Currency Trios - Using Geometric Concepts to Visualise and Interpret Relationships between Currencies

Abby Davidson

A dissertation submitted to the Faculty of Commerce, University of Cape Town, in partial fulfilment of the requirements for the degree of Master of Philosophy.

November 20, 2016

MPhil in Mathematical Finance,
University of Cape Town.

Declaration

I declare that this dissertation is my own, unaided work. It is being submitted for the Degree of Master of Philosophy in the University of the Cape Town. It has not been submitted before for any degree or examination in any other University.

Signed

November 20, 2016

Abstract

A currency trio is a set of three currencies and their respective exchange rates, which have a relationship fixed by a triangular arbitrage condition. This condition forms the basis for the derivation of a geometric interpretation of the relationships between the exchange rates. In the geometric framework, the three currencies in a currency trio are represented by a triangle, where each of the vertices represents a currency. The volatilities of the exchange rates are represented by the lengths of the sides joining the respective currencies and the cosine of each angle represents the correlation between the two exchange rates depicted by the angle's adjacent sides. The geometric approach is particularly useful when dealing with implied data as it allows the calculation of implied correlation using implied volatility. This is valuable as implied volatility is frequently quoted in the foreign exchange market; whereas, implied correlation is not directly quoted and is more difficult to extract from market data. This dissertation aims to thoroughly investigate the geometric framework and use it to visualise and interpret the relationships between currencies in a currency trio. The analysis will initially look at currency trios with realised spot data before moving on to implied data. In the implied data context, the framework will be used to extract and evaluate implied correlation estimates using implied volatility data extracted from the foreign exchange market. The framework will be extended to investigate whether an illiquid option can be proxy hedged using options on the two other currencies in a currency trio. Finally, the findings will be discussed and the feasibility of the applications of the framework will be considered.

Acknowledgements

I would like to express my thanks to my supervisor, Obeid Mahomed, for his guidance, and to my external supervisors, Daniel Polakow and Gideon Van de Linde for their input. Lastly to my family, thank you for everything, your continuous support over this period has been invaluable.

Contents

List of Figures

List of Tables

Nomenclature

USD- United States Dollar

ZAR- South African Rand

EUR- Euro

GBP- British Pound

JPY- Japanese Yen

CNY- Chinese Yuan

MXN- Mexican Peso

BRL- Brazilian Real

KES- Kenyan Shilling

NGN- Nigerian Naira

CHF- Swiss Franc

KRW- South Korean Won

AUD- Australian Dollar

INR- Indian Rupee

Chapter 1

Introduction

The foreign exchange market is one of the largest and most liquid financial markets, with participants across the world. This market facilitates the exchange of two currencies, referred to as a currency pair, at a certain rate and is fundamental to each country as exchange rates affect trade, foreign investment and interest and inflation rates. The instruments traded in this market include transactions on spot exchange rates as well as a number of currency derivatives (Levinson, 2014).

Certain exchange rates, in particular the more illiquid exchange rates, cannot be directly traded and are priced with reference to a third, more liquid currency such as the USD. This has led to the notion of a currency trio; a set of three currencies linked by the exchange rates between them. Theoretically, the three exchange rates are related by an arbitrage condition, so that the conversion between the currencies does not result in a risk free profit (Levinson, 2014). While pricing discrepancies do occur in practice, with the rise of algorithmic trading and faster systems, risk-free arbitrage opportunities in the foreign exchange spot market are infrequent and disappear quickly (Ito *et al.*, 2012).

Singer *et al.* (1998) describe how the derivation of this arbitrage condition forms the basis for a geometric interpretation of currency trios. In the geometric framework, the three currencies and exchange rates in a currency trio are represented by a triangle and each of the triangle vertices represents a currency. The volatility of the exchange rates is represented by the lengths of the sides joining the corresponding vertices and the cosine of each angle represents the correlation between the two exchange rates depicted by the angle's adjacent sides. This geometric approach does not require any assumptions about the distribution of the underlying exchange rates and it provides an intuitive way to view the relationships between the volatilities and correlations of the exchange rates in a currency trio.

The geometric approach can be extended from relationships between spot exchange rates to implied data. Implied volatility and implied correlation are forward looking measures that represent the market's expectation of future volatility and correlation (Hull, 2012). Implied volatility is easily extracted from simple market instruments and is quoted directly in the foreign exchange market. However, implied correlation is not quoted directly and is more difficult to extract from market instruments. The geometric approach gives an expression for correlation in terms of volatility, providing a method to calculate implied correlation from implied volatility using simple trigonometry.

This dissertation aims to thoroughly investigate the geometric framework and use it to visualise and interpret the relationships between the three currencies in a currency trio. Furthermore, it seeks to use the framework to calculate implied correlations between exchange rates using quoted implied volatilities. Finally this dissertation will assess the hedging applications of this framework. Chapter 2 will outline the derivation of the geometric relationship and further consider the extensions to this framework. Following this discussion, Chapter 3 will explain the concepts of volatility and correlation and the estimation of these measures, both in the historical and implied contexts. Derivative instruments such as variance swaps, which allow direct trading on implied volatility, will also be discussed. In Chapter 4, different methodologies used in hedging option positions will be explained and the application of the framework to hedging illiquid exchange rate options will be considered.

The theory developed will be applied to simulated and market data in Chapter 5. The analysis will initially consider currency trios with simulated and realised spot data before moving on to implied data. In the implied data context, the framework will be used to calculate and evaluate implied correlation estimates. Finally, the feasibility of using the framework to hedge illiquid exchange rate options will be investigated. Chapter 6 will include a brief discussion of the findings and give recommendations for extensions to this study.

Chapter 2

Geometric Interpretation

The geometric framework is derived from the triangular arbitrage condition between the three currencies in a currency trio. In this chapter, the derivation of the geometric relationship and literature on the geometric approach will be discussed. Furthermore, the benefits and possible extensions to this framework will be considered.

2.1 Triangular Arbitrage

The triangular relationship is derived from the triangular arbitrage condition. Triangular arbitrage is a strategy that takes advantage of a pricing discrepancy between the exchange rates of three different currencies (Gradojevic and Gençay, n.d.). The group of three currencies is referred to as a currency trio and the relationship between the three exchange rates of these currencies is defined by

$$S_{X/Y} = \frac{S_{X/Z}}{S_{Y/Z}}, \tag{2.1.1}$$

where $S_{X/Y}$ represents the exchange rate between currencies X and Y, i.e. the amount of currency X that can be exchanged for one unit of currency Y. Y is referred to as the base or domestic currency and X is referred to as the foreign currency.

In order to implement a triangular arbitrage strategy, one currency is successively converted to each of the other currencies in the currency trio before it is converted back to the original currency. If there are any pricing discrepancies, this could yield a risk free profit. There are two different routes that can be taken to implement a triangular arbitrage strategy, ($X \to Y \to Z \to X$ *or* $X \to Z \to Y \to X$).

While this is a relatively simple strategy, it may not be profitable in practice. The bid-ask spreads involved in the conversion process need to be considered. If the bid-ask spreads are too large, the potential profit is lost during the conversions (Gradojevic and Gençay, n.d.). Furthermore, the trades need to take place at the moment after the discrepancy occurs as discrepancies are quickly traded out and risk-free arbitrage opportunities are infrequent and fleeting (Ito *et al.*, 2012).

2.2 Geometric Interpretation of a Currency Trio

Geometric concepts have been used to analyse a number of financial problems, such as optimising portfolio selection in Markowitz (1952) and more recently to understand the relationships between currencies and exchange rates in Singer *et al.* (1998). The geometric relationship is derived from the triangular arbitrage condition, which was described in Equation (2.1.1). The derivation is demonstrated in Equation (2.2.1). The starting point for the derivation is the equation for the relationship between three currencies,

$$S_{X/Y} = \frac{S_{X/Z}}{S_{Y/Z}}.$$

This equation can be differentiated to obtain

$$dS_{X/Y} = \frac{1}{S_{Y/Z}} dS_{X/Z} - \frac{S_{X/Z}}{S_{Y/Z}^2} dS_{Y/Z}.$$

By dividing through by $S_{X/Y}$ we obtain

$$dS_{X/Y} = \frac{1}{S_{Y/Z}} \frac{S_{Y/Z}}{S_{X/Z}} dS_{X/Z} - \frac{S_{X/Z}}{S_{Y/Z}^2} \frac{S_{Y/Z}}{S_{X/Z}} dS_{Y/Z}.$$

Finally this expression is simplified to obtain

$$\frac{dS_{X/Y}}{S_{X/Y}} = \frac{dS_{X/Z}}{S_{X/Z}} - \frac{dS_{Y/Z}}{S_{Y/Z}}, \tag{2.2.1}$$

where $\frac{dS_{X/Y}}{S_{X/Y}}$ is the return on exchange rate $S_{X/Y}$.

The results of the derivation in Equation (2.2.1) shows that each exchange rate return in a currency trio can be expressed as a linear combination of the two other exchange rate returns. One of the stylised facts of the financial markets is that asset returns are stationary with respect to one another, so the correlations between the exchange rate returns should converge to the actual correlation (Cont, 2001a). Consequently, the equation for the variance of a sum of correlated variables can be applied to Equation (2.2.1). The resulting equation is

$$\sigma^2_{X/Y} = \sigma^2_{X/Z} + \sigma^2_{Y/Z} - 2\rho_{X/Z-Y/Z} \times \sigma_{X/Z} \times \sigma_{Y/Z}, \tag{2.2.2}$$

where $\sigma_{X/Z}$ is the volatility of the return of exchange rate $S_{X/Y}$ and $\rho_{X/Z-Y/Z}$ is the correlation between the returns of exchange rates $S_{X/Z}$ and $S_{Y/Z}$.

Equation (2.2.2) is easily rearranged to express the correlation in terms of the volatilities. The rearranged equation is expressed as

$$\rho_{X/Z-Y/Z} = \frac{\sigma^2_{X/Z} + \sigma^2_{Y/Z} - \sigma^2_{X/Y}}{2 \times \sigma_{X/Z} \times \sigma_{Y/Z}}. \tag{2.2.3}$$

Equation (2.2.2) can be compared to the expression for the law of cosines,

$$a^2 = b^2 + c^2 - 2 \times b \times c \times \cos(A). \tag{2.2.4}$$

Equation (2.2.2) and Equation (2.2.4) have the same form, with the correlation term between the two exchange rates ($\rho_{X/Z-Y/Z}$) equivalent to $\cos(A)$. The comparison between these equations forms the basis of the triangular framework with the three currencies each representing a vertex of a triangle and the sides representing the exchange rates between the currencies. The lengths of the triangle sides represent the volatilities of the exchange rate returns and the angles of the triangle represent the correlations between the exchange rate returns of the adjacent sides. In alignment with market practice, volatilities of the exchange rate returns are defined by the annualised standard deviations of logarithmic returns (Singer *et al.*, 1998). Figure 2.1 visually represents the geometric representation of the exchange rates between the three currencies in a currency trio.

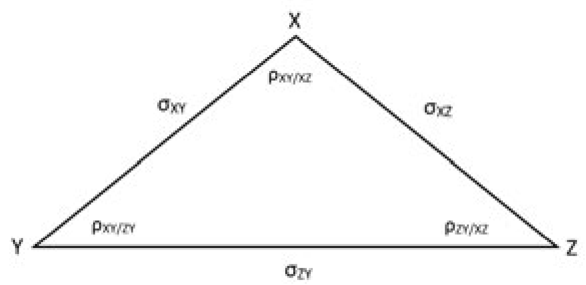

Fig. 2.1: Geometric representation of three currencies, X, Y and Z

Singer *et al.* (1998) is one of the first studies to specifically look at the geometric relationship between exchange rates. Singer *et al.* (1998) attempts to provide an intuitive explanation for this relationship using simple numerical examples. Walter and Lopez (2000a) take a more mathematically rigorous approach to explaining the relationship, by deriving the equation for the arbitrage condition as demonstrated in (2.2.1) and (2.2.2).

