The empirical distribution of implied correlation - A study of correlation implied by otc fx cross-option prices

Den empiriske fordelingen for implisitt korrelasjon

Torun Sæther Fretheim and Helene Høigaard

Norwegian University of Life Sciences
Department of economics and resource management
Master Thesis 30 credits 2012

Preface

Writing this thesis has been both academically challenging and rewarding. Many people have offered their help and advice in the process, and we are grateful to each and every one of you.

In particular we would like to thank our supervisor, Professor Espen Gaarder Haug, for providing useful practical tips on how to conduct a research project. He has been both flexible and generous with his time, and we have benefited greatly from his vast experience as an option trader.

We also wish to extend our gratitude to Thijs van den Berg, whose help and encouragement has been invaluable throughout the writing process. Thank you for your patience and the detailed explanations; your enthusiasm and positive attitude are beyond compare.

A special thanks goes to Marcos Carreira and Harel Jacobson; both professionals with extensive knowledge of the OTC FX market. Thank you both for providing us with useful information on the dynamics and distinguishing features of this marketplace. Your guidance has been greatly appreciated.

Last but not least we would like to extend a general thank you to everyone who has been willing to discuss the topic of implied correlation with us along the way.

University of Life Sciences

Ås, May 2012

Torun Sæther Fretheim Helene Høigaard

Executive summary

In this paper we will analyze the distribution and statistical behavior of implied correlation in the foreign exchange market. Implied correlation is here defined as the co-movement between two currencies as implied by the price of a combination of option contracts. We will describe the statistical properties of implied correlations derived from over-the-counter foreign exchange options at-the-money implied volatilities.

There are several reasons why one should examine how correlation evolves over time. This measure is a central input in areas like risk management, and portfolio construction and evaluation. The pricing and hedging of many derivatives are dependent on correlation levels. Our study is also relevant for large companies with cash flows in several currencies.

Market observations show that the correlation between financial assets is difficult to calculate or predict. We examine implied correlations backed out of cross-option data, which is a flexible correlation measure that can react quickly to changing market conditions. It is also attractive from a statistical point of view, as it does not contain the sampling error associated with correlation estimates calculated from logarithmic returns.

We have found that implied correlation is a more stable parameter than historical correlation, especially for short term correlations. Existing literature suggests that correlations are highly unstable. Our empirical study confirms this impression.

The core result of this thesis concerns the term structure of implied correlation. Implied correlations derived from option contracts that are nearby in terms of maturity are more variable than long-term implied correlation. Further, we find that the predictive ability of implied correlation needs to be evaluated empirically – it does not perform uniformly better in terms of forecasting, relative to a historical counterpart.

Finally, we discuss the theoretical distribution of the sample correlation coefficient, and compare this distribution with those of historical estimates. We also address several shortcomings of the theoretical framework that underpins previous research on the information content of cross-option prices, and discuss how to calculate test statistics that are valid in the context of correlation estimates.

Sammendrag

Denne oppgaven analyserer den empiriske fordelingen for implisitt korrelasjon i markedet for unoterte valutaopsjoner. Implisitt korrelasjon er her definert som samvariasjonen mellom to valutaer, definert av priser fra en kombinasjon av opsjonskontrakter. Prisene er gitt i form av *at-the-money* implisitte volatiliteter.

Korrelasjon er en interessant statistisk størrelse av flere årsaker. Samvariasjonen mellom finansielle instrumenter er en viktig parameter ved sammensetning og evaluering av porteføljer, og gode estimater er avgjørende ved beregning av hedgerater. Vår studie er også relevant for internasjonale selskaper der kontantstrømmene omfatter flere valutaer.

Korrelasjonen mellom finansielle variabler er vanskelig å beregne eller forutsi. Vi har fokusert på implisitt korrelasjon, som beregnes ved at man tar opsjonsprisene for gitt og løser med hensyn på korrelasjonen. Dette er en fremtidsrettet korrelasjonsparameter, som kan fange opp markedets forventinger om fremtidige korrelasjonsnivåer. Implisitt korrelasjon er også en attraktiv størrelse fra et rent statistisk synspunkt, da vi unngår usikkerheten som følger ved å beregne en tradisjonell form for korrelasjonsestimat.

Vi har funnet at implisitt korrelasjon er en mer stabil parameter enn historisk korrelasjon, dette gjelder spesielt på kort sikt. Eksisterende litteratur tyder på at korrelasjoner er svært ustabile, og derfor vanskelige å predikere. Vår empiriske undersøkelsen bekrefter dette inntrykket.

Hovedfunnet i vår studie gjelder terminstrukturen til implisitt korrelasjon. Korrelasjoner som er beregnet ved hjelp av opsjoner med kort tid til forfall varierer langt mindre enn langsiktig korrelasjon. Videre finner vi at prediksjonsevnen til implisitt korrelasjon må evalueres empirisk. Denne parameteren gir ikke nødvendigvis bedre prognoser enn estimater basert på tidsseriedata.

Til slutt diskuterer vi den teoretiske fordelingen for korrelasjonskoeffisienter, og sammenligner denne fordelingen med de som er basert på historiske estimater. Vi diskuterer også flere mangler ved tidligere forskning implisitt korrelasjon, og drøfter hvordan man best kan utføre statistiske tester når man studerer korrelasjoner.

Table of contents

1 Introduction

Correlation refers to a broad class of statistical relationships which measure the level of co-dependence between two random variables. In the world of finance it can be thought of as a gauge of how two (or more) securities move in relation to each other. Obtaining good estimates of the co-movement between financial securities is a crucial part of advanced portfolio management. High-quality correlation estimates are important both for assessing total portfolio risk and for making decisions regarding future trading strategies.

In this paper we will study exchange rate correlations implied by market data on a trio of option contracts. There is already a substantial literature exploring the informational content of publicly traded options, however the bulk of this literature focus on implied volatilities rather than correlations. We will examine implied correlations backed out by historical price quotes, as well as the distribution of this correlation measure over time. We will also calculate rolling historical correlation estimates, and compare the distribution of historical correlation with a frequency distribution of the sample correlation coefficient, as derived by Fisher (1915). To our knowledge we are the first to conduct such a comparison. Finally, we will evaluate the forecasting accuracy of implied versus historical correlation.

Research on exchange rate correlations implied by market data is sparse. Previous work in this area mainly evaluates the forecasting accuracy of different correlation measures. The studies closest related to our work are Bodhurta and Shen (1995), Haug (1996), Siegel (1997), Campa and Chang (1998), Lopez and Walter (2000) and Castrén and Mazotta (2005), which all focus explicitly on exchange rate correlations.

Our research will expand upon these studies in a number of ways. To our knowledge we are the first to examine the empirical distribution of implied correlation. We will also analyze implied correlation over a larger set of tenors than what has been done previously. Our study will span contracts with one month, two months, three months, six months, nine months and one year to expiration. Finally we will examine currency triangles which are not yet covered in the literature.

We will analyze data from the interbank foreign exchange options market, which is one of the largest and most liquid derivatives markets in the world (Wystup 2010b). One feature of this trading venue is of particular importance for our study, namely that the market standard is to quote implied volatilities for each contract, instead of the option invoice price.

This market convention has evolved for two reasons. First and foremost, volatility will normally change less erratically than option prices. This means that price quotes in the form of implied volatilities need to be updated less frequently than price quotes in a nominal form. Second, the OTC FX market use only one formula for quoting implied volatilities, namely Garman and Kohlhagen (1983) which is based on Merton's (1973) proportional dividend extension of the Black and Scholes (1973) model. The use of a customary option pricing formula is what makes it possible to quote option prices directly in terms of implied volatilities.

That options are quoted in terms of implied volatilities is important in the context of this study, because it means that our calculations can be made more accurate than what would be possible with data from other markets. Obtaining the implied volatility level by some iterative method would require synchronous data for options prices and its underlying, a problem we avoid when implied volatility is traded on the market.

Another beneficial property of the OTC FX market is that new options are quoted every day. This entails that time to maturity can be held constant across our analysis, and the strike price will always be at-the-money. As a result we avoid *"mixing apples and oranges",* a phenomenon which will be discussed further in the data section.

The remainder of this paper is organized as follows. In section 2 we will introduce the methodological framework of the study. The data set is described in section 3. Section 4 will provide empirical evidence of the statistical behavior of implied and historical correlation in the OTC FX market, based on the framework presented in sections 2 and 3. In section 5 we will evaluate the predictive accuracy of implied correlation. Section 6 presents an overall summary with concluding remarks.

2 Theoretical framework

2.1 Why should we be interested in implied correlation?

There are several reasons why one should examine how correlation evolves over time. This measure is a central input in areas like risk management and portfolio construction and evaluation. It is also an important factor for anyone trading financial securities. The financial crisis highlighted how correlation risk is a central parameter in the valuation of financial securities.

The textbook definition of implied volatility is that it represents the market's best estimate of future volatility (see e.g. McDonald 2006). Conversely, implied correlation can be thought of as the market's perception of the future correlation level. If implied volatility is interpreted as a measure of supply and demand for options (see e.g. Haug 2007a, Triana 2009), then a set of implied correlations could give an indication of what currencies are in excess supply or demand in the foreign exchange market.

Further, the pricing and hedging of many derivatives are dependent on correlation levels. Quantos, swaptions and outperformance options are examples of such products, which has become quite popular during the last decade. The theoretical value of an option contract can be highly sensitive to the correlation input, and changes in the correlation level can also have a large effect on the *Greeks*. There are many types of hedging techniques, but regardless of which one you rely on it will require an accurate estimate of the hedge ratio.

There are also derivatives that allow you to engage in outright bets on the future correlation level. A correlation swap is such a derivative, and its payoff will usually depend on the observed average correlation of a basket of underlying products. Knowledge about the empirical distribution of correlation levels and how this measure evolves over time will be of value to anyone who wish to evaluate the risk and profit-loss potential associated with such contracts.

This study focus on plain vanilla straddles, and it is possible to trade correlation using a combination of three such contracts. Another way of locking in a fixed correlation level involves the purchase of three *Forward Volatility Agreements* (FVA). While this is a relatively new contract, its payoff is similar to that of a forward starting straddle. A FVA is usually a forward starting contract on the future spot *implied* volatility, i.e. a pure forward *Vega* instrument. These products are usually priced close to the ATM implied volatilities we are using in this study.

Market observations show that the correlation between financial assets is difficult to calculate or predict. According to Wilmott (2006), correlation estimates extracted from financial time series are notoriously unstable. While many chose to ignore these findings and model correlation as a fixed number, the last decade has brought the introduction of a new class of option pricing models that allows the correlation coefficient to vary stochastically over time.[1]

[1] See e.g. Da Fonseca et al. (2007) and Ma (2009).

[2] As pointed out by Wilmott (2006), uncertainty is different from randomness in that it does not assume some

Kamtchueng (2010) takes the notion of unpredictability one step further and derive an uncertain correlation model.[2] If one is trying to develop (or evaluate the performance of) such models, it is important to know how correlation evolves over time. Wrong assumptions about the distribution of correlation could result in unrealistic models.

Correlations are also important from a corporate perspective. Clark (2011:225) uses the example of a hypothetical Japanese sporting goods provider, who sells merchandize both internationally and in their domestic market.

If this company expects revenues from e.g. Europe and the US, they are exposed to currency risk in both EURJPY and USDJPY. The exporter could hedge this exposure by buying puts on EURJPY and USDJPY, while another approach entails viewing their position as a portfolio of assets. If the exporter has a good forecast of future correlation between the two currencies, they can hedge their position by purchasing a multicurrency basket option instead of two separate FX contracts. This will generally be cheaper than to hedge the two exposures separately.

For concreteness, consider the case of perfectly correlated or anti-correlated currencies (Clark 2011). If the correlation between EURJPY and USDJPY is 1.0, then the hedge for this "portfolio" will indeed be the sum of two options – the assets will move in complete tandem. However, if the correlation is -1.0, then one currency will strengthen as the other one weaken. In this scenario one of the options would *always* be redundant. While perfectly correlated or anti-correlated currencies hardly exist in the market, this way of thinking about currency risk illustrates how multicurrency basket options can be a cheaper hedge than plain vanilla contracts.

In other words, there are many reasons why one should be interested in the empirical distribution of different correlation measures. Implied correlation is especially appealing because it is a forward-looking economic indicator which can incorporate both news and views concerning future market conditions. We also avoid the statistical sampling error associated with classical correlation measures calculated from historical data.

---

[2] As pointed out by Wilmott (2006), uncertainty is different from randomness in that it does not assume some probabilistic description of what may happen. Because the model is rid from probabilities we are also rid of the notion of expectations, both real and risk neutral.

2.2 Literature review

To our knowledge, Tompkins (1994) was the first to publish how to calculate implied correlation. His results are an extension of a framework presented by Margrabe (1978). The Margrabe exchange option model concerns the exchange of one risky asset for another at a prearranged strike. Tompkins draws on these results and uses the fact that the total variance of an exchange option is a function of the variances of both assets, minus 2 times their covariance. By rearranging the Margrabe volatility input Tompkins obtains a measure of implied correlation, and shows that trading correlations is nothing but an extension of trading volatility.

Bodurtha and Shen (1995) is an early paper that discusses implied correlation in an empirical context. They use exchange traded options to determine the yen-mark option implied covariance matrix. From this matrix they calculate implied volatilities (standard deviations) and correlations. They find that the implied parameter estimates adds to the forecasting accuracy of historical-based estimates, and suggests that this correlation measure provide forecast explanatory power similar to what has been documented for implied volatilities.

Haug (1996) was the first to examine implied correlations backed out from OTC cross-option price quotes. Using proprietary data from Chase Manhatten Bank (now JPMorgan Chase) Haug analyzes options on USDDEM, USDJPY and DEMJPY. The study covers options with one month and six months to expiration, and spans October 1990 through December 1994. According to Haug, implied correlation is superior to historical correlation in predicting future co-movement between USDDEM and USDJPY. The paper also note that implied correlation vary markedly less than historical correlation for this time period. Haug suggests that this could be caused by market participant's anticipation of some mean reversion in the correlation level.

Siegel (1997) examine implied correlations backed out of cross-option data from the Philadelphia Stock Exchange (PHLX). He finds that this correlation measure seems to provide improved estimates of future correlation between GBPDEM and DEMJPN relative to estimates based on historical correlations.

Campa and Chang (1998) expand upon this work by examining OTC currency options. They find that implied correlation among the DEMUSD, USDJPY and DEMJPY currency pairs from January 1989 through May 1995 shows predictive power for the future realized correlation at one and three month forecasting horizons. Castrén and Mazotta (2005) confirm this result when they examine implied correlations derived from a larger number of exchange rate pairs. Their data consist of 1 month implied volatilities from the OTC market, and their study span from 1992 to 2004.

Walter and Lopez (2000) arrive at a somewhat different conclusion. They examine data consisting of daily, one-month and three-month implied volatilities for two currency trios. For USD/DEM/JPY they analyze price quotes from October 1990 through April 1997. For USD/DEM/CHF their data span from September 1993 through April 1997. Their analysis suggests that implied correlations are not unambiguously superior in forecasting future realized correlation. Walter and Lopez find that even though it performs well for the USD/DEM/JPY triangle, implied correlations add little if any information to a forecast of future co-movement between USD/DEM/CHF.

This contradicts the findings from the other studies we have mentioned. It suggests that even though implied correlation is a forward looking economic indicator, i.e. an indicator that can incorporate the market's assessment of future events, this parameter should be interpreted with care. While evidence is found that supports the use of cross-options data to forecast future co-movement between *liquid* currency pairs, the jury is still out on the practical usefulness of this correlation measure.

Finally, to make inference about implied correlations it is vital to have a basic understanding of implied volatilities and their statistical properties. Haug et al. (2010) examine implied volatility quotes for USDJPY, USDGBP and EURUSD, with a focus on descriptive statistics and empirical distributions. The study is based on data from the OTC FX market, and the dataset cover 6 years (2000-2006) for EURUSD and 10 years (1996-2006) for the other pairs. They also report the historical correlation of implied volatility and find that the relationship between implied volatilities can vary substantially over time. In addition the study reveals that the implied volatility of options which are closer in terms of time to expiration has a higher correlation than options that differ in terms of maturities.

2.2.1 The Garman-Kohlhagen formula

We have mentioned how the use of a standard option pricing formula is what makes it possible to quote options in terms of implied volatilities. This formula is an extension of the Black-Scholes-Merton model, and commonly referred to as *Garman-Kohlhagen* (1983).[3]

[3] In passing we would like to mention that Grabbe (1983) published the exact same formula the same year. Despite this fact the model is usually referred to as the Garman-Kohlhagen formula, and we will follow this convention.

Implied volatilities always refer to the volatility level yielding a particular strike price when used in conjunction with this model.

Note that the markets adoption of a standard pricing formula should not be considered a validation of its assumptions; neither does it imply that these assumptions are believed to be realistic by practitioners (see e.g. Derman and Taleb 2005, Haug 2007b, Haug and Taleb 2011). The markets way of using this formula is nothing more than a convenient way of quoting option prices.

The standard Garman-Kohlhagen formula is given by the expression below; we use ***P*** and ***C*** to distinguish between the theoretical value of a put and a call:

$$V_0^C = S_0 e^{-r_f T} N(d_1) - K e^{-r_d T} N(d_2)$$

$$V_0^P = K e^{-r_d T} N(-d_2) - S_0 e^{-r_f T} N(-d_1)$$

$$d_{1,2} = \frac{ln(S_0 - K) + \left(r_d - r_f \pm \frac{1}{2}\sigma^2\right)}{\sigma\sqrt{T}}$$

Where $N(x)$ denotes the standard normal probability density function, S_0 denotes the spot price, K represents strike price, T represents time to maturity, σ is volatility and $r_{d,f}$ represents the domestic and foreign interest rate.

The formula above makes use of a domestic and a foreign interest rate. We would like to emphasize that the price of an FX option will be the same, no matter what currency (and hence interest rate) you consider as your domestic one. This is important because it is the only way the different sides of a trade can agree on a common price. A formal proof can be found in Clark (2011:27-28), we merely note that it does not matter which currency you take as your reference point.

2.3 The OTC market for foreign exchange

The OTC interbank market for FX options is perhaps the most actively traded derivatives market in the world. The market participants are generally large institutions like e.g. international banks, hedge funds and other investment groups. As a consequence, the dynamics of this marketplace is not something which is common knowledge – at least not among academics.

In this section we will give a brief introduction to the FX market and its trading conventions. Our outline is based on Castagna (2010), Wystup and Reiswich (2010b) and Clark (2011).

A special feature of foreign currency options is that the underlying asset consists of a pair, a duality that is unique to this marketplace. Each pair consists of two currencies, for concreteness we will call them currency *XXX* and currency *YYY*. The pair is denoted *XXXYYY* and this label dictates how many units of YYY you can by with one unit of XXX. This corresponds to the foreign currency spot rate, $S_t = XXXYYY$, which represents the required units of domestic currency needed to buy one unit of foreign currency at time t.

The first quote in a currency pair is usually referred to as the *foreign* currency and the second one is called *domestic*. These terms does not specify a geographical region, it is rather a convention to distinguish between the different sides of a trade. Alternatively one can use *base* (in place of foreign) and *numeraire* (in place of domestic). In the following we will apply the terms foreign/domestic whenever we discuss a twosome of currencies.

Most options in the FX market are European. A plain vanilla European currency option consists of both a call and a put. If the underlying is quoted as *XXXYYY*, a ***XXX call YYY put*** is defined as the buyers right to buy (sell) the notional amount of *XXX (YYY)* currency at the strike price on expiry. A ***XXX put YYY call*** gives the buyer the right to sell (buy) the amount of *XXX (YYY)* currency. The strike price is defined as the fixed price in which one can buy (sell) the relevant currencies.

Plain vanilla options are quoted for standard maturities and they are usually settled physically at maturity. The most actively traded tenors are one day, one week, one month, two months, three months, six months, nine months, and one year (Castagna 2010). When a transaction is agreed upon, the option premium can be paid in either of the currencies of the underlying pair.

2.3.1 Quoting conventions

In the OTC FX market options are quoted in terms of delta rather than premium. The standard is to ask for a delta and receive a price in the form of a Garman and Kohlhagen (1983) implied volatility in addition to the strike, given a spot reference.

Because volatilities are assigned to deltas rather than strikes, the trader is less exposed to small price movements in the underlying market. The strike price, in absolute terms, is determined only after an agreement is reached. This ensures that the option possess the “correct” features in terms of exposure to the underlying pair *and* implied volatility (Castagna 2010:15). As the spot exchange rate can move substantially during a short interval of time, this way of quoting options help maintain the efficiency of the FX OTC market.

Delta is also used to measure the degree of moneyness. An option is said to be at-the-money (ATM) if the current price of the underlying is equal to the specified strike price. In the FX market there are four different types of ATM quotations (Wystup and Reiswich 2010b). The summary below gives a quick overview of the different quotation styles:

ATM spot:	$K = S_t$
ATM forward:	$K = S_f$
ATM value-neutral:	*with K such that call* $V_{call} = V_{put}$
ATM delta-neutral:	*with K such that* $\Delta_{call} = -\Delta_{put}$

Where K represents strike price, S_t represents the spot price in the underlying market, S_f represents the forward price and V represents the theoretical value of an option contract.

In this paper we will analyze implied volatilities quoted in the form of ATM delta-neutral straddles. These structures on standard dates are the most liquid of all FX option contracts (Castagna 2010:16). Choosing the strike in the ATM-delta-neutral sense ensures that the structure has a zero spot exposure.[4] It also means that no delta hedge is needed when entering into the trade.

2.4 From implied volatility to implied correlation

2.4.1 Implied volatility

It is difficult to make inference about implied correlation without first discussing what the expression *implied volatility* really represents. Technically speaking it is the volatility that - when used in conjunction with a particular pricing formula - yields a theoretical value for an option equal to the current market price (Wilmott 2006). In other words, implied volatility is strongly related to how the market is currently pricing an option.

Textbooks often suggest that implied volatility should be interpreted as “the market’s best estimate of future volatility”, and there is some merit to this notion. By taking price quotes from the market and reverse engineer the implied volatility level, you will get a volatility measure that reflects the current (model-dependent) sentiment in the market. Wilmott (2006) also note that the supply and demand of options is an important part of what determine implied volatility levels. As long as the theoretical value is not too out of line compared with other products, there is no way of assessing whether changes in implied volatilities are driven by changed perceptions of future variance, or changed conditions in supply and demand.

[4] Formally speaking, delta is defined as the percentage of the foreign notional one should hedge in the underlying market.

Literature that discuss implied volatility and how it responds to supply and demand include Derman and Taleb (2005), Wilmott (2006), Haug (2007a), Triana (2009) and Haug and Taleb (2011).

2.4.2 Implied correlation

Correlation describes a relationship between two random variables. This relationship is usually reported in terms of the correlation coefficient, which is a measure of the linear dependence between the two variables (Tsay 2010). The correlation coefficient will always take on a value between +1 and −1. A correlation of +1 indicates that the two variables increase or decrease in lockstep, while a correlation of -1 means that the two variables always move in opposite directions. A correlation of 0 suggests that the two variables are independent.

Note that a correlation of +1 does *not* indicate that asset X will move by 1 % in response to a 1 % move in asset Y (in the same direction). Correlation really describes a *ratio* of movement (Taleb 1997:89), and to be correlated by 100 % could also imply that asset X will move up by 2 % in response to a 1 % up move in asset Y.

In this section we will take the reader from implied volatilities to implied correlation. The latter is derived by using the Garman-Kohlhagen pricing formula in conjunction with the triangle arbitrage condition on currencies. A rigorous derivation of the relationship between implied volatility and the implied correlation coefficient can be found in Castagna (2010: 269-272). We will only provide a non-technical summary to explain the reasoning behind this correlation measure, and to fix notation.

Assume that we have 3 currencies; *XXX*, *YYY* and *ZZZ*. From these currencies three pairs can be formed; *XXXYYY*, *XXXZZZ* and *YYYZZZ*. In the following we will refer to the different currencies by their first letter. The exchange rates are thus given by S_{XY}, S_{XZ} and S_{ZY}, while the implied volatilities will be denoted σ_{XY}, σ_{XZ} and σ_{ZY}.

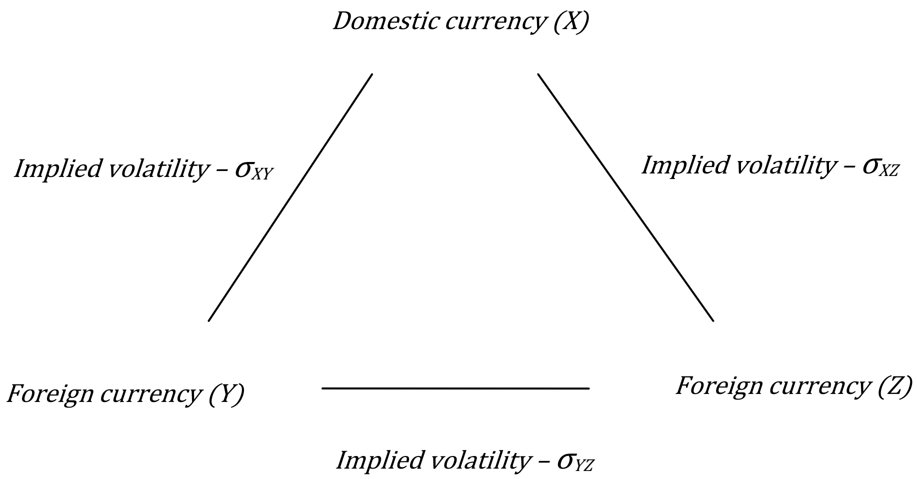

Figure 1: Given three currencies and three appropriate option contracts we are able to extract implied correlations from market prices. Note that a total of three correlations can be derived from this triangle (Siegel 1997:371).

From the no arbitrage constraint we have that:

$$S_{XZ} = \frac{S_{XY}}{S_{ZY}}$$

which implies that you can calculate the exchange rate for the pair *XZ* once the other two are known. If we assume that the logarithmic returns of the spot exchange rates are jointly normally distributed random variables, it follows that the implied variances across these exchange rates are related according to:

$$\sigma_{XY} = \sqrt{\sigma_{XY} + \sigma_{ZY} - 2\rho_{XY,ZY}\sigma_{XY}\sigma_{ZY}}$$

We see that the variance of the currency cross exchange rate is determined by the variance of the other spot rate processes, together with the correlation between them. This means that if the implied volatilities for three currency pairs can be found in the market, it is possible to infer the implied correlation among the pairs by rearranging the formula above:

$$\rho_{XY,ZY} = \frac{\sigma_{XY}^2 + \sigma_{ZY}^2 - \sigma_{XZ}^2}{2\sigma_{XY}\sigma_{ZY}}$$

In other words, implied correlation represents the degree of co-movement between two currencies using a third currency as a numeraire. If the underlying spot exchange rate is quoted as *ZX* instead of *XZ* the following expression can be derived:

$$\rho_{XY,ZY} = \frac{\sigma_{XZ}^2 - \sigma_{ZY}^2 - \sigma_{XY}^2}{2\sigma_{XY}\sigma_{ZY}}$$

While it reasonable to assume that arbitrageurs will make the triangle condition on currencies hold, it is not realistic to assume that the changes in the spot exchange rates are normally distributed (Boothe and Glassman 1987, Johnston and Scott 1999). However, the relationship of implied correlation is still valid because it is derived using a constraint, namely the "no arbitrage" condition stated above. It is difficult to make inference about the practical implications that follows from non-normal log-returns of the exchange rates, but it seems plausible that a high peak and fat tails in these distributions will increase the variability of implied correlation.

2.5 Historical correlation

To obtain a consistent estimate of the realized correlation between two currency pairs we have computed rolling historical correlations. By *rolling* it is implied that each calculation is tracking a certain period of time, and then rolls forward as time progresses.

This correlation measure allows for a time-varying relationship between variables, and has the benefit of being easy to calculate. However, it comes with certain caveats (see e.g. Engle 2002). Most importantly it gives equal weight to all observations within the chosen time frame, and zero weight to observations outside that window. It follows that extreme observations can cause jumps in the correlation level when these outliers exit the estimation window and suddenly receive zero weight. The correlation level will drop, yet this drop will be completely spurious as it is not caused by changing market conditions but rather an event in the past. Finally, assigning equal weight to all observations leads to new information being reflected rather slowly in the estimated coefficient.

From the discussion above it follows that one must be cautious when determining how many observations to include in an estimation window. While too few observations will lead to problems with sampling error, you also wish to obtain an estimate that reflects *current* market conditions. It can be argued that including observations from the distant past reduces the timeliness of the estimate; what happened three months ago might not be relevant to today.

Formally, the rolling window correlation over a T-day period is calculated by dividing the equally weighted covariance estimate by the square root of the product of the variance estimates:

$$\hat{\rho}_{(r_{XY},r_{YZ})_{t,T}} = \frac{\sum_{i=1}^{T}\left(r_{XZ,t+i}-\mu_{XZ}\right)\left(r_{YZ,t+i}-\mu_{YZ}\right)}{\sqrt{\sum_{i=1}^{T}\left(r_{XZ,t+i}-\mu_{XZ}\right)^2}\sqrt{\sum_{i=1}^{T}\left(r_{YZ,t+i}-\mu_{YZ}\right)^2}}$$

where r_{XY} and r_{ZY} represents the natural logarithm of the price changes, and μ_{XY} and μ_{ZY} are the corresponding sample means over a T-day period.

2.5.1 Detecting non-stationarity

Before we can calculate historical correlations it is necessary to confirm that the log-returns of the exchange rates are stationary. A stationary process has constant mean and variance, and the covariance is independent of time. For non-stationary variables these values are not constant, thus without stationarity one cannot be sure to obtain consistent estimators (Hamilton 1994).

If the variables used to calculate the sample correlation are non-stationary, the sample correlation will converge to a random variable rather than the true correlation. This can be shown mathematically by calculating the asymptotic distribution of the correlation constructed from non-stationary variables using the functional central limit theorem (Hamilton 1994:483-485). The sample correlation will in this case converge to integrals of Brownian motions instead of the true correlation. In contrast, the sample correlation of stationary variables will converge to a constant.

In sum, if the mean and variance of a time series are not well defined, then neither is its correlation with other variables. For this reason one should be cautious about trying to estimate correlations based on non-stationary variables.

The standard way of detecting non-stationarity involves testing for unit roots (Wooldrigde 2003). The most common non-stationary processes in financial data have one unit root, i.e. they are integrated of order one. While visual inspection of the process can provide a basis for insight, it is necessary to test for unit roots in a more formal way. A common approach takes a generalized auto-regressive model as reference point:

$$y_t = \alpha + \rho y_{t-1} + \varepsilon_t, t = 1, 2, \ldots \quad (1)$$

where $y_0$is the observed initial value.

If the sequence $\{y_t\}$ follows the process modeled in (1), it has a unit root if $\rho = 1$.The α-term indicates whether or not the price series has drift.

A convenient way of carrying out the unit root test is to subtract y_{t-1} from both sides of (1) and define $\theta = \rho - 1$:

$$\Delta y_t = \alpha + \theta y_{t-1} + e_t \quad (2)$$

It follows that when we test H_0: $\theta = 0$ we indirectly test if ρ is equal to unity. $\rho = 1$ indicates that the price series follows a random walk and that the time series is non-stationary. This test is known as the Dickey-Fuller test for unit roots (Wooldrigde 2003). Note that when testing the significance of the θ-parameter we cannot use the ordinary t-statistics, much stricter values need to be tabulated. The asymptotic distribution of the t-statistic under H_0 has come to be known as the Dickey-Fuller distribution, after Dickey and Fuller.

It is also possible to use an extended version of this test, namely the augmented Dickey-Fuller test (ADF-test):

$$\Delta y_t = \alpha + \theta y_{t-1} + \sum_{i=1}^{n} \gamma_i \Delta y_{t-i} + e_t \quad (3)$$

As we can see from (3), this version of the test allows for more complex model dynamics. Most notably it controls for serial correlation in the error term by including the higher-order autoregressive terms in the regression (Hamilton 1994).

The result of an ADF-test can be very sensitive to the number of lags included. The formal way of determining the appropriate number of lags is to examine some information criteria such as the Akaike or the Schwarz-Baysian criterion (Tsay 2010). An alternative approach is to test down from high orders and examine the t-values of the lagged parameters. We will follow the latter approach.

2.6 The distribution of the sample correlation coefficient

Above we have argued that estimates of the correlation coefficient will vary substantially as we vary the number of observations used inside the estimation window. Figure 2 display the density function of the sample correlation coefficient, which confirms that the distribution indeed gets more peaked as *n* increases.

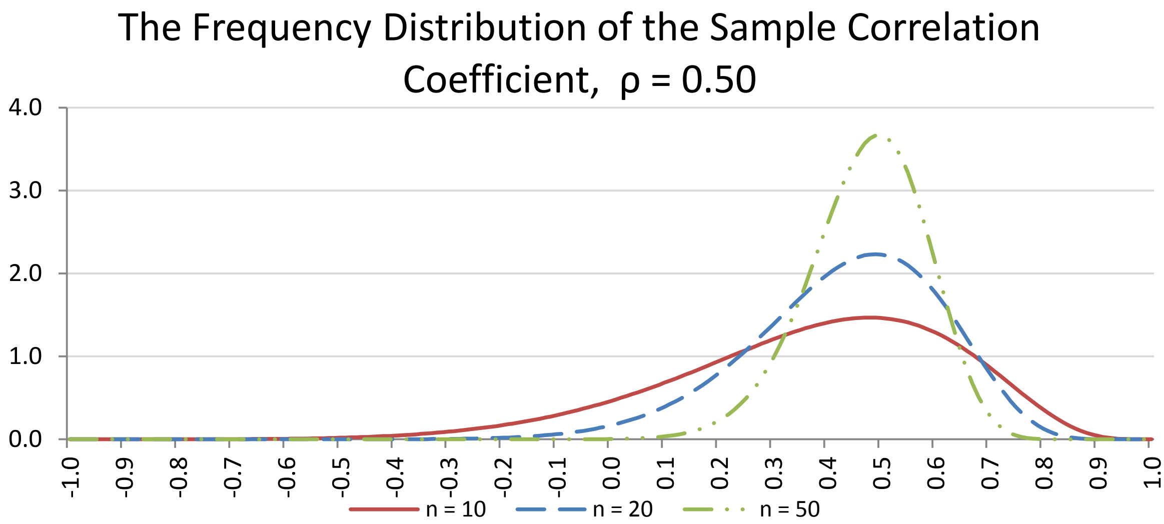

Figure 2: The distribution of the sample correlation coefficient gets more peaked as the sample size get larger. The horizontal axis represents observed correlation.

The frequency distribution of the correlation coefficient was first derived by Fisher (1915). He obtains a general distribution for non-zero correlation coefficients for any sample size, assuming a bivariate normal distribution. Our implementation of this distribution is based on Stuart and Ord (1994:559-664) and Van den Berg (2012).

The density function of the correlation coefficient is given by:

$$P(r) = \frac{(n-2)\Gamma(n-1)(1-\rho^2)^{\frac{n-1}{2}}(1-r^2)^{\frac{n-4}{2}}}{\sqrt{2\pi}\Gamma(n-\frac{1}{2})(1-\rho r)^{n-\frac{3}{2}}} \times {}_2F_1\left(\frac{1}{2},\frac{1}{2},\frac{2n-1}{2},\frac{\rho r+1}{2}\right)$$

Where ρ represent the population correlation coefficient, r is the sample correlation coefficient and n is the sample size.

$\Gamma(\cdot)$ is the standard gamma function, defined as:

$$\Gamma(p) = \int_0^\infty x^{p-1} e^{-x} dx$$

${}_2F_1(\,)$ is a hyper-geometric function, defined as:

$${}_2F_1(a,b,c,z) = \frac{1}{n}\sum_{n=0}^{\infty}\frac{(a)_n(b)_n}{(c)_n}\frac{z^n}{n!} = 1 + \frac{ab}{c}\frac{z}{1!} + \frac{a(a+1)b(b+1)}{c(c+1)}\frac{z^2}{2!} + \cdots$$

Where $n! = 1 \times 2 \times 3 \ldots .n$ is a factorial and $(a)_n = a(a+1)(a+2)\ldots(a+n-1)$ is the raising factorial (the Pochhammer symbol).

