the arbitrageur has established a long-short position, she unwinds this position when the magnitude of the mispricing falls below a level which compensates for the costs associated with the trade, plus a profit margin E (in dollars). Otherwise, the position remains open. Formally, at each subsequent time-step $(t, t+1]$, the arbitrageur's position in the stock satisfies:

$$(6) \qquad N_{t+1}^{Stock} = \begin{cases} 0 & \text{if } \|R_{t+1} - F_{t+1}\| \leq C_{t+1} + E \\ N_t^{Stock} & \text{otherwise,} \end{cases}$$

and her position in the ADR satisfies:

$$N_{t+1}^{ADR} = \begin{cases} 0 & \text{if } \|R_{t+1} - F_{t+1}\| \leq C_{t+1} + E \\ N_t^{ADR} & \text{otherwise.} \end{cases} \tag{7}$$

Intuitively speaking, equations (4)-(7) describe a simple threshold pairs trading strategy. If the magnitude of the mispricing exceeds a pre-set level, the arbitrageur shorts the expensive security and buys the cheap security. This position is held open until the mispricing narrows just enough to compensate for trading costs plus a small pre-set profit margin per unit of long/short security traded. The next section presents the empirical results of the pairs trading exercise.

4.4 - Results

This section presents the results and analysis of the returns to the pairs trading strategy. An arbitrageur is chosen to represent the proprietary trading desks at large financial institutions with direct access to the CFD market. One would reasonably expect trading costs to be a major impediment to arbitrage in markets with near-perfect substitutes (see section 4.3). Further, this choice is supported by Werner and Kleidon's (1996) suggestion that stock-ADR arbitrage is most heavily undertaken by large institutional investors.

First, this section presents the returns to arbitrage across each of the 25 stock-ADR pairs. Then, the entry/exit arbitrage bounds are varied. This process helps assess the sensitivity of arbitrage returns to the profit target set by the arbitrageur. It also helps understand the risks inherent in pairs trading, with regards to holding costs. Third, this section analyzes the risk characteristics of arbitrage returns, by way of calculating trade half-lives. Finally, this section discusses the wider risks inherent in stock-ADR arbitrage trading, and explain why mispricings persist in such a developed market.

4.4.1 Individual Returns (The $0.12/$0.01 case)

The arbitrageur sets her entry level at a point whereby the magnitude of absolute mispricing exceeds $0.12. The position is then unwound when the mispricing falls below a level which covers the bid-ask spread and exchange fees (see Table 4.3) plus a $0.01 profit margin per long/short unit.

Table 4.4 shows the results of the trading exercise (Equations 3-7) for each individual stock-ADR pair. Overall, 646 long-short positions across all pairs were initiated and completed. The median daily profit is 0.59bp, and the median time a position remains open is 507 seconds. These figures suggest that stock-ADR arbitrage are characterized by the incidence of small and quickly-disappearing windows of opportunity.

Around 5% of arbitrage positions were held for 87 seconds or less, illustrating the fact that arbitrage opportunities in markets for close substitutes disappear quickly. On the other hand, one trade (Lloyds Plc) was held for the entirety of the trading window. This may be in part due to the simplistic nature of the threshold strategy employed, or characteristic of the risks inherent in this type of arbitrage (see section 4.4.4).

Figure 4.2 illustrates the cumulative daily wealth at the portfolio level (i.e. averaging across the 25 stock-ADR pairs). Interestingly, cumulative wealth exhibits low volatility and rarely decreases at the daily level. This fact illustrates the mitigation of idiosyncratic risks from individual stock-ADR pairs across the exhaustive sample.

Company Name	long stock short adr	long adr short stock	Mean Daily Return (bp)	Max Daily Return (bp)	Min Daily Return (bp)	Median secs	5th Percentile secs	95th Percentile secs
ARM Holdings	33	0	0.38	16.13	-10.10	524	96	8,398
Astrazeneca	0	0	0.00	0.00	0.00	N/A	N/A	N/A
Aviva	2	0	0.21	12.98	0.00	107	105	109
Barclays Bank	0	0	0.00	0.00	0.00	N/A	N/A	N/A
BP	14	0	0.27	11.33	-4.06	187	18	41,594
Brit American Tobacco	11	0	0.22	13.74	-7.86	339	55	439,368
BT Group	49	0	0.90	17.15	-6.45	1,085	106	18,308
Carnival	5	13	0.60	12.38	-3.97	824	1	4,373
Diageo	1	19	0.44	8.94	-11.57	7,470	241	62,434
GlaxoSmithKline	2	2	0.11	5.79	0.00	1,292	87	6,343
HSBC	0	100	1.90	13.58	-5.51	457	37	26,899
InterContinental Hotels	4	0	0.36	22.10	-2.79	224	3	5,228
Lloyds Banking Group	0	1	-0.23	30.00	-16.77	604,793	604,793	604,793
National Grid	21	2	0.75	19.66	-11.68	2,537	49	90,451
Pearson	0	13	1.07	38.75	-11.85	4,466	524	155,951
Reed Elsevier PLC	0	20	0.59	9.66	-8.70	6,360	1,457	34,668
Rio Tinto	83	18	1.29	12.26	-4.38	199	7	7,674
Royal Bank of Scotland	0	17	1.68	49.67	-9.70	879	132	30,586
Royal Dutch Shell - A	0	115	2.11	19.08	-3.27	295	21	3,185
Royal Dutch Shell - B	11	0	0.30	6.76	-4.50	530	24	500,213
Shire	3	29	0.67	10.11	-7.46	2,517	7	42,100
Smith & Nephew	15	4	0.17	35.82	-20.24	1,413	4	217,351
Unilever	0	3	0.11	7.50	-2.15	4,766	3,225	14,426
Vodafone Group	6	7	0.46	12.60	-2.12	1,552	92	11,314
WPP	0	23	0.32	11.90	-9.04	3,733	361	84,974
Aggregates/Average	**260**	**386**	**0.59**	**15.91**	**-6.57**	**507**	**87**	**30,586**

Table 4.4 Trading results. Trading results for the \$0.12/\$0.01 case. The first column shows the number of positions initiated whereby the stock traded at a discount. Similarly the second column illustrates ADR discount. The next three columns show return statistics for each stock-ADR pair. The last 3 columns show the median, 5th, and 95th percentiles of seconds each trade was kept open.

A natural question to ask is whether the incidence of exploitable mispricings can be predicted. Are there certain times each day where mispricings are more likely to occur? To answer this question, let us split the daily trading window into 1-minute intervals, totalling 120 minutes per day. A plot is presented of how many times within each 1-minute interval an arbitrage position is initiated. The results are presented in Figure 4.3.

Besides a mild concentration of exploitable mispricings occuring in the first and last minutes of each trading day, there does not seem to be an apparent pattern as to when mispricings occur.

The next section examines the sensitivity of arbitrage returns to the profit target and entry thresholds.

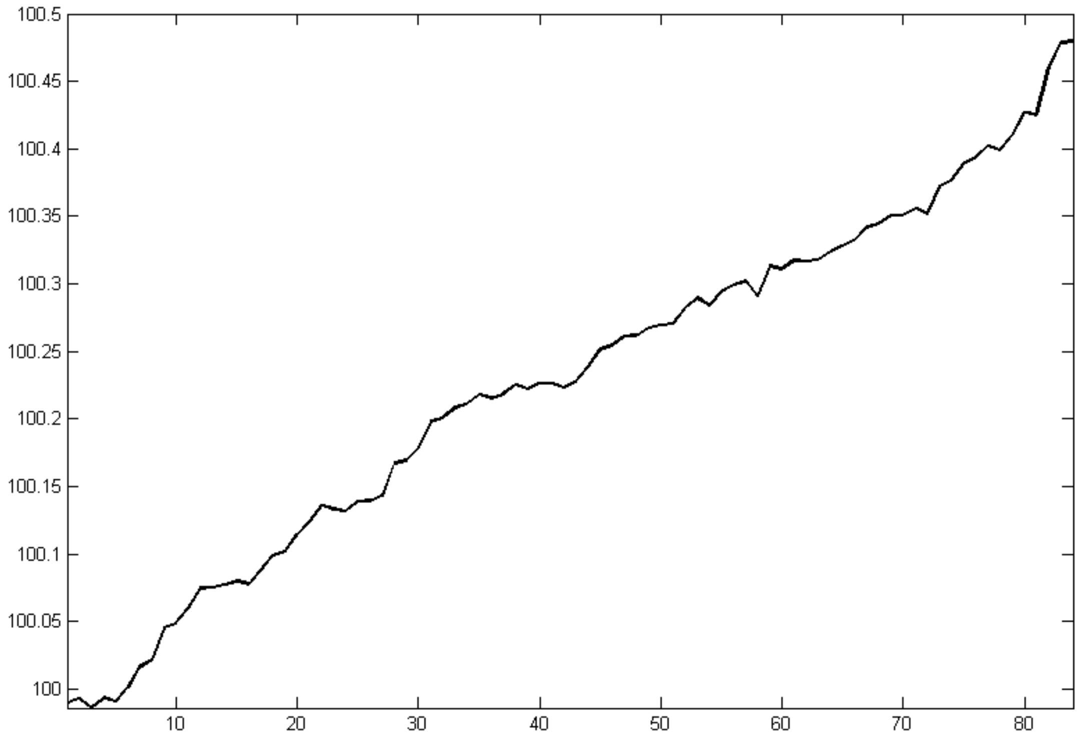

Figure 4.2: Cumulative Portfolio Wealth. This figure shows the cumulative daily wealth at the overall portfolio level, assuming an initial wealth $W_0 = \$100$ on the vertical axis. The horizontal axis denotes the number of trading days.

4.4.2 Robustness to Varying the Entry/Exit Levels

Levels at which arbitrage positions are initiated and unwound are set arbitrarily. How sensitive is the trading performance to the values at which these bounds are set?

This section presents an analysis to assess the degree to which trading performance is affected by these set levels. In doing so, the trading exercise is repeated but considers 9 combinations of entry/exit bounds. The entry are set to where the mispricing exceeds \$0.12, \$0.15, and \$0.18 away from parity. For each of these levels, it is imposed that the arbitrageur demands \$0.01, \$0.02, and \$0.03 as a profit margin per long/short unit, once all costs have been recovered. Table 4.5 presents the trading results.

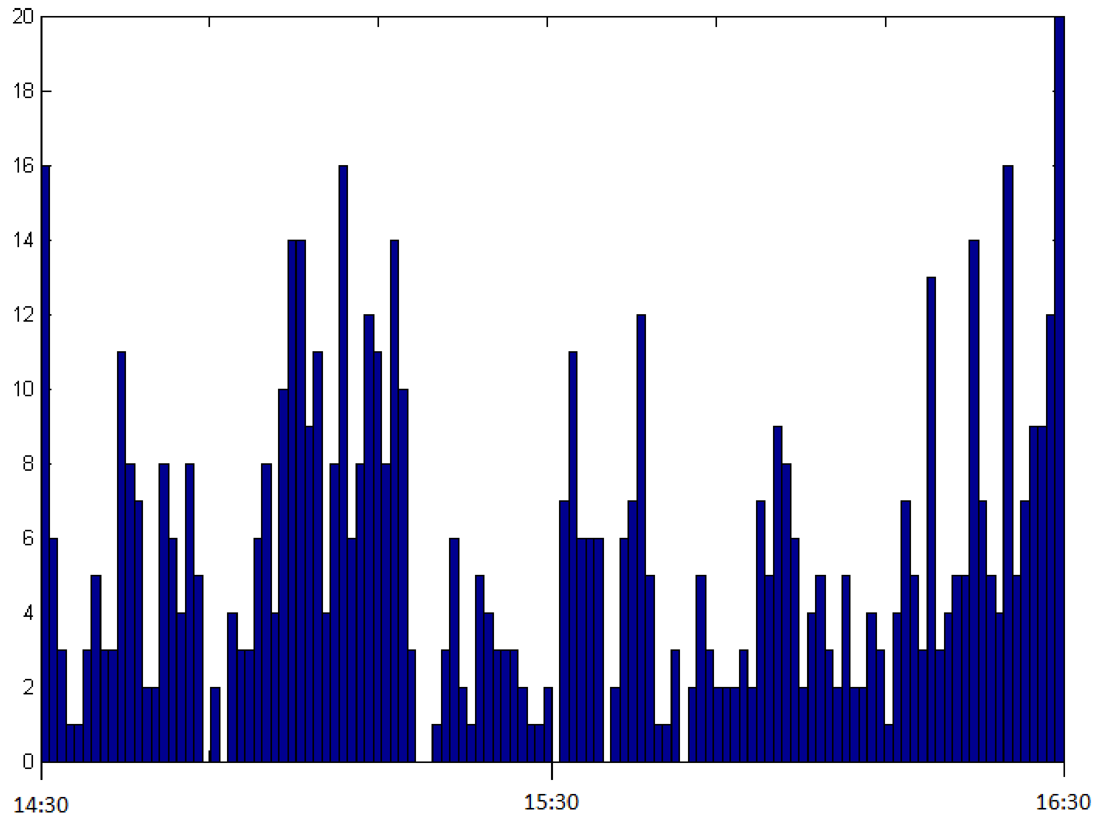

Figure 4.3: Incidence of Mispricings. Total number of arbitrage positions initiated across each 1-minute contemporaneous interval of the trading day. The horizontal axis shows the GMT interval during which both US and UK exchanges trade.

What is immediately striking to us in Table 4.5 is the large number of arbitrage positions generated. Excess returns are economically small but statistically significant. One can posit that this type of arbitrage may be lend itself well to the utilization of leverage to magnify returns. However, it can also seen from Table 4.5 that the expected trade duration is highly variable, and increases directly with the profit target set by the arbitrageur. The implications of this uncertainty toward the persistence of mispricings and the general price-correction mechanism is discussed in detail in section 4.4.4. Further, Figure 4.3 above suggests that the incidence of profitable disequilibria is scattered randomly throughout the trading day. Having to keep capital readily available to exploit quickly-disappearing opportunities may act as a deterrent to arbitrageurs (Shleifer and Vishny (1997)).

Profit Target	Annual Excess Return (%)	% of Daily Returns < 0	Median Trade Time (s)	St.Dev Trade Time (s)	# Trading Positions
			Entry Bound = $0.12		
$0.01	1.45%	19.28%	561	41,921	646
$0.02	1.14%	22.89%	1,569	58,923	345
$0.03	1.02%	30.12%	4,041	77,824	217
			Entry Bound = $0.15		
$0.01	0.84%	27.71%	614	24,799	465
$0.02	0.77%	34.94%	1,224	34,825	267
$0.03	0.73%	36.14%	2,734	54,107	166
			Entry Bound = $0.18		
$0.01	0.46%	30.12%	476	25,242	328
$0.02	0.46%	33.73%	1,125	32,910	203
$0.03	0.44%	30.12%	1,733	39,144	148

Table 4.5: Returns for Different Investors. This table documents the effect of varying the entry/exit bounds on the returns to stock-ADR arbitrage. The first column shows the annual return in excess of the risk-free rate. The second column shows the percentage of negative returns. The third and fourth columns show the median and standard deviation of trade duration (in seconds). The final column shows the total number of long/short positions established and closed.

Overall, Table 4.5 shows that larger disequilibria disappear more quickly, which is consistent with the notion that in the absence of fundamental risk factors, larger pricing anomalies are inherently less risky to exploit (De Jong et al. (2009)) and act as greater incentives for arbitrage capital (Grossman and Stiglitz (1976, 1980)).

4.4.3 The Dynamics of Mean-Reversion

At this point, it has been shown that pairs trading is an alternative mechanism to direct arbitrage, by which stock-ADR pricing misalignments can be profitably exploited. However, in contrast to direct arbitrage, pairs trading invariably involves keeping a long/short position open for a period of time. This clearly involves the risk that the stock-ADR prices may take a very long time to converge, or may even not converge at all. This section explores the dynamics of price convergence between stock-ADR pairs.

Parity in this chapter's stock-ADR relationship can be expressed in the form of a real exchange rate, $\frac{FX_t \cdot S_t}{A_t} = 1$, where the price levels are the prices of the stock and ADR respectively. This is useful because there is an established literature on mean-reversion dynamics for real exchange rates, e.g. Osbat et al. (2003), Cheung et al. (2004), Nam (2011), Astorga (2012). These papers use cointegration and an error correction model (ECM) to derive the half-life of Purchasing Power Parity. This method is adopted to compute the half-life of stock-ADR price parity.

The specific variant of ECM that is used for the analysis presented here is from Brooks (2008). An Engle-Granger cointegration (Engle and Granger (1987)) model is employed across all 25 stock-ADR pairs. Let us begin with the following transformations:

$$s_t = \log(FX_t \cdot S_t)$$
$$a_t = \log(A_t)$$

Now, an augmented Dickey-Fuller (ADF) test is used to establish that all s_t and a_t are highly significantly integrated of order $\sim$ I(1). One can then estimate the cointegrating regression using OLS:

$$s_t = \alpha a_t + u_t. \tag{8}$$

The estimates for alpha $\hat{\alpha}$ are displayed in the first column of Table 4.6. As one would expect, all are very close to unity.

The residuals, $\hat{u}_t$, are highly significantly integrated of order $\sim$ I(0) across all 25 stock-ADR cointegrating relationships. Thus let us proceed to estimate an Error-Correction Model (ECM) as follows:

$$\Delta s_t = \beta_1 \hat{u}_{t-1} + \beta_2 \Delta a_t + v_t, \tag{9}$$

where $\hat{u}_{t-1} = s_{t-1} - \hat{\alpha} a_{t-1}$ is the error correction term. The coefficients in (9) are estimated using OLS. In particular, this section is interested in the speed of mean-reversion, β_1, whose values and corresponding t-statistics across all 25 stock-ADR pairs are provided in the second and third columns, respectively, of Table 4.6. The magnitude of all t-statistics is very large, indicating that all stock-ADR pairs exhibit strong mean-reverting behaviour.

The fourth column of Table 4.6 documents the R^2 for each of the 25 ECM regressions, with average values of around 80%. The final column of Table 4.6 reveals the half-life (speed of reversion), β_1, of these mispricings which are derived using the specification of Osbat et al. (2003):

$$\text{Half-life} = \frac{\log(0.5)}{\log(1-\beta_1)}.$$

The half-life numbers show that an average stock-ADR mispricing reduces its size by half in under 7 minutes. This result is consistent with earlier results regarding the expected duration of trades (e.g. Tables 4 and 5). Overall, high mean-reversion rates, coupled with short half-lives, suggest that the risk faced by pairs traders of a widening mispricing is negligible.