Singer *et al.* (1998) and Zerolis (1998) extend the geometric framework to include four currencies, creating tetrahedrons that display the volatility and correlation relationships. Singer *et al.* (1998) explain that the extension loses its intuitive nature as more dimensions are added; however, this extension has important implications for pricing foreign exchange options involving more than three currencies as consistent implied correlation estimates are required in the pricing (Shevchenko, 2009). The extension of the framework to four currencies is derived from Equation (2.2.1),

$$\begin{aligned} \frac{dS_{Z/Y}}{S_{Z/Y}} &= \frac{dS_{X/Y}}{S_{X/Y}} - \frac{dS_{X/Z}}{S_{X/Z}} \\ \frac{dS_{X/W}}{S_{X/W}} + \frac{dS_{Z/Y}}{S_{Z/Y}} &= \frac{dS_{X/W}}{S_{X/W}} + \frac{dS_{X/Y}}{S_{X/Y}} - \frac{dS_{X/Z}}{S_{X/Z}}. \end{aligned} \tag{2.2.5}$$

The formula for the sum of correlated variables can now be applied to the left and right hand side of Equation (2.2.5),

$$\begin{aligned} LHS: \sigma^2 &= \sigma^2_{X/W} + \sigma^2_{Z/Y} - 2\rho_{X/W-Z/Y}\sigma_{X/W}\sigma_{Z/Y} \\ RHS: \sigma^2 &= \sigma^2_{X/W} + \sigma^2_{X/Y} + \sigma^2_{X/Z} + 2\rho_{X/W-X/Y}\sigma_{X/W}\sigma_{X/Y} \\ &\quad -2\rho_{X/W-X/Z}\sigma_{X/W}\sigma_{X/Z} - 2\rho_{X/Z-X/Y}\sigma_{X/Z}\sigma_{X/Y}. \end{aligned} \tag{2.2.6}$$

All of the correlation terms in the right hand side expression of Equation (2.2.6) can be replaced with the corresponding volatility terms using Equation (2.2.3). The left and right hand side expressions are then equated to get the equation for $\rho_{X/W-Z/Y}$,

$$\rho_{X/W-Z/Y} = \frac{\sigma^2_{X/Y} + \sigma^2_{W/Z} - \sigma^2_{X/Z} - \sigma^2_{W/Y}}{2\sigma_{X/W}\sigma_{Z/Y}}. \tag{2.2.7}$$

The geometric framework can be extended to more currencies; however, the extension to four different currencies allows the relationship to be applied to any exchange rates, so the relationship is not constrained to the exchange rates in a currency trio. Furthermore, the relationship loses its intuitive meaning as more dimensions are added.

The advantages of the geometric approach are described by Singer *et al.* (1998), who argue that the geometric framework provides an intuitive way to portray information as it graphically conveys concepts such as the relationship between volatility and correlation and how these values interact. Another benefit of this approach is described by Fretheim and Høigaard (2012). Fretheim and Høigaard (2012) explain that the arbitrage condition maintains the geometric relationship so distributional assumptions are not required. This is advantageous as the assumption of lognormality is particularly problematic for exchange rates, as they often exhibit larger than expected movements in the market as a result of changes in monetary policies or other factors (Hull, 2012).

2.3 Extensions of the Geometric Framework

The geometric framework is particularly useful in the interpretation of implied data from foreign exchange options, as it allows the conversion from implied volatility to implied correlation. Implied volatilities and correlations can be extracted from different derivatives which are traded in the market, which will be explained further in Section 3. These values represent the market's expectation of future volatilities and correlations for a specific maturity and strike price (Hull, 2012). The majority of the literature on implied data focusses on extracting implied volatilities, and implied volatilities are quoted daily in the foreign exchange market. In contrast, implied correlations are difficult to extract from market data. However, Equation (2.2.3) can be used to calculate implied correlation using implied volatility as shown in Equation (2.3.1),

$$\rho_{X/Z-Y/Zimplied} = \frac{\sigma^2_{X/Zimplied} + \sigma^2_{Y/Zimplied} - \sigma^2_{X/Yimplied}}{2 \times \sigma_{X/Zimplied} \times \sigma_{Y/Zimplied}}. \tag{2.3.1}$$

Forward looking estimates of correlations are important as they are used in the pricing of exotic foreign exchange options such as quanto options, which is an option on an asset denominated in one currency is where the payoff is denominated in a different currency (Shevchenko, 2009). The extraction and characteristics of the implied correlation measures will be discussed further in Section 3.3.

Walter and Lopez (2000b) extend their study of the framework to include an analysis of volatility and correlation changes over time. The geometric representation makes it easy to see the impact of economic changes and the stability of the volatility and correlation. A further application of this approach is to look at the trading ideas that arise from this representation. The framework gives a clear representation of when correlations or volatilities display unusual signals, which could inform trading strategies.

In other financial applications, Singer (1999) uses the geometric framework in the context of portfolio and risk management. A portfolio is a group of financial assets such as stocks and bonds held by an investor. The performance of a portfolio is generally measured against a benchmark, such as a broad market index. The risk of a portfolio (portfolio volatility) relative to a benchmark can be decomposed into the benchmark risk (benchmark volatility) and a tracking error risk (tracking error volatility). These three volatility components can be represented using the triangular relationships. This aids in the intuitive understanding of how the portfolio risk changes when the correlation between the portfolio and the benchmark or the benchmark volatility change.

Chapter 3

Volatility and Correlation

The concepts of volatility and correlation are intrinsically linked to the triangular framework. Volatility is an estimate of the uncertainty of the returns of an asset, while correlation is a measure of how two assets move in relation to one another (Hull, 2012).

Volatility and correlation are important concepts in financial markets. These concepts are key metrics in risk management and portfolio selection as volatility gives an estimate of the risk of asset returns and correlation shows how the returns of an asset change in relation to other assets in the portfolio. Volatility and correlation also play an important role in the pricing and hedging of derivative instruments as they are primary inputs in the pricing calculations of exotic options. While volatility has historically been an important financial metric, with the increase in correlation dependent derivatives, such as quantos, correlation is becoming an increasingly significant measure (Fretheim and Høigaard, 2012).

There are many different methods and models used in the estimation of volatility and correlation. This study will only consider realised and implied volatility and correlation. Realised volatility and correlation are estimates of volatility and correlation measures determined from historical data, while implied volatility and correlation are the market's expectation of the future value of these measures (Hull, 2012).

In this chapter, the notion of realised and implied volatility and correlation and the methods used to evaluate them will be introduced. The concept of implied volatility will be investigated further and the foreign exchange volatility surface will be briefly discussed. Following this discussion, variance swaps, which are market instruments used to trade future volatility, will be analysed. The theory behind the pricing of variance swaps provides a methodology to calculate a fair implied volatility estimate for a particular maturity across all strike prices. Finally, implied correlation and the literature on this measure will be considered.

3.1 Evaluating Volatility and Correlation

There are differing methods employed by financial practitioners to measure volatility and correlation. The simplest method of evaluating volatility is by taking the standard deviation of the log returns of an asset. This standard deviation is effectively the average daily volatility which must be converted to an annual volatility using an annualisation factor determined by the square root of time rule. This rule accounts for the idea that standard deviation is proportional to the square root of time (Alexander, 1998). Consequently, simple realised volatility is defined by

$$vol_X = \sigma_x \times \sqrt{N}, \tag{3.1.1}$$

where σ_x is the standard deviation of the log returns of asset X and N denotes the number of observations in the evaluation period.

Correlation is generally more unstable than volatility and there is significantly less literature on estimating and evaluating correlation in comparison to the literature considering volatility (Vonhoff, 2006). Furthermore, the returns of two assets must be jointly stationary for their correlation to be a coherent measure (Alexander, 1998). A generally accepted stylised fact of financial assets is that the log returns are stationary so this will not be investigated further (Cont, 2001b). Simple realised correlation is defined by

$$corr_{X,Y} = \frac{cov_{X,Y}}{\sigma_Y \times \sigma_X}, \tag{3.1.2}$$

where $cov_{X,Y}$ denotes the covariance between the log returns of asset X and Y.

There are different methods used to calculate realised correlation and volatility. Simple realised estimates of these measures are problematic as they assign an equal weight to all observations in the time period. This can lead to slow reflection of new information and large changes in volatility and correlation levels due to distant observations (Fretheim and Høigaard, 2012). To compensate for this issue, methods of determining realised volatility that assign a higher weight to recent observations, such as the exponentially weighted moving average (EWMA) method are often used (Hull, 2012).

Predictions of future of volatility and correlation are central to managing portfolios and risk. They are also particularly important in the derivative market, as these future values are primary inputs to derivative pricing formulas. The best estimates of the future levels of volatility and correlation are commonly regarded to be implied volatility and correlation. These measures will be discussed further in Section 3.2 and 3.3.

3.2 Implied Volatility

Implied volatility is an estimate of future volatility derived from vanilla instruments traded in the foreign exchange market (Hull, 2012). It is considered to be the market's best estimate of future volatility levels, taking into account both future expectations and historical information (Poon and Granger, 2005). There are numerous studies focussing on predicting future volatility, which predominantly conclude that implied volatility is the best predictor as it incorporates the most information (Poon and Granger, 2005).

Implied volatilities are extracted from simple market instruments such as call or put options, as volatility is the only unobservable variable in the pricing formula (Hull, 2012). The Garman-Kohlhagen equation, which extends the Black-Scholes formula to foreign exchange derivatives, is considered to be the standard pricing formula in the foreign exchange market (Clark, 2011). The Garman-Kohlhagen equation for the price V_t of a call or put option on an exchange rate is given as

$$\begin{aligned} V_t &= \omega S_0 N(\omega d_+) e^{-r_f(T-t)} - KN(\omega d_-)e^{-r_d(T-t)} \\ d_+ &= \frac{ln(\frac{S_0}{K}) + (r_d + r_f + \frac{\sigma^2}{2})(T-t)}{\sigma\sqrt{(T-t)}} \\ d_- &= d_+ - \sigma\sqrt{(T-t)}, \end{aligned} \tag{3.2.1}$$

where S_0 is the spot exchange rate, K is the strike rate, σ is the implied volatility of the exchange rate, T is maturity of the option, r_f and r_d are the foreign and domestic interest rates respectively and ω is +1 for call options and -1 for put options.

The implied volatility, σ, which is a primary input for Equation (3.2.1) is quoted directly as a function of the strike price, maturity of an option and time and can be obtained from sources such as Bloomberg (Vonhoff, 2006). The implied volatilities on Bloomberg are quoted in the form of a volatility surface; made up of five implied volatility values, which are generally quoted against five strike prices for a certain maturity. The method of extracting implied volatility estimates from the foreign exchange market will be discussed further in Section 3.2.1.

In order to price instruments such as variance swaps, which are a pure trade on future volatility, a single estimate of implied volatility for a particular maturity is required, as opposed to the five quoted volatility-strike pairs. There are differing methods used to obtain a single estimate of implied volatility. The simplest approach is to use the at-the-money (ATM) delta-neutral volatility quote, as ATM volatilities imply the most information about future volatility levels (Vonhoff, 2006). While this is generally a fairly accurate prediction, it does not take all the information inherent in the volatility surface into account. The theory of variance swap pricing suggests a method of extracting a single model-free estimate of implied volatility. Demeterfi, Derman, Kamal and Zou (1999) derive a method to determine the fair volatility strike, which is a prediction of future volatility levels, using replication arguments. This will be discussed further in Section 3.2.2.

3.2.1 The Foreign Exchange Volatility Surface

Volatility surfaces for different assets are quoted daily on sources such as Bloomberg. The quoting conventions differ across asset classes and between individual assets. The foreign exchange market differs from most other financial markets as the volatility surface is quoted as implied volatilities against corresponding option deltas, as opposed to corresponding strike prices. The delta of an option is the rate of change of the option price with respect to the price of the underlying exchange rate. It represents the percentage of the foreign notional which must be bought or sold to hedge the option position. The delta equation is derived from the Garman-Kohlhagen formula and is defined as

$$\begin{aligned} \delta_S &= \omega e^{-r_f T} N(\omega d_+) \\ d_+ &= \frac{ln(\frac{S}{K}) + (r_d + r_f + 0.5\sigma^2)T}{\sigma\sqrt{T}}. \end{aligned} \tag{3.2.2}$$

A further difference between the foreign exchange market in comparison to other markets is the shape of the volatility surface. One of the stylised facts of the foreign exchange volatility surface is that is is in the shape of a smile. This shape accounts for the fact that far out-the-money and far-in-the-money options are often more expensive than predicted by the Black-Scholes model. This discrepancy in theory and practice arises as these options are often in high demand from investors, which drives up the price and therefore the implied volatility of these options (Clark, 2011). In contrast to the foreign exchange market, other markets, such as the equity market have a volatility skew. In this case, far out-the-money puts and far in-the-money calls options are in particular demand (Hull, 2012).