The moments are (Stuart and Ord 1994:566-567):

$$\bar{r} = \rho \times \left[1 - \frac{\rho(1-\rho^2)}{2n} + \cdots\right]$$

$$var(r) = \frac{(1-\rho^2)^2}{n} \times \left[1 + \frac{11\rho^2}{2n} + \cdots\right]$$

$$\gamma_1 = \frac{6\rho}{\sqrt{n}} \times \left[1 + \frac{77\rho^2 - 30}{12n} + \cdots\right]$$

$$\gamma_2 = \frac{6}{n} \times [12\rho^2 - 1] + \cdots$$

Note that the mean and variance are dependent on both ρ and n, a detail that makes this distribution very different from e.g. the normal distribution.

In figure 3 below we have plotted the density function of the sample coefficient for $n = 10$ and ρ = 0, 0.25 and 0.50. We can see that as the value of the correlation coefficient increases, the density function is skewed more heavily towards the right. There is also a slight change in the curvature of the distribution; it gets more peaked as the value of the population correlation coefficient increases. In figure 2 we could see that the distribution also gets more peaked for larger *n*.

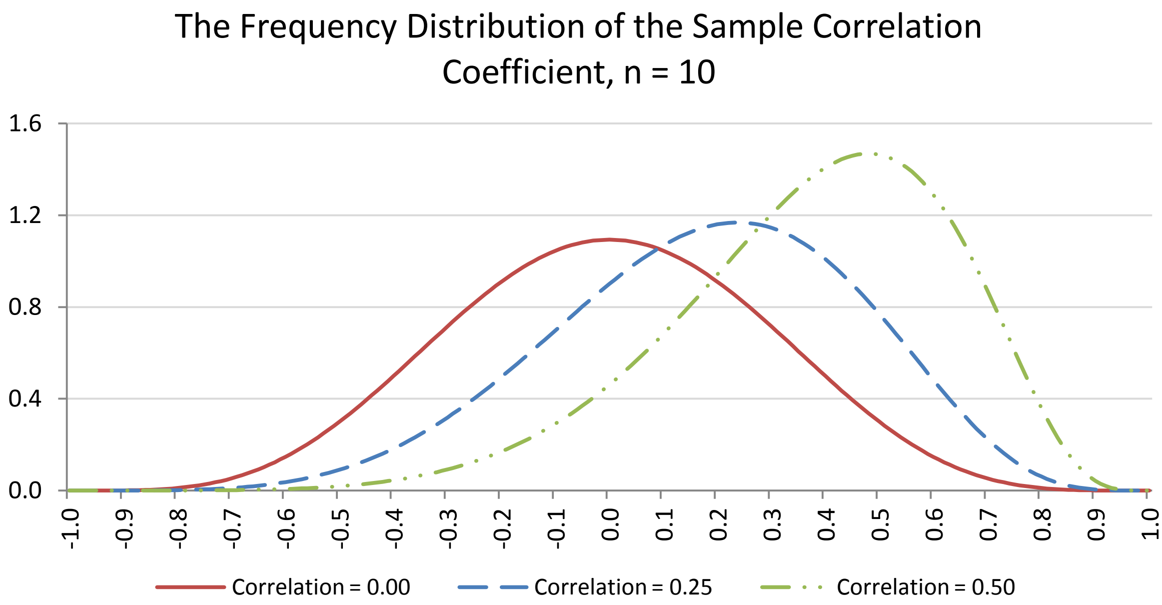

Figure 3: The frequency distribution of the sample correlation coefficient for different values of the population correlation coefficient.

Although this distribution was derived as early as 1915, we have not seen many references to a sample correlation distribution in the existing literature. Furthermore, we have not found any existing literature that compares the accuracy of the sample correlation distribution with results obtained from empirical data. We will apply this distribution in chapter 4, to evaluate how the theoretical density function aligns with empirical data.

We have implemented this distribution using VBA in Excel. Our implementation only works up to *n* of approximately 150. The reason for this is that as we increase *n*, the gamma function takes on extremely large values - larger than Excel can handle.

2. 7 The geometry of correlation coefficients

In this section we will give a brief rundown of why the correlation coefficient cannot be treated as just any other metric. While some quantities like e.g. lengths, angles and variances are arithmetically additive, correlation coefficients are not (Garcia 2011).

For an arithmetic average expression to be valid, an additive metric is required. The reason why correlation coefficients are not additive stems from a geometric interpretation of correlation as the dot product of the centered data vectors normalized by the lengths of the data vectors. When variables are defined in this way, the correlation between the original variables equates to the cosine of the angle between their corresponding vectors (Tu et al. 2006, Fisher 1915). This entails that the angles between pairs of a triplet of data vectors must form a triangle in a spherical geometry, and hence be subject to some form of triangle inequality. You can also see this by looking directly at the definition of the Pearson correlation coefficient, which obviously defines a cosine:

$$r = \frac{\sum_{i=1}^{n}(X_i - \bar{X})(Y - \bar{Y})}{\sqrt{\sum_{i=1}^{n}(X_i - \bar{X})^2}\sqrt{\sum_{i=1}^{n}(Y_i - \bar{Y})^2}}$$

We see that the *r*-value is defined as the ratio of the covariance between two variables, normalized by the product of their standard deviations. As this defines a cosine it follows that computing an average correlation coefficient as

$$\bar{r} = \frac{1}{n}\left(\frac{Cov(x_1, y_1)}{s_{x_1}s_{y_1}}\right) + \left(\frac{Cov(x_2, y_2)}{s_{x_2}s_{y_2}}\right) + \cdots + \left(\frac{Cov(x_n, y_n)}{s_{x_n}s_{y_n}}\right)$$

is essentially the same as attempting to average

$$\bar{r} = \frac{1}{n}(\cos\theta_1 + \cos\theta_2 + \cdots + \cos\theta_n)$$

a computation that cannot be valid as the sum of a cosine is not a cosine (Garcia 2010). Rather, the sum of two cosines is related according to the following formula:

$$\cos\alpha + \cos\theta = 2 \times \cos\frac{\alpha+\theta}{2} \times \cos\frac{\alpha+\theta}{2}$$

2.7.1 Fisher's Z transformation

Since the discussion above indicates that it is inappropriate to average "raw" correlation coefficients, we calculated $\bar{r}$ after employing Fisher's *r to Z transformation*. This conversion was first introduced by Fisher in 1915 (see also Fisher 1921), in order to stabilize the variance of the sampling distribution of correlation coefficients:

$$Z = \frac{1}{2} ln\left[\frac{1+r}{1-r}\right] = \text{arctanh}\ (r)$$

where ln denotes the natural logarithm, r is the sample correlation and $arctanh(\cdot)$ is the inverse hyperbolic tangent function. The new variable Z will be approximately normally distributed with $\mu = \bar{Z} = \frac{1}{2} ln\left[\frac{1+\rho}{1-\rho}\right]$ and $\sigma = (N-3)^{-0.5}$.

This transformation aims to make the variability of r values which are close to ± 1 comparable to that of r-values in the mid-range. While this bias is different from what is discussed above, the *Z*-values obtained by means of the transformation are additive. This suggests that a proper average of the correlation coefficient can be obtained by first computing a weighted average of the Z-values, and then use the inverse of the function above to find $\bar{r}$ (Garcia 2011).

A study by Silver and Dunlap (1987) uses Monte Carlo simulation to examine the negative bias caused by averaging raw correlation coefficients. Their study lends empirical support to the procedure we have just described, and conclude that it is beneficial to transform correlation coefficients into *Z*-values prior to averaging, and then reverse the process. While an average obtained by using Fisher's *Z* transformation is slightly positively biased, it is always *less* biased than an average calculated from "raw" correlation coefficients.

While the Fisher transformation seems popular in areas such as psychology and genetics (see e.g. Garcia 2011 for a list of references), we have found no references to the non-additivety of correlation coefficients in relation to *financial* variables. This makes it difficult for us to make inference about the "correctness" of employing this remedy in our study. We will therefore briefly discuss why we believe it is appropriate to use the Fisher transformation in this context.

We have already noted that in a spherical geometry, the standard deviation is the length of a centered vector, and the correlation is the cosine of the angle between these centered vectors:

$$r = \cos\theta$$

When the angle is 0, the variable vectors fall on the same line, and $\cos\theta = \pm 1$. If the angle is 90º, the vectors are at right angles and $\cos\theta = 0$ (Fox 1997:246, Rodgers and Nicewander 1988). It follows that the magnitude of bias introduced from using raw values of r should be highest for $\bar{r}$ near ± 1. For $\bar{r}$ near 0, the bias will also be zero.

We will now get a little ahead of ourselves, and use our empirical findings for illustration. We will begin by looking at the empirical distribution of historical correlation for GBPUSD. The estimates are based on a 22 day rolling window. In figure 4 below we see the empirical distribution of these data together with the theoretical sample correlation distribution based on an average calculated from *raw* r-values. Transforming the variable into *Z*-values yields the distribution in figure 5.While this distribution has a high peak and fat tails, it is *approximately* normal. This result is general across our data set.

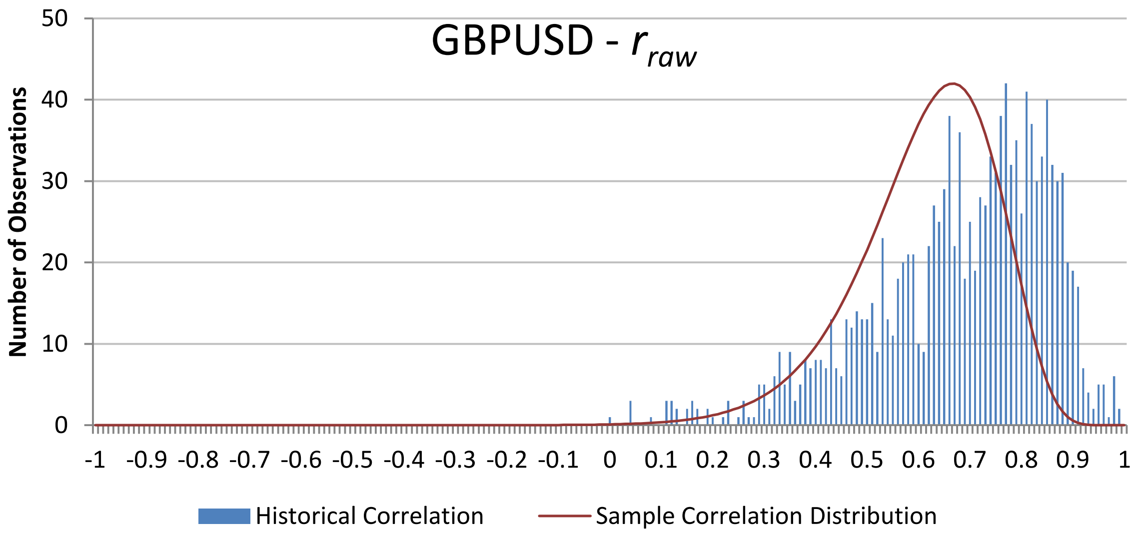

Figure 4: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPJPY. The red line represents the theoretical sample correlation distribution, and it is based on the mean of raw *r*-values.

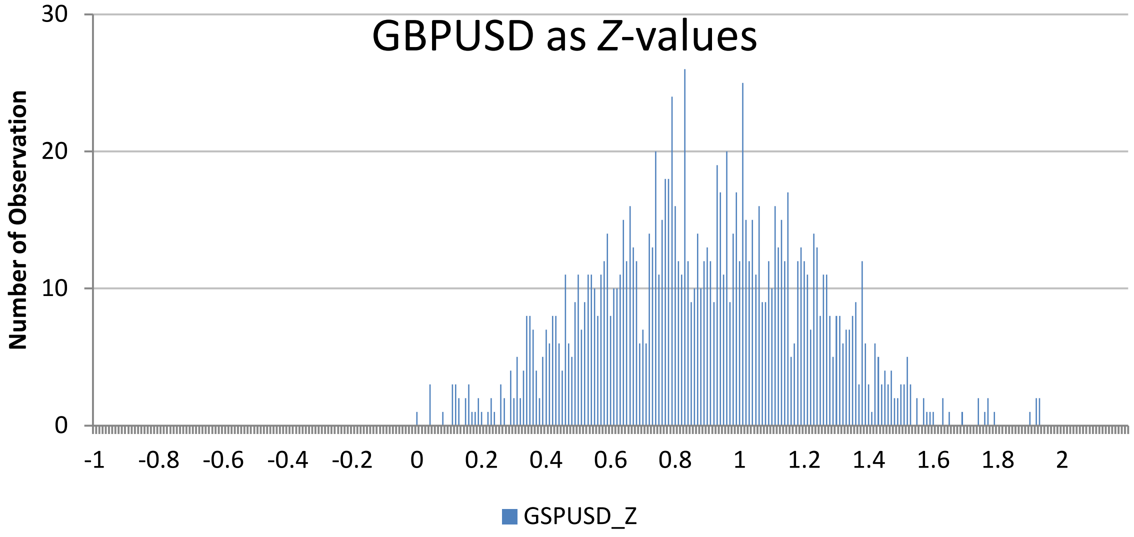

Figure 5: The empirical distribution of historical correlation, transformed from raw correlation coefficients into *Z*-values. We see that this distribution is approximately normal.

Below we see how the theoretical sample distribution changes when we use the mean obtained by transformation. The mean calculated from raw r-values was 0.675, while a transformation yields $\bar{r}$ = 0.714. This is a change of 0.039 (in absolute value), or - 5.78 %. While a difference of – 5.78 % is perhaps not dramatic, we can see that it visibly changes the shape of the theoretical sample distribution. The new distribution seems to be a better fit to the empirical distribution, a finding that holds for several currency pairs.

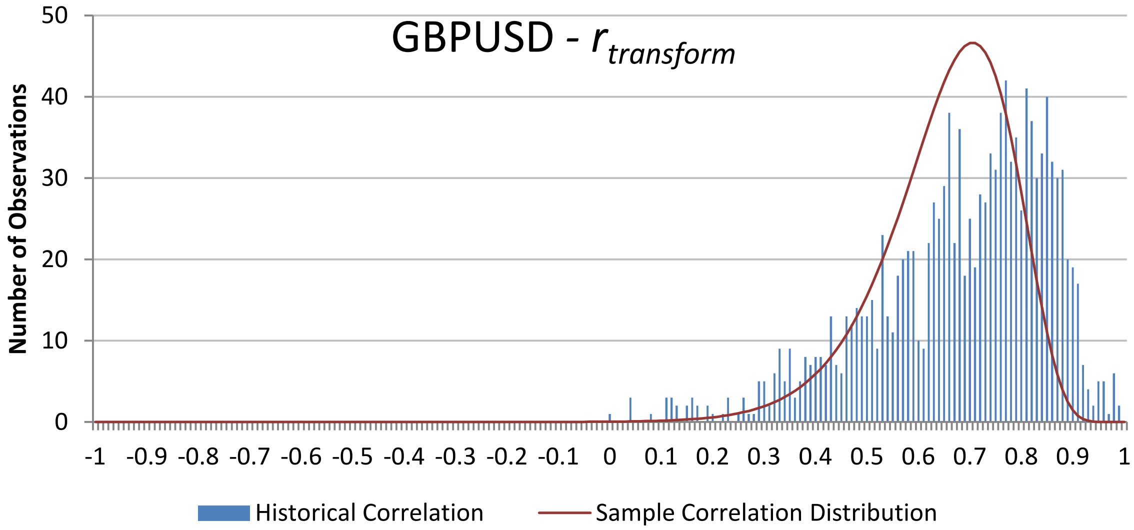

Figure 6: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPJPY. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on $\bar{r}_{\text{transform}}$.

The largest difference (in absolute value) between $\bar{r}$ calculated from raw correlation coefficients and the value calculated based on the transformation were 0.048. This 7.16 % difference was found for the historical correlation of EURUSD, based on a 22 day rolling window. Note that for both EURUSD and GBPUSD we have high $\bar{r}$. This seems to align well with the theory outlined above, i.e. the bias induced by calculating an arithmetic average from raw correlation coefficients is largest for $\bar{r}$ near ± 1. For *r* near 0, the bias should be more or less zero. This holds true for e.g. USDNOK. If we look at the historical USDNOK correlation estimates based on a 22 day rolling window we get $\bar{r}_{raw} = -0.180$ and $\bar{r}_{transform} = -0.191$.

These findings appear to be general across our data set. We have reported both $\bar{r}_{raw}$ and $\bar{r}_{transform}$ in Appendix A, but for the remainder of this paper we will only discuss $\bar{r}$-values which are calculated from transformed variables. That our empirical findings fit well with the results obtained in other sciences like psychology and genetics also support the use of the Fisher transformation.

While the interpretation of the implied correlation coefficient as a cosine value is less obvious, we believe it is appropriate to transform these variables before attempting to find an average correlation value as well (see e.g. Fox 1997:241-261 for a thorough description of the vector geometry of correlation coefficients). However, the principle of *caveat lector* applies.

We will discuss our empirical results in chapter 4. For now we merely note that historical correlation coefficients appear to be the more sensitive with respect to how we calculate $\bar{r}$, relative to implied correlation. Further we see that the difference between the two mean values decreases when the distribution narrows, and that the liquidity of the exchange rate pairs seems to have some impact on the size of the mean correlation bias.

2.7.2 Implications for hypothesis testing

While the use of a biased $\bar{r}$ might seem trivial when all you want is some rough estimate of the average value of the sample correlation, the unique properties of the correlation coefficient have important implications for hypothesis testing and the interpretation of standard deviations.

The sample standard deviation of a discrete random variable is generally defined as:

$$S_n = \sqrt{\frac{1}{n}\sum_{i=1}^{N}(x_i - \bar{x})^2}$$

However, we have already mentioned that subtracting $\bar{r}$ from individual *r*-values does not yield a cosine value. More importantly, $\bar{x}$ is an unbiased estimator of a *monovariate* distribution, while $\bar{r}$ is a biased estimator of a *skewed bivariate* distribution (see e.g. Garcia 2010, Stuart and Ord 1994).

As we saw in figures 2 and 3, the variance of a correlation coefficient is dependent on both sample size and the population correlation level. In other words it is not normally distributed, nor does it have a constant variance. To our knowledge Fisher (1915:517) was the first to discuss how these properties affected the interpretation of the mean and standard deviation of the sample correlation:

> *"The use of the correlation coefficient r as independent variable of these frequency curves is in some respects highly unsatisfactory. For high values of r the curve becomes extremely distorted and cramped, and although this very cramping forces the mean $\bar{r}$ to approach ρ, the difference compared with $1-\rho$ becomes inordinately great. Even for high values of n, the distortion in this region becomes extreme, and since at the same time the curve rapidly changes its shape, the values of the mean and standard deviation cease to have any very useful meaning. It would appear essential in order to draw just conclusions from an observed high value of the correlation coefficient, say 0.99, that the frequency curves should be reasonably constant in form."*

In other words, because *r* (and cos) is bounded by ±1 it is symmetrically distributed around 0 only for $\rho = 0$, and this will be the only place where the mean of the sampling distribution equals the population mean value (Fisher 1915:529). We also have that as ρ approaches ±1, the sampling distribution becomes skewed and the sampling variance approaches 0. It follows that standard deviations have no natural interpretation in this context, as the standard deviation is a statistical measure which aims to express to which degree individual observations within the sample differ from the sample mean.

More importantly, these features of the sample variance have implications for hypothesis testing and the calculation of standard errors. The standard error is in many ways a measure of the trustworthiness of an estimate, and tests of statistical significance hinges on their accuracy. While there are tests based on *r* that are UMPU (uniformly most powerful unbiased) given $\rho = 0$, this property no longer holds for non-zero values of *r* (Stuart et al.

1999:484). For this reason we will do all significance testing based on variables transformed by Fisher's Z transformation.

Let r be the sample correlation coefficient, and define $Z = \tanh^{-1} r$ and $\zeta = \tanh^{-1} \rho$. Assuming that Z is approximately normally distributed, this gives approximate mean

$$E(Z) = \frac{1}{2} ln \left[\frac{1+\rho}{1-\rho}\right]$$

And variance

$$var\ (Z) = \frac{1}{n-1} + \frac{2-\rho^2}{2(n-1)^2}$$

These approximations appear to be reasonably close for n as low as 11 (Stuart and Ord 1994:568). Given the assumption of bivariate normality we can calculate a standard error by taking the square root of the variance.

To test $H_0: \rho = \rho_0$ one would use the Fisher transformation and assume that under the null, Z is approximately normally distributed. It is also possible to test the composite hypothesis that the correlation coefficients of two independently sampled bivariate normal populations are identical (Stuart et al 1999:484). If the correlation parameters of these distributions are the equal, then the distributions of the two transformed statistics, Z_1 and Z_2, will also be approximately normally distributed. Further, their difference will have $mean = 0$ and variance $\left[\frac{1}{n_1-1} + \frac{2-\rho_1^2}{2(n_1-1)^2}\right] + \left[\frac{1}{n_2-1} + \frac{2-\rho_2^2}{2(n_2-1)^2}\right]$ where n is sample size and the subscripts denotes the different samples. We will use this result when we evaluate the forecasting accuracy of our correlation estimates.

3 Data description

3.1 Data set

This study makes use of implied volatility quotes for seven currency pairs, namely USDJPY, GBPUSD, GBPJPY, EURUSD, EURJPY, USDNOK and EURNOK. The data consist of close prices from the OTC FX market, collected at midmarket. The sample period covers October 2006 through November 2011 (1326 observations).

We have used daily option quotes from ATM delta neutral straddles to calculate implied correlations. These are expressed as Garman and Kohlhagen (1983) implied volatilities with fixed time to maturity, as well as fixed moneyness in terms of the Garman–Kohlhagen delta.

We have analyzed ATM volatilities for option contracts with one month, two months, three months, six months, nine months and one year to expiration. These maturities are the most actively traded tenors according to experienced option traders (see e.g. Haug et al. 2010).

To obtain a consistent and empirical estimate for the historical correlation level we have analyzed the spot exchange rates for the same currency pairs. The exchange rate data span the same time period, namely October 2006 through November 2011 (1326 observations).

All price quotes – both in the form of implied volatilities and spot exchange rates - are extracted from Bloomberg, which we consider a trustworthy source. We have used a similar data sample from Reuters to validate the accuracy of the dataset. While the data we obtained from Bloomberg were of high quality, our dataset spans five years and few occurrences of missing data are inevitable. We have used linear interpolation to fill in missing values, albeit to a limited degree. We believe this has not affected our forthcoming analysis.

As we have mentioned before, a special feature of the FX interbank market is particularly appealing when studying implied correlation: in this arena a new set of options are quoted every day. This makes these price quotes ideal for empirical studies, because an options time to maturity can be held constant across the analysis. On an exchange, new maturities are introduced once a month or less frequently, which means that an option with one month to maturity yesterday will have one month minus one day to expiration today. Because the OTC FX market introduces new contracts every day we avoid *"mixing apples and oranges"*, which is the econometrician's way of stating that you mix two totally different things. An options time to maturity has significant impact on the behavior of implied volatilities, a point we will discuss more thoroughly in chapter 4.

Another problem we would encounter in relation to exchange traded options is changes in moneyness. A move in the price of the underlying asset can make an option switch from being at-the-money, in-the-money or out-of-the-money several times each day. This last problem is also solved by using data on OTC traded options, because the interbank market quotes the strike price in terms of delta. As volatilities are assigned to deltas rather than strikes, moneyness can be held constant over time.

If one would like to compare the implied volatilities of option contracts that are quoted in the traditional way, i.e. with prices given in terms of strike, one would have to make sure that the price of the underlying asset *and* the price of the option are quoted on the *exact* same point in time. You would then have to calculate the implied volatilities using an iterative procedure, which makes this method extremely sensitive to sampling error. Not only would we face the risk of dealing with asynchronous price data, the likelihood of making mistakes increase with the number of complicated calculations to execute.

Because implied correlation is calculated from implied volatilities, it is crucial for our study that the implied volatilities are accurate and comparable. In sum, our data set will not contain the statistical sampling error we would experience if we had we tried to perform the same analysis on other option contracts.

3.2 Historical correlation

We will compare the implied correlation level with historical correlation estimates calculated as rolling windows. We will estimate historical correlations for namely 22, 43, 65, 128, 193 and 256 days. This will correspond roughly with the number of trading days of the different option contracts. As an example, the effective numbers of trading days in options with one month to maturity vary between 20 and 23 days, with a mean of 22. For the two and three month horizons the effective number of trading days ranges from 41 to 45 and 64 to 66, respectively.

3.2.1 Testing for non-stationarity

In this section we will briefly present the results from our tests for non-stationarity. We decided to start out performing ADF-tests, to capture potential problems caused by serial correlation.

We began testing with a generous number of lags, and reduced them one by one. The t-statistics on higher order autoregressive terms were invariably small and insignificant. Based on these findings we concluded that it was unnecessary to augment the test for these particular time series, and we report the statistics from ordinary Dickey-Fuller tests. The results are presented in table 1 below. However, the various ADF-tests consistently lead to a rejection of the null hypothesis of non-stationary variables. This suggests that our findings are robust.

We have decided to keep the constant term, because of the bias that would be induced if $\alpha \neq 0$. The values of this parameter are not important in this context, and neither is their level of significance.

Table 1: Dickey-Fuller test statistics **Timer period: January 2006 – November 20 11, Daily observations**		
	Constant term (α)	Coefficient (θ)
lnEURUSD	0,00 (0,42)	- 0,95** (- 37,05)
lnGBPUSD	0,00 (- 0,39)	- 0,96** (- 37,33)
lnUSDJPY	0,00 (- 1,52)	- 1,04** (- 40,53)
lnNOKUSD	0,00 (- 0,37)	- 1,03** (- 40,40)
lnGBPJPY	0,00 (- 1,29)	- 0,96** (- 37,75)
lnEURJPY	0,00 (- 0,84)	- 0,96** (- 37,03)
lnEURNOK	0,00 (- 0,11)	- 0,96** (- 37,55)

* * indicates statistical significance at the 0.01 level

H_0: α = 0, θ = 0, t-statistics in brackets, Dickey-Fuller critical value: 1 % = 3.4

From the table above we see that all θ-values are close to minus unity, which suggests a ρ close to zero. The t-statistics for all the log-returns are higher (in absolute terms) than the critical values of 3.4. We can therefore reject the null of non-stationarity for all series. High t-values suggest that these findings are very robust.

4 Empirical analysis

4.1 Implied correlation versus historical correlation

In this section we will compare the distribution of implied correlation with the distribution of historical correlation. The latter is estimated over time horizons that correspond roughly to the number of trading days in the different option contracts.

We will display the distribution of historical correlation together with the frequency distribution of the sample correlation coefficient. The sample correlation distribution is based on the correlation mean of the historical data, or more precisely the mean value obtained by using Fisher's *r* to *Z* transformation. To make the distributions comparable we have scaled the theoretical distribution from continuous to discrete time.[5]

[5] All distributions have been rescaled in the following manner: $\frac{\sum_{i=0}^{n} discrete\ sample\ points}{\sum_{i=0}^{n} continous\ time\ sample\ points} \times 2$

First we will give a detailed analysis of the USDJPY correlation and discuss not only implied correlation compared to historical correlation, but also the term structure of these correlation measures.

4.1.1 Implied and historical correlation – USDJPY

The implied correlation between GBPJPY and GBPUSD, labeled USDJPY for short, is calculated based on implied volatility quotes for GBPJPY, GBPUSD and USDJPY. Historical correlation is estimated based on the log-returns of GBPJPY and GBPUSD. We will begin by looking into the correlation level of contracts with one and three months to expiration.

Below we have plotted the level of both implied and historical correlation against time. The figures span January 2007 through October 2011, and all calculations are based on daily observations. The first graph displays implied correlation for option contracts with one month to maturity against historical correlation estimates based on 22 days, i.e. the average number of trading days during one month. Figure 8 shows implied correlation for option contracts with three months to expiration against historical correlation estimates based on a time period of 65 days.

Visual inspection of these graphs indicates that historical correlation varies substantially relative to implied correlation. We note that both correlation measures seem to vary less erratically when the time horizon increase. This could be caused by sampling error in the estimates of historical correlation. However, as implied correlation is based on a uniform number of sampling points, this explanation does not apply for this correlation measure. We will discuss this phenomenon further below.

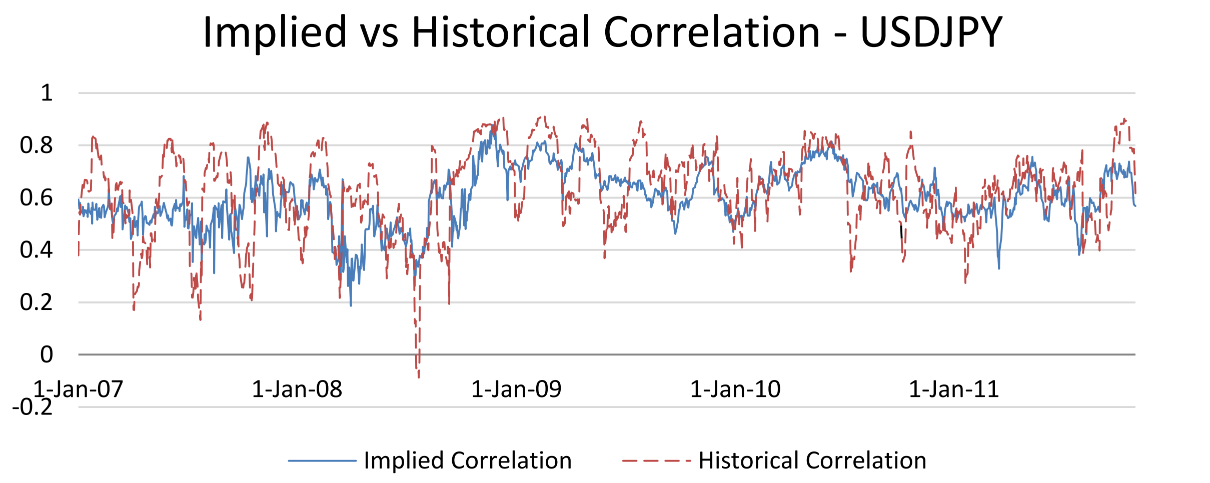

Figure 7: This graph displays implied correlation for option contracts with one month to expiration against a rolling historical correlation estimate based on a time period of 22 days.

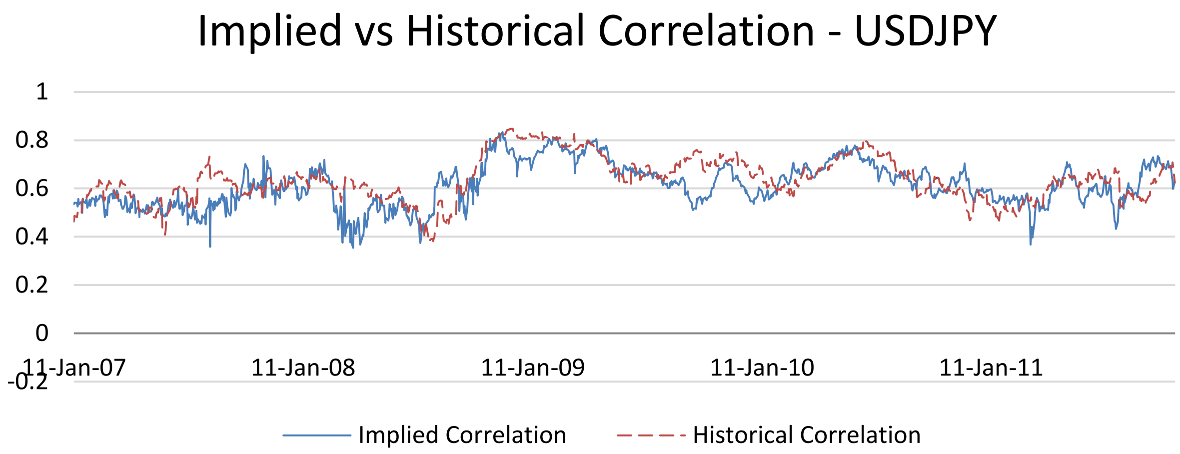

Figure 8: Implied correlation for option contracts with three months to expiration plotted against a rolling historical correlation estimate based on a time period of 65 days.

Table 2: Descriptive statistics USDJPY – Implied vs Historical Correlation

	1 month		*2 months*		*3 months*	
	Implied Correlation	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*
Mean	0.607	0.655	0.611	0.648	0.614	0.647
Mode	0.57	0.69	0.57	0.64	0.56	0.65
Minimum	0.186	- 0.125	0.231	0.245	0.255	0.378
Maximum	0.879	0.917	0.842	0.888	0.834	0.848
Observations	1326	1296	1326	1275	1326	1253

Table 2 presents a comparison of implied and historical correlation in the form of descriptive statistics. We see that the sample means are similar across both correlation measures - it is also similar across the different time horizons. The mode hardly changes as we move from one to three months.

The minimum and maximum values indicate a wide distribution for both implied and historical correlation. The spread between the minimum and maximum correlation estimates decreases slightly as the time horizon increases.

Our findings so far suggest that the historical correlation level vary more than that of implied correlation, especially for short time horizons. This could mean that the market participants expect some mean reversion in the correlation level and incorporate this view in their price quotes. It is also likely that sampling error is causing more erratic behavior of historical correlation estimates based on short time periods. Sampling error is commonly understood to be the error caused by observing a sample, rather than the whole population (Tsay 2010).

Note that there are two sources of uncertainty associated with our historical estimates. Correlation is not a visible statistical measure; like all higher moments of time series data it is unobservable (Walter and Lopez 2000:8). It is possible that the way we have chosen to calculate the sample correlation does not properly incorporate the variance dynamics of the data. Pearson's *r* is a coefficient of linear dependence, and does not capture more complex forms of interdependence.

Second, rolling correlations based on small time windows are extremely sensitive to outliers entering and exiting the estimation window. We believe that the latter phenomenon is causing a large part of the variability in our historical estimates based on a small number of observations. The fact that the level of implied and historical correlation seems to align with increased time horizon supports this notion.

Below we have plotted historical correlation estimates based on 22 and 65 days, with a 95 % confidence interval around the estimates. These graphs only display data for 2007, in order to clearly illustrate how increasing *n* affects the intervals. While the reliability of the sample correlation increases substantially with larger *n*, changing *n* does not by itself explain the lower *turbulence* in the historical correlation estimates based on 65 days. We will not discuss confidence intervals in further detail; we merely note that outlying observations entering and exiting the rolling estimation window seem to cause a lot of the turbulence we see in the historical estimates. This result is general across all currency pairs.