4.4.4 On the Limits of Arbitrage

This section considers why exploitable price disparities arise in the first place, and why they persist in such a highly developed market with no capital controls and no restrictions on cross-border arbitrage. In the context of stock-ADR pairs trading, this issue essentially boils down to two main questions: Firstly, are the trading strategies presented in this chapter implementable in practice? Secondly, are there risks which act as disincentives to arbitrage?

Company Name	Cointegrating Vector (α)	Reversion Speed (β_1)	t-Statistic for β_1	R^2	Half-Life (secs)
ARM Holdings	0.9986	-0.0031	-30.97	76.63%	222
Astrazeneca	0.9996	-0.0028	-28.92	87.99%	250
Aviva	0.9967	-0.0040	-34.13	72.77%	173
Barclays Bank	0.9985	-0.0023	-25.88	84.21%	303
BP	0.9993	-0.0013	-19.99	82.50%	529
Brit American Tobacco	0.9990	-0.0005	-12.41	83.68%	1,354
BT Group	0.9983	-0.0042	-35.34	81.98%	166
Carnival	0.9992	-0.0021	-25.12	78.46%	326
Diageo	1.0000	-0.0069	-45.86	80.70%	101
GlaxoSmithKline	0.9995	-0.0018	-23.35	81.25%	379
HSBC	1.0001	-0.0098	-54.46	84.81%	71
InterContinental Hotels	0.9980	-0.0017	-22.53	78.09%	412
Lloyds Banking Group	1.0098	-0.0010	-16.67	79.78%	711
National Grid	0.9990	-0.0008	-15.80	78.46%	817
Pearson	1.0003	-0.0013	-19.77	72.21%	526
Reed Elsevier PLC	1.0000	-0.0241	-85.19	81.63%	29
Rio Tinto	0.9994	-0.0010	-17.09	88.30%	719
Royal Bank of Scotland	0.9994	-0.0025	-27.21	81.55%	280
Royal Dutch Shell - A	1.0000	-0.0108	-57.65	85.83%	65
Royal Dutch Shell - B	0.9990	-0.0018	-23.54	87.29%	383
Shire	1.0000	-0.0095	-53.63	84.53%	73
Smith & Nephew	0.9992	-0.0008	-15.09	80.51%	908
Unilever	0.9991	-0.0013	-20.03	77.17%	521
Vodafone Group	0.9987	-0.0010	-17.49	80.20%	667
WPP	1.0000	-0.0106	-56.64	83.29%	66

Table 4.6: Dynamics of mean-reversion. For all stock-ADR pairs, the first column shows the cointegrating vector α from equation (8). The second and third columns show, respectively, the speed of mean-reversion β_1 from equation (9), and the corresponding t-statistic. The fourth column shows the R^2 value for model (9), and the final column shows the half-life of deviations from parity.

Regarding the first question, Kozhan and Tham (2012) point out the significance of execution risk. Market orders are not guaranteed to be executed at the quoted price, if a quote is either removed or has already been filled. The question then is, what happens if an order gets filled at the next best price? The data sample relies on volume-weighted average price quotes rebuilt every second of the trading overlap (see section 4.3.1). As a result, the orders submitted in the context of this analysis are never executed at the best available price. This conservative approach mitigates execution risk in the analysis. Furthermore, CFD contracts (used in this study) are widely-traded instruments which circumvent UK SDRT.

Some 30% of UK equity trades are done through CFDs, and the primary user of CFDs are large financial institutions and hedge funds (see section 4.3.5).

Liquidity is a source of risk for which arbitrageurs need compensation. The sample includes 25 of the most liquid UK companies, with liquidity on both the stock and ADR side averaging around $12,000 per second. Furthermore, the process of volume-weighting quotes also factors in market impact costs. An arbitrageur can execute the entire depth of the order book each second at the prices presented in the analysis. This work conjectures that relaxing this process would magnify returns, but leave arbitrageurs open to liquidity risk and excess inventory. Overall, it is posited that the mispricings presented in this chapter are implementable in practice.

Now, this section turns its attention to the risks associated with arbitrage trading. As Shleifer and Vishny (1997) and De Long et al. (1990) point out, pairs trading is inherently susceptible to noise trader risk. Unlike the direct arbitrage approach employed by Suarez (2005), pairs trading involves keeping open a long/short position for some time. The mispricing may worsen significantly, resulting in a capital loss to the investor, and possibly a portfolio liquidation at a highly unfavorable point in time. However, stock-ADR pairs trading has a unique way of mitigating this risk entirely. How? If an arbitrageur initiates a long/short position in a mispricing, and the mispricing worsens, the arbitrageur can exit this position by converting the security held long into the security held short to close out the position. The loss in that case is limited to the costs associated with the conversion process (see table 4.3). Note that this feature is unique to stock-ADR pairs trading, and is not possible in pairs trading cointegrated stocks (Gatev et al. (2006)) or even dual-listed companies (De Jong et al. (2009)).

Pontiff (1996, 2006) argues that there are significant holding costs associated with convergence trading. Of these, the most prevalent in the context of the work presented here is the uncertainty toward price convergence. Looking at Table

4.5 reveals some interesting insights: When arbitrageurs increase their profit target per long/short unit even slightly from \$0.01 to \$0.02, the median duration they would expect to keep positions open at least doubles. Adding another 1 cent to the profit target at least triples the expected trade duration. Also, the standard deviation of expected trade durations increases 20% for each additional cent demanded. In the extreme, one single position in Lloyds TSB remained open throughout the entire sample period. The mispricing process in that single case never reached its profit target of \$0.01. On the other hand, widening the trade entry threshold results in shorter trade durations but a decreased number of trades.

Taken together, these observations suggest that a major disincentive to arbitrage stems from the uncertainty regarding trade duration. This is not surprising in an environment with near-perfect substitutes. The uncertainty relates directly to the per-trade profit target demanded by arbitrageurs. To this end, Abreu and Brunnermeier (2002, 2003) and Kondor (2009) argue that synchronization problems and the hope of "even better opportunities tomorrow" can lead to the persistence of mispricings. Note however, that these observations are consistent with Grossman and Stiglitz (1976, 1980) who argue that mispricings are necessary to incentivise arbitrageurs to correct pricing inefficiencies.

One final explanation regarding the existence of mispricings is inspired by Schultz and Shive (2010), who ask "Do pricing disequilibria actually constitute mispricings?". The authors examine dual-listed companies whose shares carry different voting rights. In such a context, a deviation from price parity need not necessarily indicate a mispricing, as different classes of shares carry premiums associated with the conditions of ownership. In the context of this chapter, besides the currency denomination, stocks and ADRs are identical securities in every regard (see Pulatkonak and Sofianos (1999)). It is therefore likely that a deviation from parity constitutes a mispricing.

4.5 - Conclusion

This chapter reveals an important price-correcting mechanism in the stock-ADR market that had been overlooked by the existing literature. In so doing, this chapter essentially achieves two things. First, it provides compelling evidence that actual parity enforcement mechanism in UK stock-ADR parity is pairs trading and not the direct arbitrage approach, as is ubiquituously assumed in this literature. Second, it establishes that the risk associated with pairs trading in this context is very low, is bounded by the ability to implement direct arbitrage on an existing pairs trade and is thus limited to concerns about slow or non-convergence of prices, rather than fears of their divergence.

When applied to a comprehensive sample of 25 UK stock-ADR pairs, pairs trading returns 1.45% in excess of the risk-free rate. Exploitable opportunities appear randomly throughout the interval in which the UK and US markets trade contemporaneously, with a mild concentration at the US open and UK close. In contrast, only a single profitable trading opportunity can be observed using the direct arbitrage approach over the same period.

Stock-ADR arbitrage entails no fundamental risk since the two securities are identical and are two-way convertible. The dynamics of mispricings are shown to be highly mean-reverting, exhibit minimal volatility and have short half-lives. However, this work find that the main disincentive to stock-ADR arbitrage stems from uncertainty about the duration of the pairs trade. For each additional cent demanded by arbitrageurs in per-share profit, the expected trade duration more than doubles, and the standard deviation of trade duration increases by 20%. Finally, arbitrageurs who set their trade entry bounds higher will experience less duration uncertainty but will observe fewer trading opportunities.

Overall, the results of this chapter show that stock-ADR arbitrage is characterized by a high incidence of small, short-lived, exploitable mispricings. Arbitrage is an important element in enforcing stock-ADR parity, and is demonstrably successful at doing so. This analysis constitutes evidence in support of Grossman and Stiglitz (1976, 1980) in that pricing anomalies incentivise arbitrageurs to restore price efficiency.

The work presented here could easily be extended to examine the role of arbitrage in other settings. Of particular interest to us for future work is the implementation of the pairs trading strategy to ADRs originating in emerging markets. Pairs-trading in this scenario circumvents many barriers to the direct arbitrage approach, especially restrictions stock-ADR convertibility which render direct arbitrage infeasible.

Chapter 5

Ultra High Frequency Statistical Arbitrage Across International Index Futures

Abstract

This chapter shows that exploitable lead-lag relations of the order of a few hundred milliseconds exist in the three pairings between the S&P 500, FTSE 100, and DAX futures contracts. These relations exhibit clear intra-daily patterns, particularly around the US open, the European close, and the announcement of country data. Using this information, this chapter forecasts mid-quote changes in lagging contracts with a directional accuracy in excess of 85%. A simple statistical arbitrage strategy exploiting these relations yields economically significant profits which are robust to market impact costs and the bid-ask spread. However, returns are sensitive to the risk of slippage, and the most profitable trading opportunities rarely exist for longer than 300 milliseconds. Hence, it is highlighted that price slippage and infrastructure costs are the most significant limits to arbitrage in this market setting. Overall, the results presented here accord with the view that informational inefficiencies incentivise arbitrageurs to appropriate pricing anomalies.

Keywords: Lead-Lag Relationships, Futures Markets, Hayashi-Yoshida Cross Correlation Estimator, Statistical Arbitrage.

JEL classification: F36, G14, G15.

5.1 -Introduction

Employing a recent advance in the statistical measurement of lead/lag relationships (Hayashi and Yoshida (2005), Hoffman et al (2010)), this chapter's principal contribution is to uncover the existence of sub-second exploitable disequilibria which occur across international index futures. The original methodology of Hayashi and Yoshida (2005) for measuring cross-correlation was extended by Hoffman et al. (2010) to measure lagged correlations, and it is on the latter work that the research presented here is based.

This chapter is motivated by evidence of increasing international financial integration in recent years (Ayuso and Blanco (2001), Kearney and Lucey (2004), Evans and Hnatkovska (2005)). If on the one hand, markets were perfectly integrated, efficient, and frictionless, then returns on closely related securities would exhibit perfect simultaneity and contemporaneous correlation, so as to preclude the possibility of cross-market arbitrage (De Jong and Donders (1998)). On the other hand, there is wide evidence that frictions impede the flow of information across markets (Lo and Mackinlay (1990)), and give rise to lead/lag patterns.[1] Now, if the price adjustment process is not instantaneous, then the question arises: Do lead/lag patterns across segmented markets induce cross-market predictability? Does information about future price adjustments admit the possibility of arbitrage?

This is the first work to provide answers to these questions within the important but overlooked context of international stock index futures. Specifically, this work focusses on lead/lag relations in the three pairings between the S&P 500, FTSE 100, and DAX futures contracts. While previous works focus on the relation between futures contracts and their underlying stock indices, the vast majority

[1] Examples include the S&P 500 futures vs spot in Kawaller et al. (1987), the Major Market Index futures vs spot in Stoll and Whaley (1990), and the FTSE 100 futures vs spot in Brooks et al. (2001).

highlight factors such as trading costs, price staleness, and illiquidity in the index constituents which render spot-futures arbitrage infeasible (see e.g. Brennan and Schwartz (1990)). In contrast, futures contracts are highly liquid instruments with low upfront capital requirements and trading costs. Therefore, one could argue that pairs of futures contracts constitute an ideal setting under which to examine lead/lag relationships.

The extant literature is complemented in four ways. Firstly, this work employs the Hayashi-Yoshida (Hayashi and Yoshida (2005), hereafter HY) cross-correlation estimator to measure the speed of price adjustment. The HY estimator readily allows for irregular and non-synchronous quote arrival times without the need for coarse resampling. This feature allows us to examine microscopic lead/lag patterns, which is particularly useful as it is found that the majority of price adjustments take place at sub-second intervals. Secondly, Best Bid and Offer (BBO) quote data is employed, time-stamped to the millisecond. BBO quotes form the top level of the Central Limit Order Book, thus constituting a continuous price series free of stale quotes. The use of ultra-high frequency data is crucially important, because the speed of price adjustment is likely to be high in electronically traded markets. Thirdly, this chapter places emphasis on both the forecasting accuracy, and arbitrage returns, in assessing the evidence against cross-market efficiency. Finally, this chapter documents the risks and costs which constitute the greatest limits to arbitrage in this market setting.

The results of this chapter show evidence of highly significant lead/lag relationships between the three pairs. On average, the S&P 500 leads the FTSE 100 and DAX contracts by 98 and 349 milliseconds respectively, consistent with the notion that international price discovery originates in the US. The FTSE 100 leads the DAX contract by 33 milliseconds on average. Some evidence of bi-directional causality is found in the FTSE-DAX pair, but very little between either transatlantic pair. Further, the FTSE-DAX pair is more strongly correlated than either transatlantic pair, in line with the notion of imperfect international integration (Jorion and Schwartz (1986)).

The lead/lag relations across all pairs exhibit strong intraday seasonality: Correlations among all pairs peak around daily economic announcements at 13:30 GMT, and fall sharply when the FTSE 100 and DAX cash markets close at 16:30 GMT, rebounding in evening trade. The S&P 500 contract increases its lead over both FTSE 100 and DAX contracts at the announcement of data, at 14:30 GMT when the S&P 500 cash market opens, and at the European close. It is important to note that none of these results are the product of statistical artifacts arising from differences in liquidity or quote arrival frequency.

Based on simple threshold triggers, mid-quote moves of leading contracts forecast subsequent moves of lagging contracts with a directional accuracy of over 85%. This accuracy relates directly to the sensitivity of the threshold: The more conservative the threshold, the “cleaner” the trigger signal, and the more accurate the forecast.

Finally, a simple trading exercise provides evidence against auto-efficiency across international markets. Leading mid-quote moves are used to generate buy/sell signals in lagging contracts, and the returns aggregated net of all costs. Arbitrage profits are economically and statistically significant - a strategy that trades only 5 contracts per trading signal generates aggregate profits of around GBP 100,000 per month. However, returns are sensitive to the risk of price slippage: Around half the profits disappear if 20% of trades execute at the next best price. Further, the windows of arbitrage opportunity are narrow, with the most profitable trades rarely existing for more than 300 milliseconds. This suggests that technology costs are a significant consideration for arbitrageurs in this setting.

Overall, it is concluded that information flows across international markets are not instantaneous, but rather exhibit brief delays. Importantly, these delays last long enough, and induce pricing anomalies large enough, to compensate arbitrageurs for appropriating pricing disequilibria. These results accord with the Grossman and Stiglitz (1976, 1980) suggestion that temporary disequilibria incentivise arbitrageurs to correct pricing anomalies. Though clearly, the authors did not have in mind the kind of trading frequency studied in this chapter, it is interesting that the results of this chapter supports their view in the high-frequency domain.

5.2 - Related Literature

Introduced in February 1982 by the Kansas City Board of Trade, the Value Line contract was the world's first index futures contract, followed shortly by the S&P 500 contract (Gulen and Mayhew (2000)). As futures and other derivative contracts become more pervasive, research exploring the link between these and their underlying markets, as well as across markets, becomes increasingly focal. The work presented here is related to two broad streams of literature which constitute a substantial body of work documenting inter-market pricing relationships.

The first stream explores price discovery and the temporal pricing relationship between futures contracts and the underlying cash index. The seminal research of Zeckhauser and Niederhoffer (1983) documents the possibility that futures prices can predict spot prices.[2] The early work involving intraday data of Kawaller et al. (1987) provides evidence that US futures prices lead spot prices with a time lag of around 45 minutes. The authors attribute this lag to the inertia of stock trading relative to futures trading, which implies that the futures market contributes most to price discovery. Similarly, Herbst et al. (1987) and Stoll and

[2]The term *spot* refers to the underlying index cash market - e.g. the constituent stocks of an index, gold bullion, or barrels of oil.

Whaley (1990) document that futures prices lead spot prices by around 8 minutes, whereas Wang and Wang (2001) show that price discovery is bi-directional when volatility is high.

The spot-futures relationship has also been explored for a number of different countries. Wahab and Lashgari (1993), Abhyankar (1998), Brooks et al. (2001), and Brooks and Garrett (2002) analyze UK spot-futures data, Tse (1995) explores the Japanese spot-futures relationship, while the case of Greece is examined in Kenourgios (2004) and Andreou and Pierides (2008). These works largely confirm the notion that futures prices lead spot prices, suggesting that price discovery takes place in the futures market. But what is the significance of these pricing relationships?

This question is explored in the second stream of literature to which this work belongs. On the one hand, an important implication of a lead/lag relationship is the potential for cross-market return predictability and arbitrage. On the other hand, the authors in Brooks et al. (2001) point out that while futures markets are capable of responding to new information quickly, the cash index can only fully respond once every constituent stock price updates. An important question then arises: Is a persistent lead/lag relationship evidence of market inefficiency?

While works such as Figlewski (1984, 1985), Brennan and Schwartz (1990), Thomas (2002), Richie et al. (2008), and Cummings and Frino (2011) document evidence of significant disequilibria in the spot-futures relation, there are two commonalities in these works: First, the nature of the pricing relation examined therein is not a temporal one based on lead/lag effects, but rather a mispricing approach based on a *cost of carry* model. Second, disequilibria in these works are frequently attributed to transaction costs, market (im)maturity, and liquidity effects that give rise to disequilibria while precluding arbitrage.

The literature exploring arbitrage based on lead/lag relationships is relatively sparse. Stoll and Whaley (1990) suggest that although futures generally lead spot prices, the effect is not uni-directional, making spot-futures arbitrage difficult. Brooks et al. (2001) compare several forecasting models and demonstrate that the FTSE 100 futures contract can achieve over 65% accuracy when forecasting the FTSE 100 index. However, when a trading strategy based on this is tested, profits are not robust to trading costs.