The quoting conventions in the foreign exchange market are different for various currency pairs, so the extraction of data and creation of a corresponding volatility surface can become complicated. Clark (2011) gives a thorough overview of the conventions in the foreign exchange market and details the specific rules for the most commonly traded exchange rates. In contrast, (Reiswich and Wystup, 2012) take a more general approach; however, they present a clear summary of the simplified formulae used to create a tractable foreign exchange volatility surface.

The first step in creating a tractable volatility surface is transforming the quoted pair of implied volatility and corresponding delta to a pair of implied volatility and corresponding strike. The different quoting conventions must be taken into account for the transformation of each currency pair. The two primary considerations are the type of delta and whether the delta is premium adjusted. The two types of deltas are forward and spot deltas. Spot deltas correspond to the rate of change of the option price with respect to the price of the underlying currency, while forward deltas correspond to the rate of change of the option price with respect to the forward price of the underlying currency. Premium adjustment arises if the option premium is conventionally paid in the foreign currency. Consequently, the investor with the short foreign currency position has to hold less of the foreign currency to hedge the position as the premium is received in the foreign currency (Reiswich and Wystup, 2012). The relationship between spot deltas and premium adjusted spot deltas is described by

$$\delta_s = \delta_{s-p.a} - \frac{V}{S}, \tag{3.2.3}$$

where δ_s is the spot delta, $\delta_{s-p.a}$ is the premium adjusted delta and $\frac{V}{S}$ is the option premium quoted in the foreign currency.

There are five quoted volatility-delta pairs for each maturity; an ATM delta-neutral volatility (σ_{ATM}), two risk reversal volatility quotes with corresponding deltas of 0.25 and 0.1 (σ_{25D-RR} and σ_{10D-RR} respectively) and two butterfly spread volatility quotes with corresponding deltas of 0.25 and 0.1 (σ_{25D-BF} and σ_{10D-BF} respectively). Each of the quoted volatilities gives information about the volatility smile. The ATM volatility quote shows the level of volatility, the risk reversal quotes give information about the skewness of the volatility smile and the butterfly spread quotes give information about the curvature of the volatility smile (Clark, 2011). The ATM delta neutral volatility is easily converted to the corresponding strike using Equation (3.2.4),

$$\begin{aligned} K &= Se^{(r_d-r_f)T+\frac{1}{2}\sigma_{ATM}T} \quad \text{(Regular Pairs)} \\ K &= Se^{(r_d-r_f)T-\frac{1}{2}\sigma_{ATM}T} \quad \text{(Premium Adjusted Pairs).} \end{aligned} \tag{3.2.4}$$

For the remaining quotes, the implied volatilities for the risk reversals and butterfly spreads must first be converted to implied volatilities for calls and puts before they are converted to volatility strike pairs. Reiswich and Wystup (2012) provide a clear algorithm for obtaining call and put volatilities from the quoted butterfly spreads and risk reversals, defined by

$$\sigma_{10D-P} = \sigma_{ATM} + \sigma_{10D-BF} - 0.5\sigma_{10D-RR}$$
$$\sigma_{25DP} = \sigma_{ATM} + \sigma_{25D-BF} - 0.5\sigma_{25D-RR}$$
$$\sigma_{10D-C} = \sigma_{ATM} + \sigma_{10D-BF} + 0.5\sigma_{10D-RR}$$
$$\sigma_{25D-C} = \sigma_{ATM} + \sigma_{25D-BF} + 0.5\sigma_{25D-RR},$$
(3.2.5)

where σ_{10D-P} is the volatility for a put option with a delta of 0.1 and σ_{10D-C} is the volatility for a call option with a delta of 0.1.

The new volatility-delta pairs are then used to calculate the corresponding strikes. The strike of non premium adjusted rates is found by rearranging Equation (3.2.2), with the delta-volatility pairs from Equation (3.2.5) used as inputs. The rearranged equation is described by

$$K = Se^{\left((r_d - r_f)T + +\frac{1}{2}\sigma_{ATM}T - w\sigma\sqrt{T}N^{-1}(w\delta e^{(rf)T})\right)}. \quad (3.2.6)$$

Note that the delta for a put option is negative, i.e. for the volatility σ_{10D-P}, the delta is -0.1.

The conversion from volatility-delta pairs to volatility-strike pairs is more complex for premium adjusted currency pairs. Equation (3.2.3) for premium adjusted pairs cannot be rearranged to solve for the strike so it must be solved using a root finding algorithm Castagna (2010). The algorithm used is included in Appendix A.

3.2.2 Variance Swaps

Variance swaps are over-the-counter financial instruments that allow an investor to trade the realised volatility over a certain time period. The final pay-off of a variance swap is the difference between the strike variance and the actual realised variance over the life of the swap. The fair strike for a variance swap is the strike that makes the value of the swap contract zero at inception (Bossu *et al.*, 2004).

The fair strike can be determined using replication arguments. Demeterfi, Derman, Kamal and Zhou (1999) show that a variance swap can be replicated by a long position in the underlying asset and a static short position in a log contract on the asset. The log contract has a pay-off which is the logarithm of the final asset value divided by the initial asset value. Log contracts are not traded on the market, so the log contract must also be replicated. The pay-off of the log contract can be decomposed using the findings of Breeden and Litzenberger (1978), which lead to the representation of the log contract pay-off in terms of vanilla put and call options.

The resulting formula for the fair strike of a variance swap is given as

$$K_{var} = \frac{2}{T}\Bigg(rT - \Big(\frac{S_0}{S_*}e^{rT} - 1\Big) - log\frac{S_*}{S_0} + e^{rT}\int_0^{S_*}\frac{1}{K^2}P(K)dK + e^{rT}\int_{S_*}^{\infty}\frac{1}{K^2}C(K)dK\Bigg), \tag{3.2.7}$$

where K_{var} is the fair variance strike, T is the time to maturity, r is the interest rate, S_0 is the initial spot price, P and C are the Garman-Kohlhagen formulas for put and call options respectively and S_* is the ATM strike.

In practice, only a finite number of implied volatilities are quoted, so the integrals in Equation (3.2.7) must be approximated. Demeterfi, Derman, Kamal and Zhou (1999) show that an approximation is given by

$$\begin{aligned} K_{var} &= \frac{2}{T}\Bigg(rT - \Big(\frac{S_0}{S_*}e^{rT} - 1\Big) - log\frac{S_*}{S_0} + e^{rT}P\Bigg) \\ P &= \sum_i w_P(K_{i,P})P(S, K_{i,P}) + \sum_i w_C(K_{i,C})C(S, K_{i,C}) \\ w_{P/C}(K_{i,P/C}) &= \frac{f(K_{i+1,P/C}) - f(K_{i,P/C})}{K_{i+1,P/C} - K_{i,P/C}} - \sum_{i=0}^{i-1} w_P/C(K_{i,P/C}), \end{aligned} \tag{3.2.8}$$

where $S_* = K_{0,P} = K_{0,C}$ is the ATM strike and $K_{i,P}$ are the $K_{i,P}$ strikes below and above the ATM strike respectively.

There are some issues with the approximation as the small range of strikes can lead to an underestimation of the actual fair variance; however the approximation is sufficient for the purposes of this application.

3.3 Implied Correlation

Implied correlation is a future estimate of correlation derived from instruments traded in the foreign exchange market (Hull, 2012). The extraction of implied correlation from market instruments can be problematic, as instruments involving correlation are generally more complex. In addition this, exotic and complicated instruments are often not liquid or exchange traded and accordingly cannot be used to extract consistent implied correlation estimates.

In contrast to the numerous studies focussing on implied volatility, there are limited studies concerning implied correlation. However, with the rise of exotic derivatives and correlation products, implied correlation is becoming an increasingly important measure. Future correlation estimates are principal inputs in portfolio and risk management and in the pricing of financial instruments, such as basket options, best of options and quanto options. Furthermore, instruments such as correlation swaps which allow investors to trade pure correlation are traded.

The need for instruments such as correlation swaps arose as participants in the financial market began to increase their correlation exposures due to the trading of more exotic options such as basket options (Jacquier and Slaoui, 2007). Correlation swaps are instruments similar to the variance swaps discussed in Section 3.2.2, where the pay-off of the swap is the difference between the realised correlation over a specific period and a specifies strike (Jacquier and Slaoui, 2007). Another way to gain exposure to correlation is by dispersion trading. Dispersion trading is a strategy where a short position in an index option and a long position in options on the constituents of that index are taken. This gives the investor exposure to the correlation between the constituents of the index as the volatility of the index is dependent on the correlation of the constituents, while the volatility of the individual options is not affected by correlation (Cayetano, 2007).

While implied correlations are not easily extracted from the market, implied volatilities are quoted daily, so the triangular framework can be used to calculate implied correlation from the three implied volatilities of the exchange rates in a currency trio, using equation (2.3.1). Furthermore, the implied correlation of currency pairs that are not in the same currency trio can be calculated using Equation (2.2), which extends the relationship to four currencies. This allows the extraction of consistent forward looking correlation estimates between currencies that are not confined to a currency trio. This will aid the pricing of exotic options involving a number of currencies, such as outside barrier options and best of options (Shevchenko, 2009).

Fretheim and Høigaard (2012), Walter and Lopez (2000a) and Campa and Chang (1998) evaluate implied correlation in the foreign exchange market extracted using the triangular framework. Each of these papers use ATM forward straddles to extract a single implied volatility measure. A straddle is an option trading strategy where an investor takes the same position in a call and put option with the same strike and maturity. In this trade, the investor makes a profit proportional to how much the price of the underlying moves, regardless of whether the price increases or decreases. For this reason, a straddle is often considered to be a trade on volatility (Hull, 2012).

Theses papers then calculate implied correlations using the extracted volatilities and rearranging Equation (2.2.2). While these papers use similar methods to extract implied volatility, they consider varied currency trios and use differing methods to evaluate the implied correlations. Fretheim and Høigaard (2012) consider implied correlation extracted from seven different currency trios over a period spanning from October 2006 to November 2011. Their analysis includes a visual inspection of the implied correlation structure and a statistical analysis of the properties and forecasting ability of the implied correlation series. The visual analysis suggests that realised correlation series vary erratically in comparison to implied correlation series. This is partly due to the rolling period over which realised correlation is calculated, which creates large moves in the correlation as data points enter or exit the rolling period. The analysis also reveals that both series appear more stable as the time interval over which the correlation is calculated increases. Fretheim and Høigaard (2012) suggest that the increased stability of the implied correlation could be due to the market expecting the series to be mean reverting. Finally, Fretheim and Høigaard (2012) conclude that implied correlation is generally a more stable measure than historical correlation. However, they found that this result is not consistent for every pair and for some pairs, particularly the less liquid currencies, the implied correlation is a poor predictor of realised correlation.

Walter and Lopez (2000a) examine two currency trios from October 1990 to April 1997. While the analysis periods and currencies considered differ to Fretheim and Høigaard (2012), the two papers have a number of similar findings. Walter and Lopez (2000a) also observe that implied and realised correlation series vary less erratically as the time interval over which the correlations are calculated increases. Similar to Fretheim and Høigaard (2012), they suggest that this is caused by market participants expecting implied correlation series to be mean reverting. Furthermore, Walter and Lopez (2000a) find that the forecasting ability of implied correlation is dependent on the currency trio considered. In conclusion, Walter and Lopez (2000a) find that while implied correlation includes information that historically derived data does not, the calculation of implied correlation is not consistently beneficial or economically viable.

Campa and Chang (1998) consider one currency over a period from 1989 to 1995. This study includes a more rigorous statistical analysis than the previous studies considered, however; it is limited as only one currency trio is considered and the analysis period is significantly different from the period considered in this investigation. Campa and Chang (1998) conclude that implied correlation outperforms historically based estimates in forecasting future correlation. This is consistent with the findings of Fretheim and Høigaard (2012).

In addition to the three studies mentioned above, Driessen *et al.* (2013) and Mueller *et al.* (2012) look at implied correlation from a different perspective. These studies focus on investigating whether implied correlation is priced. Driessen *et al.* (2013) consider the equity market, looking at the implied correlation of the $S\&P100$. They find that implied correlation is consistently greater than realised correlation, suggesting that a considerable premium is added to correlation products. Mueller *et al.* (2012) investigate implied correlation specifically in the foreign exchange market. They calculate the implied correlation using a replication strategy derived using the triangular arbitrage condition and find evidence of priced correlation risk in foreign exchange markets. The high correlation premiums reflect the level of uncertainty in predicting future correlations and the relative illiquidity of more complex correlation products.