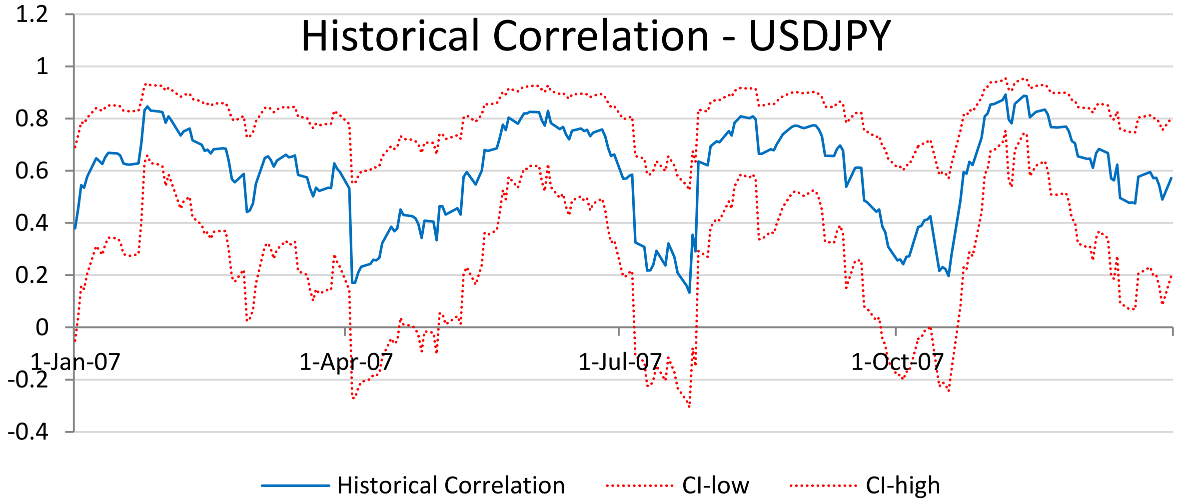

Figure 9: Rolling historical correlation estimates based on a time period of 22 days. The red lines represent a 95 % confidence interval.

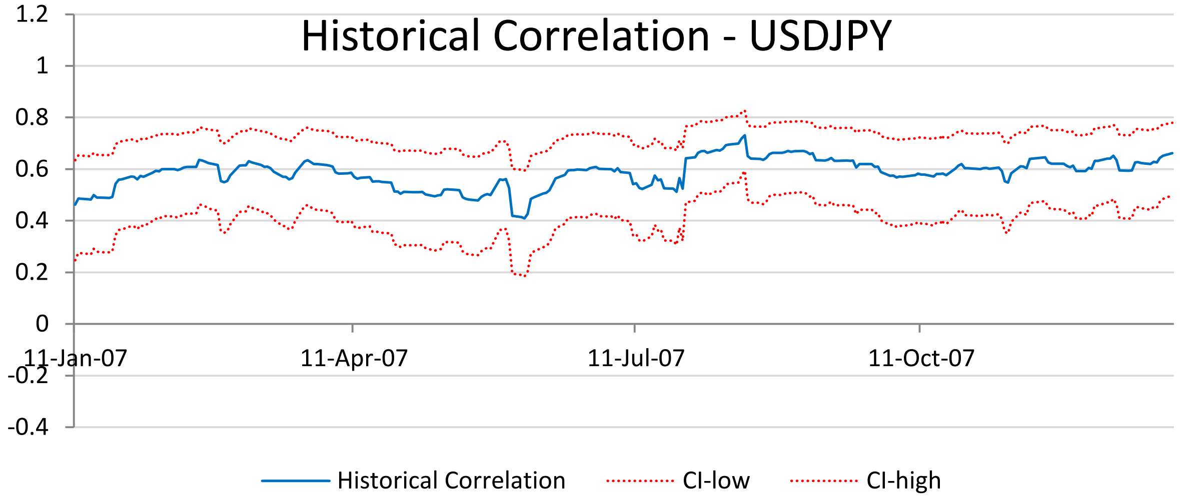

Figure 10: Rolling historical correlation estimates based on a time period of 65 days. The red lines represent a 95 % confidence interval.

The erratic behavior of both implied and historical correlation indicates that both measures have a wide distribution. This notion is confirmed by figures 11 through 14 below. Note that while the units on the x-axis are fixed, the y-axis changes across correlation measures and time horizons. The theoretical distribution of the sample correlation coefficient is based on the mean of the historical correlation estimates.

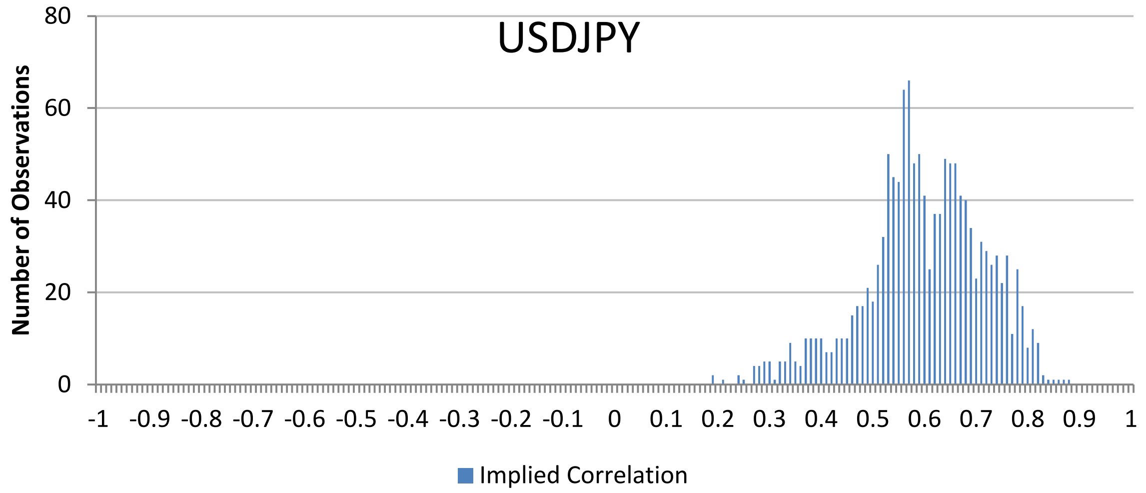

Figure 11: The empirical distribution of implied correlation based on implied volatility quotes for GBPJPY, GBPUSD and USDJPY option contracts with one month to expiration.

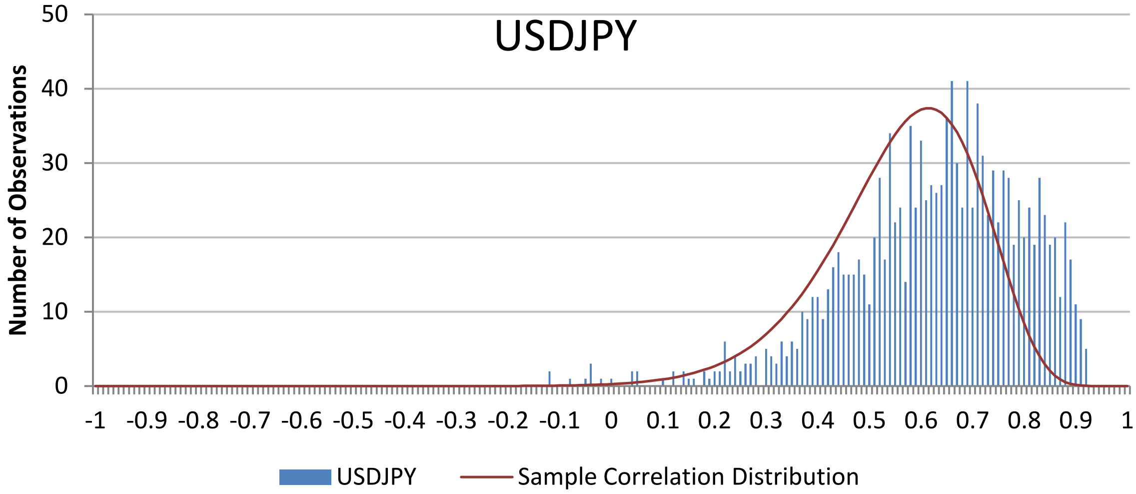

Figure 12: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for GBPJPY and GBPUSD. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

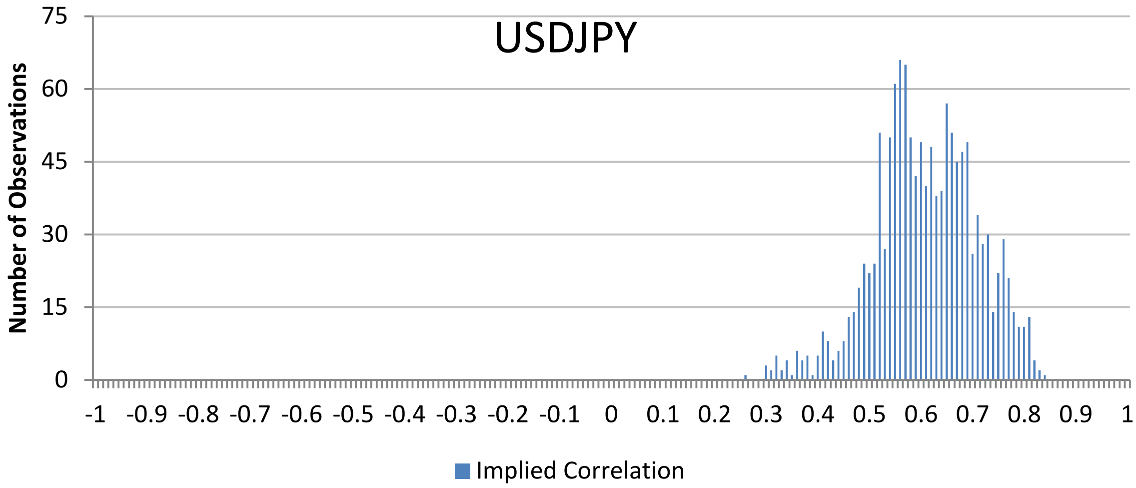

Figure 13: The empirical distribution of implied correlation based on implied volatility quotes for GBPJPY, GBPUSD and USDJPY option contracts with three months to expiration.

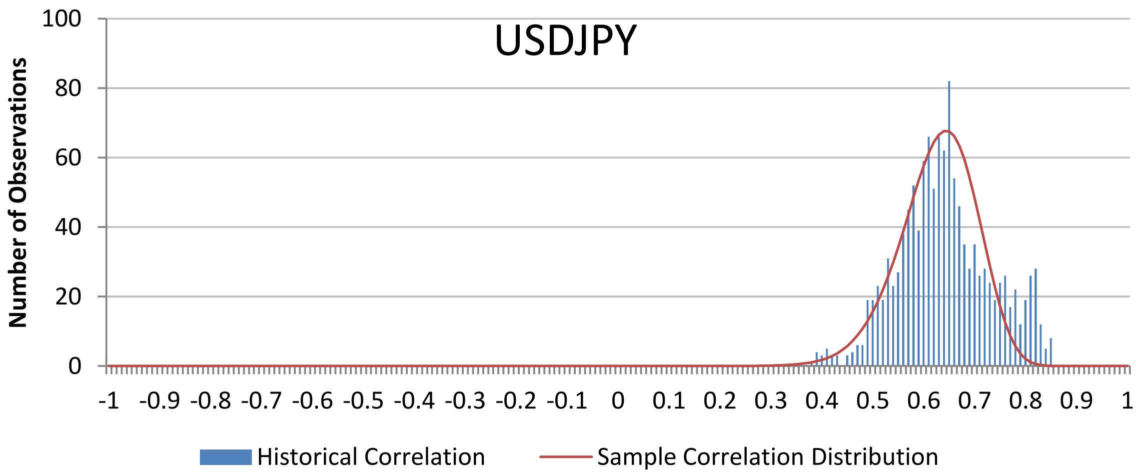

Figure 14: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for GBPJPY and GBPUSD. All estimates are based on a 65 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

The distribution of both implied and historical correlation narrows as time to maturity/the size of the estimation window increases. Figure 11 through 14 also confirm that historical correlation is the more variable of our two correlation measures. Correlation calculated from log-returns has a wider distribution relative to implied, especially when we look at a time horizon of one month. We also see that the theoretical distribution of the sample correlation is a reasonably good fit to the empirical estimates.

We will now discuss correlation for longer time horizons. Figure 15 and 16 below describe the level of both implied and historical correlation from 2007 through October 2011. Figure 15 displays implied correlation for option contracts with six months to maturity versus rolling historical correlation based on an estimation window of 128 days. Figure 16 illustrate implied and historical correlations for option contracts with one year to expiration, against correlation estimates based on a rolling window of 256 days.

Both correlation measures seem to vary less erratically as we increase the time horizon under consideration. While they do not move in complete tandem, there is a similar pattern to their behavior. We see that increasing the number of observations used to calculate the rolling historical estimate has a smoothing effect on the historical correlation level.

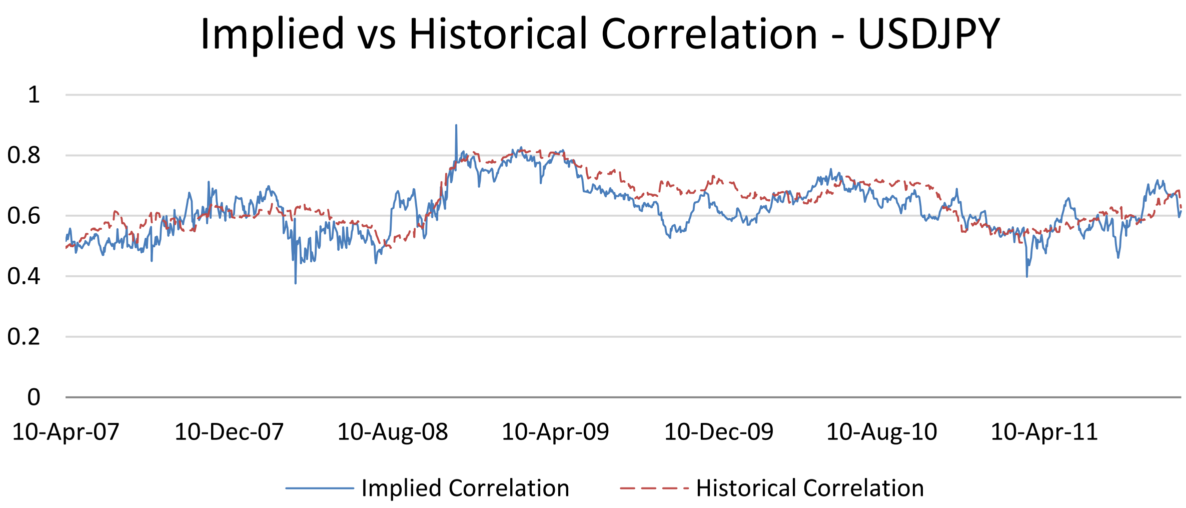

Figure 15: Implied correlation for option contracts with six months to expiration plotted against a rolling historical correlation estimate based on a time period of 128 days.

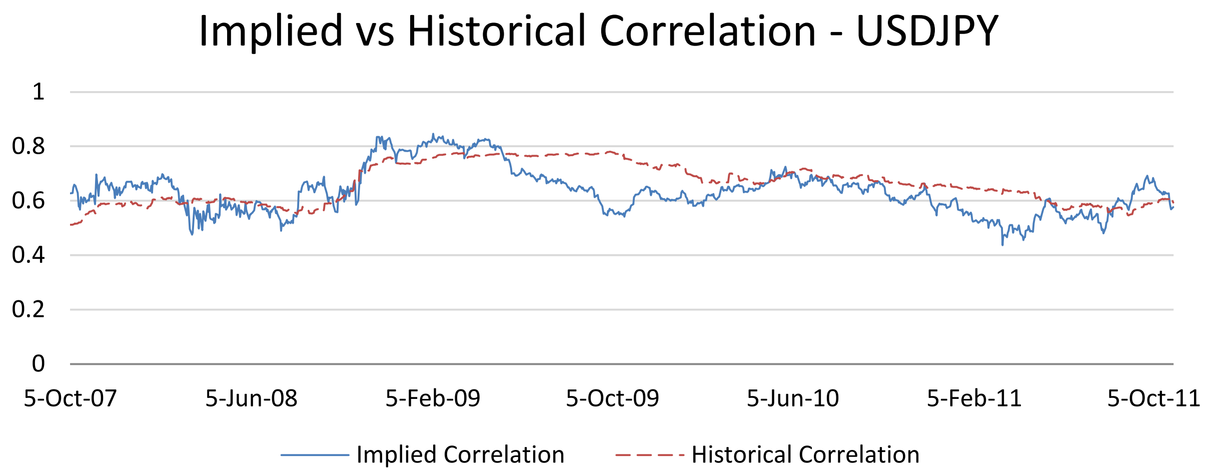

Figure 16: Implied correlation for option contracts with twelve months to expiration plotted against a rolling historical correlation estimate based on a time period of 256 days.

Table 3: Descriptive statistics USDJPY – Implied vs Historical Correlation

	6 months		*9 months*		*12 months*	
	Implied Correlation	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*
Mean	0.618	0.652	0.619	0.660	0.621	0.669
Mode	0.66	0.71	0.65/0.66	0.58	0.66	0.78
Minimum	0.317	0.493	0.337	0.503	0.350	0.512
Maximum	0.900	0.817	0.833	0.794	0.846	0.780
Observations	1326	1190	1326	1124	1326	1062

In table 3 we see that the mean values stay similar across the different time horizons. The minimum and maximum values indicate that implied correlation now has a wider distribution relative to the historical estimates. This impression is confirmed in figures 17 through 20 below.

The dispersion of the correlation measures is visibly reduced when we compare these figures to those portraying shorter time horizons (figures 11-14); the tails are markedly altered. This phenomenon is particularly pronounced when we look at the distribution of historical correlation in figure 18, in this case the tails has more or less disappeared. Finally we see that the empirical distribution of historical correlation does no longer resemble the theoretical sample distribution in terms of shape, although they are still located in the same general area.

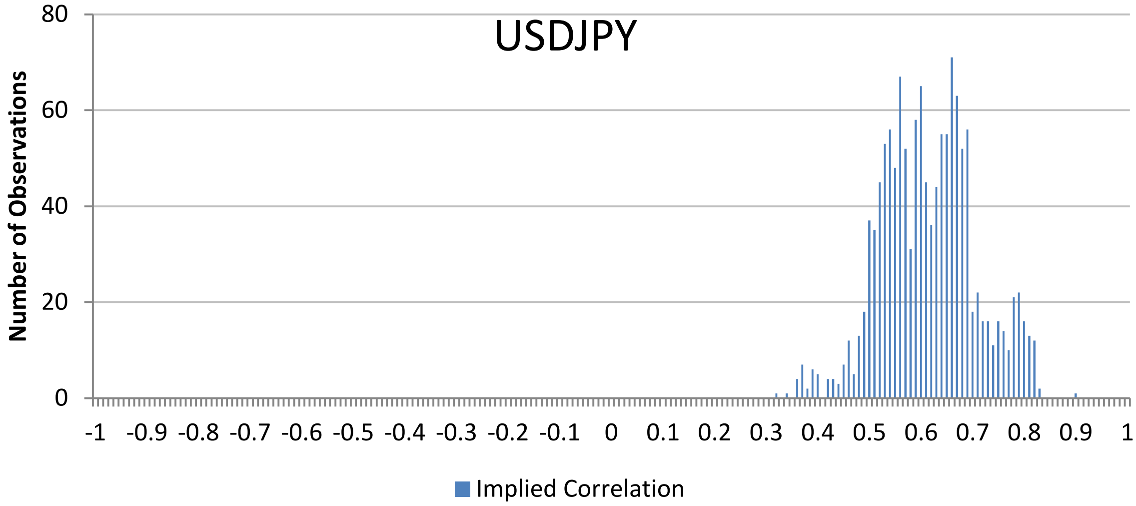

Figure 17: The empirical distribution of implied correlation based on implied volatility quotes for GBPJPY, GBPUSD and USDJPY option contracts with six months to expiration.

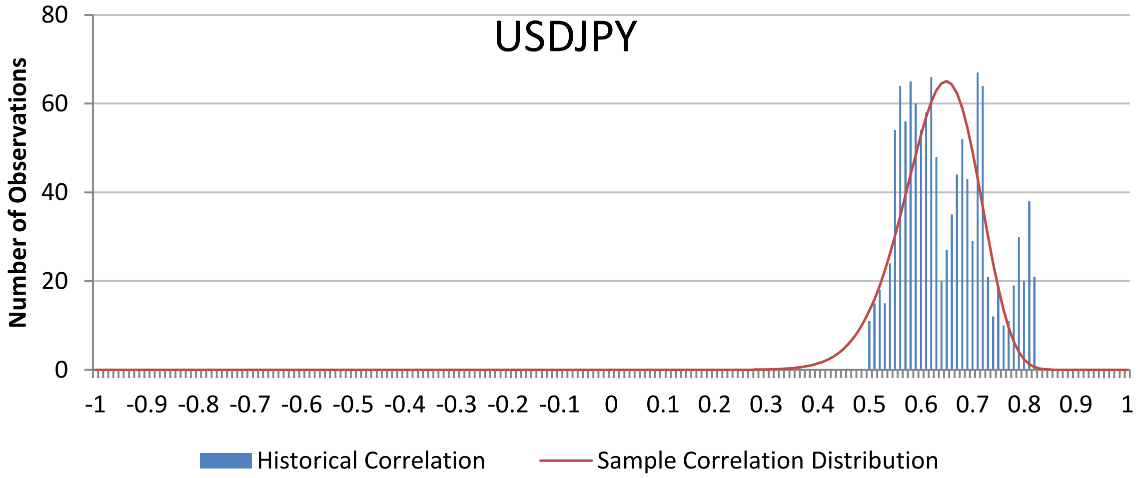

Figure 18: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for GBPJPY and GBPUSD. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

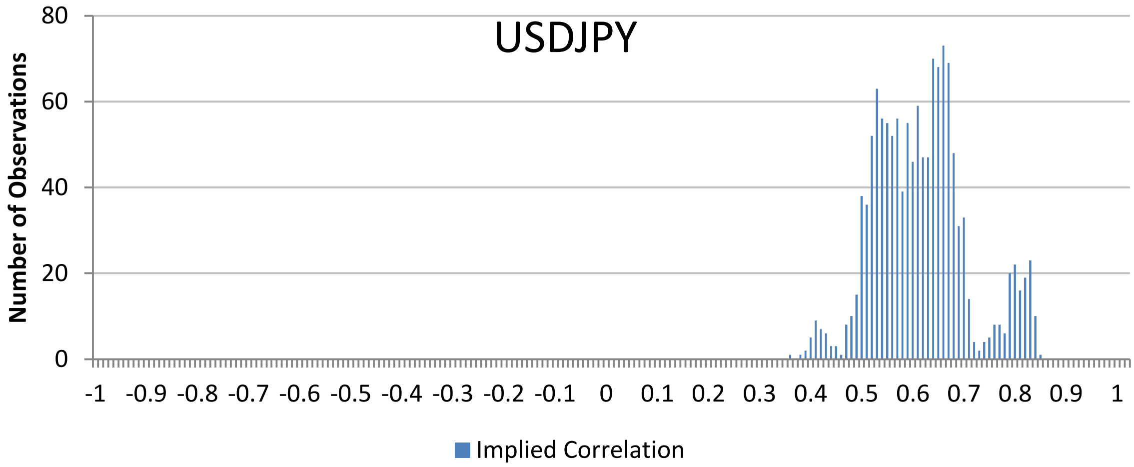

Figure 19: The empirical distribution of implied correlation based on implied volatility quotes for GBPJPY, GBPUSD and USDJPY option contracts with one year to expiration.

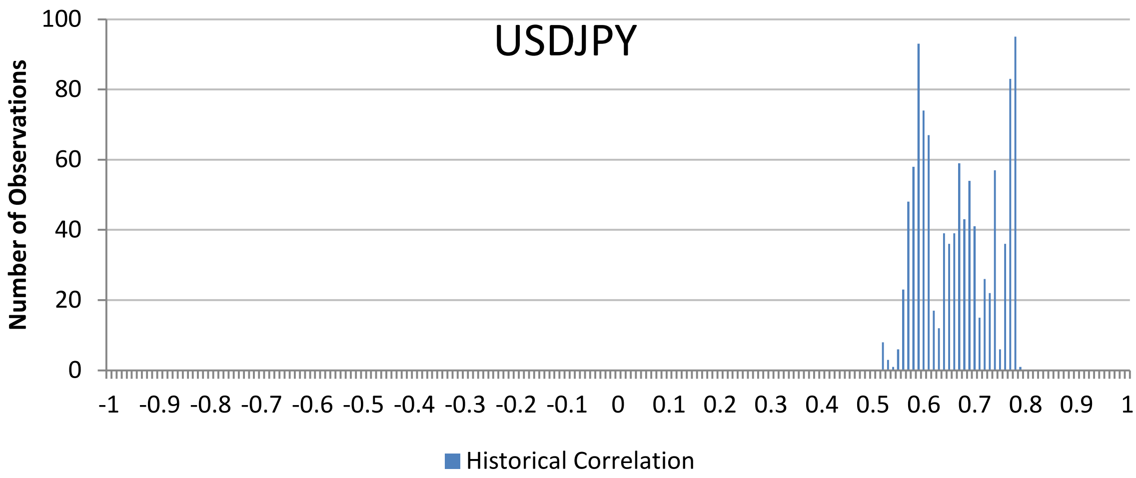

Figure 20: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for GBPJPY and GBPUSD. All estimates are based on a 256 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Our results so far suggest that the implied correlation level vary less erratically with increased time to maturity. As implied correlation is *derived* from option prices rather than *estimated*, this finding cannot be explained by sampling error – there is no uncertainty associated with implied volatility traded on the market. It is however consistent with the idea that unexpected changes in the correlation level will have a smaller impact on correlation for option contracts with one year to maturity relative to the correlation for monthly contracts.

Remember that implied correlations are calculated based on observations of implied volatilities. If some form of market shock brings on a sudden change in the implied volatility level, this will affect options with one year to maturity substantially less than an option which expires in one month. In the former case the shock will be spread out over a longer time period, and thus contribute less to the overall volatility calculation. Although we have not found a consistent way of calculating standard deviations or similar measures of dispersion for correlation coefficients, it is reasonable to expect that sudden changes in implied volatility levels will affect implied correlation.

It is also reasonable to expect strong mean reversion in implied volatility (see e.g. Wilmott 2006). After a shock that cause turbulence in the market, it is unlikely that there will be more major events for days in succession. On the contrary, it is to be expected that the market will be calm for some time while its participants digest the news. Haug et al. (2010:194) suggests that this behavior is embedded in implied volatilities, and it is plausible to expect a "spillover effect” that influences implied correlation.

4.1.2 Implied and historical correlation – USDJPY2

In the previous section we calculated the implied correlation between US dollars and Japanese yen based on implied volatility quotes for EURJPY, GBPUSD and USDJPY. Our dataset allows us to calculate a USDJPY implied correlation through an alternative route as well. In this section we will calculate what we have labeled USDJPY2, which is implied correlation backed out from implied volatility quotes for EURUSD, EURJPY and USDJPY. Conversely, historical correlation will be based on the log-returns of EURUSD and EURJPY.

Below we have plotted the level of implied and historical correlation against time. Like before, the figures span January 2007 through October 2011, and all calculations are based on daily observations. The first graph displays implied correlation for option contracts with one month to maturity versus historical correlation estimates based on 22 days. The second graph shows implied correlation for option contracts with six months to expiration, against a rolling correlation estimate based on a time period of 128 days.

These charts appear quite similar to the corresponding ones for USDJPY. However, the graphs below seem to vary even more erratically than the USDJPY correlation calculations based on the linkage through GBP.

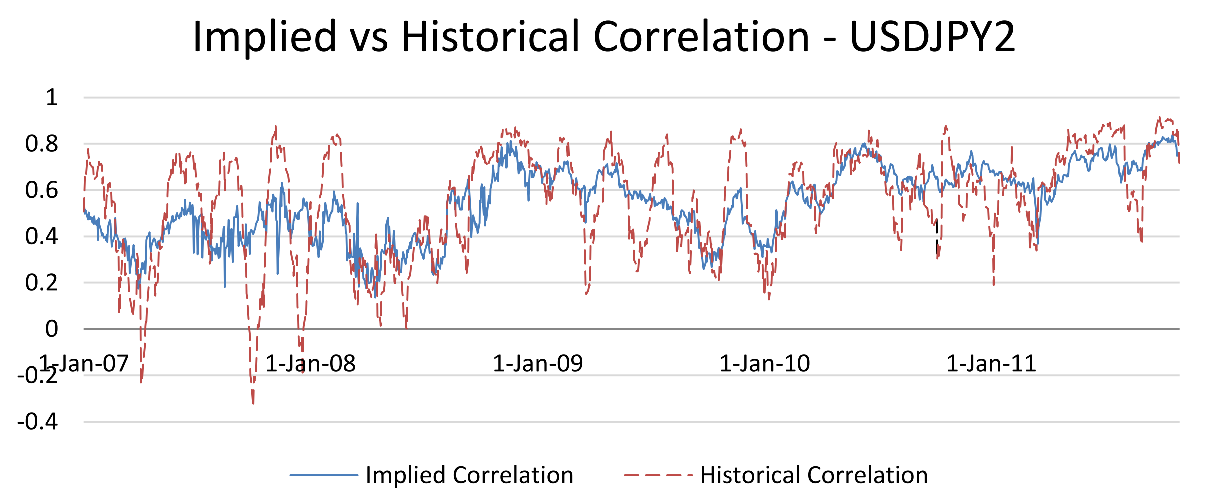

Figure 21: Implied correlation for option contracts with one month to expiration plotted against a rolling historical correlation estimate based on a time period of 22 days.

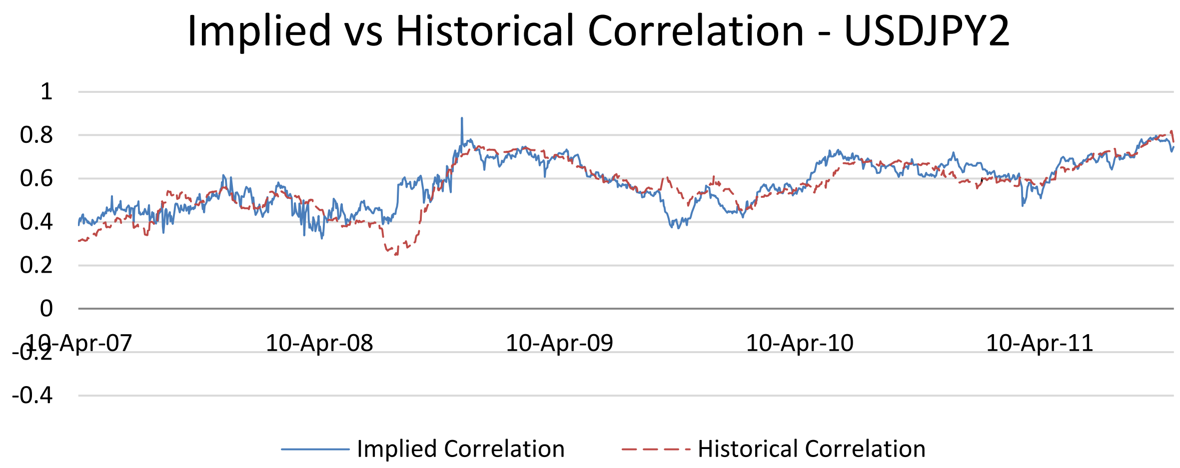

Figure 22: Implied correlation for option contracts with 6 months to expiration plotted against a rolling historical correlation estimate based on a time period of 128 days.

Table 4: Descriptive statistics USDJPY2 – Implied vs Historical Correlation

	1 month		*3 months*		*6 months*	
	Implied Correlation	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*
Mean	0.560	0.591	0.570	0.580	0.576	0.579
Mode	0.69	0.72	0.70	0.57	0.45/0.70	0.59
Minimum	0.135	-0.338	0.226	0.164	0.323	0.248
Maximum	0.844	0.924	0.820	0.852	0.881	0.819
Observations	1326	1296	1326	1253	1326	1190

Table 4 presents descriptive statistics for the USDJPY2 correlations. The sample means have dropped slightly when GBP is substituted with EUR as third leg in the currency triangle. The width of the distributions appears to be roughly the same as before. The modes are also similar, note that the six month correlation distribution appears bimodal.

Below are exhibits of the distributions of both implied and historical correlation. Figure 23 and 24 display implied correlation for option contracts with one month to maturity and historical correlation based on a 22 day window, respectively.

Both distributions are slightly wider than what we saw in the similar diagrams for USDJPY linked through GBP. The most distinct difference between these charts and the corresponding ones for USDJPY is that the distributions below are markedly less peaked, especially for implied correlation. In figure 25 we see that the empirical distribution of implied correlation is far from bell shaped. That the distribution of implied correlation takes on atypical shapes when we consider contracts with relatively longer time to expiration is an outcome we get for several currency pairs. In figure 25 this effect is very pronounced.

Finally, we note that the empirical distributions of historical correlation now have an increasing number of observations outside the tails, relative to the theoretical sample density function. This feature is noticeable in both figures 24 and 26 below.

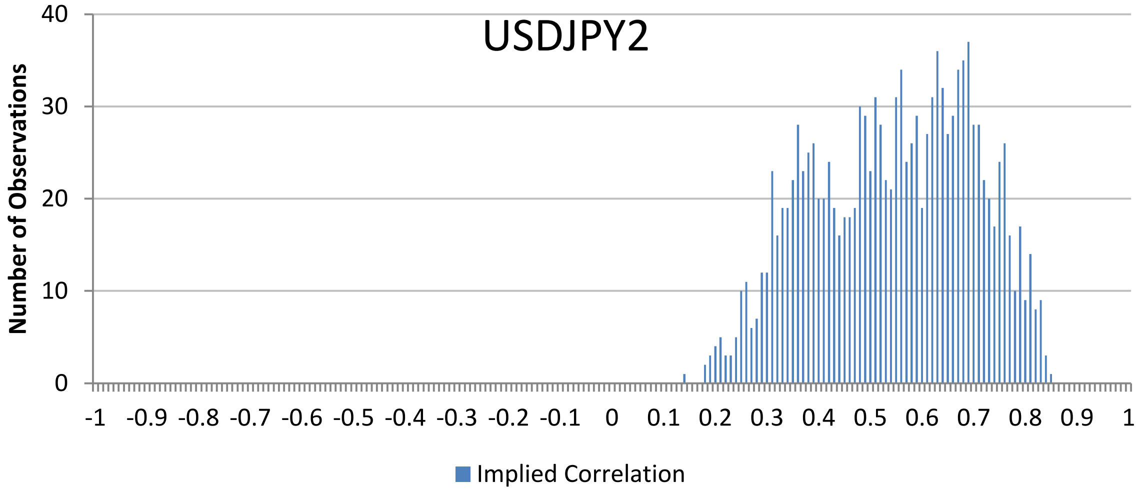

Figure 23: The empirical distribution of implied correlation based on implied volatility quotes for EURJPY, EURUSD and USDJPY option contracts with one month to expiration.

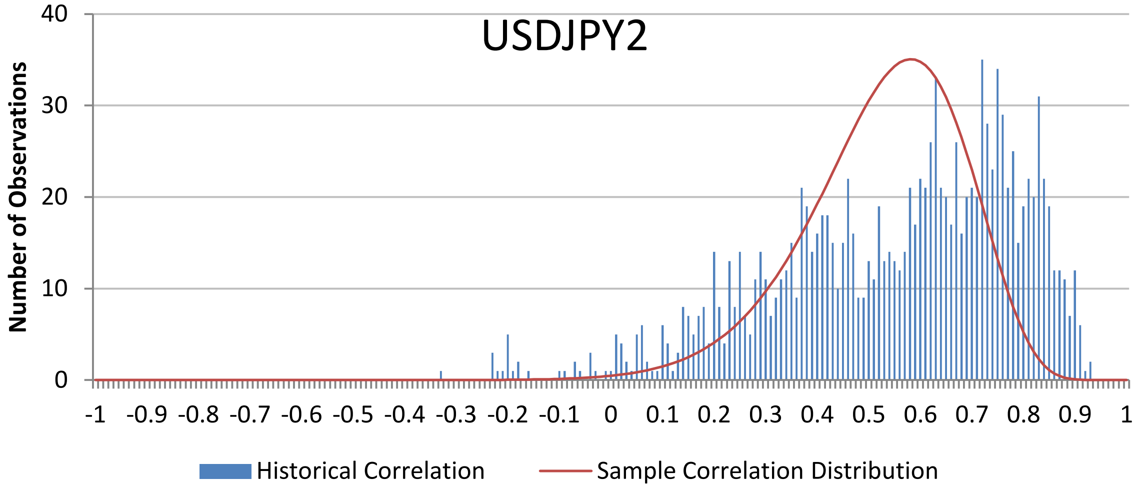

Figure 24: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURJPY and EURUSD. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, which is based on the mean of the actual data set.