This chapter relates to these works by placing emphasis on the predictability of returns, and the profitability of arbitrage. However, unlike these works, the work presented here pursues this theme in an international setting. Indeed, one major commonality linking these works is a focus on inter-market relationships within the same country. Now, recent advances in communications technology have contributed to the integration of similar but otherwise fragmented markets. One direct consequence of this integration is a trend towards the equalization of expected returns across global markets - a single "relevant" event is able to move global stock market indices jointly (Eiteman et al (1994), Medeiros et al. (2009)).

There are several works which explore temporal relationships across international markets (Eun and Shim (1989), Hamao et al. (1990), Abhyankar et al. (1997), Antoniou et al (2003), Innocenti et al. (2011)). However, few focus explicitly on the link between lead/lag effects and arbitrage, none employ high-frequency data, and none focus on futures contracts. So why is doing so important? Employing high-frequency data is particularly important to uncover temporal relations which are invisible to the discrete-price observer. Goetzmann et al. (2005) document a dramatic increase in global market correlations over the last two decades, which increases international market integration and naturally pushes evidence of lead/lag relationships deeper into the sub-minute and sub-second space. Examining the link between lead/lag effects and arbitrage has important implications towards the theory of Efficient Markets. Finally, using futures data mitigates many of the limitations of spot-futures arbitrage, such as transaction costs and illiquidity. Therefore, exploring lead/lag relationships across different futures contracts provides an important platform on which to test for inter-market predictability and arbitrage.

A major risk inherent in the type of trading explored in this chapter stems from the fact that the contracts under investigation are not perfect substitutes. Arbitrage in this context relies on exploiting a statistical relationship based on economic fundamentals. Invariably, arbitrageurs are vulnerable to noise traders, who may trade the individual contracts for reasons other than to exploit pricing disequilibria (De Long et al. (1990)), at highly unfavourable times for arbitrageurs.

In what follows, this chapter employs BBO quotes time-stamped to the millisecond for the S&P 500, FTSE 100, and DAX futures. These contracts are highly liquid financial securities based which represent the broad markets of three major developed economies. By using them, one can mitigate to the greatest extent possible the transaction cost barrier, and the spurious lead/lag relation generated by liquidity differences across the contracts. Anticipating that lead/lag patterns occur at very fine timescales, this work employs the HY estimator which is statistically robust to non-synchronous trading and irregular quote-arrival times.

5.3 - Data and Methodology

This section provides details of the data and data refinement procedures, together with details of the precise implementation of the HY estimator. Also, the performance of HY is contrasted against a common discretization procedure.

Finally, this section provides details of the forecasting and trading strategy employed.

5.3.1 Futures Contracts Data

The data consist of the three most liquid futures contracts globally, namely the DAX, FTSE 100, and S&P 500 E-mini futures contracts.[3] For each contract, BBO quotes are obtained containing price and volume information, time-stamped to the millisecond. The data span the period January 9, 2012 - February 28, 2012. Each contract is monitored for 12 hours per day, between 08:30 and 20:30 GMT. The FTSE contract trades from 01:00 - 21:00 GMT; the DAX contract trades from 07:50 - 22:00 CET, and the S&P trades 24 hours per day with the exception of times between 15:00 - 15:15, wherein quotes can be submitted by no executions are permitted. The trading period chosen in this study overlaps all these contracts' active trading hours.

Because futures contracts operate on a quarterly expiration cycle, the focus of this study is on the the current "active" contracts, namely those expiring in March 2012. In all, the dataset consists of a number of BBO quotes in excess of 100 million.

The data is sourced from the CME, Eurex, and NYSE-Liffe exchanges. Since each of these exchanges is synchronized to an atomic clock (IOSCO (2012)), the quote time-stamps reported for each exchange are precise to within one microsecond, and contemporaneous to within one millisecond (Barua (2012)). Table 5.1 provides an overview of the institutional features of the data set.

It can be seen from Table 5.1 that these contracts are highly liquid. Given that the main aim of this chapter is to ultimately establish the existence of arbitrage opportunities based on lead/lag relationships, the choice of data set

[3]The E-mini is the more liquid and electronically traded version of the S&P 500 contract.

Contract Name	Symbol	Exchange	Multiplier	Minimum Increment	Average Daily Volume	Average Daily Trade Value
DAX	DA	Eurex	x €25	0.50	137,217	€21.5 billion
FTSE 100	FT	NYSE-Liffe	x £10	0.50	104,041	£7.7 billion
E-mini S&P 500	ES	CME	x $50	0.25	1,971,484	$140 billion

Table 5.1: Futures Data. Institutional features of the data set. The last two columns show, as of January 2012, the average number of contracts traded, and the average daily turnover.

is deliberately conservative, for it is well documented that liquidity and the availability of arbitrage are inversely related.

The fourth column in Table 5.1 shows the minimum price increment in terms of index points. To obtain the monetary value of this price increment, one simply multiplies the number of index points by the multiplier. For example, the minimum price increment of the FTSE 100 contract is £5.

Inspired by Schultz and Shive (2010), a number of crucial data filtering processes are applied. Specifically, data which qualifies any of the following are excluded:[4]

- Bid Price $\geq$ Ask Price,
- Bid Volume $= 0$ and/or Ask Volume $= 0$,
- Ask Price $>$ Bid Price by more than 25%,
- Mid quote return $\geq 25\%$ or $\leq -25\%$.

Since this chapter makes use of BBO quotes, it is in effect continuously observing the current best quote within the central order book. New BBO quotes arrive when there is either a change in the current best price, or a change in the available liquidity at the current best price. In the first scenario, if a new quote arrives that is better than the current best price, or an incoming trade consumes all the liquidity available at the best price, the BBO price and volume update. In the

[4]In all, this process removes around 0.2% of the data.

second scenario, if a new quote arrives *at* the current best price, or an incoming trade does not consume all the currently available liquidity, the BBO volume updates, but the price does not.

In view of establishing lead/lag relationships, the difference between the above two scenarios is non-trivial. If one monitors changes in the mid-quote, the first scenario will always yield a non-zero return, while the second scenario will always yield a zero return. The main aim here is to establish whether mid-quote changes in one futures contract portend corresponding changes in another. This aim is best served by focussing purely on non-zero returns, and so all zero-returns from the data are removed. This process carries the benefit of allowing us to deem all non-zero returns as informative to the lead/lag relationship.[5] A similar procedure is performed in Huth and Abergel (2012).

5.3.2 Trading Costs and the Representative Investor

Costs pertaining to futures trading consist of the bid/ask spread, which is accounted for directly in the BBO quote data. Further, each exchange charges a fee for order submission and clearing. These are outlined below.

Contract Name	Exchange Fees	Clearing Fees
DAX	€ 0.50	---
FTSE 100	£0.25	£0.03
E-mini S&P 500	$0.75	$0.39

Table 5.2: Exchange Costs. This table shows the costs for order submission and clearing levied by the Eurex, NYSE-Liffe, and CME Group. These relate to trading the DAX, FTSE 100, and S&P 500 contracts, respectively.

The costs in Table 5.2 apply on a per-order basis, regardless of the size of the order. Given that the notional amount of a single futures contract is 10, 25, or 50 times its underlying value (see Table 5.1), the costs in Table 5.2 are economically

[5] Griffin and Oomen (2011) provide a thorough discussion of similar sub-sampling routines.

small. However, the strategy described in this chapter warrants the submission of a vast number of market orders aimed at exploiting minute pricing disequilibria across these markets. The costs in Table 5.2 are negligible for a long-term investor, but are a significant consideration here.

With that in mind, the representative investor is akin to a quantitative hedge fund utilizing a fully algorithmic strategy, with access to co-location services within the exchange buildings. This type of investor is chosen for two reasons: Firstly, all three contracts are 100% electronically traded. Hendershott et al. (2011) provide evidence that some 70% of NYSE trades are executed by investors of this type. Secondly, the approach in this chapter aims to exploit temporal pricing disequilibria which exist predominantly in the sub-second horizon. Clearly, such speeds lie beyond the scope of human traders.

5.3.3 The Hayashi-Yoshida Estimator

In Hayashi and Yoshida (2005),[6] the authors introduce a novel estimator of the covariance between two non-synchronous processes. Specifically, let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, and let X_t and Y_t be two correlated processes such that:

$$\begin{aligned} dX_t &= \mu^X X_t dt + \sigma^X X_t dW_t^X \\ dY_t &= \mu^Y Y_t dt + \sigma^Y Y_t dW_t^Y, \end{aligned} \tag{1}$$

where W_t^X and W_t^Y are Brownian motions with respect to $\mathbb{P}$. Assume that the correlation $\langle W_t^X, W_t^Y \rangle = \rho$. Assume further that X_t and Y_t are sampled at discrete observation times $0 = t_0^X \leq t_1^X \leq \cdots \leq t_n^X = T^X$ and $0 = t_0^Y \leq t_1^Y \leq \cdots \leq t_m^Y = T^Y$ respectively. Importantly, these observation times are assumed independent of each other and of X_t and Y_t - which in practical terms suggests

[6]For brevity, mathematical proofs in this section are omitted. However, the interested reader is directed to works in which these proofs are available.

two things. Firstly, the quote arrival frequency of one asset does not influence that of the other. Secondly, the quote arrival frequency of each asset does not depend on the value of that asset.

Let us define the following terms:

$$I_i^X = \left(t_i^X, t_{i+1}^X\right] \tag{2}$$
$$I_j^Y = \left(t_j^Y, t_{j+1}^Y\right].$$

Here, I_i^X and I_j^Y denote time intervals between the arrival of quotes in assets X and Y respectively, and this concept is illustrated in Figure 5.1.

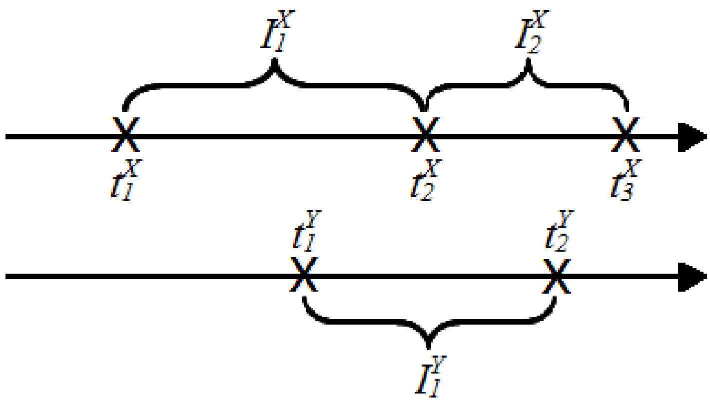

Figure 5.1. Time Intervals. A time-line illustration of quote arrivals in X (top) and Y (bottom). Quotes are marked "X". Intervals between quotes are marked I^X and I^Y.

Figure 5.1 depicts a typical scenario encountered with high-frequency data. Quotes arrive irregularly for each asset, and asynchronously across assets. Robustly estimating the covariation between data which exhibit these phenomena is the subject of HY. Given (1) and (2), the HY estimator for the covariance $C^{X,Y}$ is:

$$C^{X,Y} = \sum_{i,j} \Delta X \left(I_i^X \right) \Delta Y \left(I_j^Y \right) \mathbb{I}_{\{I_i^X \cap I_j^Y \neq \emptyset\}}, \tag{3}$$

where Δ is the difference operator, and $\mathbb{I}$ is an indicator function such that:

$$\mathbb{I} = \begin{cases} 1, & \text{if } I_i^X \cap I_j^Y \neq \emptyset \\ 0, & \text{otherwise.} \end{cases}$$

The covariance $C^{X,Y}$ in (3) yields the HY correlation coefficient ρ_{HY} as follows:

$$\begin{aligned} \rho_{HY} &= \frac{C^{X,Y}}{\sigma^X \sigma^Y} \\ &= \frac{\sum_{i,j} \Delta X \left(I_i^X \right) \Delta Y \left(I_j^Y \right) \mathbb{I}_{\{I_i^X \cap I_j^Y \neq \emptyset\}}}{\sqrt{\sum_i \left[\Delta X \left(I_i^X \right) \right]^2 \sum_j \left[\Delta Y \left(I_j^Y \right) \right]^2}}. \end{aligned} \tag{4}$$

In practice, equations (3) and (4) amount to summing the product of all returns between assets X and Y once they fully or partially share a time overlap. Because of this property, the HY estimator allows for the inclusion of all data points, without the need for regularizing (re-sampling) the data.

As it is currently stated, equation (4) measures the contemporaneous correlations between X and Y. Hoffman et al. (2010) extend the estimator to allow for leads and lags. Following from (2), let us define:

$$(I_j^Y)_\ell = \left(t_j^Y + \ell, t_{j+1}^Y + \ell \right], \tag{5}$$

where $\ell \in \mathbb{R}$ is the lag length (measured in units of time). Following the same procedure in (2)-(4) arrives at a formula for the lagged version of the HY estimator.

$$\rho_{HY}(\ell) = \frac{\sum_{i,j} \Delta X\left(I_i^X\right) \Delta Y\left(I_j^Y\right) \mathbb{I}_{\{I_i^X \cap (I_j^Y)_\ell \neq \emptyset\}}}{\sqrt{\sum_i \left[\Delta X\left(I_i^X\right)\right]^2 \sum_j \left[\Delta Y\left(I_j^Y\right)\right]^2}}. \tag{6}$$

As it is now defined, equation (6) yields the entire cross-correlation curve (hereafter HY curve) between X and Y. In practice, evaluating (6) amounts to shifting the time-stamps of Y by amount ℓ and re-evaluating the HY correlation ρ_{HY}. Doing this for all ℓ yields the HY curve.

In practical applications, it is important to establish the lag length $\hat{\ell}$ that maximizes the correlation $\rho_{HY}(\hat{\ell})$. Doing so allows conclusions to be drawn about the temporal relationship between X and Y. For example, "X leads Y by $\hat{\ell}$ seconds". The lag $\hat{\ell}$ is defined as a solution to the following equation:

$$|\rho_{HY}(\hat{\ell})| = \max_{\ell \in \mathcal{G}} |\rho_{HY}(\ell)|, \tag{7}$$

over a time-grid $\mathcal{G}$. In practice, this amounts to finding the peak of the HY curve. Equivalently, evaluating (6) for different values of ℓ until a maximum is obtained.

Importantly, if it is found that $\hat{\ell} \neq 0$, this would imply that the relationship between X and Y is not contemporaneous. Specifically, knowledge of one can be used to explicitly forecast future movements in the other. It is precisely this condition that this work relies on to examine lead/lag relations across international futures contracts.

Finally, Huth and Abergel (2012) define the notion of a lead/lag ratio (henceforth LLR) which is employed in the subsequent analysis. It is defined as follows:

$$\text{LLR} = \frac{\sum_u \rho_{HY}^2(-\ell_u)}{\sum_u \rho_{HY}^2(\ell_u)}, \tag{8}$$

where u is chosen such that $\ell_u \geq 0$.

The quantity $\rho_{HY}(\ell)$ in (6) provides the correlation coefficient when Y leads X by ℓ units of time, and similarly, the quantity $\rho_{HY}(-\ell)$ provides the same measure when Y lags X (i.e. X leads Y) by ℓ units of time. With this in mind, the purpose of the LLR is simply to compare the evidence of these two phenomena. It is established in the literature that lead/lag relationships are often bi-directional (Wang and Wang (2001) among others). The LLR is informative in decoupling bi-directional relationships into relative strengths at different lags. Put simply, it is a metric which is useful in assessing the strength and direction of the lead/lag relationship.

5.3.4 Robustness to Spurious Lead/Lag Relations

This section evaluates the robustness of HY to differences in liquidity. Further, it contrasts the performance of HY against a commonly employed method of measuring covariation.

It is well known that liquid[7] assets tend to lead less liquid ones, due to their ability to impound information faster (Lo and Mackinlay (1990) and Brooks et al. (2001)). However, liquidity differences could also give rise to spurious conclusions toward lead/lag relationships. For example, data that are known to be contemporaneously correlated but contain differences in trading activity must not exhibit any lead/lag effects (Voev and Lunde (2007)). Discovery of a non-zero lag in this case would by definition be spurious. It is therefore important that

[7]In this case, liquidity refers to trading/quoting activity.

any employed estimator be robust to *artificial liquidity* effects, particularly when dealing with high-frequency data.

This section examines the robustness of HY to spurious lead/lag relations by way of an exercise involving synthetic data. The benefit of using synthetic data is simple - the input parameters are known, hence it is clear what output to expect.

This exercise involves sampling from the stochastic processes X and Y in (1) along two independent Poisson time-grids with intensities λ_X and λ_Y respectively.[8] The data span 24 hours, and samples are observed to the nearest millisecond. Further, it is imposed that X leads Y by 400 milliseconds (0.4 seconds). At this lag length, a correlation between X and Y of 90% is imposed.

This exercise measures lead/lags between -10 and 10 seconds, with 10-millisecond increments. Specifically, a grid $\mathcal{G}$ is chosen from (7) such that:

$$\ell \in \{-10, -9.99, -9.98, \ldots, -0.01, 0, 0.01, \ldots, 9.98, 9.99, 10\}.$$

Further, the exercise is repeated for different liquidity combinations between assets X and Y. This is achieved by varying the ratio $\frac{\lambda_Y}{\lambda_X} \in \{0.5, 1, 2, 5\}$. This step is aimed at examining how the estimated lead/lag relation is affected by differences in liquidity. For example, the ratio $\frac{\lambda_Y}{\lambda_X} = 5$ suggests that the average interval between price updates for asset Y is 5 times that for asset X. Figure 5.2 shows the results of the HY exercise.

It is clear from Figure 5.2 that the HY estimator is robust to spurious temporal relations induced by liquidity differences between assets - it correctly recovers the true lead/lag relationship between X and Y under all combinations of $\frac{\lambda_Y}{\lambda_X}$.

[8]Here, λ_X and λ_Y can be thought of as proxies for illiquidity: The lower their value, the shorter the average interval between price updates.