Chapter 4

Hedging Applications

The triangular framework can be extended to practical applications where the relationship is used to inform trading and hedging strategies. Hedging is a common practise implemented by financial institutions with large positions in different assets in order to reduce their exposure. In particular, exposures to illiquid instruments need to be carefully managed as these positions cannot be quickly and easily closed out. In this chapter, the the application of the triangular relationship to hedging illiquid foreign exchange options will be investigated. The theory behind the hedging strategy and similar studies in literature will be discussed. Lastly, the algorithm for the hedging strategy will be presented.

4.1 Hedging Illiquid Foreign Exchange Options

Institutions with exposure to different foreign exchange assets hedge by investing in liquid foreign exchange instruments which offset their current positions. Illiquid foreign exchange instruments are difficult to hedge as they cannot be traded easily, they trade at a premium and reliable volatility surfaces for illiquid assets are difficult to construct (Pesjak, 2014). Furthermore, hedging strategies which utilise the same illiquid instruments as those originally held are not feasible. Holding additional illiquid instruments for hedging purposes further decreases the liquidity of the portfolio. It may be more beneficial to hedge illiquid foreign exchange positions using instruments with different underlying foreign exchange rates. This investigation will consider a method of hedging illiquid foreign exchange options using options on more liquid exchange rates.

A hedging strategy is implemented by calculating sensitivities, known as the Greeks, of an instrument and taking the opposite position in instruments with equivalent sensitivities. The two hedging strategies that will be considered in this application are delta and vega hedging. Delta is the sensitivity of the option value to changes in the spot of the underlying rate and vega is the sensitivity of the option value to changes in the implied volatility of the underlying rate. A perfect hedge would offset the losses and gains of the instrument exactly; however, the sensitivities are constantly changing and portfolios cannot be continuously rebalanced. Furthermore, transaction costs of hedging may limit the trades that are implemented. In practice, portfolios are generally hedged daily so hedging strategies will not completely offset changes in portfolio value (Hull, 2012).

The delta (δ) and vega (v) of a foreign exchange option are derived from the Garman-Kohlhagen formula. The expressions for these sensitivities are defined by

$$\begin{aligned} \delta &= \frac{dV}{dS} = \omega e^{-r_f T} N(\omega d_+) \\ v &= \frac{dV}{d\omega} = S_0\sqrt{T} e^{-r_f T} N(\omega d_+) \\ d_+ &= \frac{ln(\frac{S}{K}) + (r_d + r_f + 0.5\sigma^2)T}{\sigma\sqrt{T}}, \end{aligned} \tag{4.1.1}$$

where V is the value of the option.

Delta hedging is executed by taking an equivalent position in the underlying that offsets the delta of the instrument, i.e. for an option with a delta of +1, one of the underlying must be sold. Vega hedging is more complex than delta hedging and there are different methods that can be implemented to vega hedge a position. Any instrument which has an exposure to the volatility of the underlying can be used, from plain vanilla options to more complex instruments such as variance swaps. In the same way as delta hedging, vega hedging is executed by taking an opposite position in derivatives on the underlying, so that the two vega exposures cancel out. This study will only consider vega hedging using ATM straddles, which is a hedging instrument comprising of an ATM put option and an ATM call option. These instruments are useful for vega hedging purposes as they are close to delta neutral, so the overall delta of the hedge does not need to be significantly adjusted Pesjak (2014).

The simple delta-vega hedging strategy described above is generally effective and simple to execute; however, it it problematic to implement for illiquid options. The hedging of illiquid options becomes complicated as vega hedging requires continual trading of options on the underlying. This is not realistic as these options are also illiquid. In contrast, the foreign exchange spot market is liquid for most currency pairs, so it will be assumed that delta hedging using the underlying can still be implemented.

The triangular relationship between three currencies in a currency trio can be used to derive an expression to vega hedge an illiquid option using options on the other two exchange rates, provided that these two exchange rates are relatively liquid. The triangular relationship in this context is defined by

$$\sigma^2_{illiquid} = \sigma^2_{liquid_1} + \sigma^2_{liquid_2} - 2 \times \rho_{liquid_1 - liquid_2} \times \sigma_{liquid_1} \times \sigma_{liquid_2}. \tag{4.1.2}$$

Equation (4.1.2) shows that volatility of the illiquid exchange rate can be expressed in terms of the volatilities of the other, more liquid exchange rates in the currency trio. This suggests that the vega of the illiquid option can alternatively be interpreted as the sensitivity of the option price to the changes in volatility of the other two exchange rates. These equivalent vega positions in the more liquid exchange rate options are referred to as 'dephased vegas' (Jewitt, n.d.). The derivation of the dephased vegas is shown in more detail as follows,

$$\begin{aligned} Dephasedv_{liquid_1} &= \frac{dV}{\sigma_{liquid_1}} \\ Dephasedv_{liquid_1} &= \frac{dV}{d\sigma_{illiquid}} \times \frac{d\sigma_{illiquid}}{d\sigma_{liquid_1}} \\ Dephasedv_{liquid_1} &= v_{illiquid} \times \frac{\sigma_{liquid_2} - \rho \times \sigma_{liquid_1}}{\sqrt{u}}, \end{aligned} \tag{4.1.3}$$

where $u = \sigma^2_{liquid_1} + \sigma^2_{liquid_2} - 2 \times \rho_{liquid_1/liquid_2} \times \sigma_{liquid_1} \times \sigma_{liquid_2}$.

Note that this equation is for an illiquid exchange rate where $S_{illiquid} = \frac{S_{liquid_1}}{S_{liquid_2}}$. For an illiquid exchange rate where $S_{illiquid} = S_{liquid_1} \times S_{liquid_2}$, the expression for u becomes: $u = \sigma^2_{liquid_1} + \sigma^2_{liquid_2} + 2 \times \rho_{liquid_1/liquid_2} \times \sigma_{liquid_1} \times \sigma_{liquid_2}$ and consequently, the sign before the ρ terms will change from negative to positive.

A dephased vega is calculated for each of the two liquid exchange rates in the currency trio. These sensitivities represent the equivalent position required in liquid exchange rate instruments to hedge the vega sensitivity of the illiquid option. The hedging algorithm using the dephased vegas to hedge the illiquid option will be described in more detail in Section 4.2.

The dephased vegas get larger as the correlation between the underlying rate and the liquid rate increases. If the illiquid exchange rate and one of the liquid exchange rates have a very high correlation it may be beneficial to hedge using only that liquid exchange rate as this will reduce transaction costs.

There are some complications which arise in implementing the hedging strategy described above. Firstly, the instruments used to implement the vega hedge in the liquid exchange rates will also have to be delta hedged, which will increase the transaction costs of the strategy. Furthermore, the Profit and Loss ($P\&L$) of the hedge using the liquid exchange rates may not have the same 'numeraire' as the illiquid option. A numeraire is an asset which denominates the price of other assets, i.e. assets are priced with reference to the numeraire. In this case, the numeraire is taken to be the domestic currency of the illiquid option. The hedge $P\&L$s must be converted to the chosen numeraire daily, which represents a further spot rate exposure (Pesjak, 2014).

Jewitt (n.d.) and Pesjak (2014) consider a similar strategy for hedging illiquid foreign exchange options. Jewitt (n.d.) describes the triangular relationship between currencies and gives a qualitative description of how multi asset foreign exchange positions are managed. However, this paper does not show the actual derivation of the dephased vega expression and furthermore does test the proposed hedging strategy.

In contrast, Pesjak (2014) includes a derivation of the dephased vega expressions and provides a more thorough explanation of the hedging strategy. Pesjak (2014) consider a number of different hedging strategies, including the strategy derived from the triangular relationship. Two currency trios, AUD-JPY-USD and EUR-KRW-USD (where AUD/JPY and EUR/KRW are taken to be the illiquid rates) over a period from 30/12/2011 to 06/05/2013 are considered. The effectiveness of each hedging strategy is evaluated by looking at the root mean square error (RMSE) and $P\&L$ of the hedged portfolio over the period. Pesjak (2014) concludes that the hedging strategy derived from the triangular relationship performs better than the simple delta hedging strategy. However, if it is possible to implement a delta-vega hedging strategy using options on the underlying, this will be the optimal hedging strategy.

4.2 Hedging Algorithm

The hedging strategy described above can be explained using an algorithmic approach. The hedge is rebalanced daily and transaction costs are not considered.

There are some additional complexities that are introduced when hedging with different currencies. The illiquid option is valued in its domestic currency, which will be considered to be the numeraire. In order to value the hedged portfolio, all the positions must be converted to have the same numeraire. This conversion is implemented by multiplying the hedging positions in the liquid rates by the relevant spot rate to convert these positions to this numeraire. This in turn means that the vega of the hedging instruments is also multiplied by the relevant spot rate as this term is retained in the derivation of the instrument value in terms of the underlying volatility. In addition to these considerations, this conversion represents a further exposure to the conversion spot rate, which must also be hedged.

For each timestep, t, the hedging procedure is as follows:

1. Extract the implied volatility surface of the illiquid and liquid underlying for the maturity date in terms of implied volatility versus strike, using the methodology described in Section 3.2.1.
2. Extract the current spot price of the illiquid and liquid underlyings and the risk free interest rate curves. Note that the relevant OIS (Overnight Indexed Swap) curve will be used to approximate the risk free interest rate curve in this study. This curve will be explained further in Section 5.2.
3. Using the inputs obtained in the previous steps, determine the current Garman-Kohlhagen value, delta and vega of the illiquid option.
4. Delta hedge the position in the illiquid option using spot contracts on the underlying illiquid rate.
5. Calculate the two dephased vegas for the liquid exchange rates using Equation (4.1.3).
6. Calculate the delta and vega for the liquid hedging instruments.
7. Calculate the corresponding weights required in the options on the liquid exchange rates, i.e.

$$weight = \frac{DephasedVega}{HedgingInstrumentVega}.$$

8. If the liquid rate has a different numeraire to the illiquid rate, multiply the hedging instrument vega by the relevant spot rate, i.e.

$$weight = \frac{DephasedVega}{HedgingInstrumentVega \times Spot}.$$

   This step accounts for the spot terms that is retained in the derivation of the hedging instrument vega and ensures that all instruments are valued using the same 'numeraire'.
9. Take a position equal to the negative weight in the liquid exchange rate instruments.
10. Delta hedge the positions in both the liquid exchange rate options using spot contracts on the underlying liquid rates.
11. If the liquid rate has a different numeraire to the illiquid rate, the liquid instruments and the delta hedge for the liquid instruments must be multiplied by the relevant spot rate when calculating the portfolio value.
12. Delta hedge the additional positions in the spot rates which arise from the conversion to a common numeraire.

In the case that the illiquid option is denoted by X/Y, and the liquid rates are denoted by Z/Y and X/Z, Y is taken to be the numeraire, so the vega and hedging instrument value for instruments with underlying X/Z must be multiplied by the spot rate Z/Y to convert them to have numeraire Y. The portfolio value of the hedging strategy described above is then given by

$$\begin{aligned}
V_t^{Portfolio} &= V_t^{X/Y} - \delta_t^{X/Y} * S_t^{X/Y} \\
&- w_t^{X/Z} \times (V_t^{X/Z} - \delta_t^{X/Z} \times S_t^{X/Z}) \times S_t^{Z/Y} \\
&+ w_t^{Z/Y} \times (V_t^{Z/Y} - \delta_t^{Z/Y} \times S_t^{Z/Y}) + b \times S_t^{Z/Y} \\
w_t^{X/Z} &= \frac{DephasedVega_t^{X/Z}}{HedgingInstrumentVega_t^{X/Z} \times S_t^{Z/Y}} \\
w_t^{Z/Y} &= \frac{DephasedVega_t^{Z/Y}}{HedgingInstrumentVega_t^{Z/Y}} \\
b &= w_t^{X/Z} \times (V_t^{X/Z} - \delta_t^{X/Z} \times S_t^{X/Z}),
\end{aligned} \tag{4.2.1}$$

where, V is the value of the option, δ is the option delta and S is the spot exchange rate.

Chapter 5

Analysis and Results

The theory developed in the previous chapters is applied to simulated and historical data to further investigate the geometric relationship. The relationship will be considered with both historical spot data and historical implied data. The application of the framework to spot data and implied market data differs. In the spot data context, the data is used to evaluate how well the triangular relationship holds in various currency trios, using the historical correlation and volatilities. The validity of the relationship can be compared between currency trios to identify where the relationship fails and corresponding inefficiencies in the market.