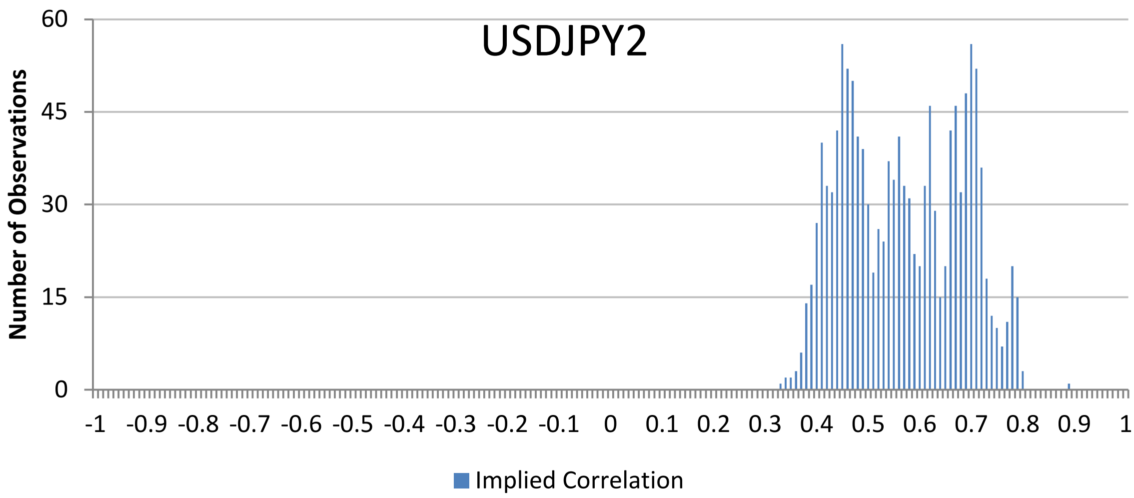

Figure 25: The empirical distribution of implied correlation based on implied volatility quotes for EURJPY, EURUSD and USDJPY option contracts with six months to expiration.

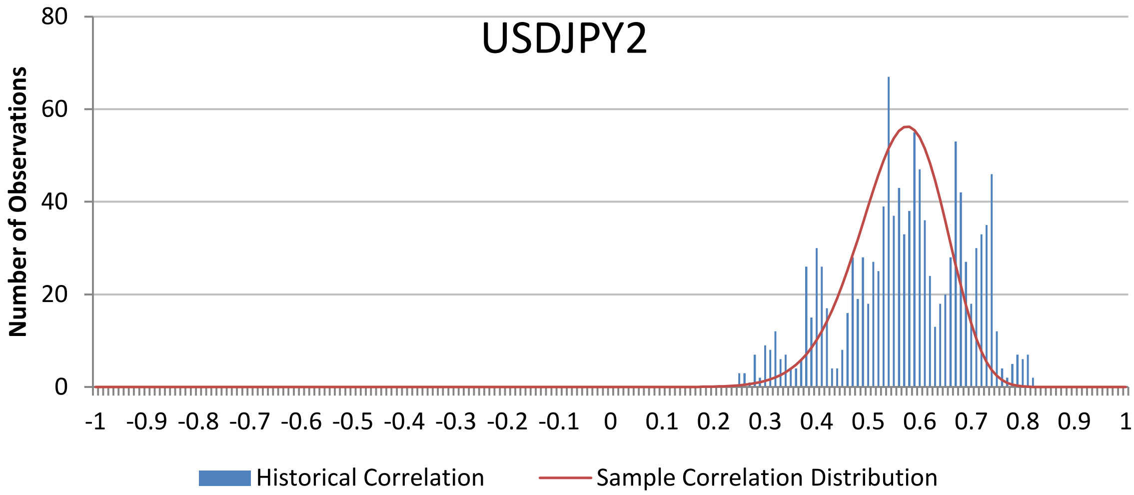

Figure 26: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURJPY and EURUSD. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

4.1.3 Implied and historical correlation – GBPUSD

In the following our calculations of implied correlation are based on implied volatility quotes for USDJPY, GBPJPY and GBPUSD. Historical correlation is estimated from the log-returns of USDJPY and GBPJPY.

Figure 27 and 28 below display the correlation level for implied and historical correlation against time. Figure 27 shows implied correlation for option contracts with one month to maturity together with historical correlation estimates based on 22 days, while figure 28 displays correlations for a time frame of six months.

Again we see that historical correlation vary substantially relative to implied correlation when we consider contracts with one month to expiration versus historical correlation based on a 22 day rolling window. Both correlation measures also vary less erratically as the time horizon increases. This suggests that what we found for USDJPY and USDJPY2 - that the variability of the two correlation measures seems to align when we increase the time frame under consideration - is something we can expect to see when we examine the other currency pairs as well.

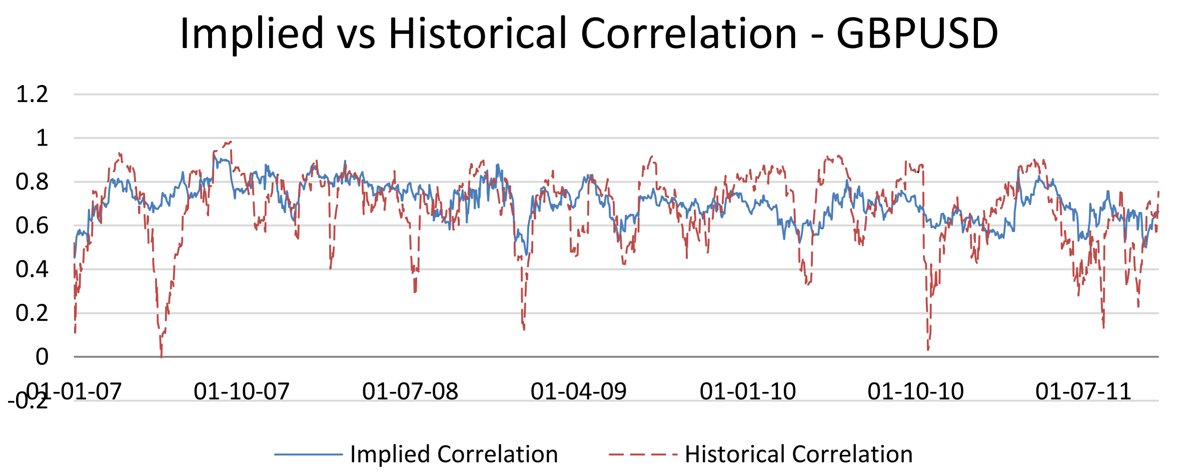

Figure 27: Implied correlation for option contracts with one month to expiration plotted against a rolling historical correlation estimate based on a time period of 22 days.

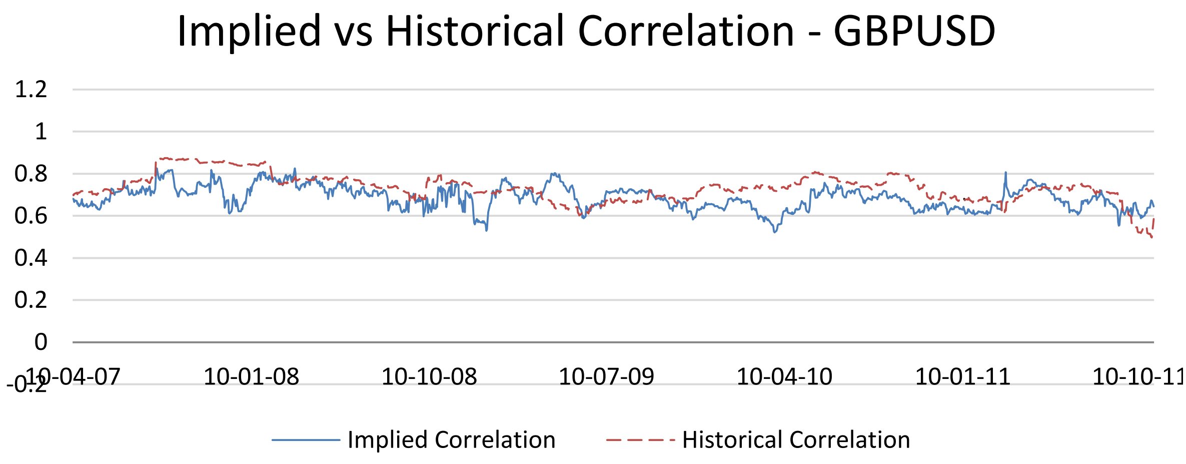

Figure 28: Implied correlation for option contracts with six months to expiration against a rolling historical correlation estimate based on a time period of 128 days.

Table 6: Descriptive statistics GBPUSD – Implied vs Historical Correlation						
	1 month		*3 months*		*6 months*	
	Implied Correlation	*Historical Correlation*	*Implied Correlation*	*Historical correlation*	*Implied Correlation*	*Historical Correlation*
Mean	0.709	0.714	0.692	0.727	0.681	0.740
Mode	0.71	0.77	0.73	0.79	0.72	0.74
Minimum	0.385	-0.002	0.403	0.340	0.392	0.497
Maximum	0.921	0.984	0.875	0.927	0.826	0.878
Observations	1326	1296	1326	1253	1326	1190

Table 6 presents descriptive statistics for GBPUSD. The minimum and maximum values confirm that both correlation measures vary less as the time frame under consideration increases. As mentioned previously, we believe it is likely that an estimation window of only 22 days cause noisy estimates of historical correlation.

Below are charts with the distribution of implied and historical correlation. Figure 29 display implied correlation for contracts with one month to expiration and figure 30 shows the distribution of a rolling correlation estimate based on 22 days. The empirical distribution of historical correlation is wider than its implied counterpart, and it has a longer (and heavier) left tail. The empirical distribution is also markedly more skewed towards the right. Again the sample correlation distribution appears to be a reasonable fit to the empirical estimates.

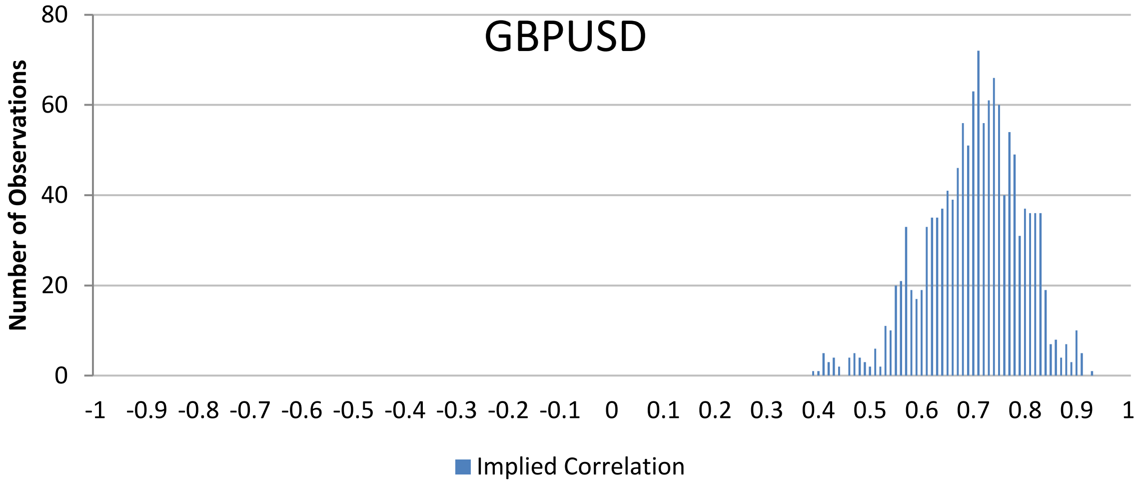

Figure 29: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPJPY and GBPUSD option contracts with one month to expiration.

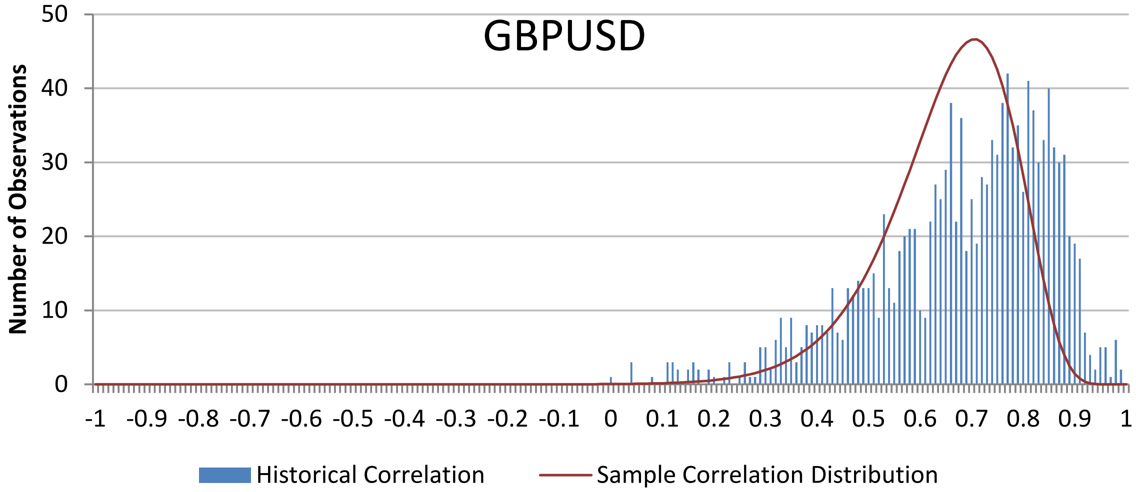

Figure 30: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPJPY. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Figure 31 and 32 below shows implied correlations for contracts with six months to expiration and historical correlation estimates based on a rolling window consisting of 128 days, respectively. Again we see that these distributions are more peaked, compared to the distributions of shorter time horizons. The theoretical sample correlation distribution is a fairly good fit to the empirical estimates.

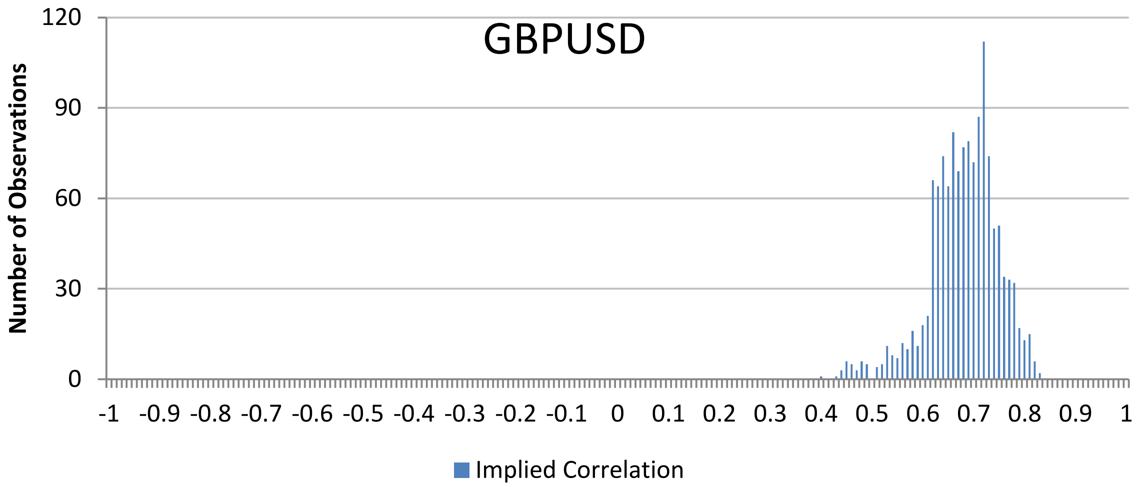

Figure 31: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPJPY and GBPUSD option contracts with six months to expiration.

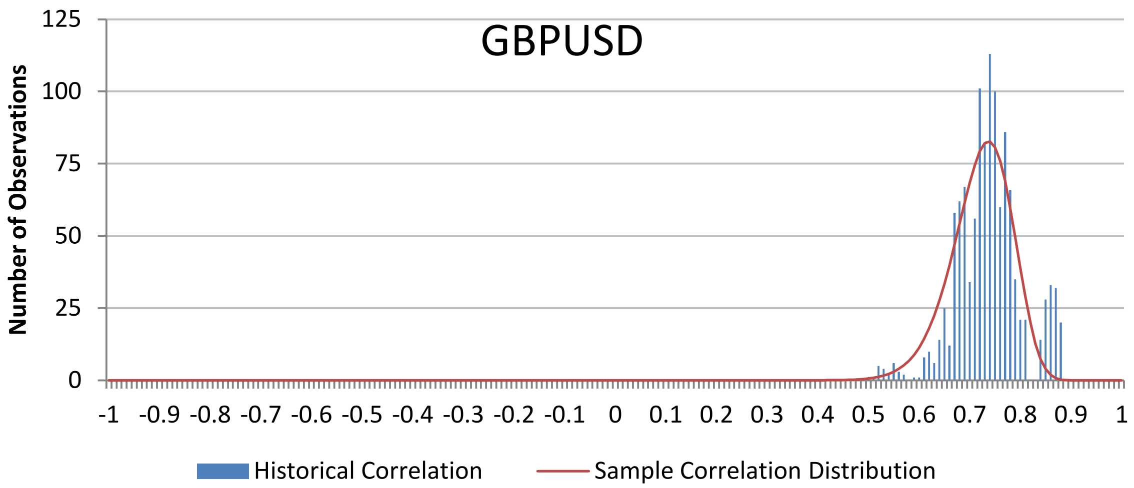

Figure 32: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for GBPJPY and USDJPY. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

4.1.4 Implied and historical correlation – GBPJPY

The implied correlation for the GBPJPY is calculated based on implied volatility quotes for USDJPY, GBPUSD and GBPJPY. The historical sample correlation is calculated from log-returns of the exchange rates USDJPY and GBPUSD.

The figures below display implied and historical correlation against time. Figure 33 shows implied correlation for option contracts with one month to maturity together with historical correlation estimates based on 22 days. Figure 34 presents correlation estimates for a time frame of six months.

The most striking feature in these figures is the "spread" between the historical and implied correlation level. The historical sample correlation consistently takes on higher values relative to implied, for both time horizons. Note that the difference between the two measures is far from constant. Several places we see that historical correlation is decreasing while implied correlation is increasing and vice versa.

Finally it looks as if implied correlation varies more erratically than historical correlation for both time horizons. This contrasts our findings so far in this study.

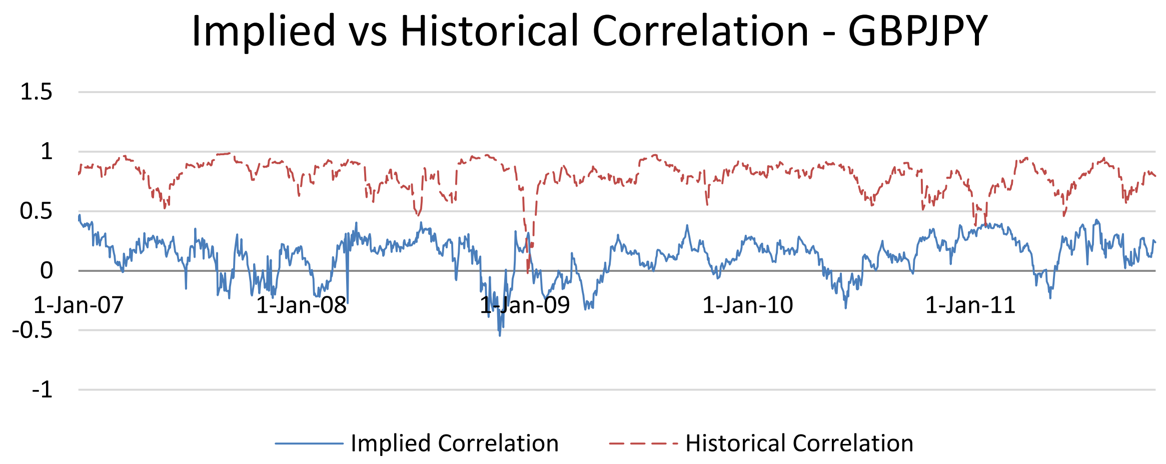

Figure 33: Implied correlation for option contracts with one month to expiration plotted against a rolling historical correlation estimate based on a time period of 22 days.

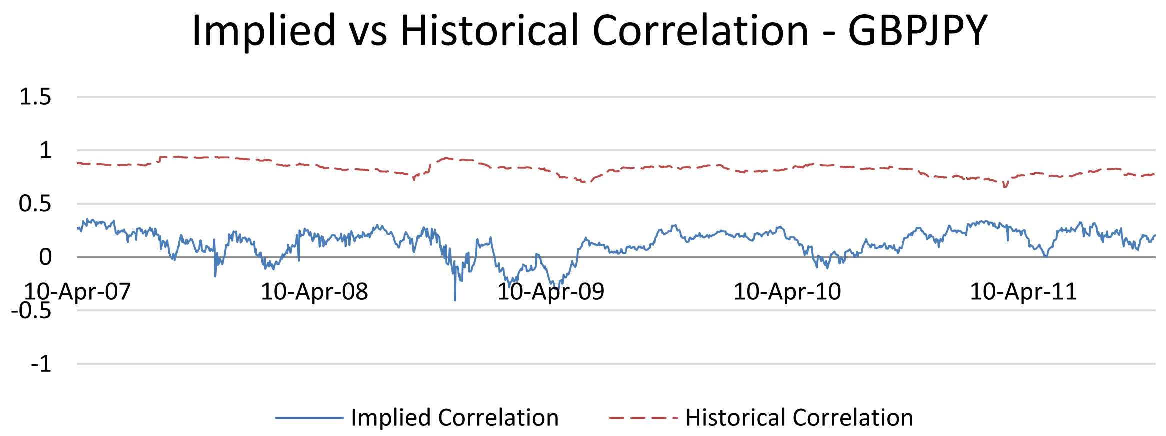

Figure 34: Implied correlation for option contracts with six months to expiration plotted against a rolling historical correlation estimate based on a time period of 128 days.

Table 7: Descriptive statistics GBPJPY – Implied vs Historical Correlation

	1 month		*3 months*		*6 months*	
	Implied Correlation	*Historical Correlation*	*Implied Correlation*	*Historical correlation*	*Implied Correlation*	*Historical Correlation*
Mean	0.150	0.828	0.159	0.834	0.166	0.838
Mode	0.17	0.89	0.24	0.86	0.21	0.83
Minimum	- 0.548	- 0.032	- 0.314	0.620	- 0.406	0.658
Maximum	0.744	0.985	0.708	0.967	0.661	0.943
Observations	1326	1296	1326	1253	1326	1190

Table 7 presents descriptive statistics. We see that the spread we discussed earlier is embodied in the mode and mean values, which differ markedly across the two correlation measures. From the minimum and maximum values we see that the distribution of historical correlation narrows as the number of data points in the rolling window estimate increase.

In figure 35 through 38 we see that while the historical correlation is skewed towards the far right of the diagram, implied correlation has a wide distribution and peak around 0.15. The implied correlation distributions also have heavy tails.

For historical correlation the tails more or less disappear as we increase the size of the estimation window, and the distribution gets very cramped. We have seen this happen for several of the currency pairs.

Finally we note that the empirical distribution of historical correlation is located further to the right, relative to the theoretical distribution.

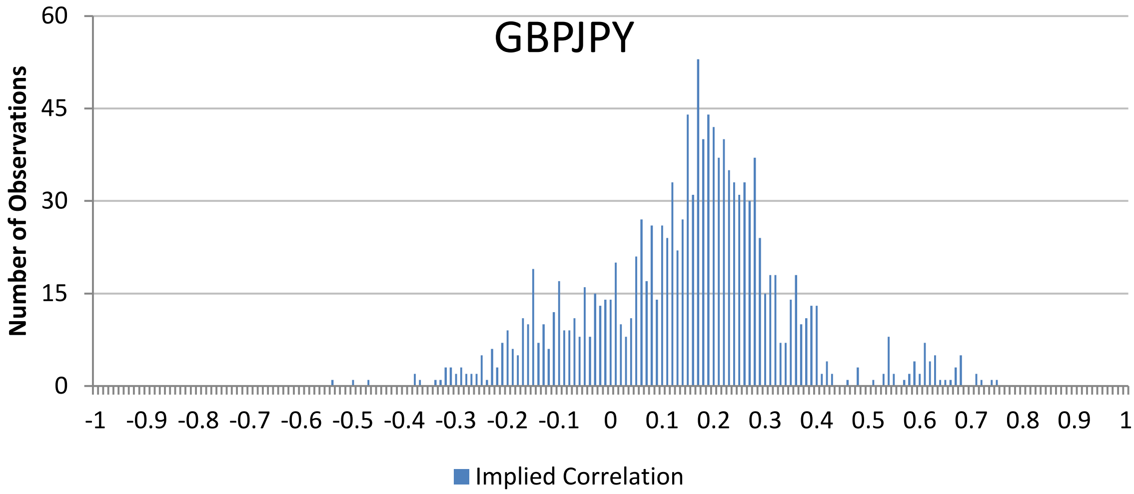

Figure 35: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPUSD and GBPJPY option contracts with one month to expiration.

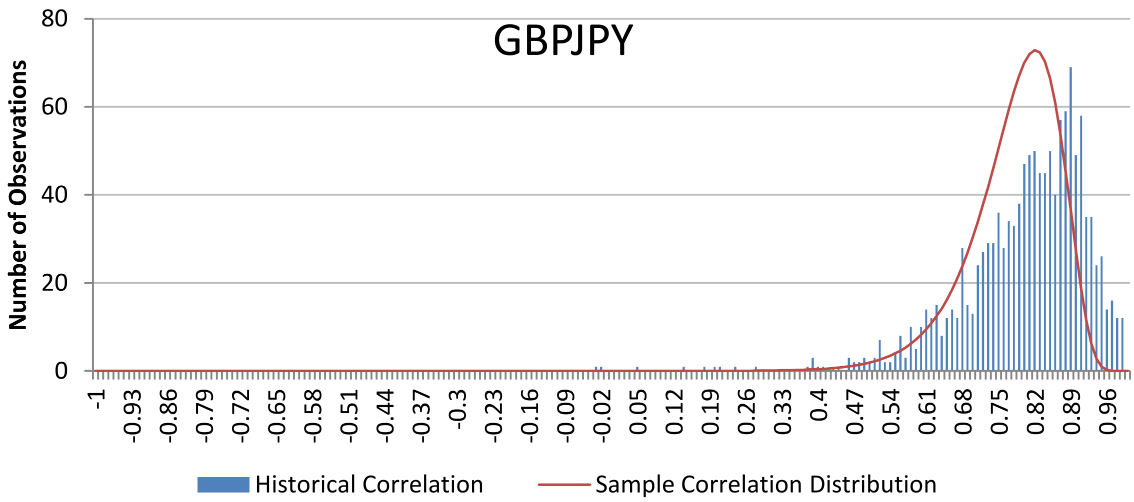

Figure 36: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPUSD. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

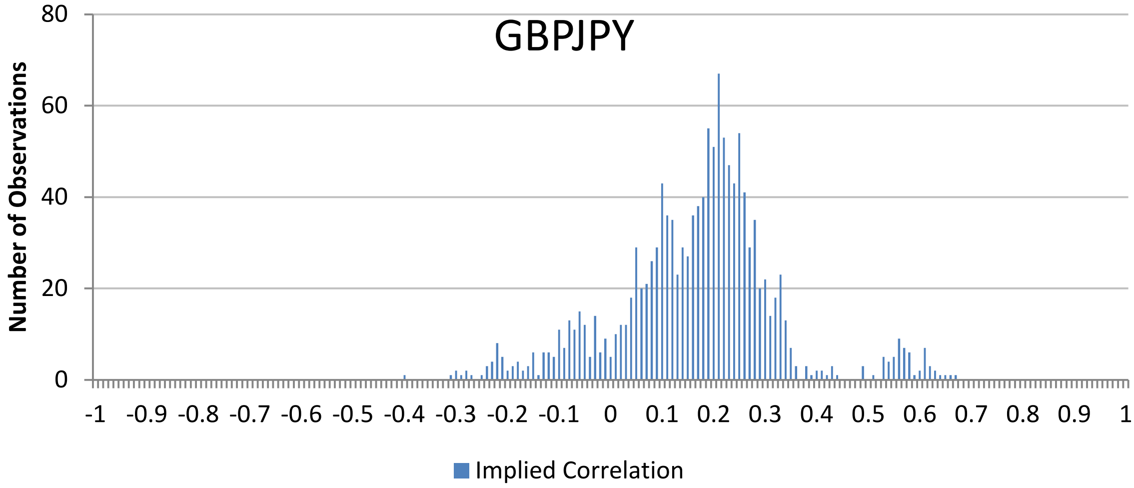

Figure 37: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPUSD and GBPJPY option contracts with six months to expiration.

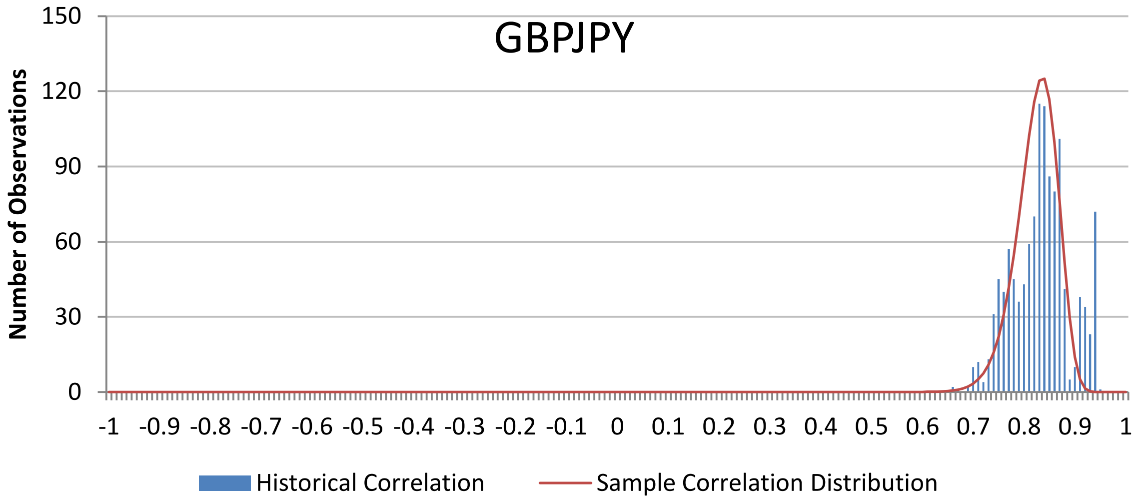

Figure 38: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPUSD. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

The final point we would like to discuss in relation to GBPJPY concerns two outlying implied correlation values, calculated from price quotes of contracts with one year to expiration. For this particular pair and contract we have two occurrences of implied correlation values larger than 1. This seems to suggest that extreme market conditions can cause the relationship governing implied correlations to break down temporarily. The two occurrences of implied correlation > 1 occurred on 28.05.07 and 29.05.07. While it is not visible in the graph, the calculated values were approximately 1.007 and 1.008.

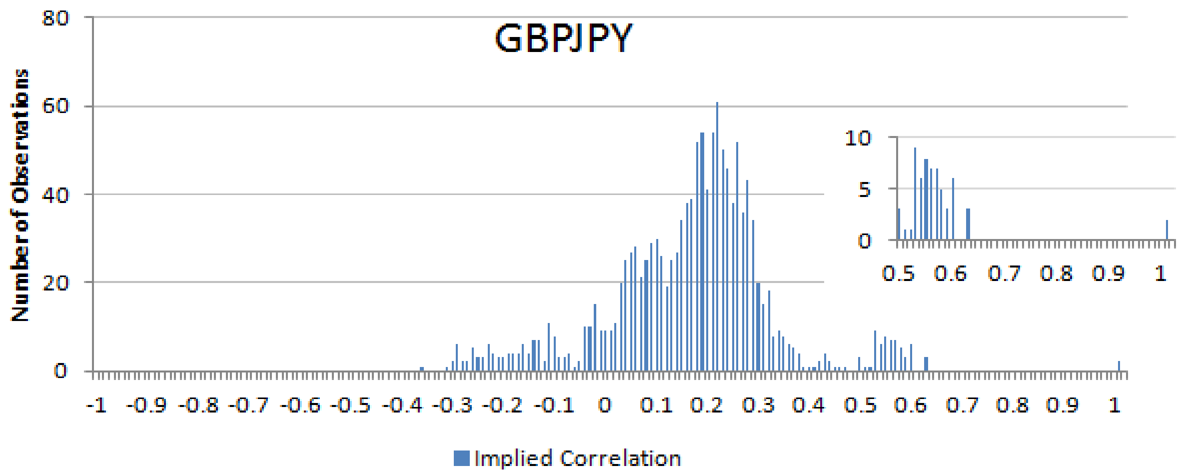

Figure 39: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPUSD and GBPJPY option contracts with 12 months to expiration. To the right we have zoomed in on the right tail of the distribution.

4.1.5 Implied and historical correlation – EURUSD

In the following our calculations of implied correlation are based on option contracts for EURJPY, USDJPY and EURUSD. Historical correlation is estimated from the log-returns of EURJPY and USDJPY.

Figure 40 and 41 below present implied and historical correlation plotted against time. Figure 40 displays implied volatility for contracts with one month to maturity against a historical correlation estimate based on 22 days. Figure 41 presents contracts with six months to expiration alongside historical correlation estimates based on a 128 day rolling window.

Once more we get the impression that historical correlation varies more than implied correlation for short time frames. They do however seem to follow the same overall trend.

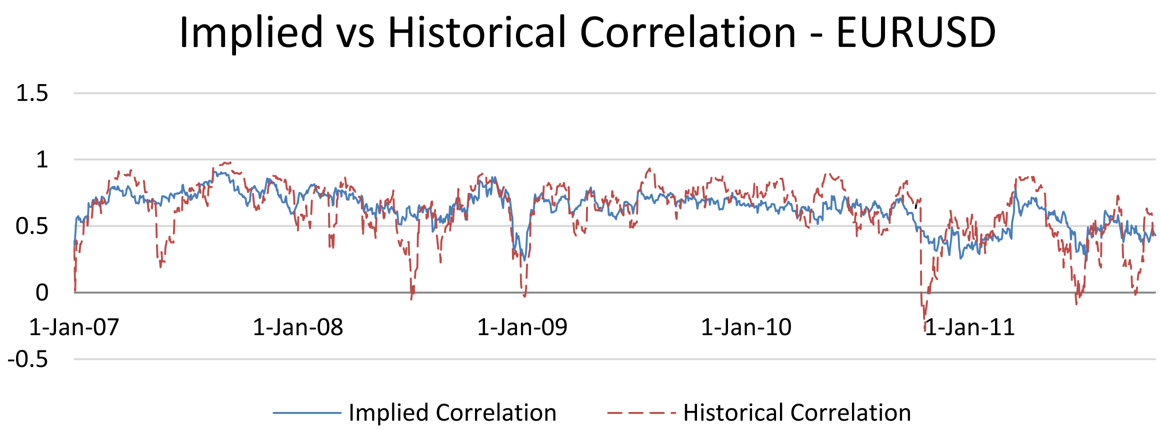

Figure 40: Implied correlation for option contracts with one month to expiration plotted against a rolling historical correlation estimate based on a time period of 22 days.