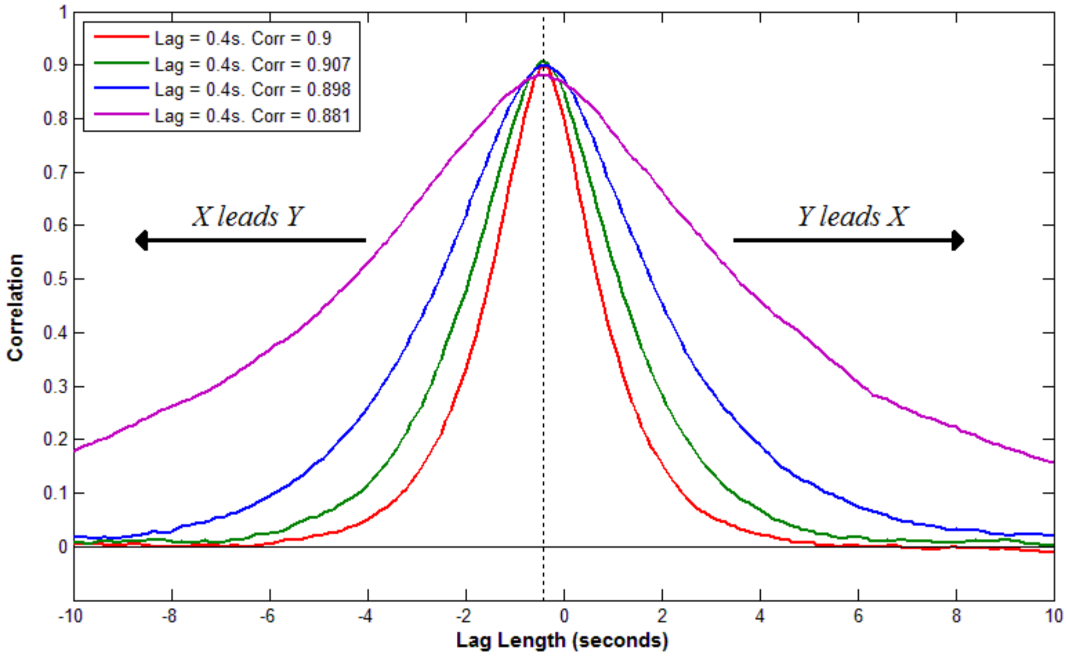

Figure 5.2. HY Curve on Synthetic Data. This figure shows the HY-estimated correlation between X and Y at various lags $\ell \in [-10, 10]$. A dotted vertical line denotes the "true" lead/lag relationship with X leading Y by 0.4 seconds ($\hat{\ell} = -0.4$). The four curves represent repetitions of the HY exercise with various liquidity ratios: $\frac{\lambda_Y}{\lambda_X} = 0.5$ shown in red, $\frac{\lambda_Y}{\lambda_X} = 1$ in green, $\frac{\lambda_Y}{\lambda_X} = 2$ in blue, and $\frac{\lambda_Y}{\lambda_X} = 5$ in purple.

This property of HY is useful to the examination of temporal relationships using high-frequency data.

It is informative to contrast the performance of HY against that of a commonly employed methodology, namely *previous-tick interpolation*, henceforth PT.[9] For completeness, the results of the same exercise using the PT method are presented.

The intuitive idea behind the PT method is to regularize irregularly-spaced and non-synchronous prices onto a synchronous time-grid with fixed and constant intervals. Where there is no price update between two or more consecutive points on the regularized grid, the previous price is interpolated forward. Formally, and using the definitions in (2), let:

[9]for a thorough discussion, see Voev and Lunde (2007), Hoffman et al. (2010), and Griffin and Oomen (2011).

$$\overline{X}_t = X(I_i^X) \text{ and, } \overline{Y}_t \ = Y(I_j^Y), \tag{9}$$

denote the PT-interpolated versions of X and Y in (1), then for a fixed-interval grid of size M and interval length h, the PT correlation estimator is given by:

$$\rho^{PT} = \frac{\sum_{i=1}^{M}\left(\overline{X}_{ih} - \overline{X}_{(i-1)h}\right)\left(\overline{Y}_{ih} - \overline{Y}_{(i-1)h}\right)}{\sigma^X \sigma^Y}. \tag{10}$$

The extension of ρ^{PT} to incorporate leads and lags $\rho^{PT}(\ell)$ is similar to that of HY in (6).

Let us now repeat the estimation exercise using the PT method. This is done for various mesh sizes $h \in \{0.1, 0.5, 1\}$ seconds. The results are given in Figure 5.3.

There are two interesting phenomena evident in Figure 5.3. First, the PT estimation procedure is adversely affected by differences in liquidity. For example, in the bottom-left panel where price updates in X occur on average 5 times as frequently as those in Y, the PT approach suggests that X leads Y by between 3 and 3.4 seconds, which is clearly not true. Second, although a smaller mesh size h yields a more accurate estimate of the true lag $\hat{\ell}$, it also results in a more biased estimate of the true correlation (Epps (1979)). Overall, the results of this section motivate the use of HY in dealing with high-frequency data.

5.3.5 The Trading Strategy

Knowledge of the temporal relationship between a pair of contracts can be used to build statistical arbitrage strategies based on exploiting the temporal disequilibria. In principle, one would expect that a higher temporal correlation would yield greater trading profits, but also more competition from other

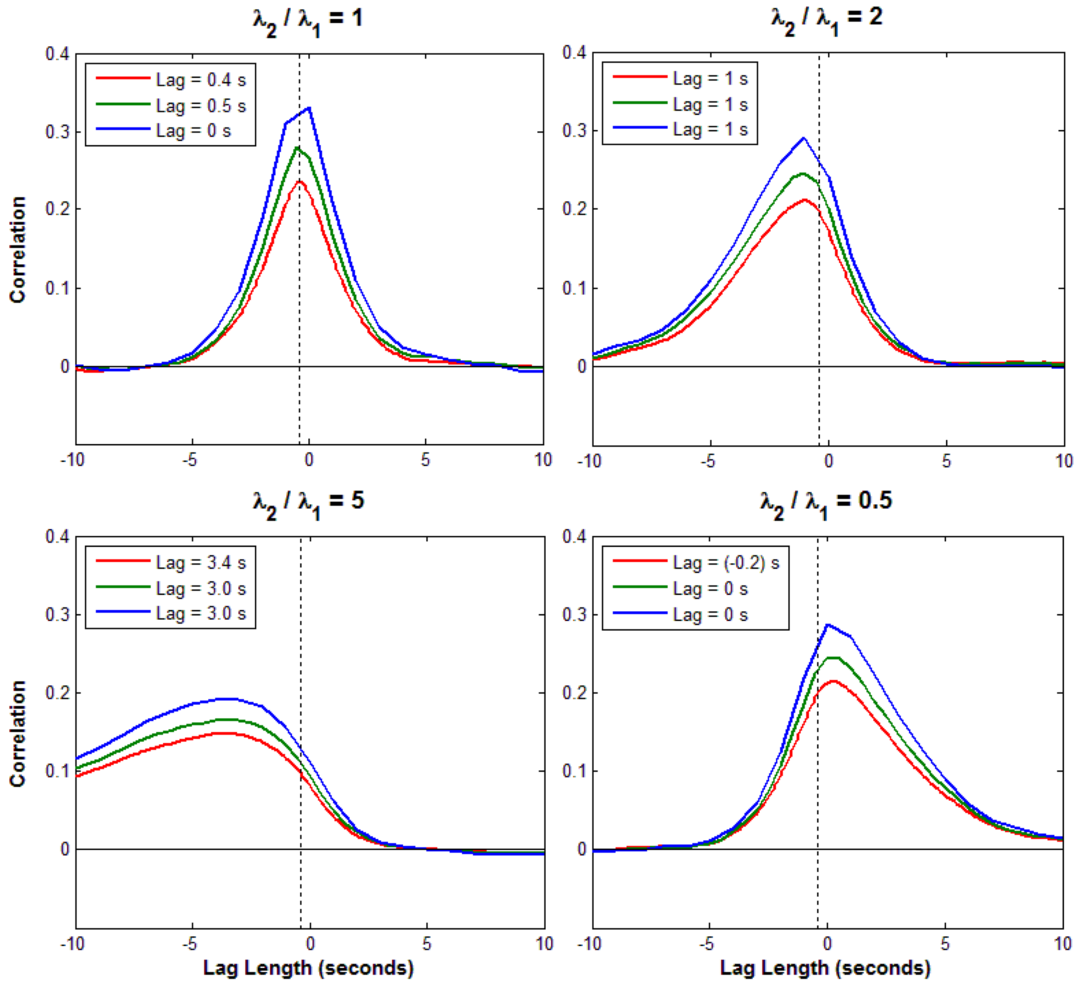

Figure 5.3. PT Curve on Synthetic Data. This figure shows the PT-estimated correlation between X and Y at various lags $\ell \in [-10, 10]$. A dotted vertical line denotes the "true" lead/lag relationship with X leading Y by 0.4 seconds ($\hat{\ell} = -0.4$). The three curves in each panel represent repetitions of the PT exercise with various mesh sizes: $h = 0.1$ shown in red, $h = 0.5$ in green, and $h = 1$ in blue. The four panels represent different liquidity combinations $\frac{\lambda_Y}{\lambda_X}$ across X and Y. Lags which maximize the correlation are given by "Lag".

arbitrageurs. On the other hand, lower temporal correlation would not intuitively yield higher profits, but arbitrageurs in the latter space would face less competition from others interested in exploiting the temporal disequilibria. This section shows how to apply the knowledge of the lead/lag relationship to forecast and exploit mid-quote changes in the lagging contract, based on mid-quote changes in the leading contract.

The forecasting and trading devices employed here are motivated by Huth and Abergel (2012) and Kozhan and Tham (2012). Given contracts X and Y, observing both time-series simultaneously amounts to observing sets of contiguous quotes in X interspersed with sets of contiguous quotes in Y. Hence let us define a *cluster* of mid-quote returns $\{C^X_{i,n} : i, n \in \mathbb{N}^+\}$ in contract X as a sequence of contiguous mid-quote returns in X uninterrupted by returns in Y. A similar definition of $C^Y_{i,n}$ holds for contract Y. Here, the subscript i refers to the cluster index (the number of clusters already observed), whereas the subscript n refers to the mid-quote return index within each cluster. This concept is illustrated graphically in Figure 5.4.

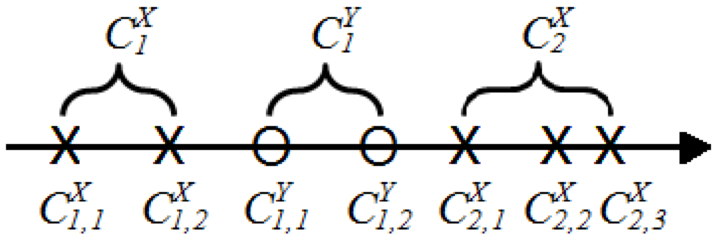

Figure 5.4. Quote Clusters. A time-line illustration of contiguous sequences of mid-quote returns. Returns in asset X are marked **X** and returns in asset Y are marked **O**, while X-clusters and Y-clusters are marked with $C^X_{i,n}$ and $C^Y_{i,n}$ respectively.

Once the leading contract is distinguished from the lagging contract within the HY framework, one can use this information to build a statistical arbitrage trading strategy. Without loss of generality, let us assume contract X leads contract Y, then:

$$N^Y_i = \begin{cases} +1, & \text{if } \max_n \left(C^X_{i,n}\right) \geq K \\ -1, & \text{if } \min_n \left(C^X_{i,n}\right) \leq -K \\ 0, & \text{otherwise,} \end{cases} \tag{11}$$

where N_i^Y is the position taken in contract Y, and where $+1$ and -1 denote long and short positions, respectively. The quantity $K > 0$ is a pre-set threshold. Intuitively, equation (11) describes a simple trading strategy: Given that X leads Y, a signal to trade Y is generated whenever the price of X moves by an amount greater than K in the current X-cluster (i.e. before the price of Y moves).[10] If Y moves in the interim, no trading signal is generated and the process restarts in the following X-cluster. The motivation for including a threshold in the trading strategy stems from the fact that exploiting microscopic price "jumps" is fruitless, unless the magnitudes of these jumps exceed the costs associated with trading (Brooks et al. (2001)). It is important to note that the trading strategy (11) depends only on present and past price observations of the leading contract.

When a signal to trade Y is generated in the current X-cluster, a position in Y is immediately opened and held for the entire duration of the following Y-cluster. The position is then closed at the start of the subsequent X-cluster. Formally:

$$N_{i+1}^Y = \begin{cases} 0, & \text{if } \{C_{i+\epsilon,n}^X : n \geq 1, \epsilon \geq 1\} \\ N_i^Y, & \text{otherwise.} \end{cases} \tag{12}$$

Taken together, equations (11) and (12) fully describe the algorithm by which temporal disequilibria are exploited between leading contract X and lagging contract Y. The calculation of profit P for each Y-cluster is as follows: First, the position taken in cluster i is multiplied by the sum of the individual mid-quote returns in cluster i, and then trading costs are subtracted:

$$P_i = N_i^Y \sum_n C_{i,n}^Y - c_i, \tag{13}$$

[10] Since all C denote mid-quote returns, it is easier to measure K in terms of ticks. Conversion between ticks and monetary value is done easily via Table 5.1

where c_i is the all-inclusive cost of trading. Calculating total profits amounts to summing the individual P_i over all clusters in the data set:

$$\text{Total Profit} = \sum_i P_i. \tag{14}$$

Alongside the trading strategy, this section measures the directional accuracy by which mid-quote changes in contract X portend mid-quote changes in contract Y. Specifically:

$$F_i = \begin{cases} 1, & \text{if } \operatorname{sgn}\left(\sum_n \left(C^Y_{i,n}\right)\right) = \operatorname{sgn}\left(\sum_n \left(C^X_{i,n}\right)\right) \\ 0, & \text{otherwise,} \end{cases} \tag{15}$$

where $\operatorname{sgn}(\cdot)$ is the signum function. Equation (15) essentially measures whether the aggregate returns in the current X-cluster portend the sign of the aggregate returns in the subsequent Y-cluster. Arriving at a final measure of directional forecasting accuracy is intuitive: sum all F_i and divide by the number of clusters in the data set:

$$\text{Directional Forecasting Accuracy} = \frac{\sum_i F_i}{\sup(i)} \times 100\%, \tag{16}$$

where $\sup(\cdot)$ denotes the supremum.

An important question related to the forecasting accuracy is whether the accuracy is affected by the choice of signal threshold K in (11). One might expect that a stronger trading signal (i.e. higher K) would intuitively yield a higher directional forecasting accuracy of X over Y. To help investigate this, equations (15) and (16) are evaluated over all clusters *in which a trade has been generated.*

5.4 - Empirical Results

This section applies the analysis of section 5.3 to the three contract pairs within the data set, namely the FTSE - DAX, S&P - FTSE, and S&P - DAX March 2012 pairs. For the remainder of this chapter, let us define the following conventions: First, the leading asset, as in the above list, is named first. Second, graphs pertaining to each pair will be coloured blue, green, and red, respectively. Third, the numerator of the lead/lag ratio (8) refers to the first-named contract; thus a highly asymmetrical left-heavy HY curve with LLR $\gg 1$ shows strong evidence that the first-named leads.

This section starts by exploring the three lead/lag relationships, and profile the intraday patterns of these relationships. Then, this section forecasts and trades lagging contracts based on signals generated by mid-quote changes in leading contracts. Finally, a discussion is presented about the limits to arbitrage across international futures contracts.

5.4.1 The Lead/Lag Relationship Between Futures Contracts

There are 35 full trading days (hereafter "days") in the sample. For each day, the entire HY cross-correlation curve is estimatee based on equation (6), then the curves are averaged for each pair over all days. Ultimately, this section obtains a single HY curve for each contract pair. This step helps quantify the relative strength and direction of the lead/lag relationships for each of the contract pairs.

For each contract pair, leads and lags of mid-quote returns are measured on a horizon of -30 to 30 seconds, with 5-millisecond increments. This section justifies this horizon by the fact that cross-correlations across all pairs diminish substantially within a few seconds. A grid $\mathcal{G}$ is chosen from (7) such that:

$$\ell \in \{-30, -29.995, -29.990, \ldots, -0.005, 0, 0.005, \ldots, 29.990, 29.995, 30\}.$$

The results of this exercise are presented in Figure 5.5.

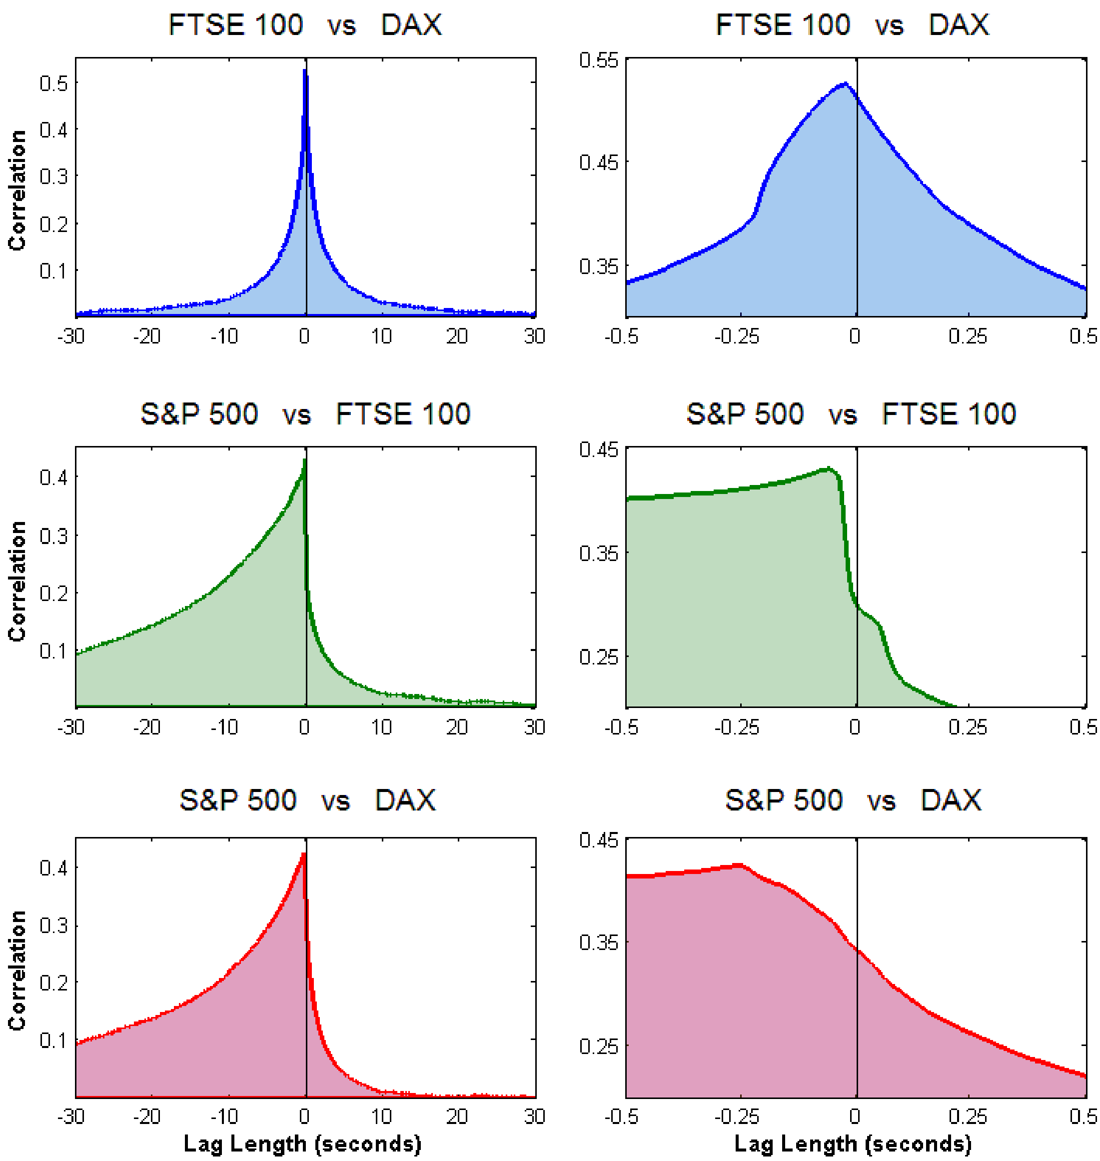

Figure 5.5. Lead/Lag Relationships. The HY curves plotted for each contract pair. Left panels show lag lengths $\ell \in [-30, 30]$ seconds. Right panels zoom in on lag lengths $\ell \in [-0.5, 0.5]$ seconds. In regions to the left of the vertical zero-line in each panel, the contract whose name appears first in the title leads, and vice versa.