The analysis for the implied context differs to the spot analysis as implied correlation is not quoted in the market, so the triangular relationship is used with implied volatility data to calculate these values. This forces the triangular relationship between implied correlation and volatility to hold, so the same analysis as performed on the spot data is immaterial. In the implied data context, the implied correlation extracted will be compared to the realised correlation over the same period for different currency trios to evaluate the accuracy of the implied correlation measure.

This chapter initially analyses the application of the triangular framework to simulated and historical spot data. Once the relationship has been thoroughly analysed in the context of spot data, implied market data is considered and implied correlations are calculated and analysed. Finally, the chapter presents and assess the effectiveness of a hedging strategy, derived from the triangular relationship, to hedge illiquid foreign exchange options.

5.1 Application of Geometric Concepts

In the spot context, the geometric relationship provides a tool to identify inefficiencies in the market and to analyse the relationships between volatility and correlation. In order to further understand the geometric relationship, described in Section 2.2, and the implications of any discrepancies in the data, the relationship is first applied to simulated data. Following the simulation, market data from different currency trios is considered.

5.1.1 Simulation

The first analysis considers simulated data to get a better understanding of the geometric relationships. The methodology is briefly described before the results of the simulation are discussed.

Simulation Methodology

The exchange rates are modelled using geometric Brownian motion, so the log returns of the exchange rates will be normally distributed (Hull, 2012). While this is not an accurate assumption for exchange rates, the framework does not depend on the distribution of the underlying exchange rates so it can be evaluated using this model. Two cases will be considered; in the first case, the processes are modelled to be arbitrage free, using the relationship described in Equation (2.1.1), while in the second case arbitrage opportunities are introduced into the simulated processes.

For each case, three processes are modelled; the first two processes are modelled using correlated Brownian motions, with stochastic differential equations given by

$$\begin{aligned} dX_t &= \mu_X X_t dt + \sigma_X X_t dW_t^X \\ dY_t &= \mu_Y Y_t dt + \sigma_Y Y_t dW_t^Y, \end{aligned} \tag{2.2.1}$$

where μ, σ and ρ (the correlation between the Brownian Motions) were arbitrarily chosen and $dW_t^X \times dW_t^Y = \rho dt$. The third process is modelled using the triangular arbitrage condition described in (2.1.1). The dynamics of the third process are then given by

$$\begin{aligned} dZ_t &= Z_t(\mu_X - \mu_Y - \rho\sigma_X\sigma_Y)dt + \sigma_Z Z_t dW_t^Z \\ \sigma_Z &= \sqrt{\sigma_X^2 + \sigma_Y^2 - 2\rho\sigma_X\sigma_Y}. \end{aligned} \tag{2.2.2}$$

For each process, 2000 points are simulated and the volatility and correlation is calculated over the whole period.

For case 2, arbitrage opportunities are introduced into the simulated processes by introducing noise into the third process in the form of normally distributed random variables which are added to the process at random points in time. The points in time were generated using uniformly distributed numbers over the whole time series.

In order to simplify the simulation, bid-ask spreads will not be considered, and the processes will be taken to represent the mid points. Ito *et al.* (2012) investigates the existence and frequency of triangular arbitrage opportunities in currency trios from 1999 to 2010 and shows that mid point indicates if there will be an arbitrage opportunity or not; however, it is important to note that if bid-ask spreads are too large, no profit can be made.

Simulation Analysis

The simulation analysis considers three different variations of cases where arbitrage opportunities are introduced, with the frequency and magnitude of the arbitrage opportunities differing for each case. The magnitude of the errors introduced into the data are described in terms of the order of magnitude of the error, i.e. for an order of magnitude of -2, the magnitude of the errors ranges between 0.01 and 0.09. It is clear that the larger the magnitude and frequency of the noise introduced, the greater the discrepancy between the standard deviation predicted by the relationship and the actual standard deviation. This is demonstrated in Figure 5.1.

If the magnitude of the random variables added to the third process is large, the actual standard deviation is always greater than the theoretical standard deviation regardless of the direction on the arbitrage opportunity. However, as the magnitude of the random variables decreases, the actual standard deviation is occasionally lower than the theoretical standard deviation. Through a number of different simulations, it became clear that the most common direction of the arbitrage opportunity, i.e. ($A \to C \to B \to A \; or \; A \to B \to C \to A$) cannot be inferred from the sign of the difference between the actual and theoretical standard deviation.

The simulation demonstrates that the relationship holds if there are no arbitrage opportunities, as shown in Case 1 in Figure 5.1. It is clear that when arbitrage opportunities are introduced into the simulation, the triangular relationship breaks. However, the point at which the arbitrage opportunity occurs only becomes clear when the discrepancy between actual and theoretical standard deviation is plotted over time. This suggests that while this approach is a good tool to assess the relationships between the higher moments of the exchange rate returns, it is not the most efficient way to look at spot arbitrage in the market. This analysis will be applied to real data to determine if the same conclusions can be made in Section

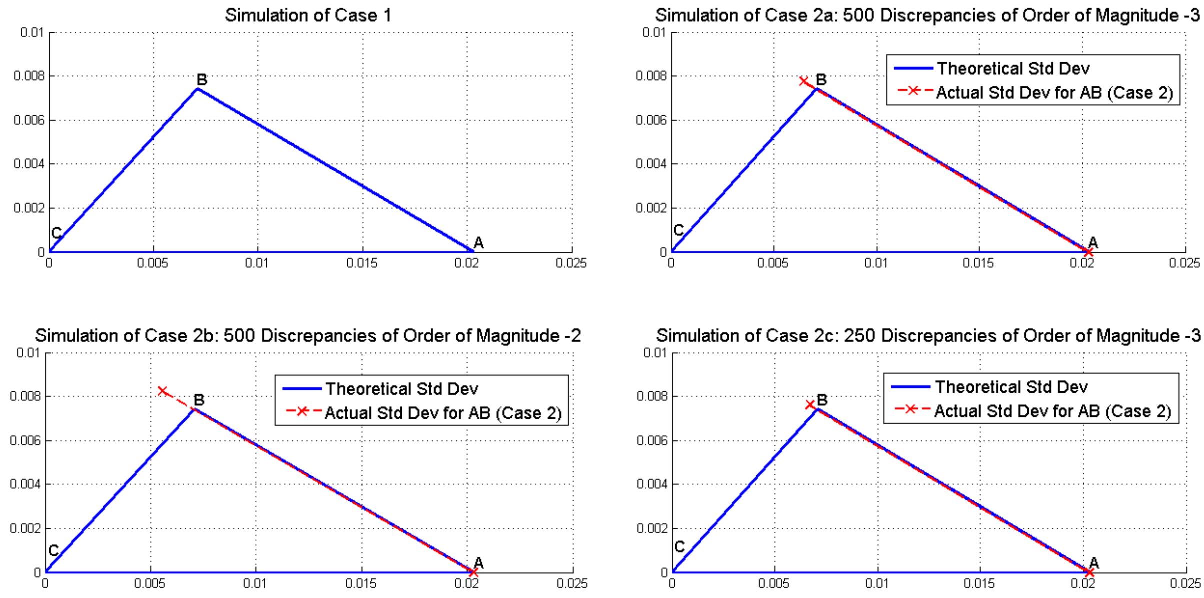

Fig. 5.1: Graphical Representation of the Relationship for Different Cases

5.1.2.

5.1.2 Market Data

Following the simulation analysis, the relationships are applied to market data. The methodology and data used in the investigation is briefly described before the results of the analysis are discussed.

Market Data Methodology

The data used in the analysis is sourced from Bloomberg and spans from January 2000 to November 2015. The spot data analysis considers nine different ZAR based currency trios, which involve the exchange rates between the ZAR, USD and various third currencies, including the EUR, GBP, JPY, CNY, MXN, BRL, KES and NGN.

The analysis in the spot contest focusses on the difference between the actual standard deviation and the theoretical standard deviation for the cross rates, where the cross rates are taken to be the exchange rates between the USD and the different third currencies. The actual standard deviation is calculated using simple historical estimates and the theoretical standard deviation is calculated using Equation (2.2.5).

Initially, end of day, mid-price data from the 17^{th} of November 2014 to the 16^{th} of November 2015 is analysed to to identify how well the triangular relationship holds. The analysis is then extended to consider how the relationships between volatility and correlation change over time. This analysis is based on the full time period of the data (January 2000 to November 2015).

Market Data Analysis

The geometric relationship holds relatively well for most currency trios considered, and discrepancy between the actual standard deviation and the theoretical standard deviation is of a magnitude of approximately -3. This is demonstrated in Figure 5.2. The relationship does not seem to hold as well for developing currencies, in particular the INR, CNY, BRL and NGN, suggesting these markets are less efficient than the more developed currencies and there may be more arbitrage opportunities present.

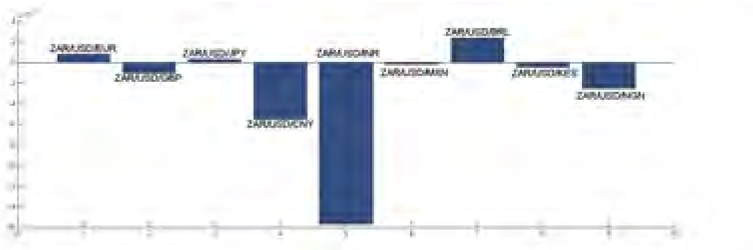

Fig. 5.2: Discrepancy between Actual and Theoretical Standard Deviation

In order to investigate how the discrepancy between the actual and theoretical standard deviation and the number of viable arbitrage opportunities present in the data are related, bid ask spreads are obtained and taken into account. The actual number of arbitrage opportunities in both directions is compared to the discrepancy between the standard deviations over the same period. Figure 5.3 substantiates the analysis of the simulation in Section 5.1.1, showing that the geometric representation does not give specific information about the direction of the arbitrage opportunities. Furthermore, the bid ask spreads need to be taken into account as if the spread is too large the arbitrage opportunity will not be profitable.

It is clear from the application to realised data, that the geometric relationship is a good indication of whether there could be arbitrage opportunities in the particular currency trio. However, once bid ask spreads are taken into account, the relationship does not always correspond directly. For example, the ZAR-USD-CNY trio has the second largest discrepancy but does not actually have many profitable arbitrage opportunities. Furthermore, as discussed in Section 5.1.1, the geometric

Fig. 5.3: Actual Number of Arbitrage Opportunities

representation does not give an indication of when the arbitrage opportunity occurred, this only becomes clear when the changes in the difference between the actual and theoretical standard deviation are plotted over time.

One of the advantages of the geometric approach is that is is a good tool to evaluate the volatility and correlation of a currency trio. This representation gives a good indication of how volatility and correlation change over time. Figure 5.4 shows how the volatility and correlation relationships in the ZAR-USD-EUR trio from 02/10/2001 to 22/03/2002.

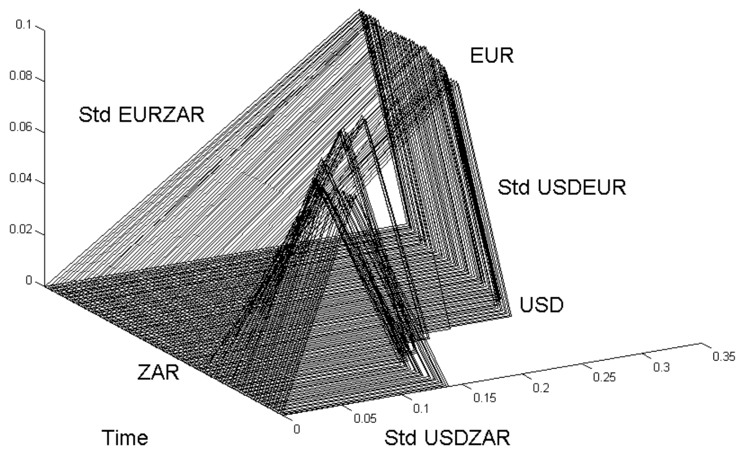

Fig. 5.4: Volatility and Correlation relationships in the ZAR/USD/EUR Trio Over Time

If is clear from Figure 5.4 that the relationships between the volatility and correlation change over time. For example, the first big jump in volatility which can be seen occurred on 21 Dec 2001 when the ZAR hit an alltime low. During this period there was a devaluation of emerging market currencies, including the ZAR, resulting in a spike in the volatility of the USD/ZAR and EUR/ZAR. The representation of the data as a triangle shows how the changes in volatility affect the correlation. When the volatility of the USD/ZAR and EUR/ZAR increased, the angle representing the correlation between the USD/ZAR and EUR/ZAR decreases, which corresponds to an increased correlation between these two exchange rates. This is consistent with the findings of Jewitt (n.d.).