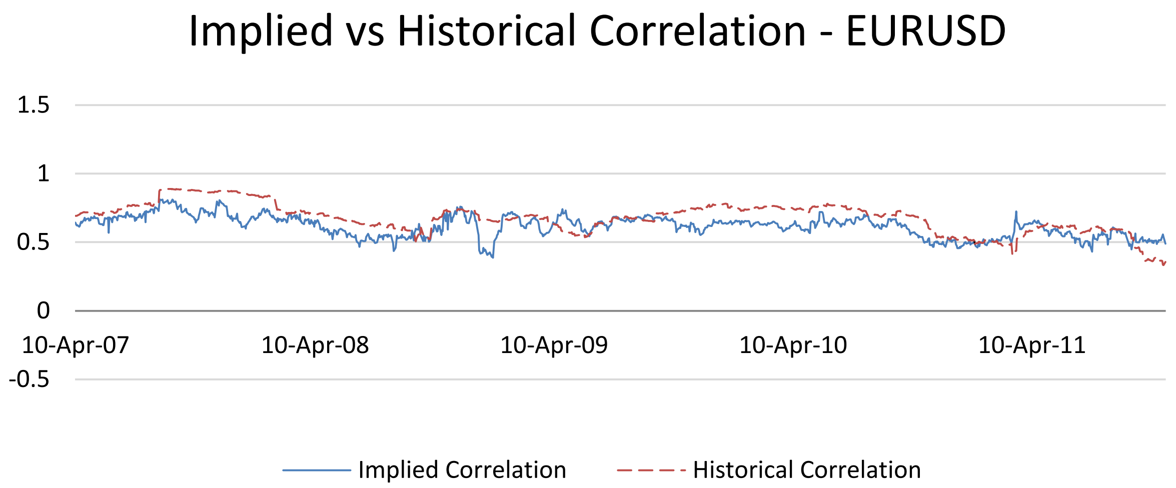

Figure 41: Implied correlation for option contracts with six month to expiration plotted against a rolling historical correlation estimate based on a time period of 128 days.

Table 8: Descriptive statistics EURUSD – Implied vs Historical Correlation

	1 month		*3 months*		*6 months*	
	Implied Correlation	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*
Mean	0.640	0.670	0.621	0.681	0.615	0.695
Mode	0.64	0.75	0.67	0.77/0.78	0.73	0.77/0.79
Minimum	0.240	- 0.288	0.290	0.198	0.332	0.333
Maximum	0.913	0.980	0.857	0.927	0.812	0.888
Observations	1326	1296	1326	1253	1326	1190

In table 8 we see that historical correlation has a wider distribution relative to implied for the shorter time frames. With increased time horizon the minimum and maximum values align.

Below are charts presenting the empirical distribution of implied and historical correlation. Figures 42 and 43 display implied and historical correlation with a monthly time horizon, while figure 44 and 45 present correlation levels for option contracts with 6 months to maturity and rolling window estimates based on 128, respectively.

When we look at figure 42 and 43 we see that the distribution of both implied and historical correlation leans heavily to the positive side, with a long tail trailing off to the left. The distribution of implied correlation appears nearly bimodal.

For historical correlation we see that the left tail is quite heavy, especially when we compare it with the theoretical sample distribution. Once more the distribution based on empirical estimates of historical correlation is located further to the right, relative to the theoretical sampling distribution.

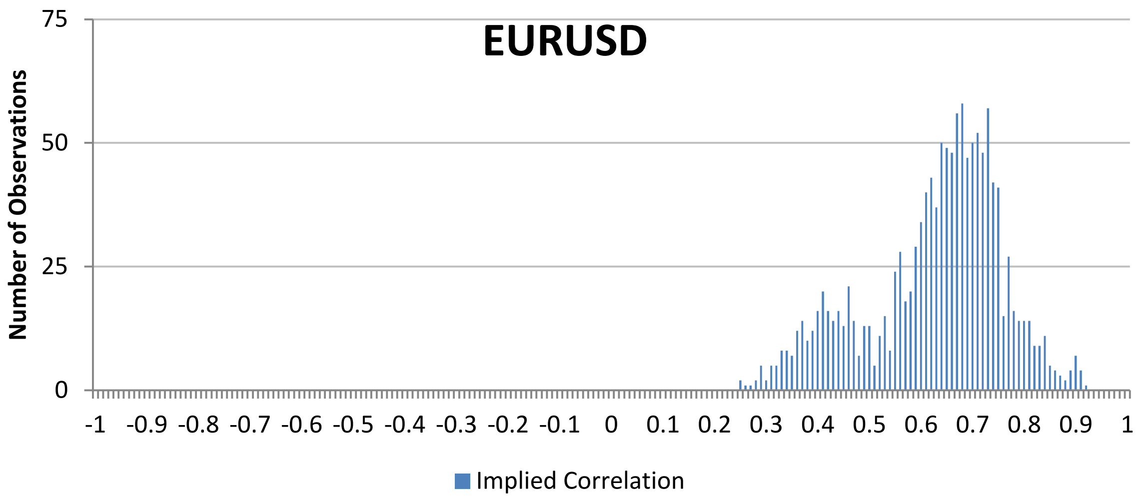

Figure 42: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURJPY and EURUSD option contracts with one month to expiration.

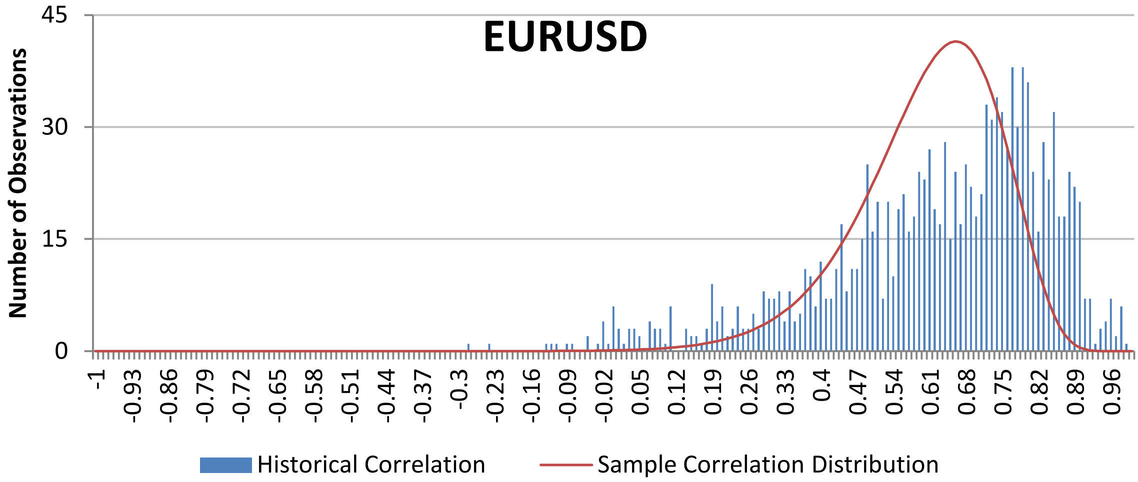

Figure 43: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and EURJPY. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

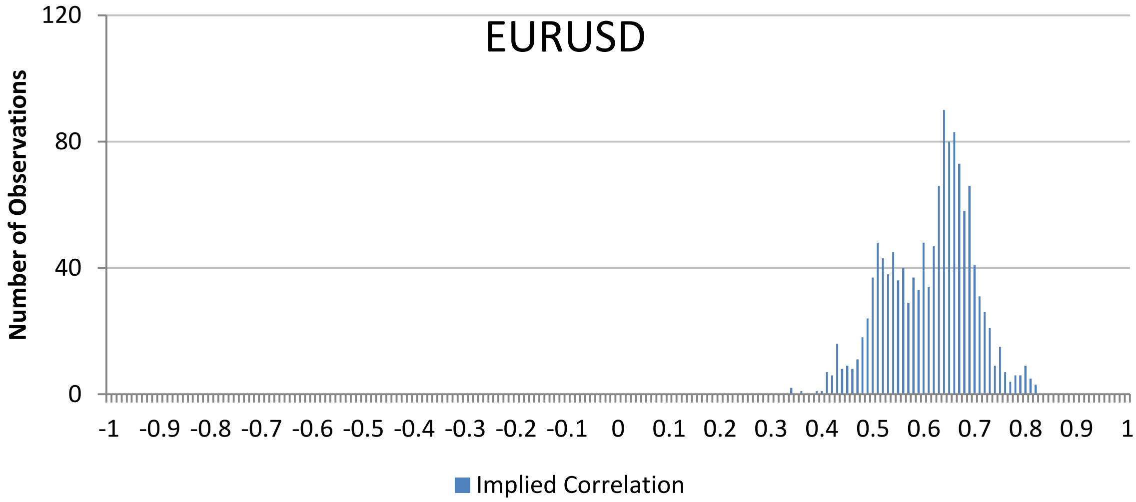

Figure 44: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURJPY and EURUSD option contracts with six months to expiration.

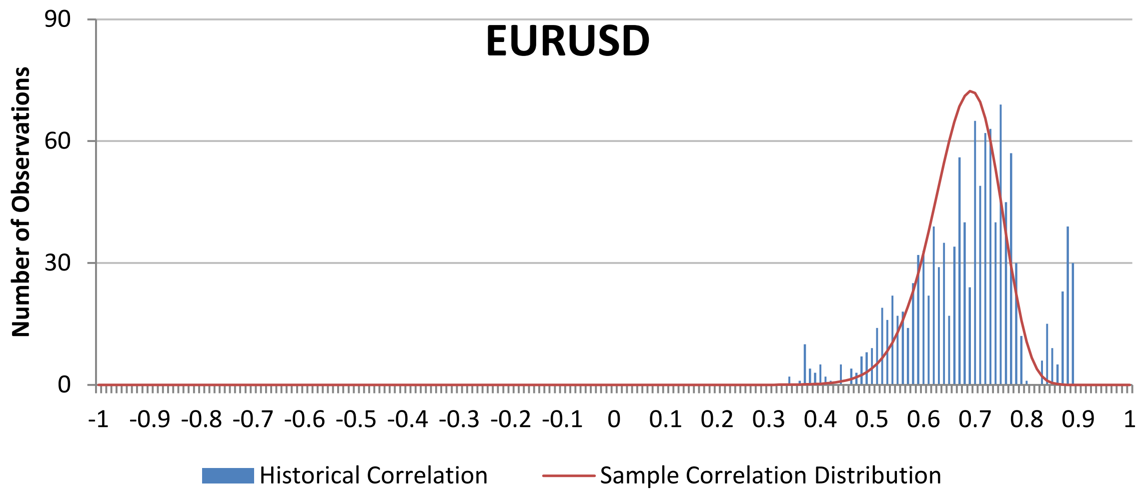

Figure 45: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and EURJPY. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

4.1.6 Implied and historical correlation – EURUSD2

Once more we have been able to calculate an implied correlation using two different currencies as numeraire. The implied correlation we discuss in this section is labeled EURUSD2, and is calculated based on implied volatility quotes for EURNOK, USDNOK and EURUSD. Historical correlation is estimated using log-returns of EURNOK and USDNOK. In other words, this section concerns the implied correlation between Euro and US dollars using Norwegian Krone as numeraire.

Figure 46 and 47 below display historical versus implied correlation for EURUSD2 against time. The first figure shows implied correlation for contracts with one month to maturity against a rolling correlation estimate based on 22 days. The second figure display implied correlation for contracts with six months to maturity against the corresponding historical correlation estimate.

In figure 46 we see that historical correlation is more variable than implied correlation. This corresponds well with what we saw in figure 40, although it looks like the historical correlation of EURUSD2 is less variable than its EURUSD counterpart. Figure 47 shows that historical correlation varies less as the time frame increases, which is consistent with our previous findings.

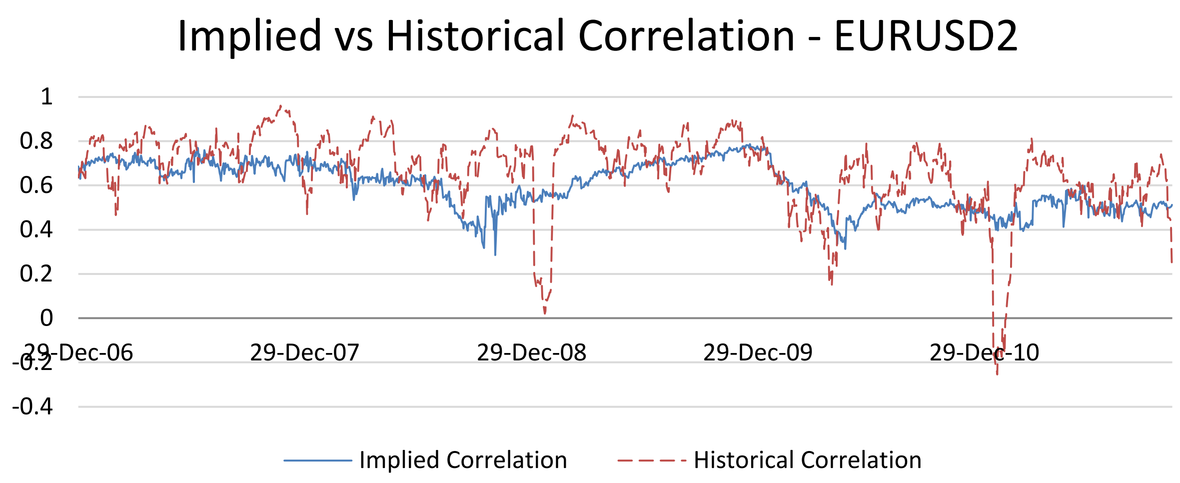

Figure 46: Implied correlation for option contracts with one month to expiration plotted against a rolling historical correlation estimate based on a time period of 22 days.

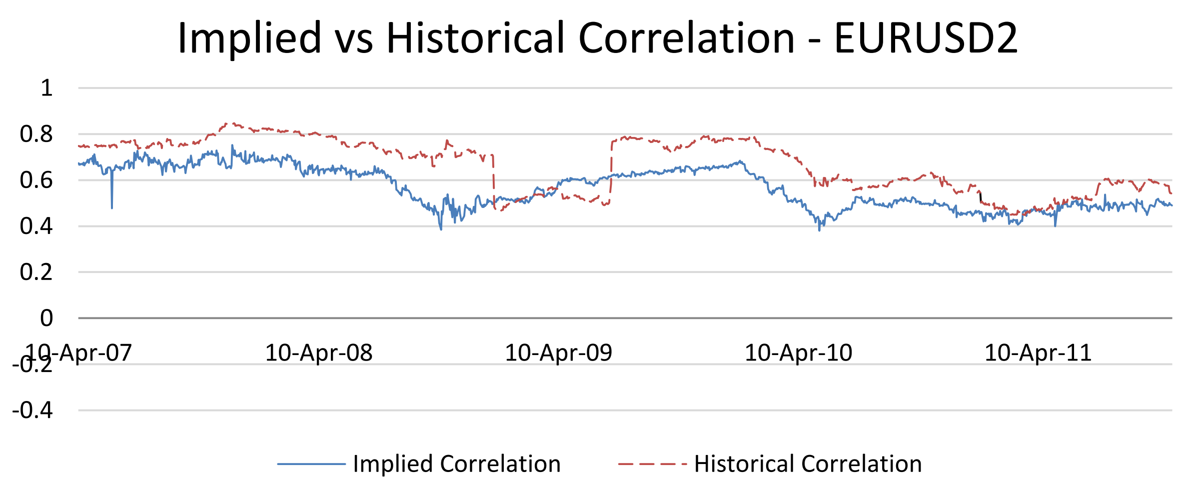

Figure 47: Implied correlation for option contracts with six months to expiration plotted against a rolling historical correlation estimate based on a time period of 128 days.

In figures 46 and 47 we can see a sharp decline in the historical correlation level in the beginning of 2009. The length of the time period with relatively lower correlation increases as we add to the number of data points used in our rolling window estimate. As previously mentioned, rolling historical correlation is tracking a certain time period. An extreme observation entering and exiting the estimation window can cause precisely such a sudden jump in the correlation level. Below is a close up taken from figure 46, which zoom in on the abrupt change in the correlation level.

The historical correlation for EURUSD2 is estimated based on the log-returns of EURNOK and NOKUSD, and the drop in correlation is likely caused by an outlier in these series. Visual inspection of the graph below reveals that the log-returns of NOKUSD indeed take on its highest *and* lowest value around the beginning of 2009. When we zoom in on the area with outliers we see that the log-returns have an all time low value of -0.094 on January 1st, before "jumping" back up to an all time high 0.083 on the next day. This corresponds perfectly with the drop in correlation. On January 1st the correlation level drops to 0.491 from 0.711 the day before. On January 2nd the historical correlation is 0.199. For reference, the rolling historical correlation estimate based on a window of 22 observations has an average value of 0.670.

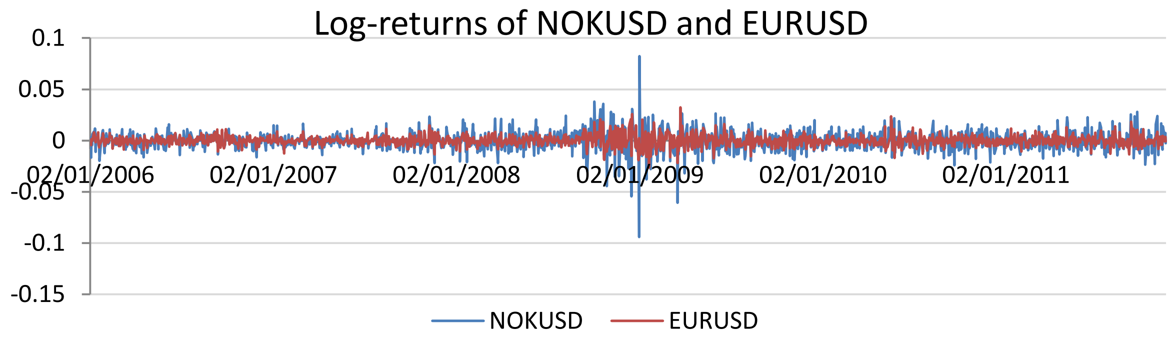

Figure 48: Log-returns of the exchange rates NOKUSD and EURUSD plotted against time.

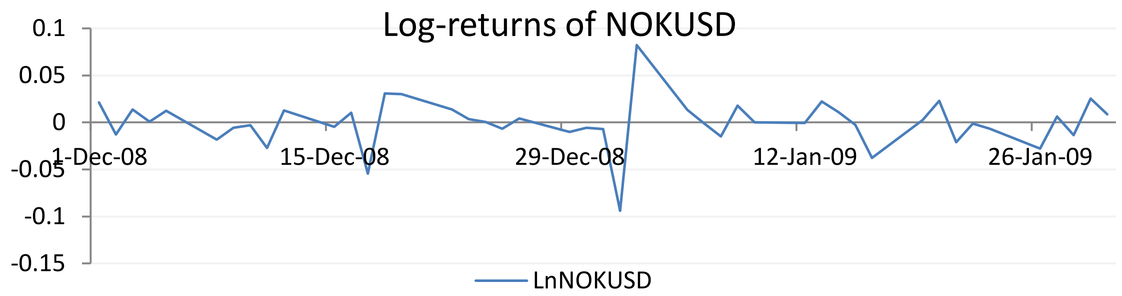

Figure 49: Log-returns of the exchange rate NOKUSD plotted against time.

If we increase the size of the rolling window from 22 days to 128, the decline in the correlation level is less pronounced. During the first days of January the correlation drops from roughly 0.70 to a level around 0.47. The average value for the correlation coefficient is 0.666 when the calculations are based on 128 samples points.

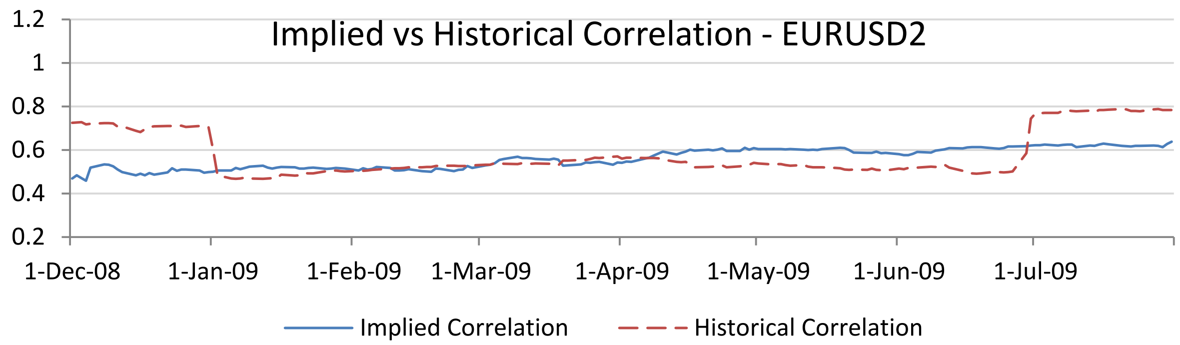

Figure 50: Implied correlation for option contracts with one month to expiration plotted against a rolling historical correlation estimate based on a time period of 128 days.

In figure 50 above we see that while the decline in correlation is less extreme, the time period with relatively lower correlation has increased. We have already discussed the reason for this phenomenon in chapter 2.5. The jump in correlation occurs when an outlying observation enters the calculation, and disappears as soon as this observation exits the estimation window. When we calculate rolling correlations for a time period of 22 days the extreme observations influence the correlation coefficient for this time period, before it suddenly receives zero weight. Conversely, when we increase the time horizon to e.g. 128 days, the outliers will be less influential on each point estimate, but the outliers will continue to affect the calculations for 128 days.

In sum, we see that the caveats we mentioned in chapter 2.5 are important to keep in mind. The previous paragraphs have demonstrated that jumps in correlation can indeed be spurious. While the decline in the correlation level occurred at the same date for both time horizons, the outliers remained influential for 22 and 128, respectively. Deciding on an appropriate number of sample points when estimating correlations will always be somewhat ad hoc, but it is definitely not trivial.

Table 9: Descriptive statistics EURUSD2 – Implied vs Historical Correlation

	1 month		*3 months*		*6 months*	
	Implied Correlation	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*
Mean	0.611	0.701	0.593	0.688	0.581	0.682
Mode	0.70	0.75	0.69	0.75	0.50	0.75
Minimum	0.286	- 0.254	0.329	0.249	0.379	0.448
Maximum	0.787	0.960	0.744	0.899	0.753	0.847
Observations	1326	1296	1326	1253	1326	1190

In table 9 we see that the correlation measures of EURUSD2 generally have similar modes and sample means as the estimates for EURUSD. The widths of the distributions also appear similar to the corresponding ones in the previous section.

In figure 51 we see that the distribution of implied correlation backed out from option contracts with one month to expiration looks similar to the one in figure 42. However, the figure below has a more pronounced bimodal shape.

When we compare the historical correlation of EURUSD linked through Japanese yen with the estimate based on Norwegian Krone, they appear similar when we consider a monthly horizon. The left tail we can see lurking in figure 43 is also present in figure 52, although it is thinner in the latter case.

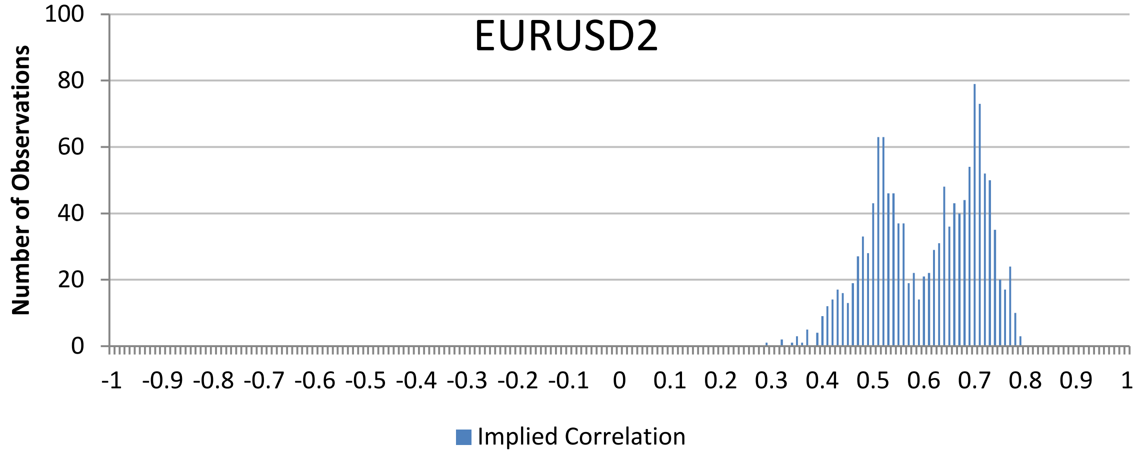

Figure 51: The empirical distribution of implied correlation based on implied volatility quotes for USDNOK, EURNOK and EURUSD option contracts with one month to expiration.

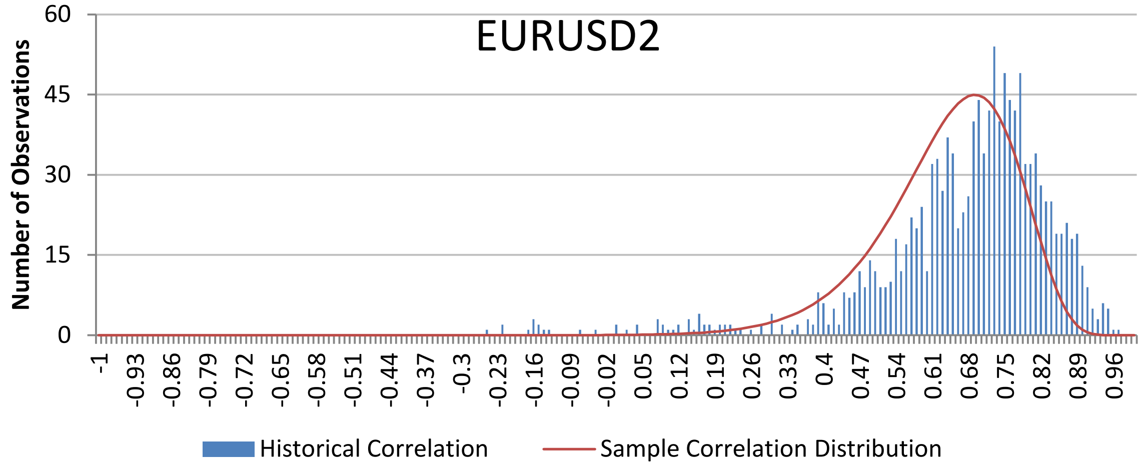

Figure 52: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDNOK and EURNOK. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

The shape of the next distributions is rather striking. They are clearly bimodal, with very distinct peaks. We cannot see any reasonable explanation for this. As far as we are aware of, the only unusual feature of this currency pair is that option contracts on the Norwegian Krone are rather thinly traded.

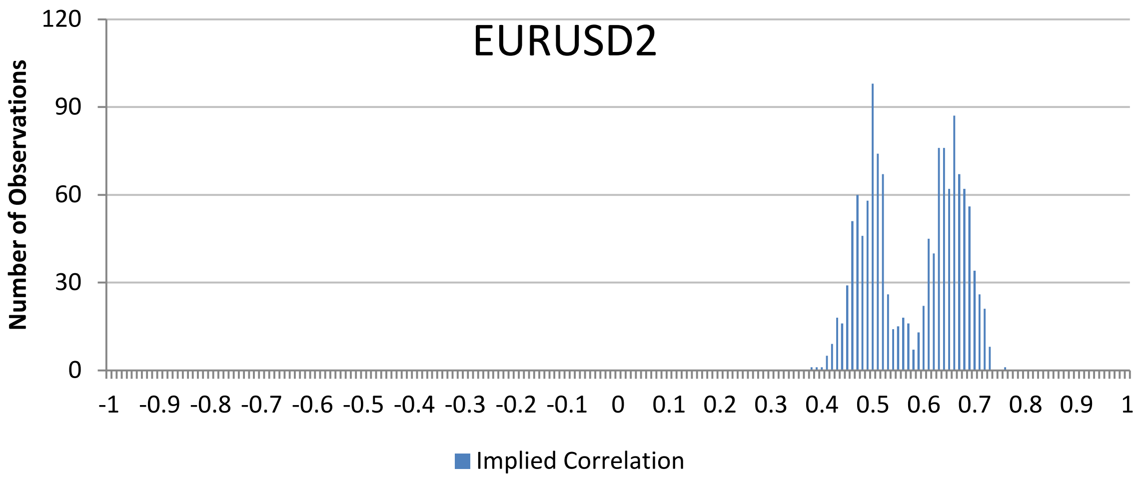

Figure 53: The empirical distribution of implied correlation based on implied volatility quotes for USDNOK, EURNOK and EURUSD option contracts with six months to expiration.

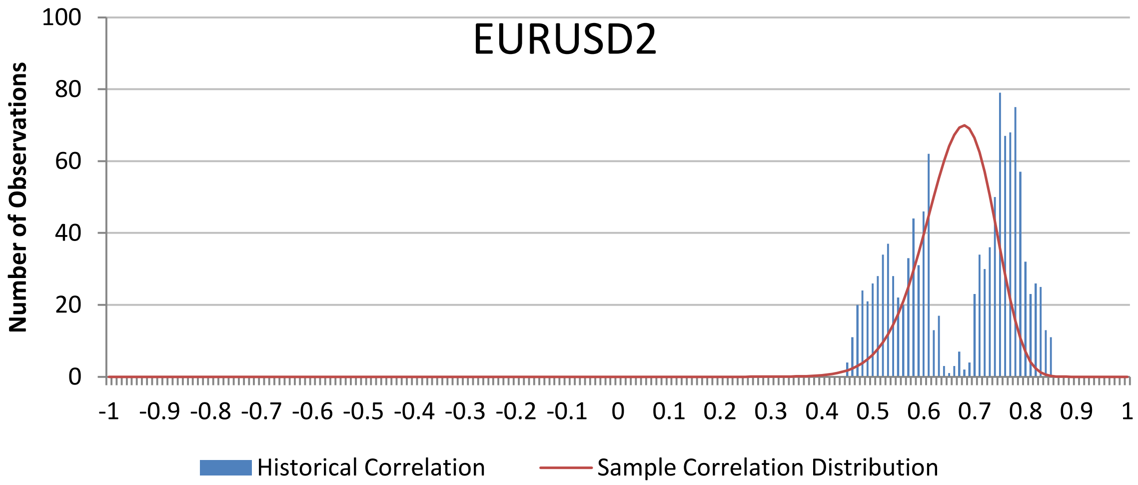

Figure 54: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDNOK and EURNOK. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

4.1.7 Implied and historical correlation – EURJPY

The implied correlation of EURJPY is calculated based on option quotes for USDJPY, EURUSD and EURJPY. The historical correlation is estimated using log-returns of the exchange rates for USDJPY and EURUSD.

Figure 55 and 56 present implied and historical correlation plotted against time. The implied correlation in figure 55 is calculated based on option contracts with one month to maturity, while the historical correlation estimate is based on 22 days. The corresponding time periods in figure 56 are six months and 128 days.

The most notable feature in these figures is the difference between implied and historical correlation. It seems as if the two correlation measures are trending in opposite directions;

when implied correlation is increasing, historical correlation is creasing and vice versa. The two measures come together occasionally, only to shoot off in different directions once more.

Further we see that the variation in both correlation measures is reduced when the time frame increases. The variability is however very high, regardless of which time horizon we look at.

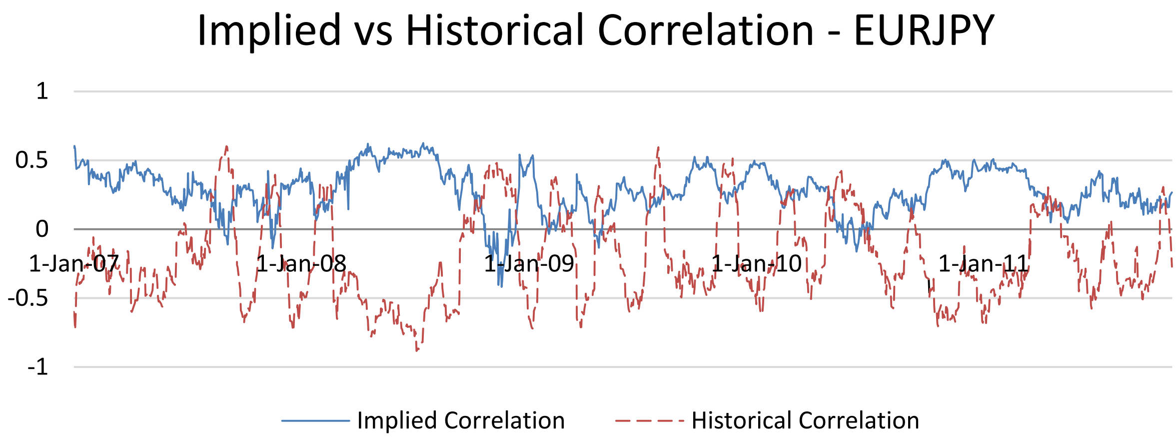

Figure 55: Implied correlation for option contracts with one month to expiration plotted against a rolling historical correlation estimate based on a time period of 22 days.

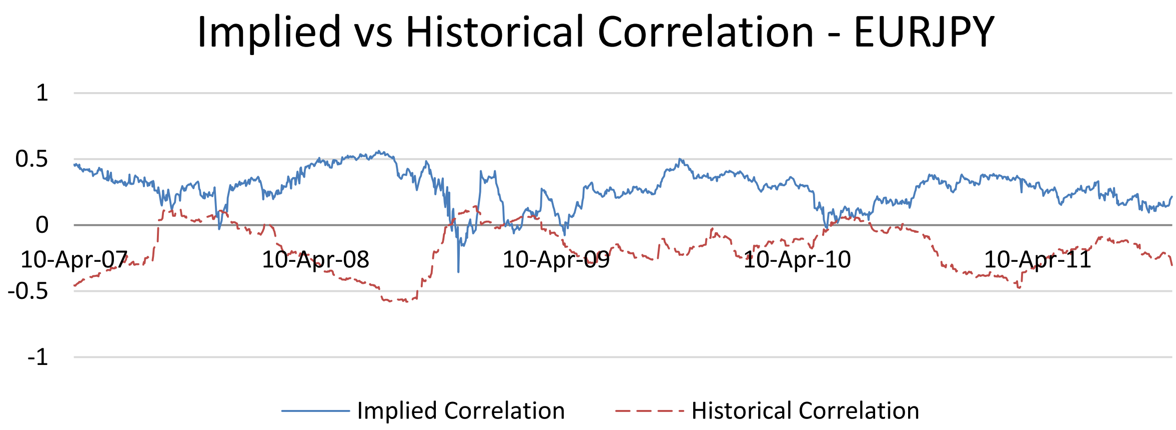

Figure 56: Implied correlation for option contracts with six months to expiration plotted against a rolling historical correlation estimate based on a time period of 128 days.

Table 10: Descriptive statistics EURJPY – Implied vs Historical Correlation

	1 month		*3 months*		*6 months*	
	Implied Correlation	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*
Mean	0.306	- 0.253	0.309	- 0.214	0.305	- 0.183
Mode	0.26	- 0.34	0.31	- 0.27/- 0.30	0.32	- 0.22
Minimum	- 0.418	- 0.882	- 0.257	- 0.697	- 0.356	- 0.586
Maximum	0.778	0.603	0.745	0.318	0.709	0.143
Observations	1326	1296	1326	1253	1326	1190

The impression we got from figure 55 and 56 is confirmed by the descriptive statistics in table 10. The modes and sample means strengthen the impression of correlation measures that tend to move in opposite directions.

Figure 57 through 60 below presents the distributions of implied and historical correlation for a one month and six month time horizon. These distributions are very wide. Looking at figure 58 we see that the correlation estimates range from -0.882 to 0.603. Figure 59 and 60 show that the distributions narrow somewhat when we consider a longer time horizon. They are however still quite wide.

The correlation sample distribution is not a good fit to the empirical distribution based on a 22 day rolling window. The empirical distribution has extremely fat tails relative to the theoretical one. In figure 60 we see that the fit does not improve when we consider a longer time frame.