Figure 5.5 displays evidence of strongly asymmetric lead/lag relationships, in which the S&P 500 contract leads both the FTSE 100 and DAX contracts.

The relationship between the FTSE 100 and DAX contracts is less pronounced, though it is clear that the FTSE leads the DAX contract. These results are qualitatively consistent with other works documenting international price linkages (Innocenti et al. (2011) among others), and the notion that international price discovery occurs in the US. Further, by looking at the left-hand panels in Figure 5.5, what is particularly striking is the speed at which the HY cross-correlation curves diminish towards zero across all three pairs. In other words, the panels in Figure 5.5 display evidence of high financial integration across the three contract pairs. This is confirmed by examining the right-hand panels, which reveal that point estimates of the lead/lag relationships, or equivalently the peaks of the HY curves, are of the order of milliseconds. This observation is striking because it has important implications for works examining international pricing relationships: by using non-granular data, it is easy to achieve the illusory effect of perfect contemporaneous correlation, as documented in Section 5.3.4 of this chapter.

Let us quantify the strength and direction of the lead/lag relationships in terms of three parameters, namely the lag length $\hat{\ell}$ which maximizes the correlation between a given contract pair, the maximum correlation itself $\hat{\rho}^{HY}$, and the lead/lag ratio LLR. These results are given in Table 5.3.

Pair Name	Leading Asset	Mean Lag Length (milliseconds)	Median Lag Length (milliseconds)	Maximum Correlation	LLR
FTSE 100 / DAX	FTSE 100	25	33	52.36%	1.20
S&P 500 / FTSE 100	S&P 500	60	98	42.80%	19.64
S&P 500 / DAX	S&P 500	250	349	42.25%	15.96

Table 5.3. Lead/Lag Relationships. Three contract pairs along with summary statistics. The third and fourth columns show, respectively, the mean and median of the lag length $\hat{\ell}$, by which the leader (second column) leads. The fifth and sixth column show, respectively, the maximum correlation $\hat{\rho}^{HY}$ and lead/lag ratios for each pair.

Table 5.3 reveals intuitive insights into the nature of the lead/lag relationships under investigation. Lag lengths, maximum correlations, and lead/lag ratios all appear to be directly related to the pairwise geographical distances between the markets.

Given that market information pertinent to any lead/lag relationship must traverse geographical distances, it is natural to wonder whether the lag lengths reported in Table 5.3 are not indicative of price discovery, but merely the result of delays in the physical transmission of data. The latter is unlikely. Built in 1999 by Global Crossing, the current transatlantic communication network (AC-1) spans 14,000 km, and links the US and major European markets via fibre-optic cables. One-way transmission between New York and London takes 32.4 milliseconds (see Johnson et al. (2012) and references therein). By extrapolating this number to the distances between the CME, NYSE-Liffe, and Eurex, it is unlikely that the lag lengths in Table 5.3 are the result of physical delays. Further, Hasbrouck and Saar (2012) suggest that the entire event-analysis-action cycle for co-located algorithmic trading systems is less than 2 milliseconds.

By examining the HY curves generated for each day in the sample, one can reject at the 90%-level the statistical hypothesis that the maximum correlation occurs at zero lag for the FTSE-DAX pair; and at the 99%-level for both transatlantic pairs. This result is not surprising, given the proximity of the NYSE-Liffe to the Eurex, relative to either transatlantic distance.

5.4.2 The Intraday Profile of Lead/Lag Relationships

It is a well-known stylized fact that financial markets produce intraday patterns (e.g. U-shaped volatility). Having established an overall lead/lag structure across the three contract pairs, let us turn attention to the intraday profile of these relationships.

This section splits each day into 24 half-hourly intervals, spanning 08:30 - 20:30 GMT. For each interval, over each day, and for each contract pair, let us measure the same three quantities as in Table 5.3, namely the Lead/Lag ratio, the maximum correlation $\hat{\rho}^{HY}$, and the lag length $\hat{\ell}$ which maximizes the correlation. Let us then average these results over the number of days. Ultimately, three 24-point curves are obtained: one for each contract pair. The results of this exercise are shown in Figure 5.6.

Figure 5.6 reveals clear evidence of intraday seasonality exhibited for all three pairs. The daily events most interesting to us are the announcement of macroeconomic news at 13:30 GMT, the US cash market open at 14:30 GMT, and the close of both UK and German cash markets at 16:30 GMT. All three panels confirm that for each pair, mid-quote returns in the first-named contracts portend mid-quote returns in the second-named contracts. Although this notion was established in the previous section, Figure 5.6 shows that this effect is consistent throughout the day.

The Lead/Lag ratio generally decreases (towards unity) throughout the day, indicating a diminishing asymmetry in the HY curves. This implies that the pairwise causal link between the leading and lagging contracts deteriorates throughout the day (Huth and Abergel (2012)). The maximum correlation remains largely range-bound, and the lag length largely constant. Notable exceptions to the above are as follows: Around the announcement of macroeconomic data, the LLR and maximum correlation both increase, while in terms of lag length, the US increases its leads over both UK and German contracts. This indicates that both UK and German traders react more attentively to cues from the US market, which impounds information first. The same effect occurs at the US cash market open. Following the close of the UK and German cash markets, the US sharply increases its lead over each, but both the

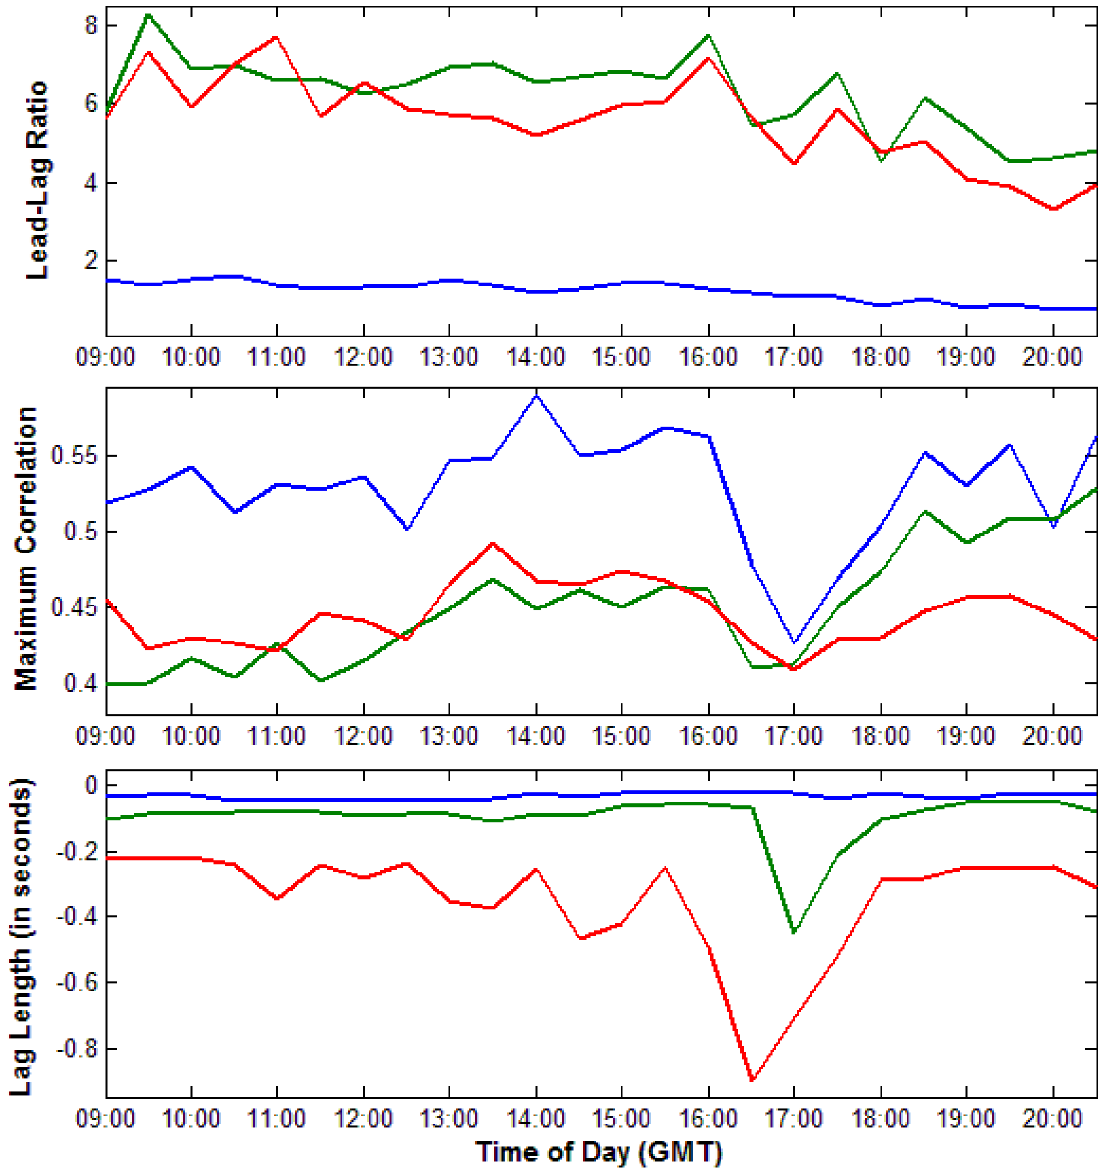

Figure 5.6. Lead/Lag Intraday Profile. Intraday patterns in the Lead/Lag Ratio (top), Maximum Correlation $\hat{\rho}^{HY}$ (middle), and Lag Length $\hat{\ell}$ (bottom). Times on the horizontal axes refer to the previous half-hour interval (e.g. 10:00 refers to the 09:30-10:00 interval).

LLR and maximum correlation across all three pairs fall sharply. Taken together, these observations suggest that at the close of the UK and German markets, the FTSE and DAX contracts are less sensitive to global factors, and more so to local factors (Jorion and Schwartz (1986)).

5.4.3 Trading Results: Exploiting Temporal Disequilibria

This section performs the statistical arbitrage exercise based on the methodology presented in section 5.3.5. Note that the trading strategy relies only on present and past mid-quote observations. Importantly, the trading strategy is simple in nature, and uses no *ex-ante* knowledge of the dynamic lead/lag relationship presented in section 5.4.2. This way, the entire data set is taken as out-of-sample for forecasting and trading purposes. Although it is discussed later how one may enhance the trading strategy by incorporating the information conveyed via the HY output, the results presented here rely chiefly on the simple strategy (11) - (12), in the interests of robustness.

Let us begin by presenting the accumulated profits gained by trading the FTSE - DAX, S&P - FTSE, and S&P - DAX pairs over the number of days in the sample. Throughout, the convention that the first-named contract leads is maintained. Initially, this analysis imposes a signal threshold $K = 5$ ticks in equation (11). That is to say, a trading signal is generated in the lagging contract whenever a cluster of leading quotes is encountered, wherein the mid-quote moves by 5 or more ticks in any aggregate direction. The position is then held open until the end of the lagging cluster. Later, the signal threshold is varied to examine the effect of the signal on arbitrage profits. Furthermore, this section initially imposes that the arbitrageur trades a single contract per signal; a step that ensures no market-impact costs, since this chapter employ BBO quotes. Figure 5.7 shows the accumulated wealth gained by trading the three contract pairs, following the strategy (11) - (12). The curves denote profits over the 35-day sample period, net of the bid-ask spread and all order submission and clearing costs.

It is worth pointing out that futures contracts are inherently leveraged, since it is not required that the investor possess the full notional amount of the contract (index points times multiplier). Typically, a clearing house will require

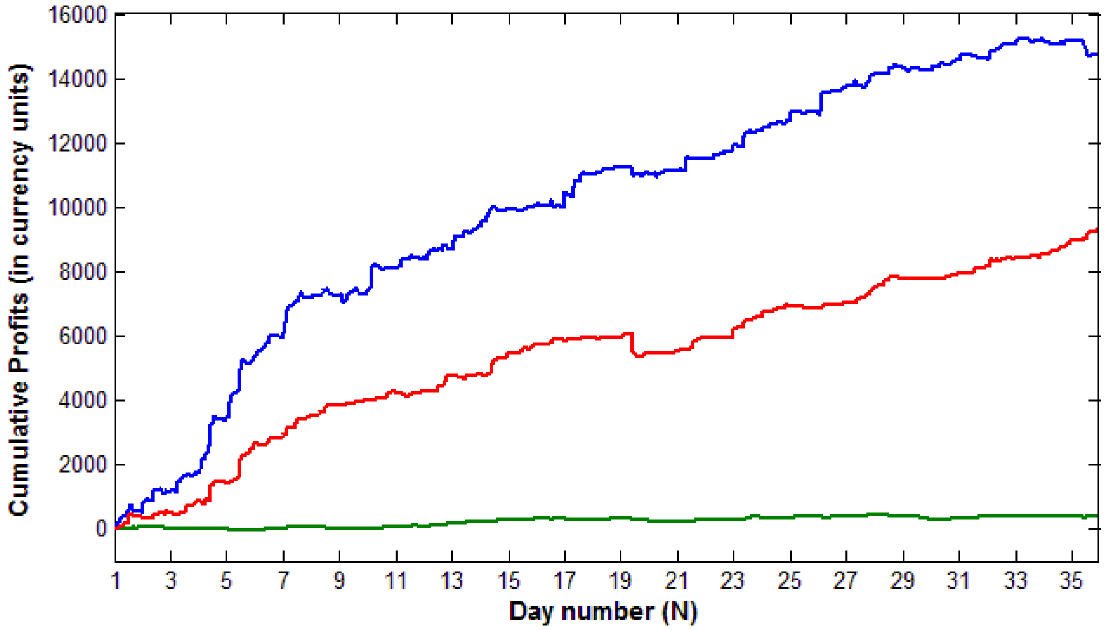

Figure 5.7. Cumulative Profits. Profits gained by applying strategy (11)-(12) to the three contract pairs over the 35 days in the sample. The vertical axis is measured in Euros for the FTSE - DAX and S&P - DAX pairs, and Pounds Sterling for the S&P - FTSE pair.

the investor to post a *margin* equivalent to a small percentage of the contract's notional value (Figlewski (1984), Reilly (1999), and Hancock (2005)). The exact amount demanded by clearing houses to initiate a futures contract trade can vary depending on the recent historical volatility of the contract. There is a risk that, given an adverse market move, the investor may receive a margin call, or face having to close part of a trading position at an unfavourable point in time. However, in the context of the trading exercise presented here, this is unlikely, since each trading position is held for a number of seconds.

Although each profit time-series trends consistently upwards, they exhibit vast differences in magnitude. The most profitable pair is the FTSE - DAX, while the least profitable is the S&P - FTSE. One reason for this may relate to the geographical proximity of London to Frankfurt: It is likely that the FTSE and DAX contracts respond to a larger set of overlapping idiosyncratic factors than either transatlantic pair, yielding a more consistent relationship. This suggestion is confirmed by a greater magnitude of correlation, shown in Figure 5.5. Further, the low profitability which characterizes the S&P - FTSE pair may be explained by Werner and Kleidon's (1996) suggestion that the US-UK market is among the most heavily arbitraged.

The profits in Figure 5.7 are measured in absolute, not percentage terms, for a number of reasons: First, the arbitrageur does not know *a priori* the number of trading opporutnities that would be presented on a given day. Works which present returns in percentage terms such as Gatev et al. (2006) and De Jong et al. (2009) rely on the full re-investment of capital per trading opportunity. This is clearly not feasible here, since one cannot purchase futures contracts in fractional amounts. Second, measuring profits in absolute terms makes this work comparable to Suarez (2005), who measures the returns to a direct arbitrage strategy in ADRs in absolute terms.

Now a study of the intraday profile of arbitrage is presented. Specifically, the analysis aggregates over the entire sample, the directional forecasting accuracy, number of generated signals, and durations of profitable disequilibria. The results are shown in Figure 5.8.