In summary, if the triangular relationship holds to a certain tolerance, it is a clear indication that there are minimal arbitrage opportunities in the currency trio. If the relationship does not hold it is an indication that pricing discrepancies exist in the data, however, it is not always indicative of the number and magnitude of the opportunities. Furthermore, considerations such as the liquidity of the currencies involved may negate any trading opportunities. It is more efficient to consider the actual data when attempting to identify triangular arbitrage opportunities. This is consistent with the findings in Section 5.1.1.

5.2 Implied Correlation

The primary application of the triangular relationship in the implied context is the expression of implied correlation as a function of implied volatility. In this section, the method followed to extract implied correlation will be discussed and the resulting implied correlation will be analysed and compared to the realised correlation over the same period. The results will be compared to the findings of Fretheim and Høigaard (2012) and Walter and Lopez (2000a).

5.2.1 Methodology

In contrast to the currency trios considered in the spot context, the implied data analysis is not centered on ZAR based currency trios. Options on ZAR based exchange rates are not as liquid as options based on more frequently traded rates such as the USD. Consequently, the trios considered all include the USD and EUR to provide a more reliable analysis. The analysis spans a period from 2^{nd} January 2007 to the 16^{th} November 2015.

Three currency trios are considered for the implied data analysis. The trios are based on the USD and EUR, with the third currency as the JPY, the CHF and the ZAR. In the implied context, in addition to the spot exchange rate data, interest rates curves for each currency and implied volatility surfaces for each exchange rate are required. This data is extracted from Bloomberg, which quotes implied volatility for all of the currency pairs analysed as volatility-spot delta pairs. A new interest rate curve and volatility surface are extracted for each time step to create a historical time series of the curves and surfaces. Table 5.1 shows the premium adjustment convention for the currency pairs used in the analysis.

Currency	Convention
EUR/USD	Regular
USD/ZAR	Premium Adjusted
EUR/ZAR	Premium Adjusted
USD/JPY	Premium Adjusted
EUR/JPY	Premium Adjusted
USD/CHF	Premium Adjusted
EUR/CHF	Premium Adjusted

Tab. 5.1: Premium Adjustment Conventions

OIS curves obtained from Bloomberg are used as an approximation of the risk free rates for each currency, with the exception of the ZAR, which does not have a quoted OIS curve. OIS curves are extracted from overnight index swaps, which are swaps where the overnight rate floating rate is exchanged for a specified fixed interest rate. These swaps have a low credit risk as they are often collateralised and between creditworthy counterparties. For this reason OIS are considered to be a good proxy for a risk fee rate (Hull and White, 2013). The OIS rates are approximated as simple rates for the purposes of the analysis. The JIBAR (Johannesburg Interbank Agreed Rate) swap curve is used as a proxy for the ZAR risk free rate. This yield curve is obtained from bootstrapping liquid swap instruments in the South African market.

In order to deal with the various working days and conventions of the different currencies, all weekdays are considered as working days. In the case of missing data due to public holidays or other reasons, linear interpolation is used to fill in the missing data. This should not affect the results of the analysis.

Implied correlation is calculated using Equation (2.3.1), which is derived from the triangular arbitrage condition. This calculation requires a single implied volatility estimate for each exchange rate in the currency trio, so a fair implied volatility estimate for each exchange rate must be calculated. Implied volatility is first extracted from the market in the form of volatility-spot delta pairs and the quoted pairs are converted to volatility-strike pair, using the methodology described in Section 3.2.1. The resulting volatility surface is the form of five volatilities against five corresponding strikes at each timestep. Demeterfi, Derman, Kamal and Zhou (1999) present a technique to determine a fair volatiliy estimate for a variance swap, described in Section 3.2.2. This technique is used to calculate a fair volatility estimate at each time step for a particular maturity. Finally, equation (2.3.1) is solved for the implied correlation at each time step.

Historical correlation is used as a comparison to the calculated implied correlation. The historical correlation is calculated over the corresponding time period to the implied correlation, i.e. for implied correlation with a horison of 1 year on day x, with a maturity day y, the corresponding historical correlation for day y is evaluated looking back from date y to x. The historical correlation is calculated using simple rolling estimates, as described in Section 3.1.

5.2.2 Analysis

The implied correlation series, generated using the method described in Section 5.2.1, are analysed by looking at the difference between the implied and realised correlation values across different maturities and currency pairs. The series are compared by visual inspection, by looking at the statistical properties of the series and by evaluating the root mean square error between the two series. The term structure of each implied correlation series is also considered.

Figure 5.5 shows one set of realised and implied correlations from each the three currency trios considered. Specifically, the correlations between the USD/ZAR and EUR/ZAR, the USD/JPY and EUR/JPY and the USD/CHF and EUR/CHF for a 1 month horison and a 6 month horison.

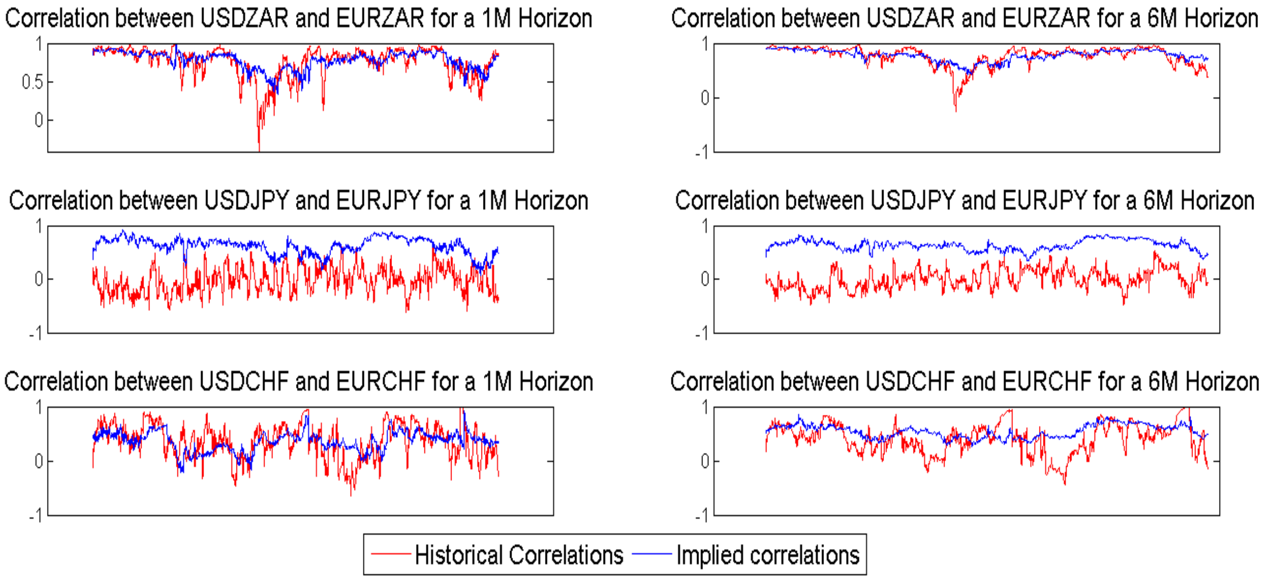

Fig. 5.5: Realised versus Implied Correlation for Three Exchange Rates Over horisons of and 1 Month and 6 Months from 2^{nd} January 2007 to the 16^{th} November 2015

Visual inspection of the correlation series suggests that implied correlation is a more stable measure than realised correlation, demonstrated by the smaller fluctuations of the implied correlation series in comparison to the realised correlation series in Figure 5.5. This effect is pronounced for correlations with shorter time horisons, such as one month. This observation is substantiated by the statistical analysis of the correlations, presented in Table 5.2. The standard deviation of the implied correlation series is consistently lower than the standard deviation of the historical correlation series, except for the correlation between the EUR/JPY and the EUR/USD. It is also clear that both the implied and historical series vary less erratically as the time horison increases. This is supported by the statistical analysis, which shows that the standard deviations for all the series decreases as the time horison increases. These observations correspond to the findings of both Fretheim and Høigaard (2012) and Walter and Lopez (2000a), as discussed in Section 3.3.

For most of the correlations analysed, the implied correlation gives a reasonable estimate of the realised correlation over the same period, as shown in the implied correlation for the USD/ZAR and EUR/ZAR and the USD/CHF and EUR/CHF in Figure 5.5. Furthermore, the mean correlation for both the implied and historical correlation series, shown in Table 5.2, is very similar for many of the correlation pairs analysed, in particular those pairs in the USD-EUR-ZAR currency trio and the USD-EUR-CHF trio.

For the USD/ZAR and EUR/ZAR correlation series, the implied correlation forecast tracks the realised correlation closely, but seems to lag the realised correlation following large moves. The lagging estimates of implied correlation may be partly due to implied correlations incorporating current correlation levels in future expectations.

Correlation Pair	1 Month		3 Months		6 Months		1 Year	
	Mean	SDev	Mean	SDev	Mean	SDev	Mean	SDev
pUSD/ZAR-EUR/ZAR Im	0.780	0.112	0.782	0.102	0.778	0.102	0.789	0.097
pUSD/ZAR-EUR/ZAR Hist	0.759	0.196	0.768	0.168	0.775	0.148	0.782	0.122
pUSD/ZAR-EUR/USD Im	-0.463	0.138	-0.457	0.106	-0.456	0.091	-0.459	0.084
pUSD/ZAR-EUR/USD Hist	-0.442	0.277	-0.445	0.216	-0.460	0.174	-0.470	0.139
pEUR/ZAR-EUR/USD Im	0.167	0.174	0.177	0.133	0.173	0.116	0.166	0.112
pEUR/ZAR-EUR/USD Hist	0.167	0.327	0.171	0.263	0.157	0.230	0.148	0.184
pUSD/JPY-EUR/JPY Im	0.621	0.152	0.621	0.125	0.634	0.097	0.644	0.077
pUSD/JPY-EUR/JPY Hist	-0.028	0.221	-0.014	0.145	-0.009	0.113	-0.007	0.094
pUSD/JPY-EUR/USD Im	-0.285	0.167	-0.280	0.143	-0.269	0.130	-0.254	0.123
pUSD/JPY-EUR/USD Hist	-0.242	0.342	-0.226	0.266	-0.204	0.224	-0.190	0.158
pEUR/JPY-EUR/USD Im	0.548	0.167	0.559	0.142	0.560	0.124	0.565	0.106
pEUR/JPY-EUR/USD Hist	-0.008	0.217	0.003	0.134	-0.002	0.104	-0.002	0.079
pUSD/CHF-EUR/CHF Im	0.360	0.179	0.544	0.135	0.542	0.120	0.312	0.151
pUSD/CHF-EUR/CHF Hist	0.368	0.312	0.404	0.266	0.437	0.259	0.486	0.235
pUSD/CHF-EUR/USD Im	-0.804	0.146	-0.510	0.132	-0.521	0.106	-0.812	0.112
pUSD/CHF-EUR/USD Hist	-0.812	0.226	-0.795	0.228	-0.766	0.235	-0.723	0.234
pEUR/CHF-EUR/USD Im	0.221	0.224	0.416	0.240	0.414	0.199	0.273	0.170
pEUR/CHF-EUR/USD Hist	0.129	0.361	0.132	0.300	0.143	0.268	0.165	0.217

Tab. 5.2: Means and Standard Deviation of Historical and Implied Correlation over Different horisons for Each Currency Pair

In contrast to the implied correlations discussed above, the implied correlation between the USD/JPY and EUR/JPY, shown in Figure 5.5, is a poor estimate of the realised correlation over the same period. The correlation forecast for the correlation between the EUR/JPY and EUR/USD also performs poorly. The statistical analysis, presented in Table 5.2 reveals that the mean of the implied and realised correlations for both these correlation series differs by more than 0.5.

In order to further evaluate the accuracy of the implied correlation forecasts, the RMSE between the implied and realised correlation for different currency pairs and different maturities is analysed. The average RMSE for each currency pair is shown in Table 5.3.