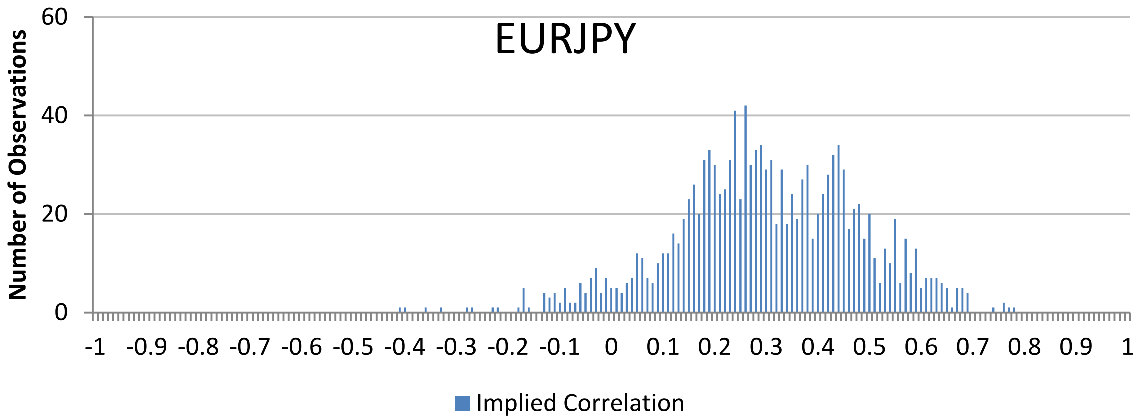

Figure 57: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURUSD and EURJPY option contracts with one month to expiration.

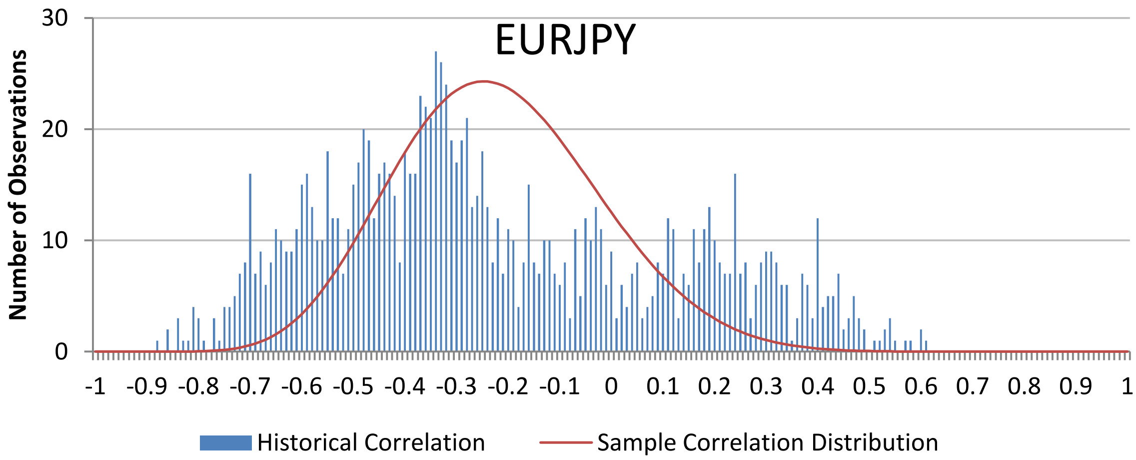

Figure 58: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and EURUSD. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

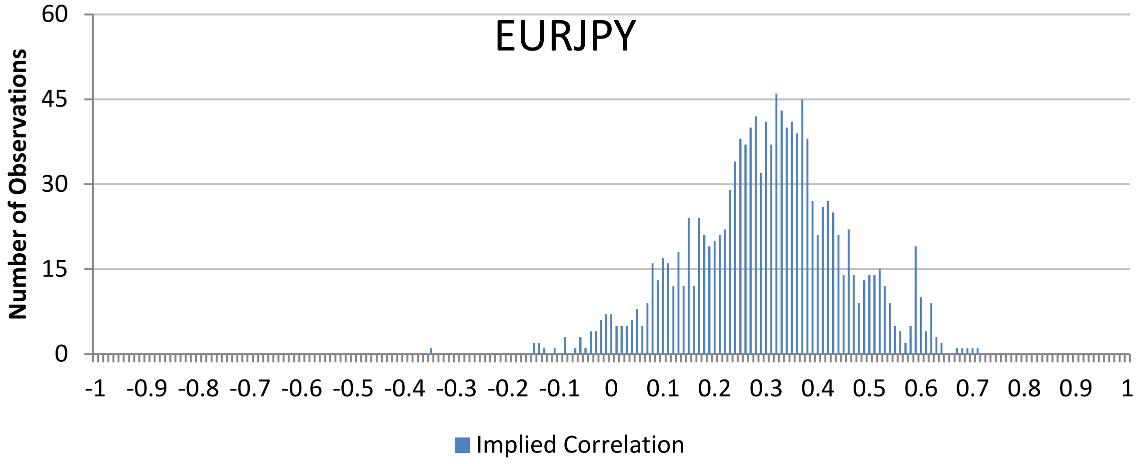

Figure 59: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURUSD and EURJPY option contracts with six months to expiration.

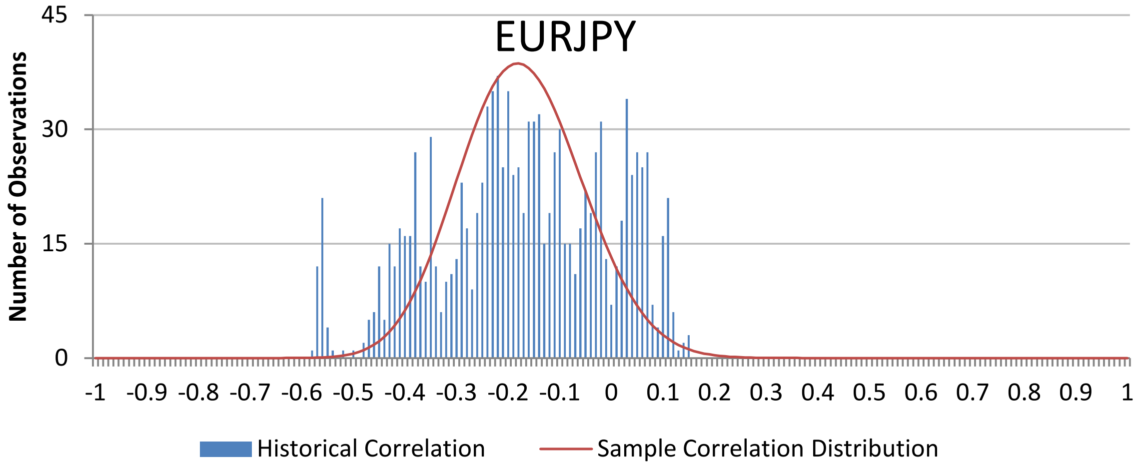

Figure 60: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and EURUSD. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

4.1.8 Implied and historical correlation – USDNOK

The implied correlation between EURUSD and EURNOK, labeled USDNOK, is calculated based on implied volatility quotes for EURUSD, EURNOK and USDNOK. Historical correlation is estimated from the log-returns of EURUSD and EURNOK.

Figure 61 and 62 display implied and historical correlation plotted against time. The implied correlation in figure 62 is calculated based on option contracts with one month to maturity, while rolling historical correlation is calculated using 22 days. The corresponding time periods in figure 62 are six months and 128 days, respectively.

Both of the correlation measure exhibit quite frenetic zig-zaging when we consider a time horizon of one month. When we look at figure 62 we get the impression of correlation coefficients that repel each other. It appears as if implied correlation is increasing when historical correlation is decreasing and vice versa. This behavior is however not as pronounced as what we saw for EURJPY in the previous section.

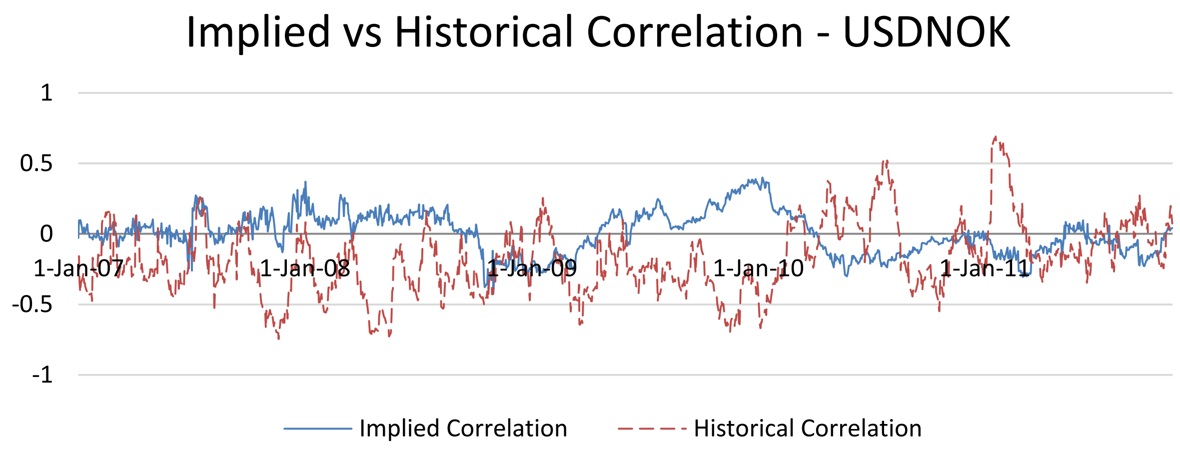

Figure 61: Implied correlation for option contracts with one month to expiration plotted against a rolling historical correlation estimate based on a time period of 22 days.

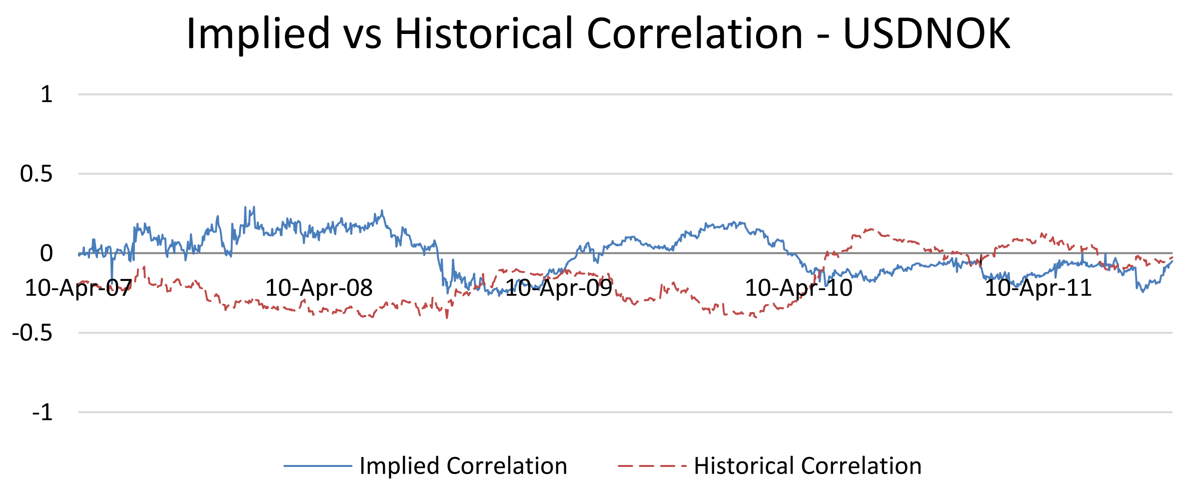

Figure 62: Implied correlation for option contracts with six months to expiration plotted against a rolling historical correlation estimate based on a time period of 128 days.

Table 11: Descriptive statistics USDNOK – Implied vs Historical Correlation

	1 month		*3 months*		*6 months*	
	Implied Correlation	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*
Mean	0.005	- 0.191	- 0.003	- 0.180	- 0.009	- 0.176
Mode	- 0.04	- 0.25	- 0.06	- 0.22	- 0.06	- 0.34
Minimum	- 0.425	- 0.745	- 0.316	- 0.564	- 0.268	- 0.407
Maximum	0.399	0.691	0.313	0.320	0.293	0.153
Observations	1326	1296	1326	1253	1326	1190

From table 11 we see that both correlation measures have wide distributions. While the sample means of implied correlation are close to zero, the sample means of historical correlation lean to the negative side. They take on similar values for all time horizons.

Evidence of this behavior can be found in the figures below. Figure 63 display implied correlation backed out from option contracts with one month to expiration. Figure 64 shows rolling historical correlation based on a time period of 22 days. Figures 65 and 66 present the distribution of implied correlations based on option contracts with six months to maturity and historical correlation estimates for a time frame of six months.

We see that when the mean correlation level is close to zero, the distribution of the correlation coefficients becomes bell shaped. Figure 64 gives the impression of bivariate normality and very heavy tails. When we look at the next figures, which consider a time frame of six months, we see that the tails are more or less chopped off. In figure 66 we see that the theoretical sample distribution is now a very poor fit to the empirical estimates.

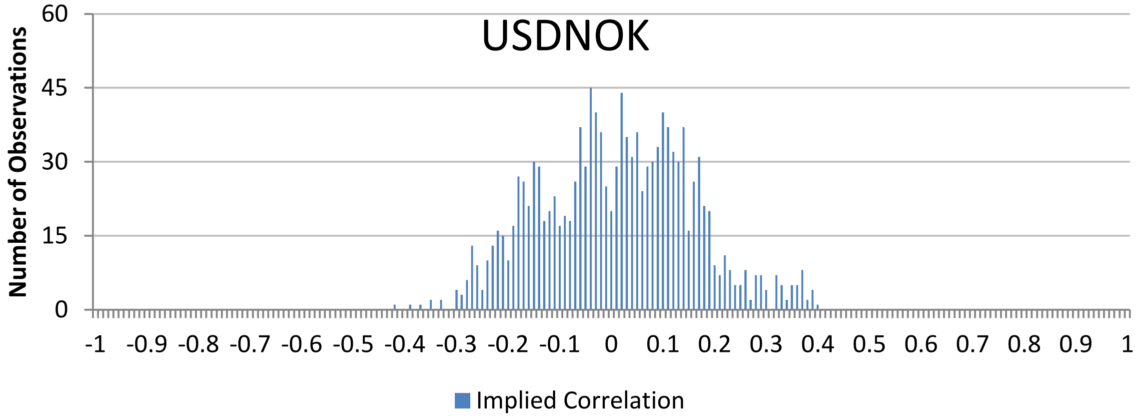

Figure 63: The empirical distribution of implied correlation based on implied volatility quotes for EURNOK, EURUSD and USDNOK option contracts with one month to expiration.

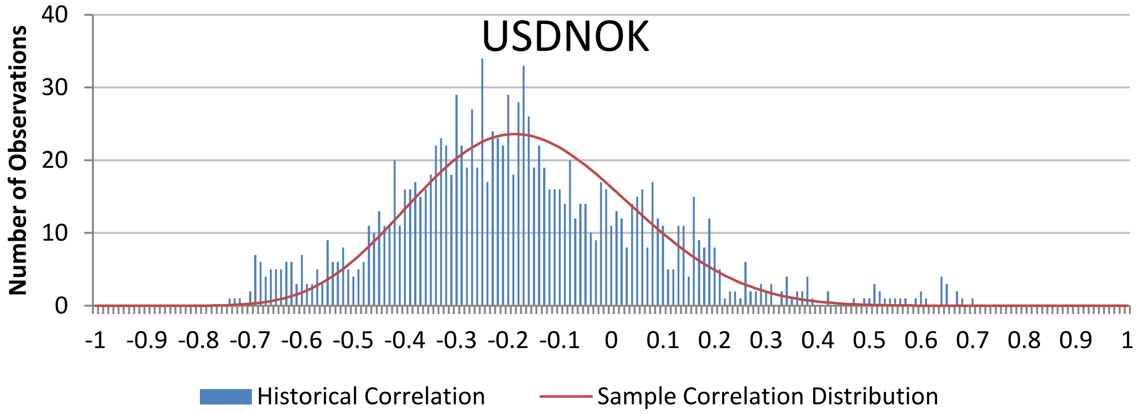

Figure 64: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURNOK and EURUSD. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

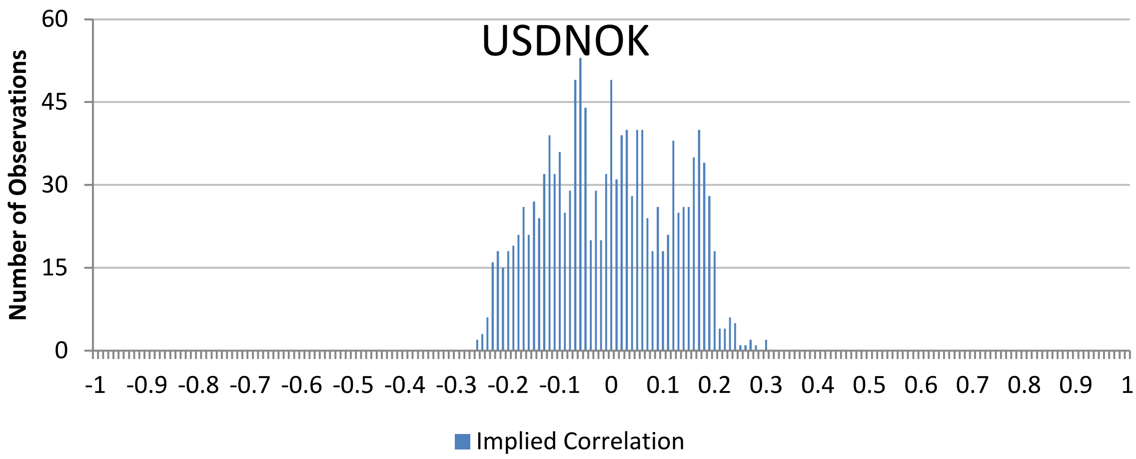

Figure 65: The empirical distribution of implied correlation based on implied volatility quotes for EURNOK, EURUSD and USDNOK option contracts with six months to expiration.

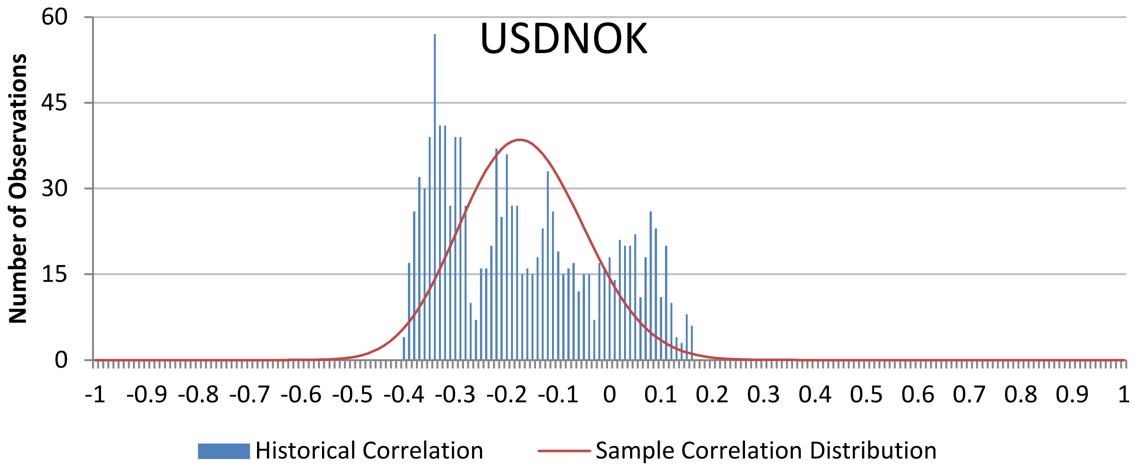

Figure 66: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURNOK and EURUSD. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

4.1.9 Implied and historical correlation – EURNOK

In the following our calculations of implied correlation are based on implied volatility quotes for USDNOK, EURUSD and EURNOK. Historical correlation is estimated from the log-returns of USDNOK and EURUSD.

Figure 67 and 68 display implied and historical correlation against time. The implied correlation in figure 67 is calculated based on option contracts with one month to maturity, while rolling historical correlation is calculated using 22 days. The corresponding time frames in figure 68 are six months and 128 days.

The figures below give the impression of correlation measures that take on similar values, albeit with different signs. This will be discussed further in relation to the descriptive statistics in table 12. We also note that both correlations vary less with increased time horizon.

In the figures below it appears as if these correlations are less "turbulent" than the other correlation series we have examined in this study. We see no apparent reason for this phenomenon. Once more, we can only think of one unusual feature of these correlations, namely the low trading volume of NOK options.

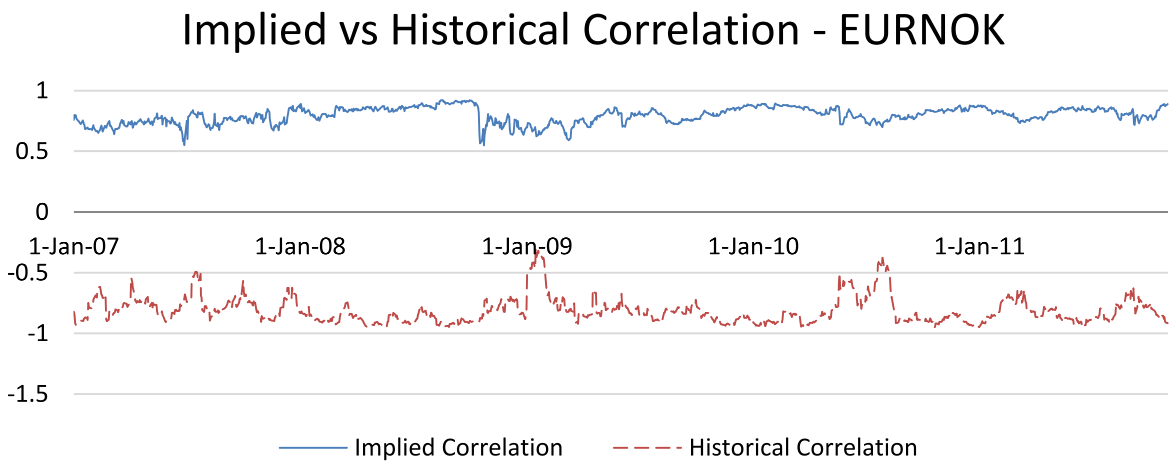

Figure 67: Implied correlation for option contracts with one month to expiration plotted against a rolling historical correlation estimate based on a time period of 22 days.

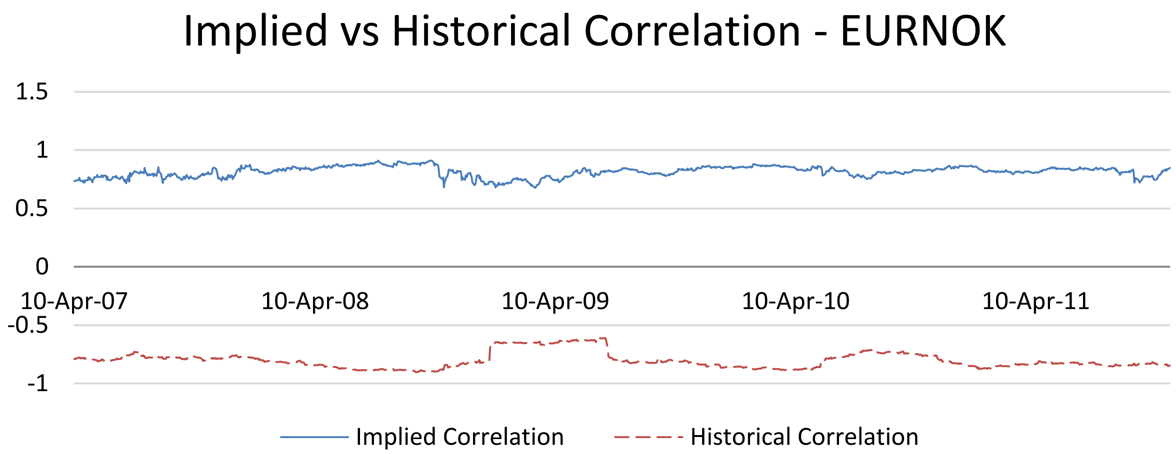

Figure 68: Implied correlation for option contracts with six month to expiration plotted against a rolling historical correlation estimate based on a time period of 128 days.

Table 12: Descriptive statistics EURNOK – Implied vs Historical Correlation						
	1 month		*3 months*		*6 months*	
	Implied Correlation	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*	*Implied Correlation*	*Historical Correlation*
Mean	0.803	- 0.834	0.810	- 0.822	0.813	- 0.812
Mode	0.84	- 0.90	0.85	- 0.88	0.82/0.84	- 0.82
Minimum	0.548	- 0.959	0.641	- 0.922	0.676	- 0.904
Maximum	0.921	- 0.293	0.921	- 0.494	0.911	- 0.609
Observations	1326	1296	1326	1253	1326	1190

Table 12 confirms that the sample means take on nearly identical values, but with different signs.

Figure 69 below display the empirical distribution of implied correlation for contracts with one month to expiration. Figure 70 shows the distribution of historical correlation based on a 22 day rolling window. We see that the two distributions take on similar shapes, but are located at the opposite sides of the diagrams. Further, the distribution of historical correlation has a longer right tail compared to that of implied correlation, which barely has any tail at all.

Figure 71 and 72 consider a time frame of six months. The shape of the distributions is similar to that of those above. The only apparent difference is that both distributions have narrowed slightly, and the tail of the historical correlation distribution has more or less disappeared.

We can find no reasonable explanation for why the implied and historical distributions look like mirror images of each other. We have repeated all relevant calculations, and found no errors. We have also calculated both implied and historical correlation from a similar dataset obtained from Reuters, and get the same results – distributions that are located at opposite sides in the diagram.

Finally, we see that the theoretical distributions are a good fit to the empirical estimates, compared to what we have seen for some of the other currency pairs. This applies to both estimations windows.

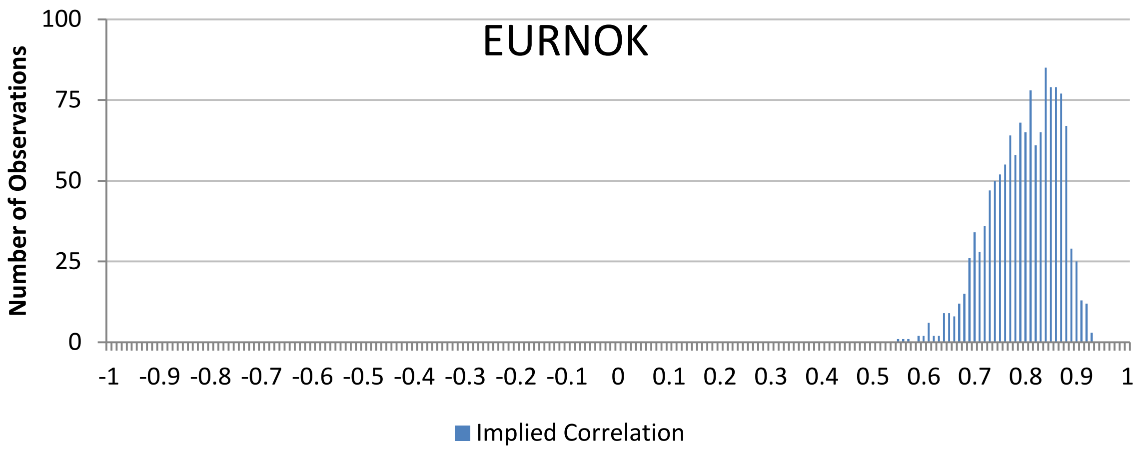

Figure 69: The empirical distribution of implied correlation based on implied volatility quotes for EURUSD, USDNOK and EURNOK option contracts with one month to expiration.

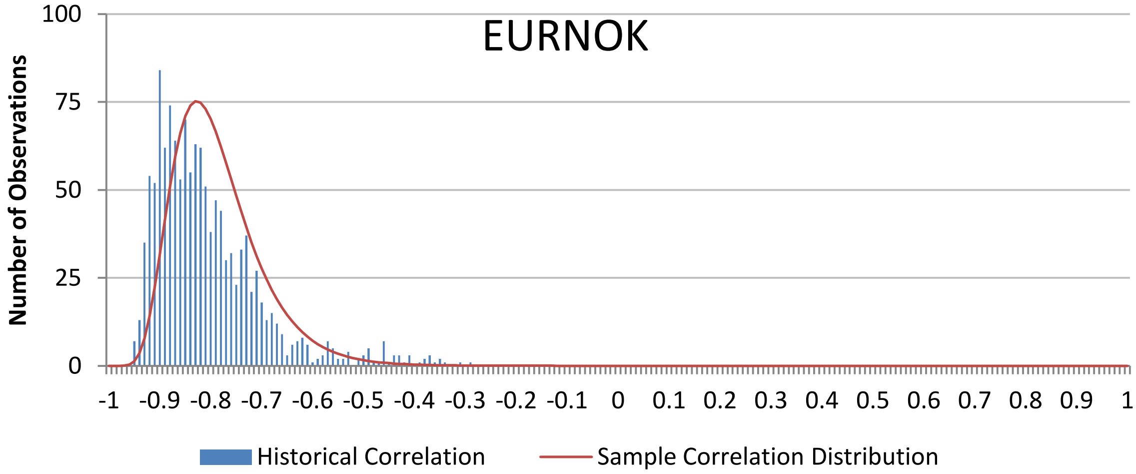

Figure 70: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURNOK and EURUSD. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

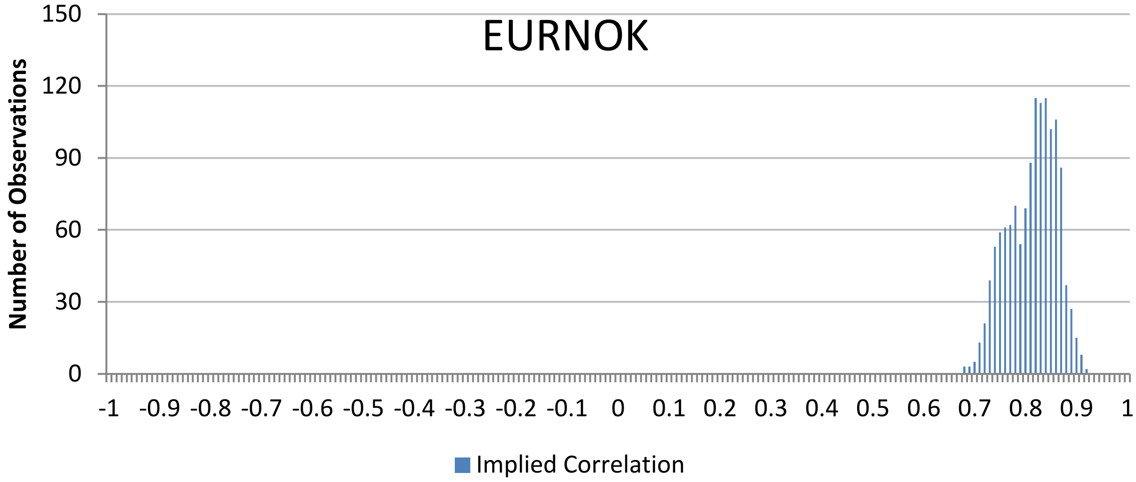

Figure 71: The empirical distribution of implied correlation based on implied volatility quotes for EURUSD, USDNOK and URNOK option contracts with six months to expiration.

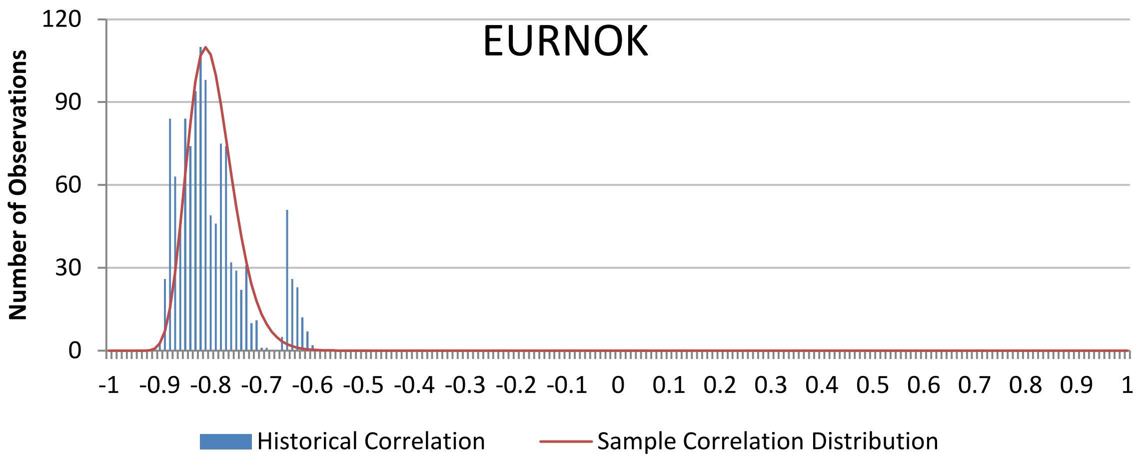

Figure 72: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURUSD and EURNOK. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

4.2 Section summary

In this section we will briefly discuss findings which appear general across the different currency pairs.

First, we find that historical correlation varies substantially relative to implied when we consider short-term correlation. This applies to nearly all of the currency pairs. Both correlation measures also tend to vary less erratically as we increase the time horizon under consideration. This effect is most pronounced for historical correlation.

We believe sampling error is causing noise in the historical correlation estimates based on a small number of observations; correlation estimates are known to be affected by sample size. We have also seen that rolling regression estimates are sensitive to outlying observations entering and exiting the estimation window.

As implied correlation is not an estimator, but rather a number calculated from option prices, sampling error cannot explain the high variability in correlations calculated from short-term option contracts. One plausible explanation for this pattern is that sudden changes in implied volatility levels influence implied correlation. It could also be that the market participants expect some mean reversion in the implied correlation level. Finally, we note that nearby option contracts are typically more heavily traded than contracts with longer time to expiration. It could be that the variability in short-term implied correlation is volume driven.

For several currency pairs we see that implied and historical correlation trend in the same direction, but this result is not universal. GBPJPY and EURNOK have a nearly constant spread between the two correlation measures. Perhaps most interesting is the pattern we see for EURNOK, where historical and implied correlation take on similar values – only with different signs. For EURJPY and USDNOK, implied correlation is high when historical correlation is low and vice versa.

In chapter 2.6 we presented the theoretical distribution of the sample correlation. We also discussed how the mean of correlation estimates is a somewhat intricate parameter, especially for high values of r. Our study confirms that the sample correlation distribution is indeed very sensitive to changes in the sample mean. $\bar{r}$ does not seem to describe the statistical properties of the curve in a consistent manner.

We also see several large discrepancies between the theoretical frequency distribution and our empirical results, which confirm that correlation is a highly unstable measure. As correlation is a higher moment, and therefore not a *visible* statistical measure, it is difficult to know if our estimates properly incorporate the dynamics of interdependence among the variables. We merely remark that although the distribution of the sample correlation coefficient is quite involved, it still gets outperformed by reality in terms of complexity.

Regarding implied correlation, it is important to note that for implied correlation to be a forward-looking indicator, i.e. an indicator that can incorporate a view of future market conditions, some level of *liquidity* is required. The information content of option prices is strongly related to trading volume. We see some evidence of this fact in our analysis; implied correlation appears inaccurate for the less liquid currency pairs, namely EURNOK and USDNOK.

In sum, we find that short-term correlations are more turbulent than long-term correlation. Both correlation measures also vary substantially over time, and most of our currency pairs exhibit wide empirical distributions.

5 The predictive accuracy of implied correlation

As previously mentioned, implied correlation can be thought of as the market's perception of future correlation. Hence, it is of interest to examine whether implied correlation can outperform the accuracy of correlation forecasts based on historical data. In the following we will carry out a simple comparative analysis, to evaluate the two correlation measures in terms of forecasting.