The top panel of Figure 5.8 shows a consistently high directional forecasting accuracy across all three contract pairs. Accuracy rates rise slightly following the US open, which accords with the idea that the UK and German markets follow the US more closely during this period. Further, accuracy rates fall following the UK and German close, which is also an intuitive result exhibiting the opposite effect. Following Huth and Abergel (2012), the analysis presented here tests the robustness of the directional forecasts against both a random forecast, and a forecast generated via an auto-correlation in the lagging contract. The auto-correlation forecast slightly outperforms the random forecast, but is not statistically significantly better. The directional forecasts presented are

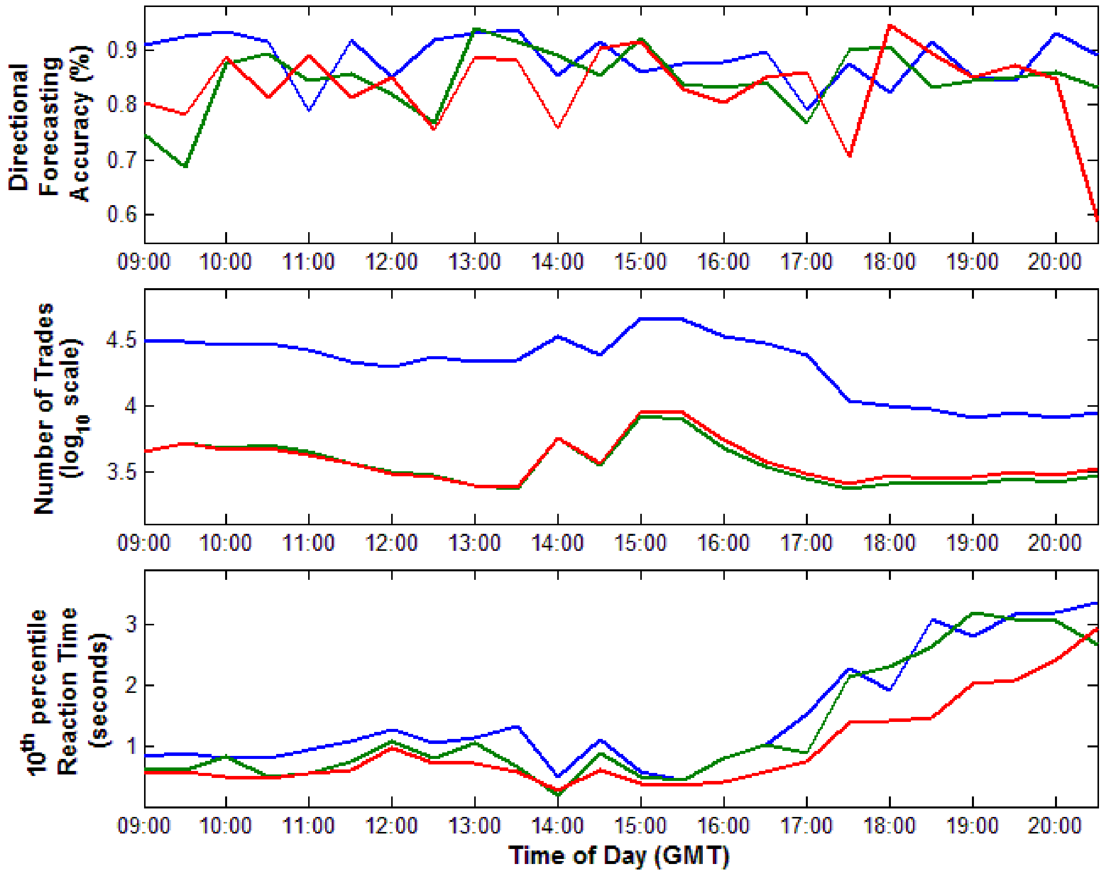

Figure 5.8. Intraday Arbitrage Profile. Intraday patterns of Directional Forecasting Accuracy (top), the total number of individual trades on a $\log_{10}$ scale (middle), and the 10th percentile duration length of profitable disequilibria (bottom).

significant at the 99% level for all contract pairs over all intervals in the sample, with the exception of the S&P - DAX pair between 20:00 and 20:30.

The middle panel of Figure 5.8 shows that the number of trades (or equivalently, trade signals) generated per interval increases sharply following the announcement of market data and again following the US open. This increase is then maintained for as long as all three cash markets remain open. The bottom panel plots the 10th-percentile *reaction times*, denoting interval durations between the point at which a trade signal is generated, and the point at which the lagging contract begins to move following the signal. This panel shows that reaction times decrease following the announcement of market data, and remain low for as long as all three cash markets are open, before gradually rising following the close of the UK and

German markets. Taken together, these results suggest that market participants who act to appropriate temporal disequilibria concentrate their activity during specific periods.

5.4.4 Varying the Number of Contracts per Trade

So far, this chapter has assumed that the arbitrageur trades one single futures contract each time a trading signal is generated. In reality, statistical arbitrage systems of this kind are only profitable if they can be scaled to yield economically significant profits.

Here, the trading exercise is repeated 20 times. At each repetition, the arbitrageur is assumed to trade a number R of contracts per trade signal, where $R \in \{1, 2, \ldots, 20\}$. Table 5.4 shows the results of this exercise.

# Contracts	FTSE - DAX	S&P500 - DAX	S&P500 - FTSE	Total Profits
1	14,028	9,647	592	24,266
2	30,265	20,059	1,546	51,870
3	46,503	30,472	2,500	79,474
4	62,740	40,884	3,454	107,078
5	78,978	51,297	4,408	134,682
6	95,215	61,709	5,362	162,286
7	111,453	72,122	6,316	189,890
8	127,690	82,534	7,270	217,494
9	143,928	92,947	8,224	245,098
10	160,165	103,359	9,178	272,702
11	176,403	113,772	10,132	300,306
12	192,640	124,184	11,086	327,910
13	208,878	134,597	12,040	355,514
14	225,115	145,009	12,994	383,118
15	241,353	155,422	13,948	410,722
16	257,590	165,834	14,902	438,326
17	273,828	176,247	15,856	465,930
18	290,065	186,659	16,810	493,534
19	306,303	197,072	17,764	521,138
20	322,540	207,484	18,718	548,742

Table 5.4. Number of Contracts per Trade. This table shows the effect on profits of varying the number of contracts per trade from a single contract to 20. Profits are converted to Euros for ease of comparison, and totalled in the final column.

Interestingly, as the arbitrageur trades a larger number of contracts, the strategy benefits from economies of scale with regard to trading costs. To see this, note that the profits from trading 10 contracts per signal exceed ten times the profits gained by trading 1. This is because exchange-related costs are charged per market order, not per contract. Clearly, one cannot trade a large number of contracts simultaneously without incurring market impact costs. This is discussed in section 5.4.6.

5.4.5 Varying the Signal Threshold

Up to now, this work has assumed a pre-set signal threshold K. It is informative to study the effects of varying this threshold on the profitability of arbitrage.

Intuitively, a higher signal threshold yields better directional forecasts profitability, at the expense of a lower trading frequency. Studying the effects of varying the signal threshold is informative to both academics and practitioners. To academics, it establishes a link between the frequency of temporal price disequilibria and risk. To practitioners, it permits an investigation into optimal trading rules.

Let us proceed by varying the signal threshold from $K \in \{1, 2, \ldots, 10\}$ ticks. For each instance and for each contract pair, several statistics are measured, namely the aggregate directional forecasting accuracy, total number of trades, average profit per trade, and 10th percentile reaction times. The results are shown in Table 5.5.

Table 5.5 reveals several insights into the effects of varying the signal threshold on arbitrage returns. The effects are consistent across all contract pairs. Firstly, increasing the signal threshold directly increases the directional forecasting accuracy. This result is consistent with the notion that increasing the signal

Threshold (ticks)	1	2	3	4	5	6	7	8	9	10
Directional Forecasting Accuracy (%)										
FTSE / DAX	70.87%	83.31%	90.02%	92.11%	92.28%	92.21%	92.61%	90.90%	89.64%	89.45%
S&P / FTSE	66.02%	66.55%	93.42%	93.72%	94.25%	93.69%	100.00%	100.00%	100.00%	100.00%
S&P / DAX	68.54%	69.26%	92.75%	93.32%	90.79%	91.62%	84.13%	84.13%	83.33%	83.33%
Total Number of Trades										
FTSE / DAX	276,645	186,012	68,112	21,913	6,604	2,210	900	406	216	133
S&P / FTSE	47,844	44,135	595	539	67	55	14	13	5	5
S&P / DAX	49,559	46,461	799	766	70	66	15	14	4	4
Profit per Trade (EUR or GBP)										
FTSE / DAX	-€ 14.62	-€ 9.84	-€ 4.60	-€ 0.36	€ 3.41	€ 6.35	€ 11.43	€ 12.45	€ 15.20	€ 13.29
S&P / FTSE	-£7.49	-£7.36	£0.44	£0.92	£3.77	£5.26	£19.44	£20.98	£40.44	£40.44
S&P / DAX	-€ 15.26	-€ 14.71	€ 12.08	€ 12.59	€ 28.64	€ 32.14	€ 30.67	€ 25.79	€ 2.13	€ 2.13
10th Percentile Reaction Time (milliseconds)										
FTSE / DAX	1,166	1,277	1,432	1,668	1,628	1,130	769	549	275	162
S&P / FTSE	1,177	1,255	1,218	1,068	82	115	64	58	37	37
S&P / DAX	1,024	1,055	456	459	135	142	178	174	112	112

Table 5.5. Varying the Signal Threshold This table shows, for each contract pair, the effect of varying the signal threshold K from 1 to 10 ticks. For each instance, the table shows the directional forecasting accuracy, total number of trades, profit per trade, and 10th percentile reaction time.

threshold results in a *cleaner* signal with which to trade. Similarly, while the total number of trades falls sharply, the profit per trade increases. In particular, despite low signal thresholds yielding directionally accurate forecasts, trading based on low signal thresholds does not overcome costs. To practitioners, these results can be used to devise optimal trading strategies.

The final three rows of Table 5.5 are interesting. The 10th percentile reaction times decline sharply as the signal threshold increases. In other words, arbitrageurs who wait for stronger trade signals do so at the cost of facing drastically shorter durations of disequilibria. This is consistent with the view that competition among arbitrageurs acts to enforce pricing efficiency (Grossman and Stiglitz (1976, 1980)), thereby perpetuating the need to invest in technological infrastructure. Clearly, this has a limit in the form of the speed of light. Over time, the payoff to investing in infrastructure would yield diminishing marginal returns.

It is informative to compare the reaction times stated in Table 5.5 to the time it takes data to physically traverse the relevant geographical distances. Recall from section 5.4.1 that one-way transmission between New York and London takes 32.4 milliseconds via Global Crossing's AC-1 fibre-optic cable. By extrapolating this time to the distances between the CME, NYSE-Liffe, and Eurex, it is highly unlikely that any of the FTSE - DAX or S&P - DAX trades lie inside the transmission speed boundary, even for high signal thresholds. This is particularly striking for the FTSE - DAX pair, which benefits from close geographical proximity, and which has been shown to be more profitable than either transatlantic pair. It is therefore puzzling why the 10th percentile reaction time for this pair should be significantly higher than either transatlantic pair. However, looking closely at the profits per trade is informative: The FTSE - DAX pair generates a large number of trades, each of which is relatively less profitable than either transatlantic pair. This suggests that the FTSE - DAX pair contains a relatively high idiosyncratic risk component which acts as a deterrent to arbitrage.

Looking at the S&P - FTSE pair, a signal threshold $K \geq 9$ suggests that for a number of arbitrage opportunities, namely those at 37 milliseconds, the duration of disequilibria lies at or just within the data transmission speed barrier. However, it is noteworthy that in the entire sample, this phenomenon occurs for only one single trade. Overall, the S&P - FTSE result is consistent with Werner and Kleidon's (1996) suggestion that the US-UK market is among the most heavily arbitraged.

Finally, it is important to note that the computing processes behind generating a trading signal and managing a position require only the evaluation of one logical operation and one floating-point operation per observed mid-quote, as per (11) and (12). Modest desktop computing resources can perform these kinds of calculations within a small number of microseconds, akin to suggestions by Hasbrouck and Saar (2012). As for the millisecond environment in which this exercise operate, it is unlikely that the time required for calculation will significantly affect the visible duration of temporal disequilibria.

5.4.6 On the Limits of Arbitrage

The data set employed in this chapter consists of the three most liquid and widely traded futures contracts globally, with no restrictions on trading or cross-border arbitrage.

In markets such as this, professional arbitrageurs often compete over a limited supply of available arbitrage opportunities. This competition creates an imbalance stemming from the excess demand by arbitrageurs for profitable trades. As a result, competition among arbitrageurs effectively reduces the magnitude and duration of profitable disequilibria. This in turn creates a natural selection mechanism by which arbitrageurs who possess the most powerful computing resources survive. Overall, this system poses two major limits to arbitrage, namely costs and risk.

Liquidity, price slippage, market impact, and competition for scarce opportunities from other traders are all sources of risk for which arbitrageurs demand compensation. This fact may explain why even seemingly *riskless* arbitrage opportunities sometimes go unexploited. It is informative to explore the effects each of these risks have on arbitrageurs' capacity to appropriate pricing disequilibria, by examining their effects on profitability.

The subsequent analysis parsimoniously captures the effects of liquidity risk, price slippage, market impact, and competition from other arbitrageurs, following an approach by Kozhan and Tham (2012). Specifically, this section imposes that a

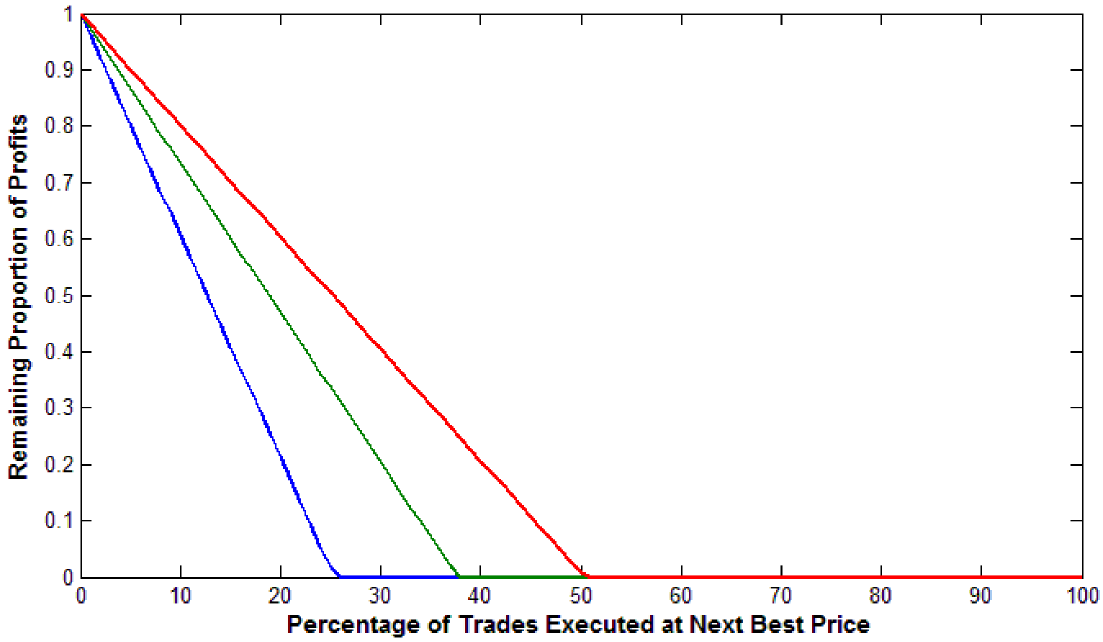

Figure 5.9. Slippage Risk. The effect of having a percentage of trades fulfilled at the next best price (horizontal) on the proportion of otherwise available profit opportunities (vertical).

certain proportion of the arbitrageur's trades are not executed at the observed (best) price, but at the next best price instead, which is assumed to be available in infinite supply. The next best price in the context of this work means a one-tick increment down the limit order book, equivalent to GBP 5 and EUR 12.50 for the FTSE 100 and DAX contracts, respectively. The objective of this exercise is to see what effect this restriction has on the profitability of the individual arbitrageur. The results are shown in Figure 5.9.

For each contract pair, Figure 5.9 reveals the sensitivity of an arbitrageur's profits to the risk of executing part of a trade at the next best price. This sensitivity is particularly high for the FTSE-DAX pair, which is an intuitive result, given the generally low per-trade profitability of this pair as seen in Figure 5.8.

Further, one can interpret the results of this section within the context of Table 5.4, and ask: How is the probability of price slippage affected by increasing the number of contracts per trade? In the data set, around 29% of FTSE 100 quotes on either the bid or ask side carry a size of 5 contracts or less, and around 46% carry 11 or more contracts. The corresponding figures for the DAX are 49% and 22% respectively. Therefore, the FTSE 100 futures contract carries a greater per-quote liquidity. Based on this, one can suggest that the risk of price slippage is relatively less profound when the contract being traded is the FTSE 100.

Besides risk, the costs associated with trading in this market setting act as a significant deterrent to arbitrage. Figure 5.8 shows the directional forecasting accuracy to be high. However, in order to profit from a trade, an arbitrageur has to not only be directionally accurate, but also overcome the bid-ask spread. Given that the majority of trades individually extract very small amounts of profit, the bid-ask spread is a major cost incurred by arbitrageurs. Furthermore, arbitrageurs face fixed rents and other costs incurred through co-location subscriptions and access to live market feeds. Hedge funds implementing these kinds of strategies must also hire analysts to create and maintain computer code. To the extent that these extra costs collectively diminish the profits reported in Table 5.4, they act as a deterrent to arbitrage. Due to the costs involved, one could suggest that arbitrage operations in this setting can most readily be undertaken by larger hedge funds and proprietary trading desks whose co-located infrastructure is not dedicated solely to this pursuit, but forms part of their overall business activity.

Overall, this section highlights that technological infrastructure costs and the risk of price slippage present significant limits to arbitrage in the market for international futures contracts.

5.5 - Conclusion

This chapter complements previous work by examining pricing relationships in the important but overlooked market setting of international index futures. This chapter documents clear evidence of consistent sub-second lead/lag patterns across the S&P 500, FTSE 100, and DAX futures contracts (in order of leadership), suggesting that the diffusion of information across these markets is not instantaneous.

Importantly, this chapter shows that these lead/lag patterns induce predictability in lagging assets, and give rise to profitable disequilibria. Taken together, these observations provide evidence against auto-efficiency across international futures markets, and suggest that arbitrage is an important component to information flow across markets. However, arbitrage profits are sensitive to the risk of price slippage, and profitable opportunities rarely exist for more than 300 milliseconds. Therefore, slippage risk and technological costs are highlighted as the most significant limits to arbitrage in this market setting.