Correlation Pair	1 Month	3 Months	6 Month	1 Year
pUSD/ZAR-EUR/ZAR	0.1558153	0.1311462	0.1048891	0.1290643
pUSD/ZAR-EUR/USD	0.2609741	0.1926106	0.1606861	0.1321190
pEUR/ZAR-EUR/USD	0.3209245	0.2589468	0.2471295	0.2267675
pUSD/JPY-EUR/JPY	0.726125	0.680098	0.667043	0.663907
pUSD/JPY-EUR/USD	0.324826	0.270916	0.272999	0.221987
pEUR/JPY-EUR/USD	0.625028	0.597613	0.596261	0.587934
pUSD/CHF-EUR/CHF	0.2789475	0.2903944	0.2640210	0.2807189
pUSD/CHF-EUR/USD	0.1192494	0.3739891	0.3504363	0.2569138
pEUR/CHF-EUR/USD	0.3133829	0.4138556	0.3672184	0.2684303

Tab. 5.3: RMSE over Different horisons for Each Currency Pair

Note that pEUR/CHF-EUR/USD refers to the correlations between the EUR/CHF and the EUR/USD.

The RMSE is below 0.5 for all pairs considered, except for the correlation between the USD/JPY and the EUR/JPY and the correlation between the USD/EUR and the EUR/JPY. Furthermore, for each currency pair, the RMSE is similar across maturities, with the exception of the RMSE for the implied correlation between USD/CHF and EUR/USD and the correlation between the EUR/CHF and the EUR/USD, where the three month and 6 month average RMSE is consistently greater than the 1 month and 1 year average RMSE. Another pattern that emerges is the decrease in RMSE as the time horison increases for most of the correlation series.

These findings are consistent with the results of Fretheim and Høigaard (2012), who showed that implied correlation is generally a good forecast of realised correlation; however, this does not hold for every currency trio. Walter and Lopez (2000a) also concluded that the forecasting accuracy of implied correlation varies across currency trios. Both Fretheim and Høigaard (2012) and Walter and Lopez (2000a) used RMSEs to evaluate the correlation series. They found that the RMSE is approximately between 0.1 and 0.3 for most pairs. However, Fretheim and Høigaard (2012) also find that for some correlation series, such as the correlation between the USDGBP and the GBPJPY, the RMSE is as great as 0.8.

Another aspect of the correlation series that is analysed is the term structure of the correlation series. Most of the implied correlations between exchange rate pairs have a flat or upward sloping structure, i.e. the implied correlation is higher for greater maturities. However, this varies over time and for different currency trios. The term structure is approximately flat for the pairs in the USD/EUR/ZAR currency trio and the upward sloping feature is most pronounced in the USD/EUR/JPY trio, as shown in Figure 5.6. An interesting feature in the correlation between the USDJPY and EURJPY is that when the correlation seems to change significantly or jump, then the term structure changes from upward sloping to downward sloping. This is similar to the findings of Faria and Kosowski (2015), who find that the term structure of implied correlation is predominantly upward sloping or flat, except in turbulent periods, when the term structure becomes downward sloping. The implied correlation structure between the USDCHF and the EURCHF is an exception to the trend as it is neither upward or downward sloping. For this series, the three month and and six month correlation series are predominantly greater than the one month and one year implied correlation series.

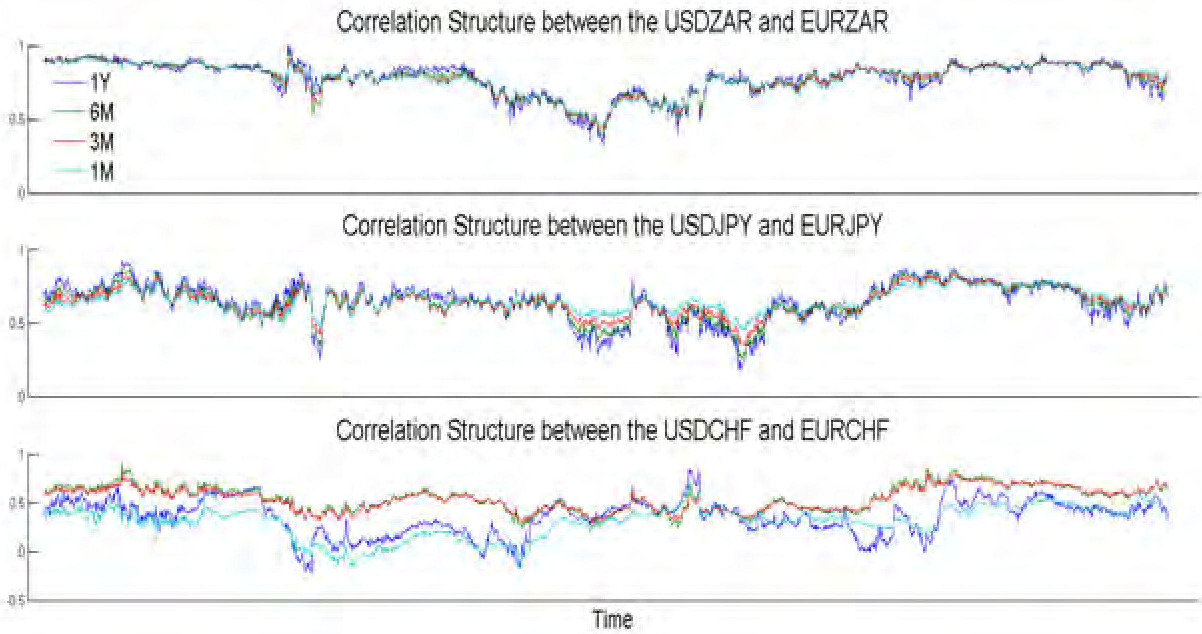

Fig. 5.6: Term Structure of Implied Correlation for Different Exchange Rate Pairs from 2^{nd} January 2007 to the 16^{th} November 2015

5.3 Hedging Applications

The application of the triangular framework to hedging illiquid options will be analysed by implementing the hedging strategy described in Section 4.2 and evaluating the effectiveness of the strategy. The results obtained using this strategy will be compared to the findings of Pesjak (2014).

5.3.1 Methodology

The three currency trios analysed in Section 5.2, the USD-EUR-ZAR, the USD-EUR-CHF and the USD-EUR-JPY trios, will be considered in the hedging analysis. The exchange rates involving the USD are generally the most liquid in the foreign exchange market, so the illiquid rate will be considered to be the exchange rate rate between the EUR and the third currencies in the respective trios. The illiquid exchange rate option that is chosen to be examined is a vanilla call option with a time to maturity of one year.

In order to evaluate the hedging strategies, a simple delta hedging strategy using the illiquid underlying of the option and a delta-vega hedging strategy using options on the illiquid underlying will be used as a benchmark. However, in practice the delta-vega hedging strategy is not feasible for illiquid options as this will decrease the liquidity of the portfolio and options on the underlying cannot be easily traded. All strategies include a delta hedge using the underlying of the illiquid option as it is assumed that all the exchange rates are sufficiently liquid in the spot market.

In order to analyse the hedging strategy, five different scenarios are considered:

- Strategy 1- No hedging.
- Strategy 2- Delta hedging using the illiquid underlying.
- Strategy 3- Delta-vega hedging using the illiquid underlying and options on the illiquid underlying.
- Strategy 4- Delta-vega hedging using the illiquid underlying and options on the the other two liquid exchange rates in the currency trio.
- Strategy 5- Delta-vega hedging using the illiquid underlying and options on the the liquid exchange rate in the currency trio which has the greatest implied correlation with the illiquid underlying over the time period of the option.

At each time step, the relevant market information is extracted, such as spot prices, interest rates and implied volatility. A new interest rate curve and volatility surface is extracted for each time step.This data is sourced in the same way as described in Section 5.2.1. If the option goes too far in or out of the money the implied volatility extrapolation gives unreasonable results so the implied volatility is flat-lined at each end of the volatility curve. The implied correlation required in the hedging calculations is determined by calculating the fair volatility estimate (using the method of Demeterfi, Derman, Kamal and Zhou (1999) described in Section 3.2.2) and inputting the volatility estimates into the triangular relationship described in Equation (2.2).

The hedging effectiveness of each strategy is analysed by calculating the RMSE between the actual and theoretical daily $P\&L$ for each strategy. Theoretically, a fully hedged portfolio should have a daily $P\&L$ of zero.

The testing procedure described above is repeated every day from the 4^{st} January 2010 to 4^{st} January 2014 to obtain a total number of 1044 tests and corresponding RMSE measures. On each day, a position in the illiquid option is taken and the hedging simulation is run for a year, before moving on to the next day. The the option was set to be ATM at inception for each simulation. While this limits the scope of the tests, the hedging strategy is largely testing the effectiveness of vega hedging and the vega effects are the largest when an option is close to ATM.

For the purposes of this study, the hedging positions will be rebalanced daily and transaction costs will not be considered.

5.3.2 Analysis

Each of the trios considered are discussed individually before the common trends are analysed. For each trio, the average RMSE of the $P\&L$ for each strategy is considered.

For the USD-EUR-ZAR currency trio, the average RMSE of the $P\&L$ for each strategy over all the simulations is shown in Table 5.4. This is further decomposed into the average RMSE over each year considered in the simulation. For this currency trio, the liquid rate with the greatest correlation with the EUR/ZAR, which will be used in Strategy 5 is the USD/ZAR.

It is clear that the RMSE for Strategy 1 (no hedging) is significantly greater than the RMSE for the other strategies, which is as expected. Furthermore, the RMSE analysis shows that Strategy 2 (delta hedging) significantly reduces the daily $P\&L$ for all of the periods. The best performing strategy over the whole period is Strategy 3 (delta-vega hedging with options on the underlying), however this strategy is not feasible in practice if the underlying is illiquid.

For the majority of the simulations and the overall average, Strategy 4 (vega hedging with options on the liquid rates) and Strategy 5 (vega hedging with op-

	2010-2014	2010-2011	2011-2012	2012-2013	2013-2014
Strategy 1	0.0727833	0.0484857	0.0654074	0.0836415	0.0937650
Strategy 2	0.0072254	0.0074853	0.0094692	0.0059879	0.0059316
Strategy 3	0.0050414	0.0044579	0.0051057	0.0046204	0.0059784
Strategy 4	0.0064846	0.0062878	0.0073962	0.0061871	0.0060520
Strategy 5	0.0064864	0.0062924	0.0073979	0.0061879	0.0060524

Tab. 5.4: Daily RMSE of the $P\&L$ for the Different Hedging Strategies for a Vanilla EUR/ZAR Call Option

tions on one of the liquid rates) perform better than Strategy 2 and worse than Strategy 3. However, this is not the case over every time period.

The average RMSE for Strategy 4 and 5 is consistently similar, as the EUR/ZAR and USD/ZAR have an implied correlation close to 1, while the EUR/ZAR and EUR/USD have an implied correlation of around zero. As discussed in Section 4.1, this results in a small EUR/USD dephased vega and consequently a negligible weight in the EUR/USD hedging instruments.

The simulation period from 2013-2014 is inconsistent with the rest of the time periods. The average RMSE for Strategy 3 is greater that that of Strategy 2. Furthermore, Strategy 4 and 5 also have a greater RMSE than Strategy 2. Upon further inspection, this period has a number of large ZAR interest rate moves, which affect the strategies involving options (Strategy 3, 4 and 5). These moves could be responsible for the discrepancies in this period. This was further substantiated by setting the ZAR rate to be constant over the whole time period, which resulted in the hedges performing as expected, with Strategy 3 performing the best and Strategy 4 and 5 performing better than Strategy 2 .

The next currency trio that is considered is the USD-EUR-CHF currency trio. The average RMSE for each strategy over all the simulations is shown in Table 5.5. This is further decomposed into the average RMSE over each year of the simulation. For this trio, the liquid rate with the greatest correlation with the EUR/CHF, which will be used in Strategy 5 is the USD/CHF.

The RMSE analysis for the USD-EUR-CHF trio shows that Strategy 2 significantly reduces the daily $P\&L$ for all of the periods. Furthermore, Strategy 3 is consistently the most effective hedging strategy. Strategies 4 and 5 consistently outperform strategy 2, suggesting that vega hedging the illiquid exchange rate using the liquid exchange rates is a viable strategy.