Implied an historical correlation will be compared in terms of their ability to forecast actual correlation over various future horizons, corresponding to specific option contract's times to expiration. We measure the actual correlation by rolling correlation estimates calculated from daily log-returns over a T-day period, labeled *"realized correlation".* Note that both implied and historical correlation are evaluated using the same measure of *ex post* correlation. This is in line with previous studies of the forecasting ability of implied correlation, and will make it easier to compare our results with similar research.

We will focus on the ability to predict the correlation level one, three and six months into the future. For implied correlation we will evaluate the forecasting accuracy of correlation estimates based on contracts with one month, three months and six months to maturity. We will compare these forecasts with the performance of rolling historical correlation estimates based on time periods of 22, 65 and 128 days.

Realized correlation will be estimated by 22, 65 and 128 day rolling windows. Our empirical findings in chapter 4 suggest that estimates based on relatively few discrete sampling points cause noisy estimates. That aside, we cannot see that there is an obvious way of selecting a "correct" *n*, and have thus decided to compare our forecasts with realized correlation based on rolling windows of different sizes. This should increase the information content of the analysis.

5.1 Forecast errors

To evaluate the forecasting accuracy of the different correlation measures we will analyze forecast errors. A straightforward way of evaluating forecasting performance is to check if the errors have zero mean. A forecast is said to be *biased* if the mean forecast error, *MFE,* is significantly different from zero (Enders 2010).

The forecast errors are here defined as:

$$e_{m,t} = \hat{\rho}_m\left(r_{xy}, r_{yz}\right)_{t,T} - \rho\left(r_{xy}, r_{yz}\right)_{t,T} \quad (4)$$

where $\hat{\rho}_m(\cdot)_{t,T}$ is the correlation forecast at time *t* generated by method *m*, and $\rho(\cdot)_{t,T}$ is the realized correlation at time *t*. The mean forecast error is calculated as:

$$MFE = \frac{\sum_{t=1}^{n} e_t}{n}$$

From (4) it follows that evaluating the unbiasedness of the forecast errors is analogous to the composite hypothesis we described in chapter 2.7.2, i.e. that the correlation coefficients of two independently sampled bivariate normal populations are identical. We hypothesize that the mean forecast errors are zero under the null, and compute the Z-statistic by dividing MFE_Z by a pooled standard error:

$$Z = \frac{MFE_Z}{\sqrt{\left[\frac{1}{n_1 - 1} + \frac{2-\rho_1^2}{2(n_1-1)^2}\right] + \left[\frac{1}{n_2 - 1} + \frac{2-\rho_2^2}{2(n_2-1)^2}\right]}}$$

where MFE_Z represents the mean forecast error of the Z-values, and *n* is the number of observations. If we reject the null hypothesis we say that the forecast errors are biased.

Further, we will evaluate the relative performance of the forecasts by comparing *root-mean squared forecast errors* (RMSFE) across the currency pairs. RMSFE-values are useful in this context, because they take into account that these variates take on both positive and negative values. Simply taking the mean of forecast errors *can* lead to false inference about the accuracy of the estimates, as errors with opposite signs will cancel each other out.

Note that we will once more use variables which are transformed using Fisher's Z transformation. First we compute RMSFE-values for the different forecasts. RMSFE is calculated as:

$$RMSFE = \sqrt{\frac{\sum_{t=1}^{n} e_t^2}{n}}$$

We proceed by identifying which forecasting method gives the lowest RMSFE score, and thus implies the best forecast in RMSFE terms.

Several methods have been suggested to determine if one RMSFE is significantly different from another. We will use a test proposed by Enders (2010:86). We have modified the test slightly to suit our purposes.

This test puts the larger of the two RMSFE in the numerator, and a standard recommendation is to us the F-statistic:

$$F = \frac{\sqrt{\sum_{t=1}^{n} e_{1t}^2}}{\sqrt{\sum_{t=1}^{n} e_{2t}^2}} = \frac{RMSFE_{Z1}}{RMSFE_{Z2}} \qquad (5)$$

where $RMSFE_{Z2}$ represents the prediction error of the forecast with the lowest RMSFE value. If we hypothesize that the two measures are equal in their ability to forecast future values, then (5) follows a standard F-distribution with (n, n) degrees of freedom. For this to hold the forecast errors need to:

1. Have zero mean and be normally distributed
2. Be serially uncorrelated
3. Be contemporaneously uncorrelated with each other

It is immediately obvious that this test is not ideal for our purposes. While our variates can be assumed to be approximately normally distributed, they might violate one or more of the other assumptions. Our main reason for employing this test is the intuitive interpretation of the test-statistics. If the two forecasts are identical, then F should equal unity. A large value of F would indicate that the mean error from the first model is substantially larger than that from the second, and vice versa. While this simple test has evident shortcomings, we find the no-nonsense way of interpreting the results appealing.

5.2 Analysis of forecast errors

In this section we will compare the predictive accuracy of implied and historical correlation. The test statistics are summarized in table 1b through 9b in appendix B.

The mean forecast errors indicate that historical correlation is an unbiased forecast of future realized correlation; we can never reject the null of $MFE = 0$. The MFE-scores of historical correlation are generally low. This can, at least partially, be explained by historical and realized correlation being calculated in the same manner.

For implied correlation we can reject the null of unbiased forecasts in several cases. This is not surprising when we consider the differences we found between historical and implied correlation in chapter 4. Again we note that our measure of realized correlation is actually lagged values of the historical estimates.

Further, we see that that the mean forecast errors of implied correlation are generally larger in absolute value, compared to the errors of historical correlation. It is however not recommended to draw conclusions based on MFE-values alone. Given that the forecasting errors take on both positive and negative values, the mean forecasting error can easily be zero even if the individual errors are substantial. We have analyzed scatterplots of the individual forecast errors to check whether this “averaging” is reducing the size of the MFE-values. We find that they do not give an accurate description of the precision of the correlation estimates, and to illustrate this point we have included two of the scatterplots below. USDJPY and GBPJPY have identical MFE-scores, namely 0.026. In figure 73 and 74 we see that the variation in the forecast errors of GBPJPY is much larger than the variation of the individual errors of USDJPY.

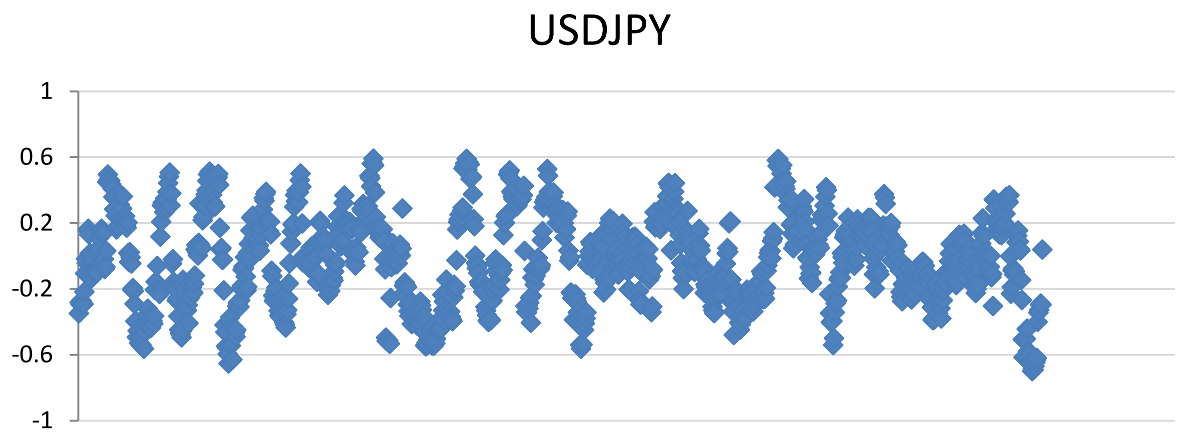

Figure 73: The forecast errors of historical correlation estimates based on a time period of 65 days. The forecasting horizon is one month.

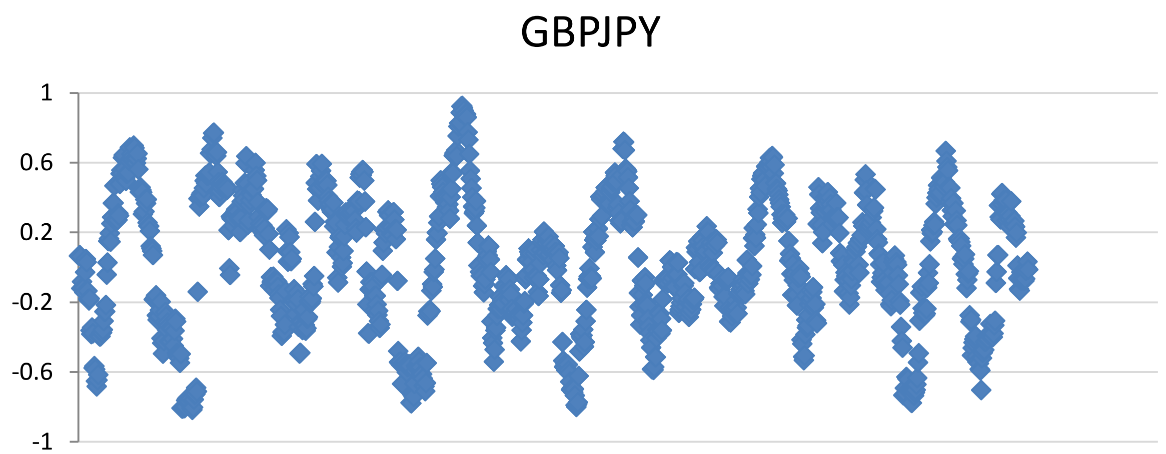

Figure 74: The forecast errors of historical correlation estimates based on a time period of 65 days. The forecasting horizon is one month.

We will now discuss the root mean squared forecast errors and the relative forecasting performance of implied and historical correlation. In table 1b through 9b (appendix B), the lowest RMSFE-scores for each forecast is marked in bold.

Historical correlation calculated from 22 days performs poorly in terms of forecasting future correlation levels. This result is general across the analysis. It seems that using only 22 observations to forecast correlation cause noisy estimates, which in turn obscure the information content of the forecasts. From now on, if we suggest that historical correlation is outperforming implied correlation in terms of forecasting, we will refer to historical correlation calculated from 65 and 128 day rolling windows.

In chapter 4 we found that there is substantial divergence between the empirical estimates of implied and historical correlation. For some currency pairs this was visible as a “spread” when we plotted the correlation levels against time. This effect was most pronounced for GBPJPY, EURJPY and EURNOK. For these currency pairs we find that historical correlation is a significantly better forecast than implied correlation. The null of no difference between the RMSFE-values is rejected, and large F-statistics indicate that these findings are robust.

For USDJPY, USDJPY2, EURUSD and EURUSD2 we find that implied correlation more often than not provide the lowest RMSFE-values. These results are often significant, i.e. we can reject the null of equal forecasting ability. Note that even though we can reject the null, the F-statistics are in most cases close to unity. This suggests that these results should be taken with a pinch of salt.

Finally, we have divided the sample into four sub-periods, to examine whether the forecasting performance of the correlation measures changes over time. The sub-samples consist of the years 2006-2007, 2008-2009 and 2010-2011.

We find that for the time period 2006-2007, historical correlation is generally the best predictor of future realized correlation. For the other sub-periods there is no such “generally best” indicator. Another interesting result is that the time period that presumably would have been most affected by the financial crisis, namely 2007-2008, does not appear different in any way compared to the other subsamples.

Overall, our analysis indicates that when either implied or historical correlation performs significantly better than the other in terms of RMSFE-scores, this relationship is relatively stable through time. We also see that the forecasting accuracy of the two correlation measures vary across the currency pairs, but much less so across the different forecasting horizons. Based on this very simple analysis it seems equally difficult to predict future correlation one month and three months ahead. The forecast errors also seem to confirm the tendencies we saw in chapter 4, namely that the information content of implied correlations vary across the different currency pairs.

Previous research (see e.g. Campa and Chang 1998, Castrén, and Mazzotta 2005) conclude that implied correlation adds to the forecasting accuracy of time-series based forecasts. These results are however based on evaluation criteria (mainly Diebold-Mariano statistics and encompassing regressions) which differ markedly from those we have used in this study. This means that their results are not directly comparable to ours.

Our results are more in line with the findings of Walter and Lopez (2000). They state that the forecasting performance of implied correlation varies substantially across different currency pairs, and conclude that the predictive accuracy of implied correlation needs to be evaluated empirically. While Walter and Lopez rely on more sophisticated statistical techniques than our simple study, we arrive at the same main conclusion.

6 Implications, limitations and discussion

We have found that implied correlation is a more stable parameter than historical correlation, especially for short term correlations. Our analysis further shows that the term structure of correlation is far from constant.

The empirical distribution of both implied and historical correlation narrows markedly for longer time horizons. As implied correlation is not an estimator, but rather a number inverted from option prices, sampling error cannot explain the high variability in correlations calculated from short-term option contracts.

We believe several factors contribute to the variation in short-term implied correlations. First, it seems plausible that sudden changes in implied volatility levels influence implied correlation. Second, it could be that market participants incorporate expectations of mean reversion in their price quotes. Finally, nearby option contracts typically have larger turnover than contracts with longer time to maturity. This can introduce volume driven variability in short-term implied correlation, as we have argued that supply and demand will influence implied correlation levels.

Existing literature suggests that correlations are highly unstable, and hence difficult to estimate and/or predict. Our empirical study confirms this notion; the behavior of correlation appears complex and varies markedly across the different currency pairs. For some pairs, like e.g. USDJPY and EURUSD, we see that implied and historical correlation trend in the same direction. For other pairs, most notably EURJPY and USDNOK, the two correlation measures appear to repel each other when we plot them against time. Most puzzling is the empirical distribution of EURNOK, where the two correlation measures cluster in opposite sides of the diagrams.

Our analysis further suggests that the forecasting performance of implied correlation varies markedly across currency pairs. We arrive at a conclusion which is similar to that of Walter and Lopez (2000), namely that the predictive ability of implied correlation needs to be evaluated empirically. The prediction power of this measure is not uniform across different currency pairs.

A large part of this thesis has been devoted to discussing the distribution of the sample correlation coefficient, derived by Fisher in 1915. This has been useful because it forces us to think carefully about what a correlation value actually represents. We have shown that the descriptive statistics commonly reported in empirical studies are not well suited to describe the statistical properties of correlation. While studies such as Campa and Chang (1998) Castrén and Mazotta (2005) find that implied correlation provides the most accurate forecasts of future realized correlation, we believe these studies would have benefited from a more thorough review of the theoretical framework underlying their analysis. The same can be said about other correlation studies as well.

We have compared the theoretical distribution of the correlation coefficient with the distribution of historical estimates, and see several large discrepancies between the theoretical frequency distribution and our empirical results. We have stressed that correlation is a highly unstable measure, and even though Fisher's distribution is a brilliant piece of mathematics, it falls short of describing the complexity of historical correlation in the OTC FX market.

Today it is commonly accepted that implied volatilities are useful in terms of forecasting the future realized volatility of the underlying financial variables. Our study indicates that it is far from obvious that the assumptions made about implied volatility as a useful economic indicator is directly transferable onto implied correlations. Regardless, while implied correlation just like implied volatility is influenced by other factors than the market's view on future market conditions, most notably supply and demand, it has the ability to incorporate additional information relative to time series forecasts.

To evaluate the predictive accuracy of implied correlation, more research is needed. While several papers have addressed the information content of cross-option prices, these studies rely on test statistics that are not entirely accurate in the context of correlations. It would be useful to find a consistent way of evaluating correlation measures in relation to financial variables, as well as a common standard of reporting standard errors. Finally, it would interesting to see if there exist a consistent way of reporting the central tendency and variability of correlation estimates.

References

Andersson, M., E. Krylova, and S. Vähämaa, 2008, Why Does the Correlation Between Stock and Bond Returns Vary Over Time? *Applied Financial Economics: February 2008, Vol. 18, No. 2, 2008.* URL: http://ssrn.com/abstract=1740666 (18.02.12)

Ayache, E., 2006, What Is Implied by Implied Volatility?, *Wilmott Magazine: January, pp. 28-35.*

Black, F., and M. Scholes, 1973, The Pricing of Options and Corporate Liabilities, *The Journal of Political Economy: May - June 1973, Vol. 81, No. 3, pp. 637-654.*

Bodurtha, J. N. and Q. Shen, 1994, Historical and Implied Measures of "Value at Risk": The DM and Yen Case, URL: http:--faculty.msb.edu-bodurthj-research-implcrr7.pdf (14.01.12)

Boothe, P., and D. Glassman, 1987, The Statistical Distribution Of Exchange Rates, *Journal of International Economics: May 1987, pp. 297-319.*

Campa J. M. and K. Chang, 1998, The Forecasting Ability of Correlations Implied in Foreign Exchange Options, *Journal of International Money and Finance: December 1998, Vol. 17, Issue 6, pp. 855-880.*

Carr, P. and L. Wu, 2007, Stochastic Skew in Currency Options, *Journal of Financial Economics 86, pp. 213–247.*

Castagna, A., 2010, FX Options and Smile Risk, West Sussex: John Wiley & Sons.

Castrén, O. and S. Mazzotta, 2005, Foreign Exchange Option and Return Based Correlation Forecasts, *ECB Working Paper No 447.* URL: http:--www.ecb.int-pub-pdf-scpwps-ecbwp447.pdf (13.01.12)

Clark, I.J., 2011, Foreign Exchange Option Pricing: A Practitioners Guide, West Sussex: John Wiley & Sons.

Da Fonseca, J., M. Grasselli and C. Tebaldi, Option Pricing When Correlations Are Stochastic: an Analytical Framework, *Review of Derivatives Research: May 2007, Vol. 10, No. 2, pp. 151-180.*

Derman, E. and N. N. Taleb, 2005, The Illusions of Dynamic Replication, *Quantitative Finance: August 2005, Vol. 5, No. 4, pp. 323–326.*

Enders, W., 2010, Applied Econometric Time Series, United States: John Wiley & Sons.

Engle, R., 2002, Dynamic conditional correlation - A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models, *Journal of Business and Economic Statistics: July 2002, Vol. 20, Issue 3, pp. 339-350.*

Faller, A. J., 1981, An Average Correlation Coefficient, *Journal of Applied Metereology: Vol. 20, pp. 203-205,* URL: http://journals.ametsoc.org/doi/pdf/10.1175/1520-0450(1981)020%3C0203:AACC%3E2.0.CO;2 (07.04.12)

Fisher, R. A., 1915, Frequency Distribution of the Values of the Correlation Coefficient in Samples From an Indefinitely Large Population, *Biometrika: May 1915, Vol. 10, No. 4, pp. 507-521.*

Fisher, R.A., 1921, On the “probable error” of a coefficient of correlation deduced from a small sample. *Metron Vol.1, no. 4, pp. 2-32.*

Fox, J., 1997, Applied Regression Analysis – Linear Models and Related Methods, United States: Sage Publications, Inc.

Garcia, E., 2010, A Tutorial on Correlation Coefficients. URL:

http://www.miislita.com/statistics/on-the-non-additivity-correlation-coefficients.pdf (01.04.12)

Garcia, E., 2011, On the Non-Additivity of Correlation Coefficients. URL

http://www.miislita.com/statistics/on-the-non-additivity-correlation-coefficients.pdf (01.04.12)

Garman, M. B., and S. W. Kohlhagen, 1983, Foreign Currency Option Values, *Journal of International Money and Finance: December 1983, Vol. 2, Issue 3, pp. 231–237.*

Grabbe, J., 1983, The Pricing of Call and Put Options on Foreign Exhange, *Journal of International Money and Finance: December 1983, Vol. 2, Issue 3, pp. 239–254.*

Haug, E. G., 1996. Implisitt Korrelasjon i Valutamarkedet, *Beta: Vol. 1, pp. 39-43.*

Haug, E. G., 2007a, Derivatives: Models on Models, West Sussex: John Wiley & Sons.

Haug, E. G., 2007b, The Complete Guide to Option Pricing Formulas, 2nd edition, United States: McGraw-Hill.

Haug, E.G., S. Frydenberg and S. Westgaard, 2010, Distribution and Statistical Behavior of Implied Volatilities, *Business Valuation Review: Winter 2010, Vol. 29, No. 4, pp. 186-199.*

Haug, E. G., N. N. Taleb, 2011, Option Traders Use (very) Sophisticated Heuristics, Never the Black–Scholes–Merton Formula, *Journal of Economic Behavior and Organization, Vol. 77, No. 2, pp. 97–106.*

Johnston, K. and E. Scott, 1999, The Statistical Distribution of Daily Exchange Rate Price Changes: Dependent vs Independent Models, *Journal of Financial and Strategic Decisions: Fall 1999 Vol. 12. No. 2, pp. 39-49.*

Joshi, M. S., 2003, The Concepts and Practice of Mathematical Finance, 2nd edition, Cambridge: Cambridge University Press.

Kamtchueng, C., 2010, Uncertain Correlation Model, *Wilmott Journal: December 2010, Vol. 2, Issue 6, pp. 317–349.*

Lopez, J.A. and C. Walter, 2000, Is Implied Correlation Worth Calculating? Evidence from Foreign Exchange Options and Historical Data, *Journal of Derivatives, Spring 2000, Vol. 7, No. 3: pp. 65-81.*

Ma, J., 2009, A Stochastic Correlation Model with Mean Reversion for Pricing Multi-Asset Options, *Asia-Pacific Financial Markets, Vol. 16, No. 2, pp. 97-109.*

Margrabe W., 1978, The Value of an Option to Exchange One Asset For Another, *The Journal of Finance: March 1978, Vol. 33, No. 1, pp. 177-186.*

McDonald, R. L., 2006, Derivatives Markets, 2nd edition, United States of America: Pearson Education, Inc.

Merton, R., 1973, Theory of Rational Option Pricing, *The Bell Journal of Economics and Management Science: Spring 1973, Vol. 4, No. 1, pp.141-183.*

Natenburg, S., 1994, Option Volatility & Pricing - Advanced Trading Strategies and Techniques, United States: McGraw-Hill.

Rodgers, J. L., and W. A. Nicewander, 1988, Thirteen Ways to Look at the Correlation Coefficient, *The American Statistician, Vol. 42, No. 1, pp. 59-66.*

Siegel, A., 1997, International Currency Relationship Information Revealed by Cross-Option Prices, *Journal of Futures Markets, June 1997, Vol. 17, Issue 4, pp. 369–384.*

Silver, N. C. and W. Dunlap, 1987, Averaging Correlation Coefficients: Should Fisher's Z Transformation Be Used?, *Journal of Applied Psychology, Vol. 72, No. 1, pp. 146-148.*

Stuart, A. and J. K. Ord, 1994, Kendall's Advanced Theory of Statistics - Volume 1: Distribution Theory, 6th edition, London: St. Edmundsbury Press Limited.

Stuart, A., J. K. Ord and Steven Arnold, 1999, Kendall's Advanced Theory of Statistics - Volume 2A: Classical Inference and the Linear Model, 6th edition, London: Arnold, a Member of The Hodder Headline Group.

Tompkins, R., 1994, Options Explained 2, 2nd edition, Cornwall: Palgrave Macmillan.

Taleb, N., 1997, Dynamic Hedging – Managing Vanilla and Exotic Options, United States: John Wiley & Sons.

Triana, P., 2009, Lecturing Birds on Flying - Can Mathematical Theories Destroy the Financial Markets?, New York: John Wiley & Sons.

Tsay, R. S., 2010, Analysis of Financial Tiime Series, New Jersey: John Wiley & Sons.

Tu, Y.K, G. T. H. Ellison, M. S. Gilthorpe, 2006, Growth, current size and the role of the 'reversal paradox' in the foetal origins of adult disease: an illustration using vector geometry, *Epidemiologic Perspectives & Innovations.* URL:

http://www.epi-perspectives.com/content/3/1/ (01.04.12)

Van den Berg, T., 2012, The Distribution of The Sample Correlation. URL: http://www.sitmo.com/article/the-distribution-of-the-sample-correlation/ (18.03.12)

Weisstein, E. W., Correlation Coefficient, Bivariate Normal Distribution, *MathWorld Wolfram.* URL: http://mathworld.wolfram.com/CorrelationCoefficientBivariateNormalDistribution.html (08.03.12)

Wilmott, P., 2006, Paul Wilmott on Quantitative Finance, 2nd edition, West Sussex: John Wiley & Sons.

Wooldridge, J. M., 2003, Introductory Econometrics - A Modern Approach, 2nd edition, United States: Thomson South-Western.

Wystup, U., 2008, Foreign Exchange Symmetries, Contribution to Encyclopedia of Quantitative Finance, Chichester: John Wiley & Sons, pp.752-759.

Wystup, U. and D. Reiswich, 2010a, A Guide to FX Options Quoting Conventions, *Journal of Derivatives: Winter 2010, Vol. 18, No. 2: pp. 58-68.*

Wystup, U. and D. Reiswich, 2010b, FX Volatility Smile Construction, *submitted for publication.* URL: http:--www.mathfinance2.com-MF_website-UserAnonymous-Company-papers-CPQF_Arbeits20.pdf (13.01.12)

Zhang, P. G., 1998, Exotic Options - A Guide to Second Generation Options, 2nd edition, Singapore: World Scientific Publishing Co. Pte. Ltd.

Zimmerman, D. W., B. D. Zumbo and R. H. Williams, 2003, Bias in estimation and hypothesis testing of correlation. Psicológica (2003), 24, pp. 133-158. URL:

http://www.uv.es/revispsi/articulos1.03/9.ZUMBO.pdf (14.03.12)

7 Appendix A

7.1 Time horizon: 1 month

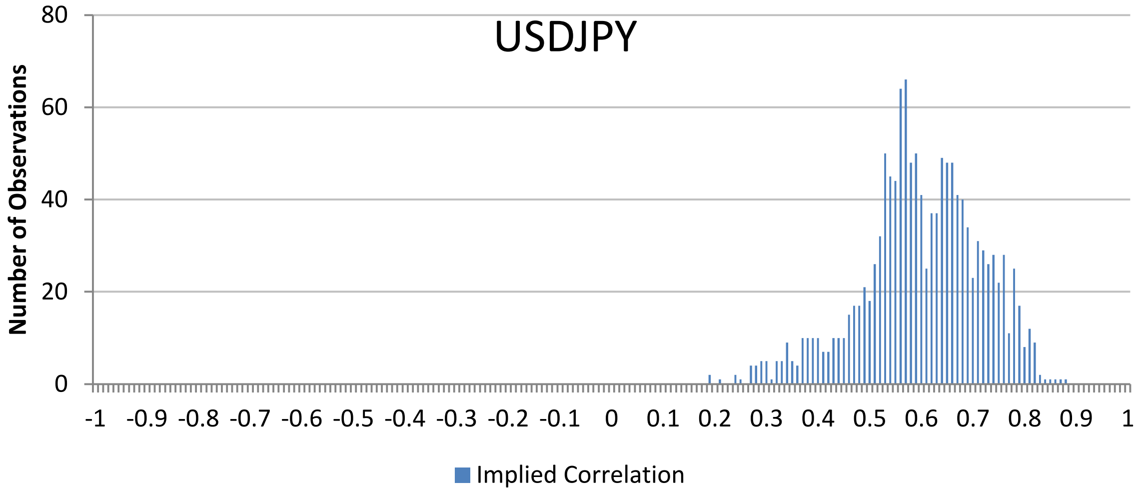

Figure 1a: The empirical distribution of implied correlation based on implied volatility quotes for GBPJPY, GBPUSD and USDJPY option contracts with 1 month to expiration.

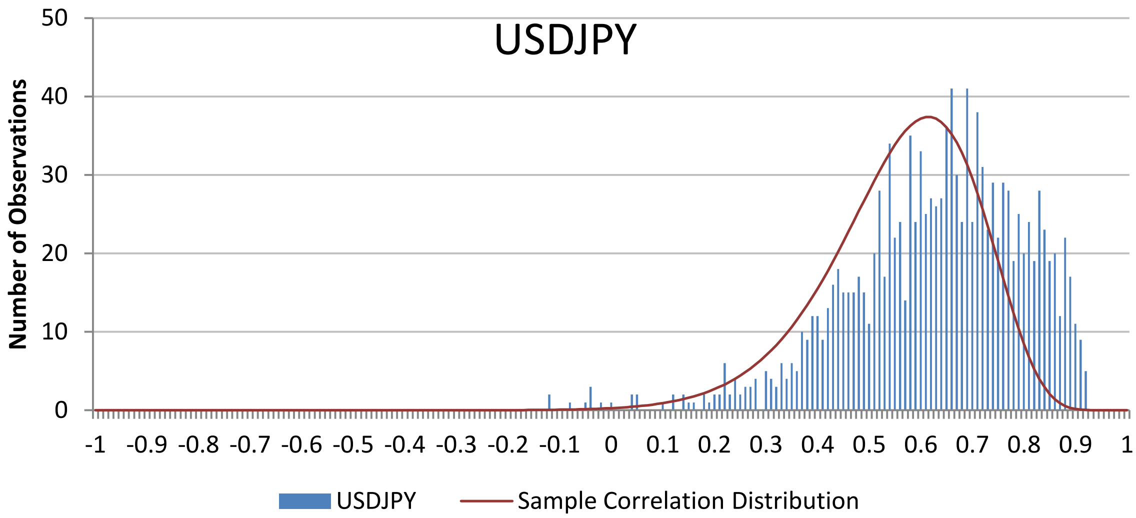

Figure 2a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for GBPJPY and GBPUSD. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.594	0.607	0.624	0.655
Mode	0.57	0.57	0.69	0.69
Minimum	0.186	0.186	-0.125	-0.125
Maximum	0.879	0.879	0.917	0.917
Number of observations	1326	1326	1296	1296

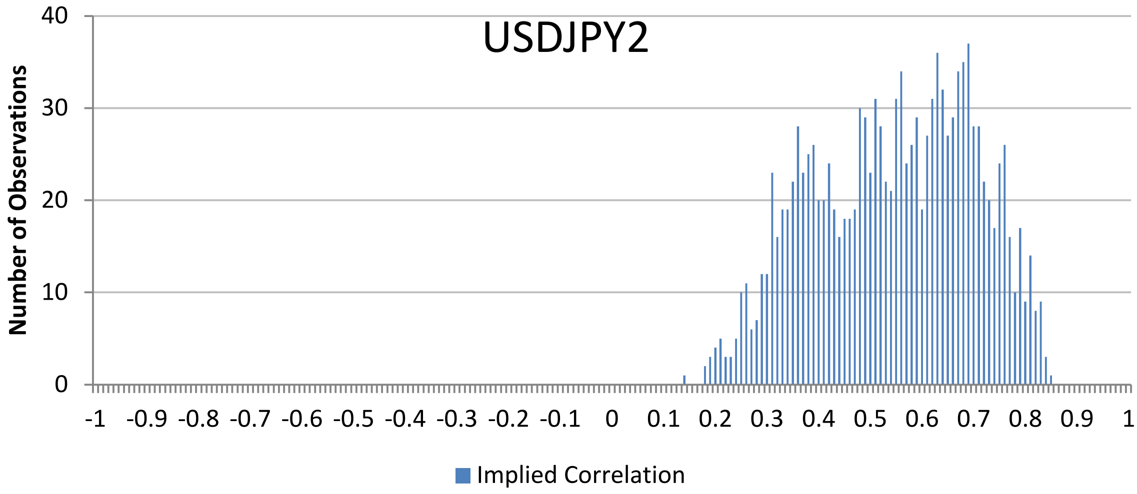

Figure 3a: The empirical distribution of implied correlation based on implied volatility quotes for EURJPY, EURUSD and USDJPY option contracts with 1 month to expiration.

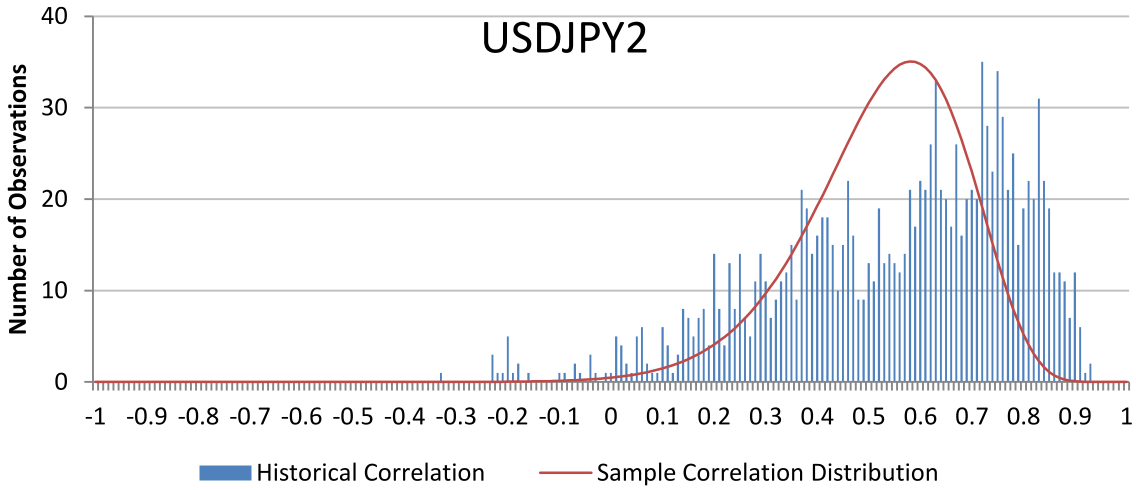

Figure 4a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURJPY and EURUSD. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, which is based on the mean of the actual data set.

Descriptive statistics: USDJPY2				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.541	0.560	0.547	0.591
Mode	0.69	0.69	0.72	0.72
Minimum	0.135	0.135	-0.338	-0.338
Maximum	0.844	0.844	0.924	0.924
Number of observations	1326	1326	1296	1296

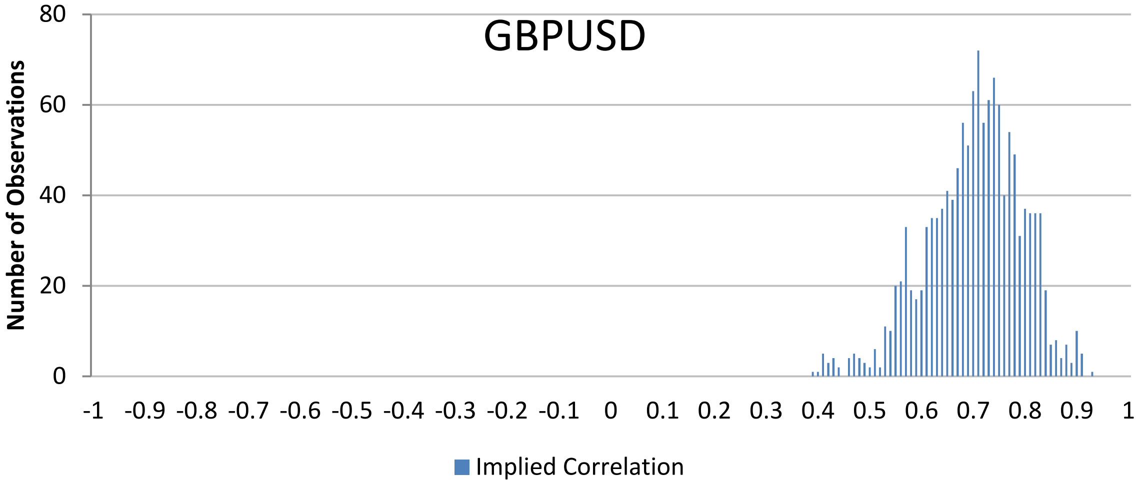

Figure 5a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPJPY and GBPUSD option contracts with 1 month to expiration.