The data employed consists of three highly liquid futures contracts, with no capital controls and no restrictions to cross-border arbitrage. Based on this, this work suggests that the results obtained in this chapter are generalizable to other global markets. Furthermore, The importance of utilizing high-frequency data is highlighted, based on compelling evidence that most of the price adjustment mechanisms in these markets operate deep into the sub-second domain.

In this chapter, a key assumption has been that the profile of the temporal correlation structure between any two futures contracts would remain so out-of-sample. It may be that the temporal relationship is seasonal through time, particularly with the FTSE-DAX pair, which exhibits bi-directional causality. However, the correlation profile was demonstrably stable through the sample period, even when the sample period was split into individual days. It could therefore be argued that this assumption is not unreasonable.

Future work could focus on extending the lead/lag analytical framework into the multi-dimensional space. This would facilitate the analysis of global flows of information and collectively transitive effects across multiple geographical markets.

Chapter 6

Conclusion

The primary goal of this thesis has been to answer three big questions in Financial Economics: 1) To what extent does aversion to fundamental risk diminish the extent to which arbitrageurs are willing to correct pricing disequilibria? 2) In the absence of fundamental risk, how effective is arbitrage at exploiting pricing anomalies? 3) To what extent does arbitrage appropriate geographical informational delays across markets?

Chapter 2 provided a review of the main relevant theoretical, empirical, and methodological bodies of literature on which the answers to the above questions are based. Chapters 3-5 then answered these questions, respectively. In addition to reflecting on the three main questions, opportunities arose along the way to make additional contributions and insights which would hopefully benefit future research.

The remainder of this chapter proceeds as follows: Section 6.1 provides a summary of the gaps in the extant literature. Section 6.2 summarizes the contributions that this thesis makes to the extant literature. Section 6.3 outlines the implications of these findings. Section 6.4 discusses the limitations of the work presented in this thesis. Finally, section 6.5 suggests potential directions for future research.

6.1 - The Gaps in the Extant Literature

The first goal of this thesis was to explore the behaviour of arbitrageurs who are averse to fundamental risk. Fundamental risk is an important component of the propensity of arbitrage to correct pricing disequilibria, and denotes the risk of deterioration in the fundamental relationship between prices of stocks which are the subject of an arbitrage trade. Chapter 2 presented a summary of the empirical literature which clearly recognizes the importance of this risk. While the intuitive and obvious effect of aversion to fundamental risk would be for arbitrageurs to trade more sceptically, the exact mechanics of this are poorly understood, because as yet no quantitative characterization of this risk has been incorporated into an arbitrageur's theoretical portfolio allocation problem. Understanding the exact relationship between fundamental risk and arbitrage is important to developing a fuller understanding of the relationship between arbitrage and price efficiency.

The second goal of this thesis was to further our understanding about the extent to which arbitrage can correct pricing anomalies that occur in an environment with virtually no fundamental risk, namely the UK stock-ADR market. There is an inconsistency in the extant literature that is presented in Chapter 2: On the one hand, the body of empirical market microstructure works has documented a clear segmentation between the US markets on which ADRs trade, and the domestic markets on which the stock trades. In theory, such an environment should constitute fertile ground for arbitrageurs, because the two securities in question are perfect substitutes. On the other hand, the extant empirical literature finds little evidence that arbitrage is profitable in this environment. The gap in the literature stems from the fact that existing works have focussed exclusively on analysing the returns to direct arbitrage, and not statistical arbitrage. The latter is important in addressing the puzzle that stock-ADR relationships appear to exhibit price efficiency despite the fact that neither arbitrageurs nor investors in each of the two markets are incentivised to bring prices back into line.

The third goal that this thesis has set out to achieve is motivated by evidence of increasing global correlations across markets. If markets were perfectly integrated and frictionless, then any non-idiosyncratic price shock to one market would be diffused instantaneously and reflected in prices across the other markets, so as to preclude the possibility of cross-border arbitrage. In reality, frictions impede the flow of information, which means that the transmission of information across correlated markets gives rise to lead/lag patterns. Whether arbitrage is successful at appropriating such geographical pricing disequilibria adds an important dimension to the debate as to whether return-predictability exists in financial markets. However, the returns to arbitrage in this setting, particularly in the high-frequency domain, have not been explored.

6.2 - Contributions to the Literature

This section summarizes the theoretical and empirical contributions that this thesis makes to the literature.

Chapter 3 is the first work to endogenously characterize the important concept of fundamental risk into a quantitative model of arbitrage trading. This chapter introduces to the literature a model that makes use of the shape of the hyperbolic tangent function to capture the idea that arbitrageurs grow increasingly sceptical towards prices which continue to diverge away from fundamental values after an arbitrage trade is initiated. This model specification has its roots in Physics, but is otherwise totally new to the literature on arbitrage. Moreover, the model presented contains the Ornstein-Uhlenbeck model, which is favoured by the extant theoretical literature, as a special case.

Based on the model, several theoretical inferences can be drawn that complement the existing empirical literature on the behaviour of arbitrageurs facing fundamental risk, and quantify the ideas presented in the empirical literature on fundamental risk discussed in Chapter 2: First, arbitrageurs who are averse to fundamental risk exhibit a diminishing propensity to exploit increasingly large mispricings away from parity. Second, aversion to fundamental risk precipitates the gradual unwinding of losing trades far sooner than is implied given a complete confidence in the pricing relationship. Empirically, it is shown that arbitrageurs' profit-seeking is enhanced by aversion to fundamental risk, by way of prudent position management in the face of what would otherwise be interpreted as highly attractive trading opportunities.

Chapter 4 carries out the first examination of a high frequency pairs trading strategy in the stock-ADR market. The main empirical contribution of this work is to show, in contrast to the prior literature, that arbitrage is successful at exploiting pricing anomalies in this market, and is therefore an important part of the price-correcting mechanism between stocks and ADRs.

Besides the main empirical result, this work also presents a number of key insights into the role of arbitrage in enforcing parity between stocks and ADRs. While there is clearly no fundamental risk inherent in stock-ADR arbitrage, arbitrageurs do face significant deterrents associated with exploiting pricing anomalies in this market: First, because stocks and ADRs are virtually the same security and are two-way fungible, mispricings in this market are small, and exist for a median duration of less than 10 minutes. It is shown that stock-ADR arbitrage can only be profitably undertaken by those with the lowest transaction costs. Second, arbitrageurs who choose to exploit the narrowest mispricings between stocks and ADRs face significant holding costs stemming from uncertainty toward the time of price convergence. Pairs trading necessitates holding trading positions in both the stock and ADR open while waiting for convergence, and so this uncertainty acts as a deterrent to arbitrage.

Chapter 5 provides an empirical contribution to the literature by way of uncovering the existence of sub-second exploitable pricing disequilibria stemming from lead/lag effects across international index futures contracts. It also contributes to the literature by employing a recent advance in the statistical estimation of lagged correlation based on Hoffman et al. (2010), and applying it to analysing the lagged cross-correlation profile of futures contract data time-stamped to the millisecond.

In so doing, this work finds, in contrast to the extant literature, that arbitrage is successful at appropriating informational delays across markets. However, such informational delays are only exploitable by those with direct access to the exchanges, and who are financially able to employ the most advanced technological infrastructure with which to trade. To put this in some context, the work presented here shows that the most profitable trading opportunities caused by geographical informational disequilibria rarely last for longer than 300 milliseconds.

6.3 - Implications of the Findings in this Thesis

This thesis provides a deeper understanding of the important relationship between arbitrage and price efficiency. Moreover, the specific contributions that this thesis makes weigh in on a number of contemporary debates in Financial Economics.

The existence of profitable disequilibria across international markets weighs in on an age-old debate surrounding the existence of predictability in asset returns. In particular, it contrasts assertions made in Fama (1965), Fama (1991), Fama and French (1992), Fama (1995), and Pontiff (1996, 2006) among many others within the deep literature on asset returns. Not only can returns can be forecast, but can be done so with accuracy in excess of 85%.

Furthermore, the work presented in this thesis reconciles a puzzle in the market for cross-listed securities. Specifically, how can the stock-ADR market exhibit a high degree of price efficiency when neither arbitrageurs (Miller and Morey (1996), Suarez (2005), Gagnon and Karolyi (2010)) nor investors in each of the individual markets (Werner and Kleidon (1996), Chen et al. (2009)) are incentivised to correct pricing anomalies? The application of a pairs trading strategy as a mechanism of parity enforcement solves this puzzle. Put simply, the extant literature had not considered the full range of tools available to arbitrageurs. Under realistic transaction costs, this thesis shows that arbitrage can profitably exploit pricing anomalies in the UK stock-ADR market, a highly developed and deeply liquid market in which by definition arbitrage opportunities are unlikely to exist. By that measure, arbitrage is an important price-correction mechanism in the market for cross-listed securities.

Overall, the conclusions presented in this thesis accord with the view asserted in Grossman and Stiglitz (1976, 1980) that arbitrage opportunities exist to incentivise arbitrageurs to correct pricing disequilibria. However, chapters 4 and 5 of this thesis find that exploitable opportunities exist in the sub-minute and sub-second space. It is unlikely that Grossman and Stiglitz (1976, 1980) had this frequency in mind. The important implication here is that it is not surprising that works which employ low-frequency data such as Kato et al. (1991), Wahab and Lashgari (1992), and Park and Tavakkol (1994) conclude in favour of price efficiency and the preclusion of arbitrage as a parity enforcement mechanism.

6.4 - Limitations of this Work

This section outlines the limitations of the work presented in this thesis.

Because pairs trading entails keeping open one long and one short position, it is therefore vulnerable as a strategy to institutional factors such as short-selling bans. In particular, the empirical period over which the model presented in chapter 3 is tested contains a 4-month period during which short-selling was banned in financial stocks in the UK. Nevertheless, the bulk of the empirical period studied in this chapter had no such restrictions. Further, it could be argued that while incorporating the short-selling ban would obviously impact the numerical result of arbitrage profits, the theoretical and qualitative insights discussed in section 6.2 which form the main contributions of chapter 3 to the literature would be unaffected.

Furthermore, short-selling involves borrowing stock in order to sell, with the promise of returning the stock to the owner in the future. In real markets, the arrangement for borrowing stock is not trivial. An investor who wishes to short stocks must request a specific amount of stock from a broker before the start of trading each day, and the broker in turn attempts to find this stock either from inventory or long-term institutional investors. It is not uncommon with less liquid stocks that the requested amount of stock is unavailable, and this would affect the empirical results presented in chapter 4 as it would hinder the ability of arbitrageurs to exploit stock-ADR mispricings. Nevertheless, it is shown that the firms upon which this study is based are highly liquid, such that the likelihood of no short-stock availability is negligible.

Finally, in order for general inferences to be valid, the length of the historical data period under empirical examination should be sufficiently long, so as to include a range of market settings relevant to the main research questions. The length of each of the data periods in chapters 4 and 5 are four and two months, respectively. However, each of these chapters is concerned with phenomena that occur primarily in the high-frequency domain, and between them, include a number of price observations in the hundreds of millions, and a total number of data points in the billions. It could be argued that these chapters analyse a richer set of historical data even than that of chapter 3, which includes 9 years of daily price observations.

6.5 - Avenues for Future Research

The applicability of the model presented in chapter 3 to future research in the study of arbitrage behaviour is twofold: First, it demonstrably admits price dynamics rich enough to simulate a wide range of price behaviour. Second, the specification of the model is simple enough to yield closed-form solutions to an arbitrageur's portfolio optimization problem. Such a model would be desirable for study in a general-equilibrium framework, wherein both the arbitrageur's trading strategy, and that strategy's effect on the wider market, is studied jointly. General-equilibrium models do not readily admit closed-form solutions (Xiong, 2001) due to the complexity of modelling the effect of trading behaviour on price formation. Thus, incorporating a rich but simple model, such as that presented in Chapter 3, would be useful to future research in this area.

It would be informative to apply the pairs trading strategy examined in chapter 4 to other markets, particularly South American stock-ADR markets. This would complement research by Hong and Susmel (2003), Rabinovitch et al. (2003), Grammig et al. (2005), and Grossman et al. (2007), who study the prevalence of stock-ADR arbitrage in a number of different countries. The appeal of applying a pairs trading strategy to these markets is threefold: First, pairs trading circumvents country-specific institutional restrictions on the conversion between stocks and ADRs (discussed in Rabinovitch et al. (2003), Puthenpurackal (2006), and Gagnon and Karolyi (2010)). Therefore, applying a pairs trading strategy would add a new and important dimension to the study of whether arbitrage is successful at eliminating pricing disequilibria in these markets. Second, South American trading hours overlap US trading hours almost entirely, and would hence provide a deep source of empirical data on which to test the prevalence of arbitrage. Third, these markets provide an excellent platform on which to test what and to what extent the risks inherent in emerging economies act as deterrents to arbitrage.

Chapter 5 provided the first implementation of the modified Hayashi-Yoshida correlation estimator (Hayashi and Yoshida (2005) and Hoffman et al. (2010)) to study correlations across international markets. This work provides a platform upon which both theoretical and empirical advances could be made. Empirically, the same methodology could readily be applied in a much wider setting, covering not only geographical informational disequilibria, but also futures contracts of different maturities, and the relationship between futures and exchange-traded funds. On the theoretical front, this work motivates the need to expand the Hayashi-Yoshida estimator to study lead/lag effects in more than two dimensions (i.e. more than one leader per lagger). One application of this is to facilitate the study of whether information stemming from two or more disparate markets is pertinent to a third. Under the premise that global markets are increasingly integrated (Ayuso and Blanco (2001), Kearney and Lucey (2004), Evans and Hnatkovska (2005)), it is important to analyse not just the pair-wise transmission of price shocks, but rather the effect of price shocks on all markets jointly.

References

Abhyankar, A. (1998). Linear and Nonlinear Granger Causality: Evidence From the UK Stock Index Futures Market. *Journal of Futures Markets* 18, pp. 519 -540.

Abhyankar, A., L.S. Copeland, W. Wong (1997). Uncovering Nonlinear Structure in Real-Time Stock-Market Indexes: The S&P 500, the DAX, the Nikkei 225, and the FTSE-100. *Journal of Business & Economic Statistics* 15(1), pp. 1-14.

Abreu, D. and M.K. Brunnermeier (2002). Synchronization Risk and Delayed Arbitrage. *Journal of Financial Economics* 66, pp. 341-360.

Abreu, D. and M.K. Brunnermeier (2003). Bubbles and Crashes. *Econometrica* 71, pp. 173-204.

Andreou, P.C. and Y.A. Pierides (2008). Empirical Investigation of Stock Index Futures Market Efficiency: The Case of the Athens Derivatives Exchange. *European Journal of Finance* 14, pp. 211-23.

Antoniou, A., G. Pescetto, A. Violaris (2003). Modelling International Price Relations and Interdependencies between the Stock Index and Stock Index Futures Markets of Three EU Countries: A Multivariate Analysis. *Journal of Business Finance & Accounting* 30, pp. 645-667.

Ayuso, J. and R. Blanco (2001). Has Financial Integration Increased During the 1990s? *Journal of International Financial Markets, Institutions and Money* 11, pp. 265-287.

Bachelier, L. (1900), Thorie de la spculation, Annales Scientifiques de lcole Normale Suprieure 3 (17): 2186. Barua, A., J. Ranganayakulu, N. Parakh, A.K. Sharma (2012).

SEBI Bulletin. *Securities and Exchange Board of India* 10(4).

Basak, S. and B. Croitoru (2000). Equilibrium Mispricing in a Capital Market with Portfolio Constraints. *The Review of Financial Studies* 13, pp. 715-748.

Bellman, R.E. (1957). Dynamic Programming. *Princeton University Press*, Princeton, NJ. Republished 2003: Dover, ISBN 0486428095.

Bjork, T. (2004). Arbitrage Theory in Continuous Time. *Oxford University Press.*

Black, F. (1986). Noise. *The Journal of Finance* 41, pp. 529-543.

Bodie, Z., A. Kane, A. Marcus (2009). Investments (8th edition). *McGraw- Hill-Irwin Limited, New York.*

Boguslavsky, M. and E. Boguslavskaya (2004). Arbitrage Under Power *RISK* 17(6), pp. 69-73.

Breiman, L. (1961). Optimal Gambling System for Favorable Games. *Proceedings of the 4th Berkeley Symposium on Mathematical Statistics and Probability.* University of California Press, Berkeley. Vol. 1, pp. 63-68.

Brennan, M.J. and E.S. Schwartz (1990). Arbitrage in Stock Index Futures. *Journal of Business* 63, pp. 7-31.

Brooks, C. (2008). Introductory Econometrics for Finance (2th edition). *Cambridge University Press, Cambridge and New York.*

Brooks, C., A.G. Rew, and S. Ritson (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, pp. 31-44.

Brooks, C., and I. Garrett (2002). Can We Explain the Dynamics of the UK FTSE 100 Stock and Stock Index Futures Markets? *Applied Financial Economics* 12, pp. 25-31.

Broumandi, S., T. Reuber (2012). Statistical Arbitrage and FX Exposure with South American ADRs Listed on the NYSE. *Financial Assets and Investing* 3(2), pp. 5-18.

Brunnermeier, M.K. and L.H. Pedersen (2005). Predatory trading. *Journal of Finance* 60, pp. 1825-1863.

Chaboud, A., B. Chiquoine, E. Hjalmarsson, C. Vega (2009). Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market. Working Paper. *Board of Governors of the Federal Reserve System.*

Chen H., G. Noronha, V. Singal (2004). The Price Response to S&P 500 Index Additions and Deletions: Evidence of Asymmetry and a New Explanation. *Journal of Finance* 59, pp. 1901-29.

Chen, J., Y. Tse, M. Williams (2009). Trading Location and Equity Returns: Evidence from US trading of British Cross-listed Firms. *Journal of International Financial Markets, Institutions and Money* 19, pp. 729-741.

Cheung, Y.W., S.L. Kon, M. Bergman (2004). Dissecting the PPP Puzzle: The Unconventional Roles of Nominal Exchange Rate and Price Adjustments. *Journal of International Economics* 64, pp. 135-150.

Chordia, T., R.W. Roll, A. Subrahmanyam (2011). Recent Trends in Trading Activity and Market Quality. *Journal of Financial Economics* 101, pp. 243-263.