In contrast to the strategies for hedging the EUR/ZAR option, the performance of Strategy 4 and Strategy 5 are not similar. The average RMSE of Strategy 5 is consistently greater than that of Strategy 4 over each time period. This is attributed

	2010-2014	2010-2011	2011-2012	2012-2013	2013-2014
Strategy 1	0.0021858	0.0024935	0.0029703	0.0021580	0.0011210
Strategy 2	0.0008046	0.0008228	0.0013576	0.0006112	0.0004271
Strategy 3	0.0003161	0.0004029	0.0005525	0.0001562	0.0001530
Strategy 4	0.0006501	0.0006034	0.0010596	0.0005537	0.0003841
Strategy 5	0.0007340	0.0007376	0.0012100	0.0005863	0.0004022

Tab. 5.5: Daily RMSE of the $P\&L$ for the Different Hedging Strategies for a Vanilla EUR/CHF Call Option

to the correlation between the EUR/CHF and the EUR/USD, which is greater than zero, so the weight in the EUR/USD hedging instruments is not negligible.

The last currency trio that is analysed is the USD-EUR-JPY currency trio. The average RMSE for each strategy over all the simulations is shown in Table 5.6. This is further decomposed into the average RMSE over each year of the simulation. For this trio, the liquid rate with the greatest correlation with the EUR/JPY, which will be used in Strategy 5 is the USD/JPY.

	2010-2014	2010-2011	2011-2012	2012-2013	2013-2014
Strategy 1	0.545909	0.468635	0.318762	0.837089	0.560253
Strategy 2	0.077570	0.102455	0.088081	0.049055	0.070765
Strategy 3	0.035203	0.041730	0.035005	0.029248	0.034922
Strategy 4	0.067492	0.089213	0.074248	0.047373	0.059188
Strategy 5	0.062257	0.081116	0.072556	0.039857	0.055520

Tab. 5.6: Daily RMSE of the $P\&L$ for the Different Hedging Strategies for a Vanilla EUR/JPY Call Option

The performance of the hedging strategies are in-line with the results of the EUR/CHF option, with Strategy 3 having the lowest RMSE. Furthermore, Strategy 4 and 5 consistently outperform Strategy 2 but are not as effective as Strategy 3. This further suggests that vega hedging the illiquid exchange rate using the liquid exchange rates is a feasible strategy.

In contrast to both the EUR/ZAR and the EUR/CHF hedging strategies, Strategy 5 performs consistently better than Strategy 4. This is unexpected as Strategy 5 only uses options on only one of the liquid exchange rates into account. However, for the USD-EUR-JPY currency trios, the implied correlation estimates were shown to be poor estimates of the realised correlation over the same period, as discussed in Section 5.2. Implied correlation is used to calculate the dephased vegas for the hedging instruments, so if these measures are incorrect this will lead to discrepancies in the hedging results.

The RMSEs of the hedging strategies cannot be compared directly across currency trios as the scale of the daily $P\&L$'s for the different trios vary considerably; however, the overall trends can be analysed. The three currency trios show a number of similarities. The overall trends show that delta hedging removes a large amount of the variance in the daily $P\&L$. Furthermore, the delta-vega hedging strategy using the illiquid underlying and options on the illiquid underlying is the most effecting hedging strategy, with the lowest average RMSE over the whole period. However, this strategy is not possible to implement in practice for illiquid options. Strategy 4 and 5, which are practically viable as they do not involve options on the illiquid underlying, generally perform better than simple delta hedging. These strategies have a lower average RMSE over the whole period than Strategy 2 for all of the currency trios considered.

While the trend differs for whether Strategy 4 or 5 performs best, the RMSE for these strategies are very similar across all the currency trios. If transaction costs are taken into account, it may be better to implement Strategy 5 as options on only one liquid underlying are traded which could reduce the transaction costs.

The RMSE values cannot be directly compared to those of Pesjak (2014) as different currency trios and time periods are considered. However, the overall conclusions are very similar. Pesjak (2014) also concludes that the hedging strategy derived from the triangular relationship performs better than the simple delta hedging strategy. However, if it is possible to implement a delta-vega hedging strategy using options on the underlying, this will be the optimal strategy.

Chapter 6

Discussion and Conclusion

This dissertation interrogates the applications of the geometric framework derived from the no arbitrage condition for the three exchange rates in a currency trio. The theoretical investigation initially looks at the derivation and extensions to the triangular relationship before moving on to a study of volatility and correlation, focussing on extracting and interpreting implied volatility and the literature on implied correlation. Finally, the theory includes a discussion on hedging strategies, in particular for hedging illiquid options.

The developed theory is applied in three different applications. Firstly, the relationship is applied in the spot context, where the robustness of the triangular relationship for various currency trios is considered. Following this analysis, the implied context is examined, where the accuracy of using the relationship to calculate implied correlation is considered. Finally, the application of the relationship in hedging illiquid foreign exchange options is investigated.

The analysis of the relationship in the spot context reveals that while the triangular relationship does provide an intuitive method of interpreting the relationships between volatility and correlation in a currency trio, it is not the most effective tool to look at arbitrage opportunities. The analysis suggests that if the relationship does not hold it is an indication that pricing discrepancies exist in the series. However, this does not correspond exactly to the number and magnitude of the opportunities and it is ultimately more efficient to consider the actual series when attempting to identify arbitrage opportunities.

In contrast, the analysis of the triangular relationship in the implied context shows that the relationship has some important applications in calculating implied correlation. There are a number of interesting properties that arise from the analysis of the implied and realised correlation series. The implied correlation series is a more stable measure than the realised correlation and sometimes lags the realised correlation series. Furthermore, the term structure of the implied correlation is generally upward sloping. The forecasting ability of the implied correlation varies between the currency trios analysed. It is a fairly good predictor of future correlation levels for the USD-EUR-ZAR currency trio and the USD-EUR-CHF trio, but not for the USD-EUR-JPY trio. These results are similar to those in literature, particularly those of Walter and Lopez (2000a) and Fretheim and Høigaard (2012). While the forecasting ability of implied correlation does show potential to be a good predictor, the varied results for different currency trios show that this is not a robust result and requires further analysis across other currency trios.

The application of the triangular relationship to hedging illiquid options reveals that the hedging strategy derived from the relationship performs better than the simple delta hedging strategy. However, the delta-vega hedging strategy using options on the underlying is the most effective strategy if it can be implemented in practice. These findings are consistent with those of (Pesjak, 2014). These results suggest that the proposed hedging strategy is feasible, however, factors such as hedging costs have not been taken into account and further currency trios need to be considered before the strategy can be validated.

In conclusion, the triangular relationship has some interesting and useful applications in computing implied correlation and in hedging illiquid options. However, the relationship should be analysed with respect to the particular currency trio before it is used in any practical applications. Further extension of this research should include an analysis of the forecasting ability of implied correlation and the effectiveness of the hedging strategy across more currency trios and time periods.

Bibliography

Alexander, C. (1998). Volatility and correlation: measurement, models and applications, *Risk management and analysis* **1**: 125–171.

Bossu, S., Gu, Y., Marketing, E. D. I. and London, J. S. L. (2004). Fundamental relationship between an indexs volatility and the correlation and average volatility of its components, *JP Morgan Equity Derivatives* .

Breeden, D. T. and Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices, *Journal of business* pp. 621–651.

Campa, J. M. and Chang, P. K. (1998). The forecasting ability of correlations implied in foreign exchange options, *Journal of International Money and Finance* **17**(6): 855–880.

Castagna, A. (2010). *FX options and smile risk*, John Wiley & Sons.

Cayetano, G. (2007). Studying the properties of the correlation trades.

Clark, I. J. (2011). *Foreign exchange option pricing: A practitioners guide*, John Wiley & Sons.

Cont, R. (2001a). Empirical properties of asset returns: stylized facts and statistical issues.

Cont, R. (2001b). Empirical properties of asset returns: stylized facts and statistical issues.

Demeterfi, K., Derman, E., Kamal, M. and Zhou, J. (1999). More than you ever wanted to know about volatility swaps. goldman sachs quant, *Res. Stratég. Res. Notes, March* .

Demeterfi, K., Derman, E., Kamal, M. and Zou, J. (1999). A guide to volatility and variance swaps, *The Journal of Derivatives* **6**(4): 9–32.

Driessen, J., Maenhout, P. J. and Vilkov, G. (2013). Option-implied correlations and the price of correlation risk.

Faria, G. and Kosowski, R. (2015). The correlation risk premium: Term structure and hedging.

Fretheim, T. S. and Høigaard, H. (2012). The empirical distribution of implied correlation: a study of correlation implied by otc fx cross-option prices.

Gradojevic, N. and Gençay, R. (n.d.). Predictors of triangular arbitrage opportunities: Interdependence and order book indicators.

Hull, J. (2012). *Options, Futures, and Other Derivatives, 7/e (With CD)*, Pearson Education India.

Hull, J. C. and White, A. (2013). Libor vs. ois: The derivatives discounting dilemma, *The Journal of Investment Management, Forthcoming* .

Ito, T., Yamada, K., Takayasu, M. and Takayasu, H. (2012). Free lunch! arbitrage opportunities in the foreign exchange markets, *Technical report*, National Bureau of Economic Research.

Jacquier, A. and Slaoui, S. (2007). Variance dispersion and correlation swaps, *Available at SSRN 998924* .

Jewitt, G. (n.d.). Vanilla fx derivatives trading topics, *FX Derivatives Trader School* .

Levinson, M. (2014). Guide to financial markets, *Journal of financial Economics* **6**: 14–36.

Markowitz, H. (1952). Portfolio selection*, *The journal of finance* **7**(1): 77–91.

Mueller, P., Stathopoulos, A. and Vedolin, A. (2012). International correlation risk, *Available at SSRN 2049809* .

Pesjak, D. (2014). Hedging illiquid fx options: An empirical analysis of alternative hedging strategies.

Poon, S.-H. and Granger, C. (2005). Practical issues in forecasting volatility, *Financial Analysts Journal* **61**(1): 45–56.

Reiswich, D. and Wystup, U. (2012). Fx volatility smile construction, *Wilmott* **2012**(60): 58–69.

Shevchenko, P. V. (2009). Implied correlation for pricing multi-fx options, *arXiv preprint arXiv:0904.4822* .

Singer, B. D. (1999). Risk analysis: A geometric approach, *AIMR Conference Proceedings*, Vol. 1999, CFA Institute, pp. 73–79.

Singer, B. D., Terhaar, K. and Zerolis, J. (1998). Maintaining consistent global asset views (with a little help from euclid), *Financial Analysts Journal* **54**(1): 63–71.

Vonhoff, V. (2006). Dispersion trading, *Available at SSRN 1889147* .

Walter, C. A. and Lopez, J. A. (2000a). Is implied correlation worth calculating? evidence from foreign exchange options, *The Journal of Derivatives* **7**(3): 65–81.

Walter, C. A. and Lopez, J. A. (2000b). The shape of things in a currency trio.

Zerolis, J. (1998). Picturing volatility and correlation, *AIMR Conference Proceedings*, Vol. 1998, CFA Institute, pp. 49–64.

Appendix A

Premium Adjusted Currency Pairs

The conversion from volatility-delta pairs to volatility-strike pairs is more complex for premium adjusted currency pairs. (Castagna, 2010) give the outline of an algorithm to solve for the strike of premium adjusted currency pairs, given the volatilities and corresponding deltas. The algorithm used in this study is detailed below.

1. Set the initial strike (K_1). The equation for the strike is given by

$$K = Se^{\left((r_d - r_f)T + + \frac{1}{2}\sigma_{ATM}T - w\sigma\sqrt{T}N^{-1}(w\delta e^{(rf)T})\right)}, \tag{A.1}$$

   with the premium adjusted delta ($\delta_{s-p.a}$) and the corresponding strike as inputs.

2. Calculate the value (V_1) of the option and the delta of the option (δ_1) using the initial strike.

3. Calculate the derivative of the premium adjusted delta with respect to the strike. This is done numerically using the forward difference method. The new strike, ($\hat{K}$) is set to $\hat{K} = 1.01K$ and the new value (V_2) of the option and the new delta of the option (δ_2) are calculated using $\hat{K}$. This method is described as follows,

$$\begin{aligned} \hat{K} &= 1.01K \\ \hat{\delta} &= \frac{(\delta_2 - \frac{V_2}{S}) - (\delta_1 - \frac{V_1}{S})}{\hat{K}}. \end{aligned} \tag{A.2}$$

4. The new strike is then given by

$$K_2 = K_1 - \frac{(\delta_1 - \frac{V1}{S}) - \delta_{s-p.a}}{\hat{\delta}}. \tag{A.3}$$

5. Set the initial strike to the strike calculated in Step 4 ($K_1 = K2$).

6. Return to Step 2 and repeat the process until the difference between the new and initial strike is below a specified tolerance ($|K_{i+1} - K_i| < \epsilon$). A tolerance of 0.001 was used for this investigation.