Cummings, J. and A. Frino (2011). Index Arbitrage and the Pricing Relationship Between Australian Stock Index Futures and Their Underlying Shares. *Accounting and Finance* 51, pp. 661-683.

D'Avolio, G. (2002). The Market for Borrowing Stock. *Journal of Financial Economics* 66, pp. 271-306.

De Jong, A., L. Rosenthal, M. van Dijk (2009). The Risk and Return of Arbitrage in Dual-Listed Companies. *Review of Finance* 13, pp. 495-520.

De Jong, F. and T. Nijman (1998). Intraday Lead-Lag Relationships Between the Futures, Options, and Stock market. *European Finance Review* 1, pp. 337-359.

De Long, J.B., A. Shleifer, L.H. Summers, R.J. Waldmann (1990). Noise Trader Risk in Financial Markets. *Journal of Political Economy* 98, pp. 703-738.

Easley, D. and M. O'Hara (1987). Price, Trade Size and Information in Securities Markets. *Journal of Financial Economics* 19, pp. 69-90.

Edwards, F.R. (1999). Hedge Funds and the Collapse of Long-Term Capital Management. *Journal of Economic Perspectives* 13, pp. 189-210.

Eiteman, D., A. Stonehill, M.H. Moffett, and M. Moffet (1994). Multinational Business Finance. *Addison Wesley*. ISBN: 9780201845532.

Elliott, R.J., J. van der Hoek, W.P. Malcolm (2005). Pairs trading. *Quantitative Finance* 5(3), pp. 271-276.

Engle, R.F. and C.W.J. Granger (1987). Co-Integration and Error Correction: Representation, Estimation and Testing. *Econometrica* 55, pp. 149-158.

Epps T. (1979). Co-movements in Stock Prices in the Very Short Run. *Journal of the American Statistical Association* 74, pp. 291-298.

Eun, C. and S. Sabberwal (2003). Cross-border Listings and Price Discovery: evidence from U.S.-listed Canadian stocks. *Journal of Finance* 58, pp. 549-577.

Eun, C. and S. Shim (1989). International Transmission of Stock Market Movements. *Journal of Financial and Quantitative Analysis* 24(2), pp. 241-256.

Evans, M.D. and V. Hnatkovska (2005). International Capital Flows, Returns and World Financial Integration. *NBER Working Paper* 11701.

Fama, E. (1965). The Behavior of Stock-Market Prices. *The Journal of Business* 38(1), pp. 34-105.

Fama, E. (1991). Efficient Capital Markets: II. *Journal of Finance* 46(5), pp. 1575-617.

Fama, E. Random Walks in Stock Market Prices. *Financial Analysts Journal* 51(1), pp. 75.

Fama, E. and K. French (1992). The Cross-Section of Expected Stock Returns. *Journal of Finance* 47(2), pp. 427-65.

Figlewski, S. (1984). Hedging Performance and Basic Risk in Stock Index Futures. *Journal of Finance* 39, pp. 657-670.

Figlewski, S. (1985). Hedging with Stock Index Futures: Theory and Application in a New Market. *Journal of Futures Markets* 5, pp. 183-200.

Friedman, M. (1953). The Case For Flexible Exchange Rates. *Essays in Positive Economics, University of Chicago Press.*

Froot, K.A. and E.M. Dabora (1999). How are Stock Prices Affected by the Location of Trade? *Journal of Financial Economics* 53, pp. 189-216.

Gagnon, L. and G.A. Karolyi (2010). Multi-Market Trading and Arbitrage. *Journal of Financial Economics* 97, pp. 53-80.

Gatev, E., W.N. Goetzmann, K.G. Rouwenhorst (2006). Pairs Trading: Performance of a Relative Value Arbitrage Rule. *Review of Financial Studies* 19, pp. 797-827.

Goetzmann, W., L. Li, and K.G. Rouwenhorst (2005). Long-Term Global Market Correlations. *Journal of Business* 78(1), pp. 1-38.

Grammig, J., M. Melvin, C. Schlag (2005). Internationally Cross-Listed Stock Prices During Overlapping Trading Hours: Price Discovery and Exchange Rate Effects. *Journal of Empirical Finance* 12, pp. 139-164.

Griffin, J.E. and R.C.A. Oomen (2011). Covariance Measurement in the Presence of Non-Synchronous Trading and Market Microstructure Noise. *Journal of Econometrics* 160(1), pp. 58-68.

Gromb, D. and D. Vayanos (2002). Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs. *Journal of Financial Economics* 66, pp. 361-407.

Grossman, S. and J. Stiglitz (1976). Information and Competitive Price Systems. *American Economic Review* 66, pp. 246-252.

Grossman, S. and J. Stiglitz (1980). On the Impossibility of Informationally Efficient Markets. *American Economic Review* 70, pp. 393-408.

Grossmann, A., T. Ozuna, M. Simpson (2007). ADR Mispricing: do Costly Arbitrage and Consumer Sentiment Explain the Price Deviation? *Journal of International Financial Markets, Institutions and Money* 17, pp. 361-371.

Gulen, H. and S. Mayhew (2000). Stock Index Futures Trading and Volatility in International Equity Markets. *Journal of Futures Markets* 20(7), pp. 661-685.

Hamao, Y., R. Masulis, and V. Ng (1990). Correlations in Price Changes and Volatility Across International Stock Markets. *Review of Financial Studies* 3(2), pp. 281-307.

Hancock, G. D. (2005). A Text Book Treatment of Calculating Returns on Non-Traditional Portfolios. *Review of Financial Economics* 14(2), pp. 173–186.

Hasbrouck, J. and G. Saar (2012). Low-Latency Trading. *Journal of Finance* 16(4), pp. 646-679.

Hayashi, T. and N. Yoshida (2005). On Covariance Estimation of Non-Synchronously Observed Diffusion Processes. *Bernoulli* 11(2), pp. 359-379.

Hendershott, T., C.M. Jones, and A.J. Menkveld (2011). Does Algorithmic Trading Improve Liquidity? *Journal of Finance* 66, pp. 1-33.

Herbst, A.F., J.P. McCormack, and E.N. West (1987). Investigation of a Lead-Lag Relationship between Spot Stock Indices and Their Futures Contracts. *Journal of Futures Markets* 7, pp. 373-381.

Hoffmann, M., M. Rosenbaum, and N. Yoshida (2010). Estimation of the Lead-Lag Parameter from Non-Synchronous Data. *Working Paper.*

Hong, G. and R.Susmel (2003). Pairs-Trading in the Asian ADR Market. *working paper, Available at: http://www.bauer.uh.edu/rsusmel/Academic/ptadr.pdf.*

Hume, D. (1741). Of the Balance of Trade, in Essays: Moral, Political and Literary. Oxford University Press, (1963 reprint).

Hurn, A.S., K.A. Lindsay, V.L. Martin (2003). On the Efficacy of Simulated Maximum Likelihood for Estimating the Parameters of Stochastic Differential Equations. *Journal of Time Series Analysis* 24(1), pp. 45-63.

Huth, N. and F. Abergel (2012). High Frequency Lead/lag Relationships: Empirical facts. *Working Paper, available at arXiv:1111.7103.*

Innocenti, A., P. Malpenga, L. Menconi, and A. Santoni (2011). Intraday Anomalies in the Relationship between U.S. Futures and European Stock Indexes. *Journal of Index Investing* 1(4), pp. 40-52.

IOSCO (2012). Technological Challenges to Effective Market Surveillance Issues and Regulatory Tools. *The Board of the International Organization of Securites Commissions.* August 2012.

Jegadesh, N. and S. Titman (1993). Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *Journal of Finance* 48, pp. 65-91.

Johnson, N., G. Zhao, E. Hunsader, J. Meng, A. Ravindar, S. Carran, B. Tivnan (2012). Financial Black Swans Driven By Ultrafast Machine Ecology. *University of Miami working paper, available at arXiv:1202.1448.*

Jorion, P. and E. Schwartz (1986). Integration vs. Segmentation in the Canadian Stock Market. *Journal of Finance* 41, pp. 603-616.

Jurek, J. and H. Yang (2007). Dynamic Portfolio Selection in Arbitrage. *Harvard University working paper.*

Kargin, V. (2003). Optimal Convergence Trading. *Working Paper, available at arXiv:math/0302104.*

Kato, K., S. Linn, J. Schallheim (1991). Are There Arbitrage Opportunities in the Market for American Depositary Receipts? *Journal of International Financial Markets, Institutions and Money* 1, pp. 73-89.

Kawaller, I.G., P.D. Koch, and T.W. Koch (1987). The Temporal Price Relationship Between S&P 500 Futures and the S&P 500 Index. *Journal of Finance* 42(5), pp. 1309-1329.

Kearney, C. and B.M. Lucey (2004). International Equity Market Integration: Theory, Evidence and Implications. *International Review of Financial Analysis* 13(5), pp. 571-583.

Kenourgios, D.F. (2004). Price Discovery in the Athens Derivatives Exchange: Evidence for the FTSE/ASE-20 Futures Market. *Economic and Business Review* 6, pp. 229-243.

Kim, S.J., J.A. Primbs, S. Boyd (2008). Dynamic Spread Trading. *Stanford University Working Paper.*

Kim, T.S. and E. Omberg (1996). Dynamic Nonmyopic Portfolio Behavior. *The Review of Financial Studies* 9, pp. 141-161.

Klibanoff, P., O. Lamont, T. Wizman (1998). Investor Reaction to Salient News in Closed-End Country Funds. *Journal of Finance* 53(2), pp. 67399.

Kloeden, P.E. and E. Platen (1999). Numerical Solution of Stochastic Differential Equations. Springer, Berlin. ISBN 978-3-540-54062-5.

Knight, F.H. (1921). Risk, Uncertainty, and Profit. *Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Company.*

Kondor, P. (2009). Risk in Dynamic Arbitrage: Price Effects of Convergence Trading. *Journal of Finance* 64, pp. 638-658.

Korn, R. and H. Kraft (2002). A Stochastic Control Approach to Portfolio Problems with Stochastic Interest Rates. *SIAM Journal on Control and Optimization* 40(4), pp. 1250-1269.

Kozhan, R. and W.W. Tham (2012). Execution Risk in High-Frequency Arbitrage. *Management Science* 58(11), pp. 2131-2149.

Kyle, A.S. and W. Xiong (2001). Contagion as a Wealth Effect. *Journal of Finance* 56, pp. 1401-1440.

Lamont, O. and R. Thaler (2003a). The Law of One Price in Financial Markets. *Journal of Economic Perspectives* 17(4), pp. 191-202.

Lamont, O. and R. Thaler (2003b). Can the Market Add and Subtract? Mispricing in Tech Stock Carve-outs. *Journal of Political Economy* 111, pp. 227-268.

Liu, J and F.A. Longstaff (2004). Losing Money on Arbitrage: Optimal Dynamic Portfolio Choice in Markets with Arbitrage Opportunities. *Review of Financial Studies* 17, pp. 611-641.

Liu, J. (2007). Portfolio Selection in Stochastic Environments. *Review of Financial Studies* 20(1), pp. 1-39.

Lo, A. and C. MacKinlay (1990). When are Contrarian Profits Due to Stock Market Overreaction? *Review of Financial Studies* 3, pp. 175-205.

Lv, Y. and B.K. Meister (2009). Application of the Kelly Criterion to Ornstein-Uhlenbeck Processes. *Renmin University of China working paper, available at arXiv:0903.2910.*

MacKenzie, D. (2003). Long-Term Capital Management and the Sociology of Arbitrage. *Economy and Society* 32, pp. 349-380.

Maldonado, R. and A. Saunders (1983). Foreign Exchange Restrictions and the Law of One Price. *Financial Management* 12, pp. 19-23.

Malkiel, B. (1977). The Valuation of Closed-End Investment-Company Shares. *The Journal of Finance* 32, p. 847-859.

Marshall, B.R., N.H.and Nguyen, N. Visaltanachoti (2010). ETF Arbitrage. *Working Paper.*

Medeiros, O., G. de Oliveira, B. van Doornik (2009). Testing for Lead-Lag Effects Between the American and the Brazilian Stock Markets. *University of Brasilia Working Paper.*

Merton, R.C. (1969). Lifetime Portfolio Selection: The Continuous-Time Case. *Review of Economics and Statistics* 51, pp. 247-257.

Merton, R.C. (1971). Optimum Consumption and Portfolio Rules in a Continuous-Time Model. *Journal of Economic Theory* 3, pp. 373-413.

Miller, D. and M. Morey (1996). The Intraday Pricing Behavior of International Dually Listed Securities. *Journal of International Financial Markets, Institutions and Money* 6, pp. 79-89.

Mitchell, M., T. Pulvino, E. Stafford (2002). Limited Arbitrage in Equity Markets. *Journal of Finance* 57, pp. 551-584.

Mudchanatongsuk, S., J. Primbs, W. Wong (2008). Optimal Pairs Trading: A Stochastic Control Approach. *Proceedings of the American Control Conference, Seattle, WA* pp. 1035-1039.

Nam, D. (2011). The Roles of Nominal Exchange Rate and Relative Price Adjustments in PPP Reversion. *Journal of Money, Credit and Banking* 43, pp. 775-785.

Osbat, C., R. Ruffer, B. Schnatz (2003). The Rise of Yen Vis--vis The “Synthetic” Euro: Is It Supported By the Economic Fundamentals? *Working paper No. 224, European Central Bank.*

Park, J.W. and A. Tavakkol (1994). Are ADRs a Dollar Translation of their underlying Securities?: The Case of Japanese ADRs. *Journal of International Financial Markets, Institutions and Money* 4, pp. 77-87.

Picchini, U. (2007). SDE Toolbox: Simulation and Estimation of Stochastic Differential Equations with MATLAB. . *http://sdetoolbox.sourceforge.net.*

Pontiff, J. (1996). Costly Arbitrage: Evidence From Closed-End Funds. *Quarterly Journal of Economics* 111, pp. 1135-1151.

Pontiff, J. (2006). Costly Arbitrage and the Myth of Idiosyncratic Risk. *Journal of Accounting and Economics* 42, pp. 35-52.

Pulatkonak, M. and G. Sofianos (1999). The Distribution of Global Trading in NYSE-Listed non-U.S. Stocks. *NYSE Working Paper No. 99-03.*

Puthenpurackal, J. (2006). Explaining Premiums in Restricted DR Markets and Their Implications: the Case of Infosys. *Financial Management* 35, pp. 93-116.

Rabinovitch, R., A.C. Silva, R. Susmel (2003). Returns on ADRs and Arbitrage in Emerging Markets. *Emerging Markets Review* 4, pp. 225-247.

Reilly, F. K., K. C. Brown (1999). Investment Analysis and Portfolio Management, (7th edition), South Western Educational Publishing.

Richie, N., R. Daigler, and K. Gleason (2008). The Limits to Stock Index Arbitrage: Examining S&P 500 Futures and SPDRs. *Journal of Futures Markets* 28(12) pp. 1182-1205.

Rosenthal, L. (1983). An Empirical Test of the Efficiency of the ADR Markets. *Journal of Banking and Finance* 7, pp. 17-29.

Schultz, P. and S. Shive (2010). Mispricing of Dual-Class Shares: Profit Opportunities, Arbitrage, and Trading. *Journal of Financial Economics* 98, pp. 524-549.

Shaw, W. and M. Schofield (2012). A Model of Returns for the Post-Credit-Crunch Reality: Hybrid Brownian Motion with Price Feedback. *Quantitative Finance*, forthcoming.

Shleifer A. (1986). Do Demand Curves for Stocks Slope Down? *Journal of Finance* 41, pp.579-90.

Shleifer, A. and L.H. Summers (1990). The Noise Trader Approach to Finance. *The Journal of Economic Perspectives* 4(2), pp. 19-33.

Shleifer, A. and R. Vishny (1997). The Limits of Arbitrage. *Journal of Finance* 52, pp. 35-55.

Stoll, H.R. and R.E. Whaley (1990). The Dynamics of Stock Index and Stock Index Futures Returns. *Journal of Financial and Quantitative Analysis* 25, pp. 441-468.

Suarez, E. (2005). Arbitrage Opportunities in the Depositary Receipts Market: Myth or Reality? *Journal of International Financial Markets, Institutions and Money* 15, pp. 469-480.

Thomas, S. (2002). The Saga of the First Stock Index Futures Contract: Benchmarks, Models, and Learning. *Journal of Money, Credit, and Banking.* 34(3), pp. 767-808.

Tse, Y.K. (1995). Lead-Lag Relationships between Spot Index and Futures Price of the Nikkei Stock Average. *Journal of Forecasting* 14, pp. 553-563.

Vasicek, O. (1977). An Equilibrium Characterization of the Term Structure. *Journal of Financial Economics* 5, pp. 177-188.

Voev, V. and A. Lunde (2007). Integrated Covariance Estimation Using High Frequency Data in the Presence of Noise. *Journal of Financial Econometrics* 5(1), pp. 68-104.

Wachter, J. (2002), Portfolio and Consumption Decisions under Mean-Reverting Returns: An Exact Solution for Complete Markets. *The Journal of Financial and Quantitative Analysis* 37(1), pp. 63-91.

Wahab, M., and M. Lashgari (1993). Price Dynamics and Error Correction in Stock Index and Stock Index Futures Markets: A Cointegration Approach. *Journal of Futures Markets* 13, pp. 711-742.

Wahab, M., M. Lashgari, R. Cohn (1992). Arbitrage Opportunities in the American Depositary Receipts Market Revisited. *Journal of International Financial Markets, Institutions and Money* 2, pp. 97-130.

Wang, P., and P. Wang (2001). Equilibrium Adjustment, Basis Risk and Risk Transmission in Spot and Forward Foreign Exchange Markets. *Applied Financial Economics* 11, pp. 127-136.

Werner, I. and A. Kleidon (1996). U.K. and U.S. Trading of British Cross-Listed Stocks: An Intraday Analysis of Market Integration. *Review of Financial Studies* 9, pp. 619-664.

Wong, E. (1964), The Construction of a Class of Stationary Markov Processes. *Proceedings of the American Mathematical Society Symposium on Applied Mathematics* 16, pp. 264-276.

Xiong, W. (2001). Convergence Trading with Wealth Effects: An Amplification Mechanism in Financial Markets. *Journal of Financial Economics* 62, pp. 247-292.

Zeckhauser, R. and V. Niederhoffer (1983). The Performance of Market Index Futures Contracts, *Financial Analysts Journal* 39(1), pp. 59-65.