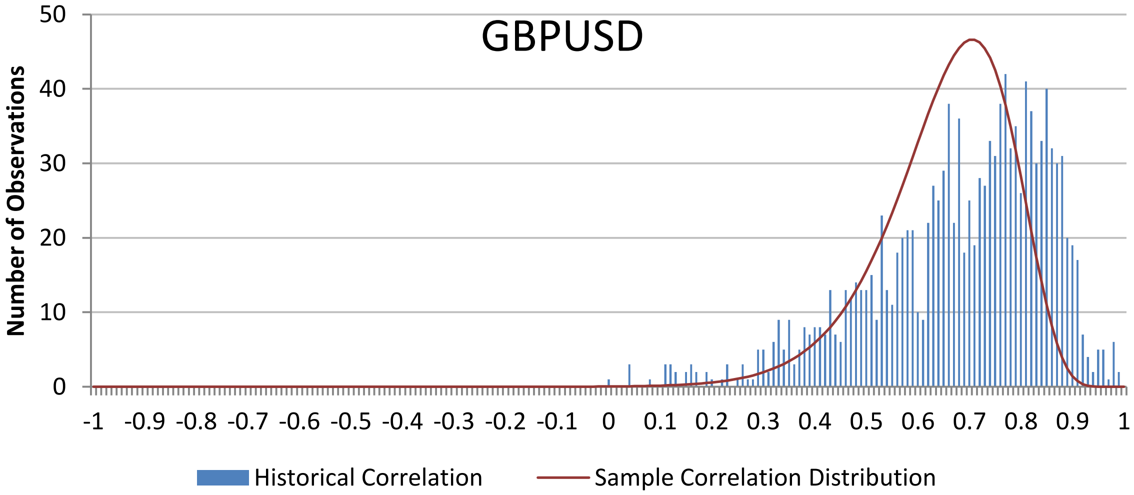

Figure 6a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPJPY. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: GBPUSD				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.697	0.709	0.675	0.714
Mode	0.71	0.71	0.77	0.77
Minimum	0.385	0.385	-0.002	-0.002
Maximum	0.921	0.921	0.984	0.984
Number of observations	1326	1326	1296	1296

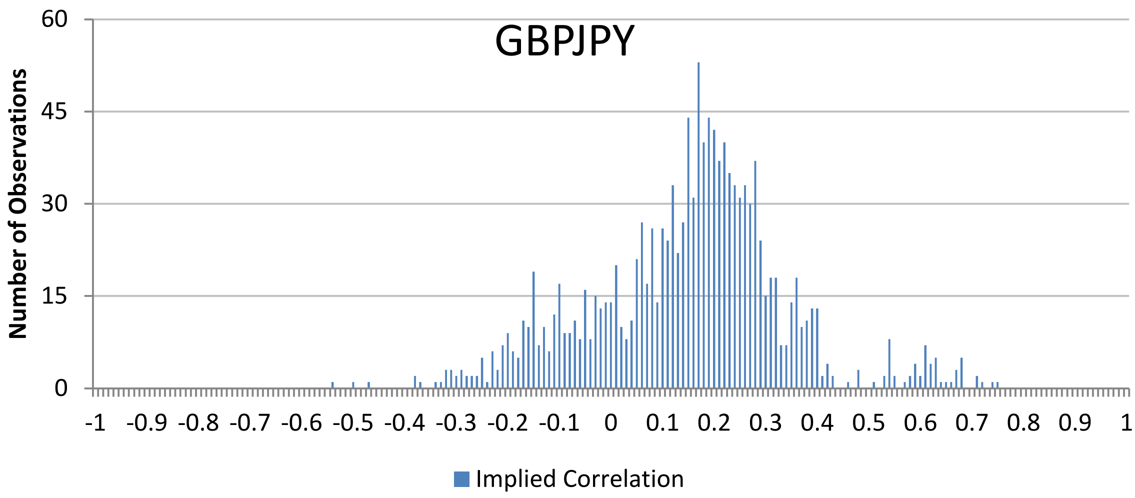

Figure 7a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPUSD and GBPJPY option contracts with 1 month to expiration.

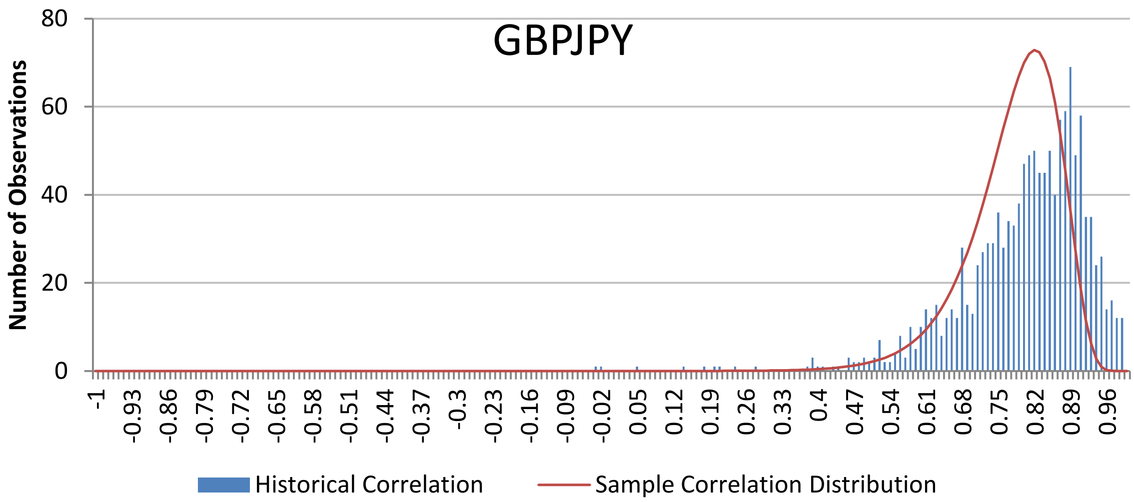

Figure 8a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPUSD. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: GBPJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.144	0.150	0.797	0.828
Mode	0.17	0.17	0.89	0.89
Minimum	-0.548	-0.548	-0.032	-0.032
Maximum	0.744	0.744	0.985	0.985
Number of observations	1326	1326	1296	1296

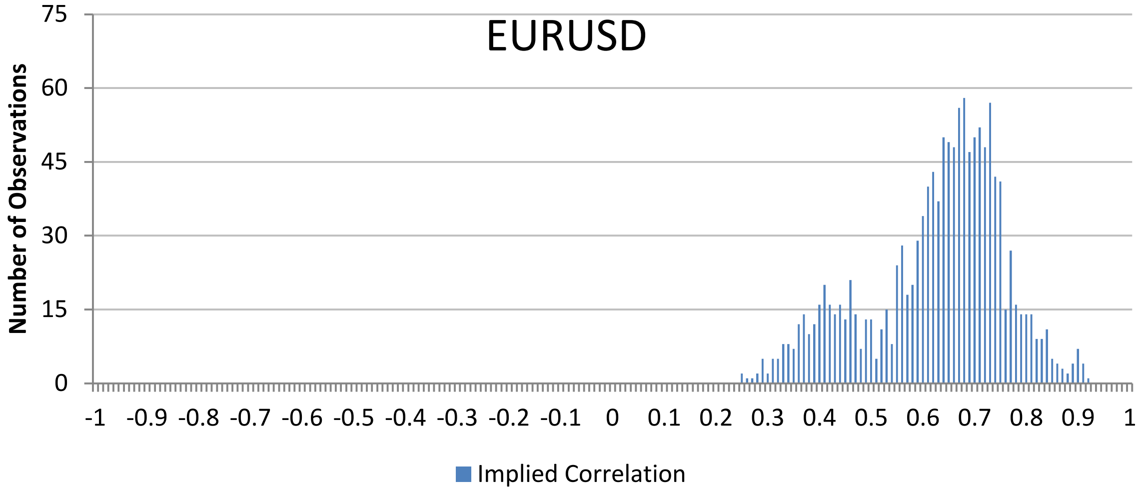

Figure 9a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURJPY and EURUSD option contracts with 1 month to expiration.

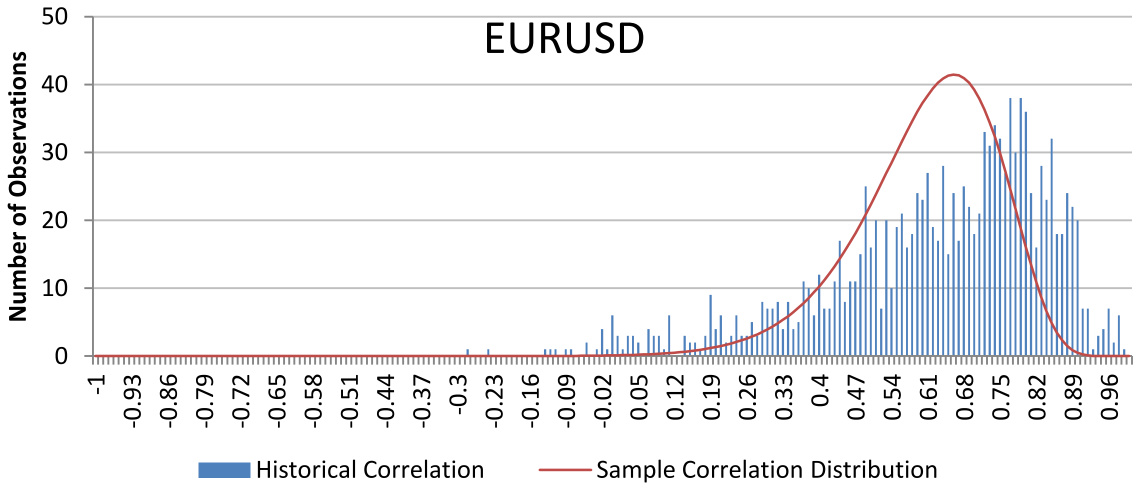

Figure 10a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and EURJPY. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURUSD				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.623	0.640	0.622	0.670
Mode	0.73	0.73	0.77/0.79	0.77/0.79
Minimum	0.240	0.240	-0.288	-0.288
Maximum	0.913	0.913	0.980	0.980
Number of observations	1326	1326	1296	1296

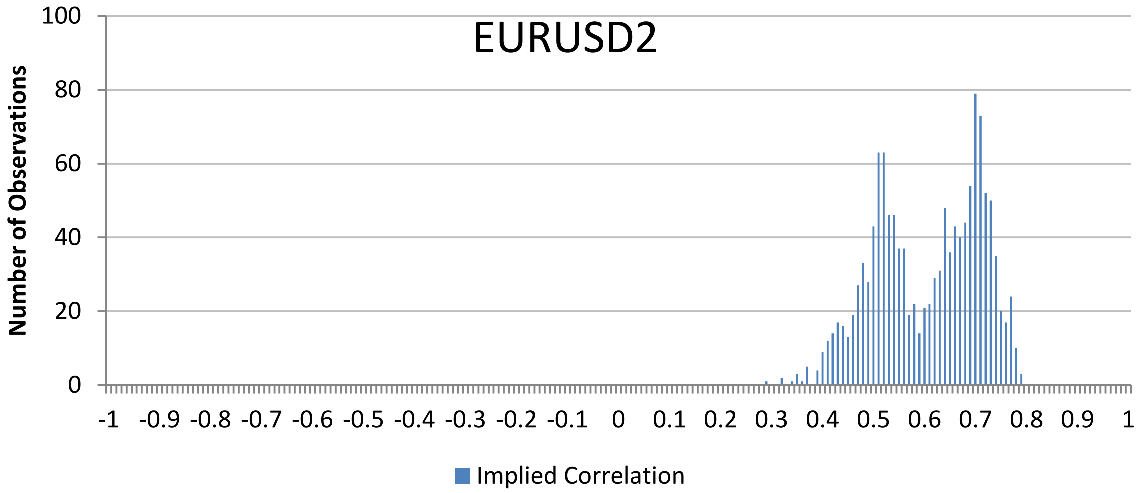

Figure 11a: The empirical distribution of implied correlation based on implied volatility quotes for USDNOK, EURNOK and EURUSD option contracts with 1 month to expiration.

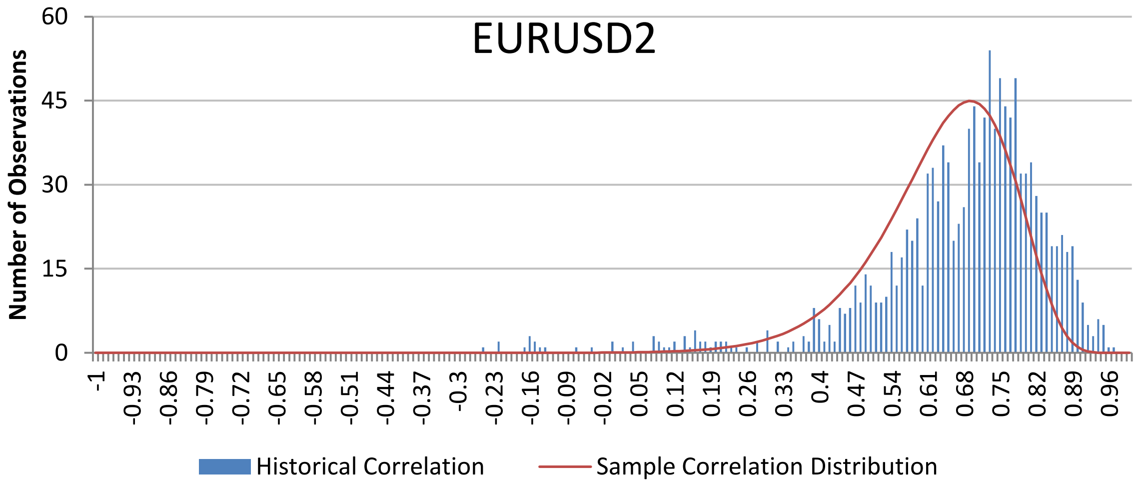

Figure 12a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDNOK and EURNOK. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURUSD2				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.601	0.611	0.670	0.701
Mode	0.70	0.70	0.75	0.75
Minimum	0.286	0.286	-0.254	-0.254
Maximum	0.787	0.787	0.960	0.960
Number of observations	1326	1326	1296	1296

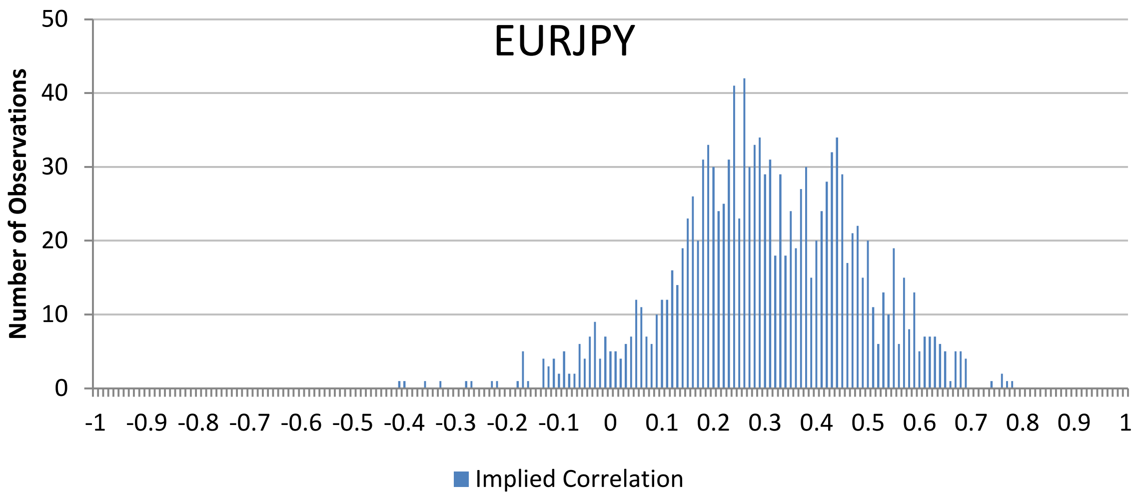

Figure 13a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURUSD and EURJPY option contracts with 1 month to expiration.

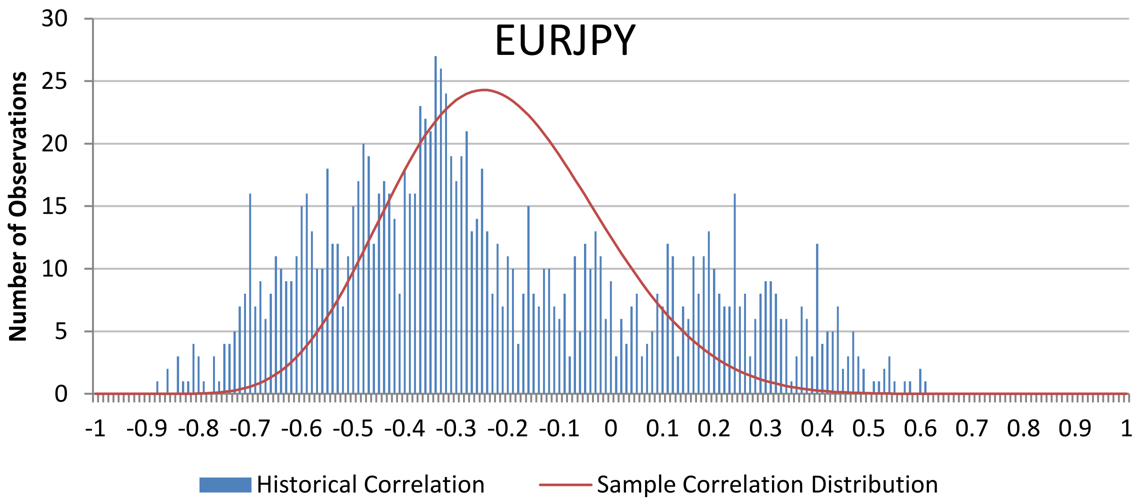

Figure 14a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and EURUSD. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.295	0.306	-0.226	-0.253
Mode	0.26	0.26	-0.34	-0.34
Minimum	-0.418	-0.418	-0.882	-0.882
Maximum	0.778	0.778	0.603	0.603
Number of observations	1326	1326	1296	1296

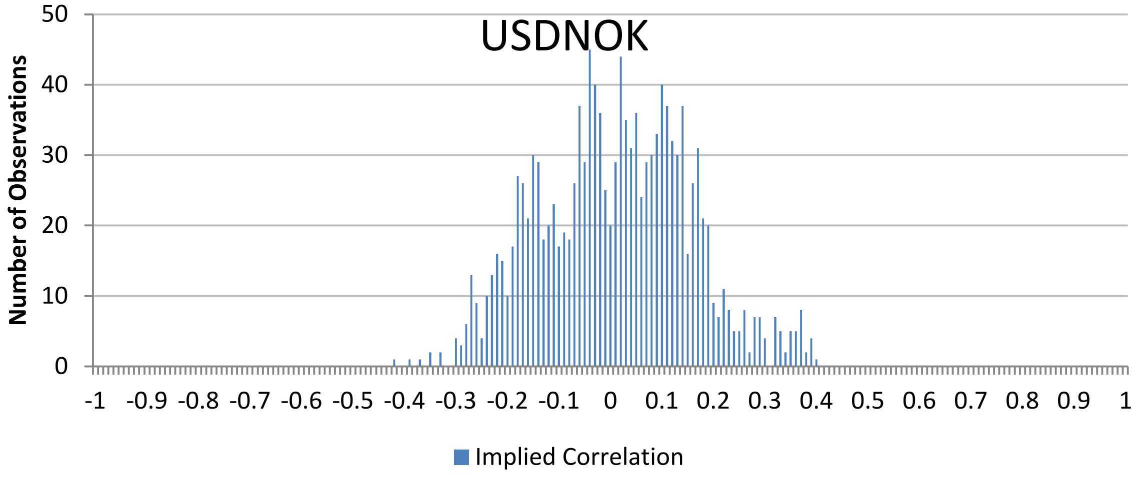

Figure 15a: The empirical distribution of implied correlation based on implied volatility quotes for EURNOK, EURUSD and USDNOK option contracts with 1 month to expiration.

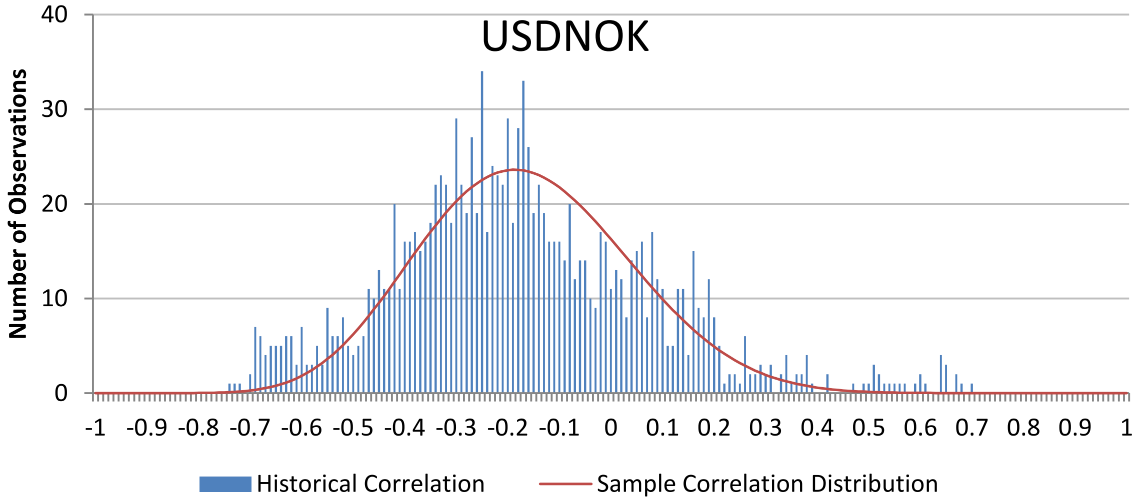

Figure 16a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURNOK and EURUSD. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDNOK				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.005	0.005	-0.180	-0.191
Mode	-0.04	-0.04	-0.25	-0.25
Minimum	-0.425	-0.425	-0.745	-0.745
Maximum	0.399	0.399	0.691	0.691
Number of observations	1326	1326	1296	1296

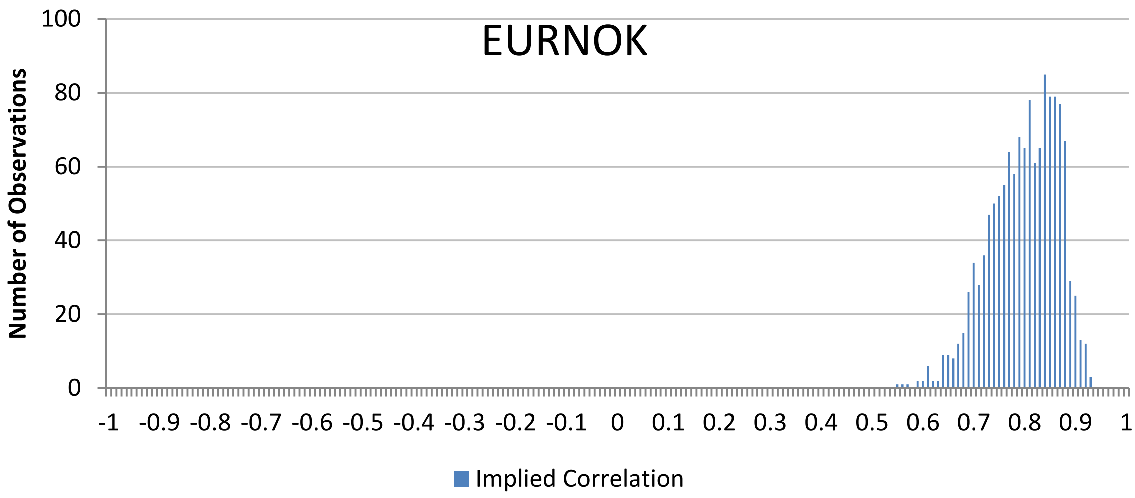

Figure 17a: The empirical distribution of implied correlation based on implied volatility quotes for EURUSD, USDNOK and EURNOK option contracts with 1 month to expiration.

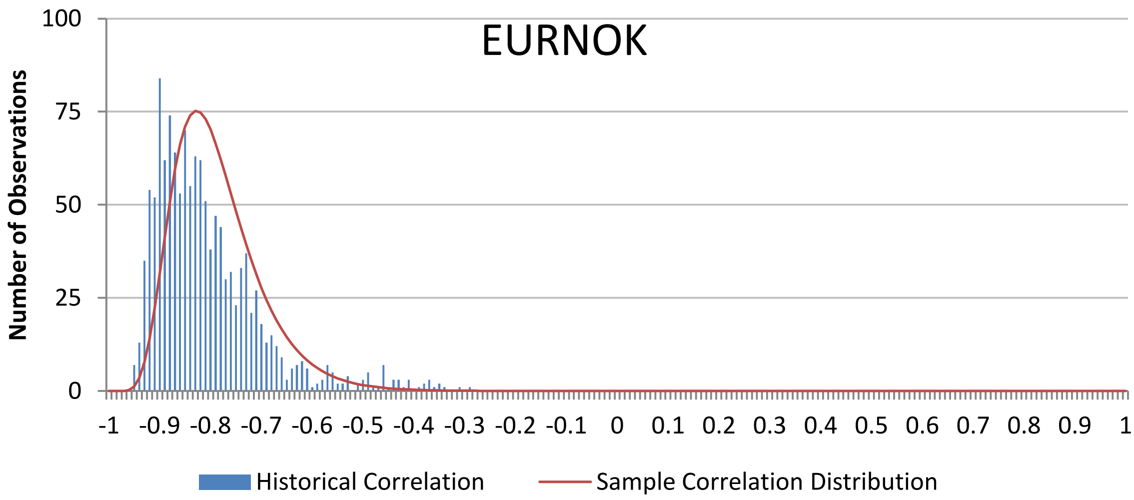

Figure 18a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURNOK and EURUSD. All estimates are based on a 22 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURNOK				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.793	0.803	-0.811	-0.834
Mode	0.84	0.84	-0.90	-0.90
Minimum	0.548	0.548	-0.959	-0.959
Maximum	0.921	0.921	-0.293	-0.293
Number of observations	1326	1326	1296	1296

7.2 Time horizon: 2 months

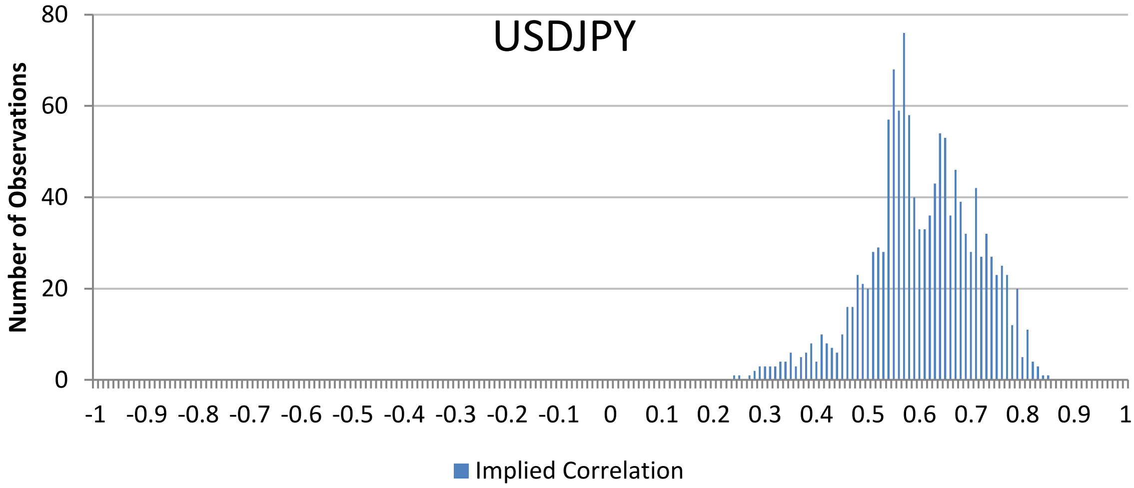

Figure 19a: The empirical distribution of implied correlation based on implied volatility quotes for GBPJPY, GBPUSD and USDJPY option contracts with 2 months to expiration.

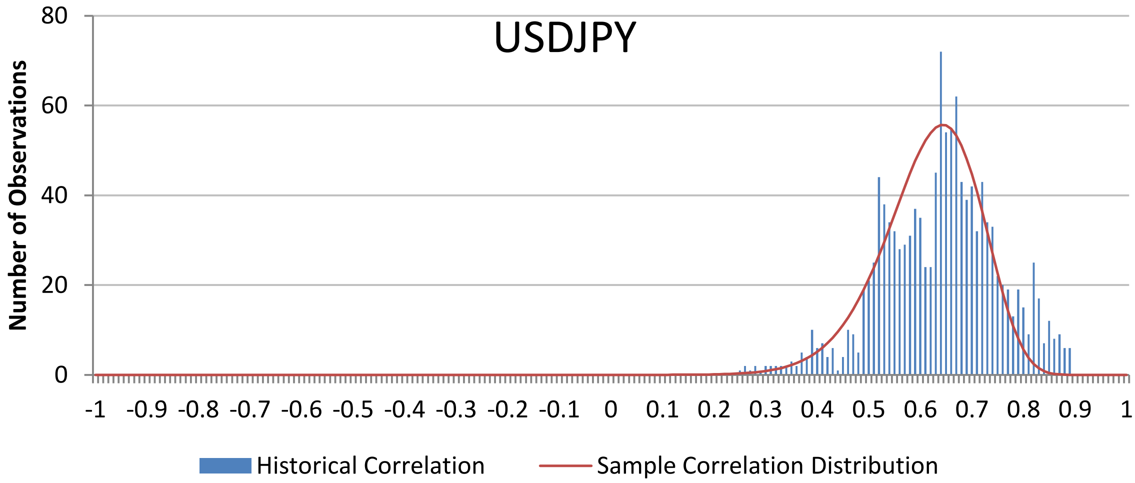

Figure 20a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for GBPJPY and GBPUSD. All estimates are based on a 43 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.600	0.611	0.633	0.648
Mode	0.57	0.57	0.64	0.64
Minimum	0.231	0.231	0.245	0.245
Maximum	0.842	0.842	0.888	0.888
Number of observations	1326	1326	1275	1275

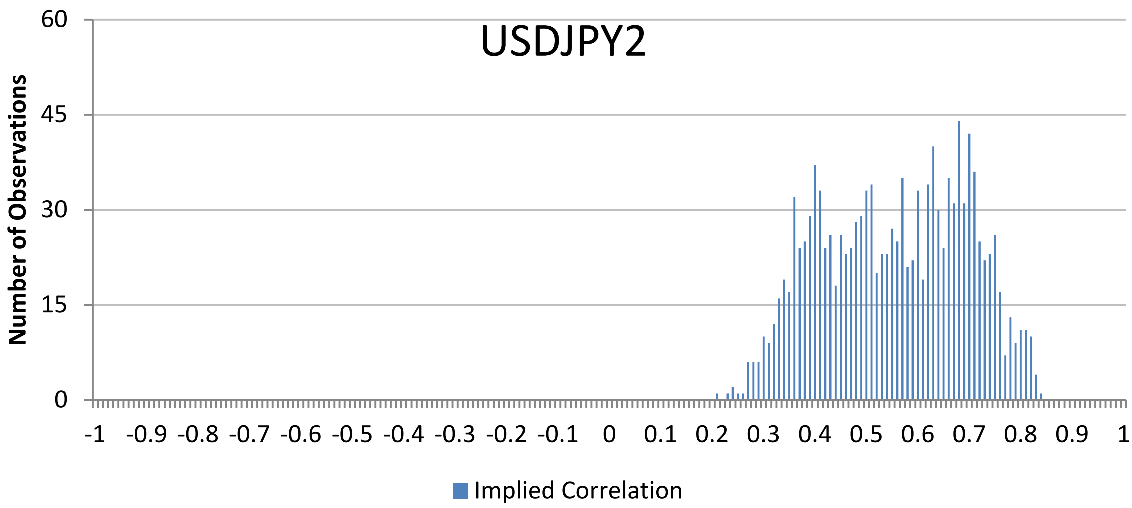

Figure 21a: The empirical distribution of implied correlation based on implied volatility quotes for EURJPY, EURUSD and USDJPY option contracts with 2 months to expiration.

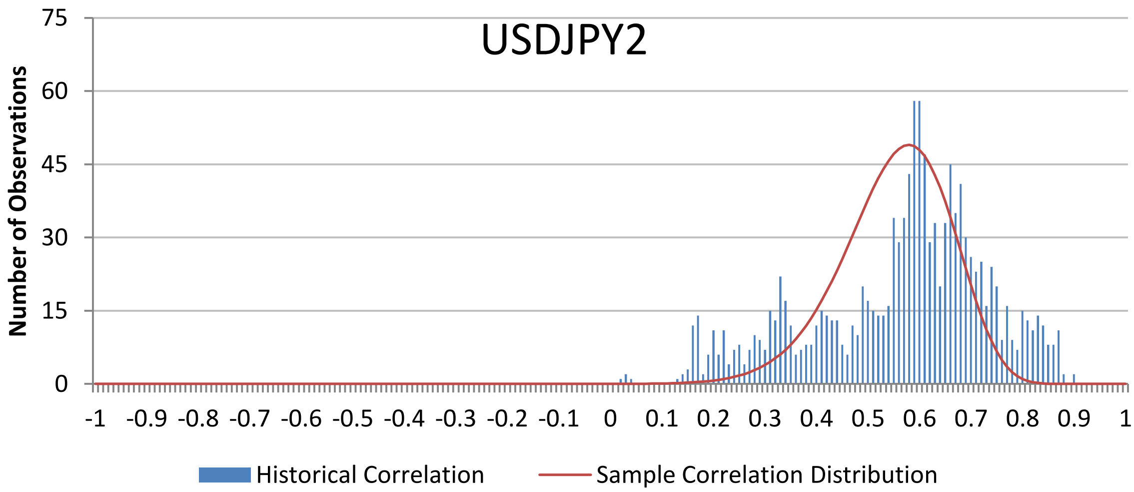

Figure 22a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURJPY and EURUSD. All estimates are based on a 43 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDJPY2				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.550	0.566	0.559	0.584
Mode	0.69	0.69	0.59/0.60	0.59/0.60
Minimum	0.200	0.200	0.016	0.016
Maximum	0.831	0.831	0.896	0.896
Number of observations	1326	1326	1275	1275

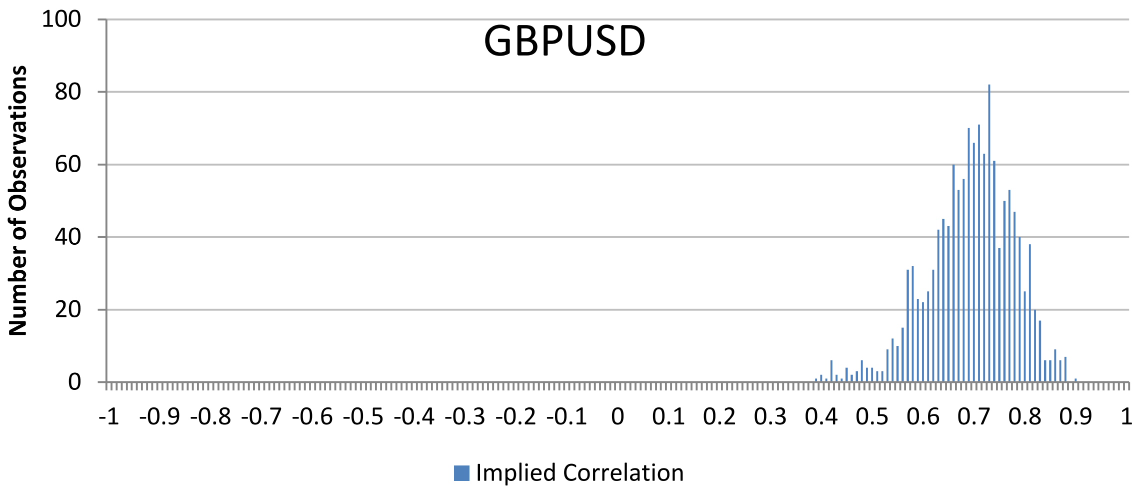

Figure 23a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPJPY and GBPUSD option contracts with 2 months to expiration.

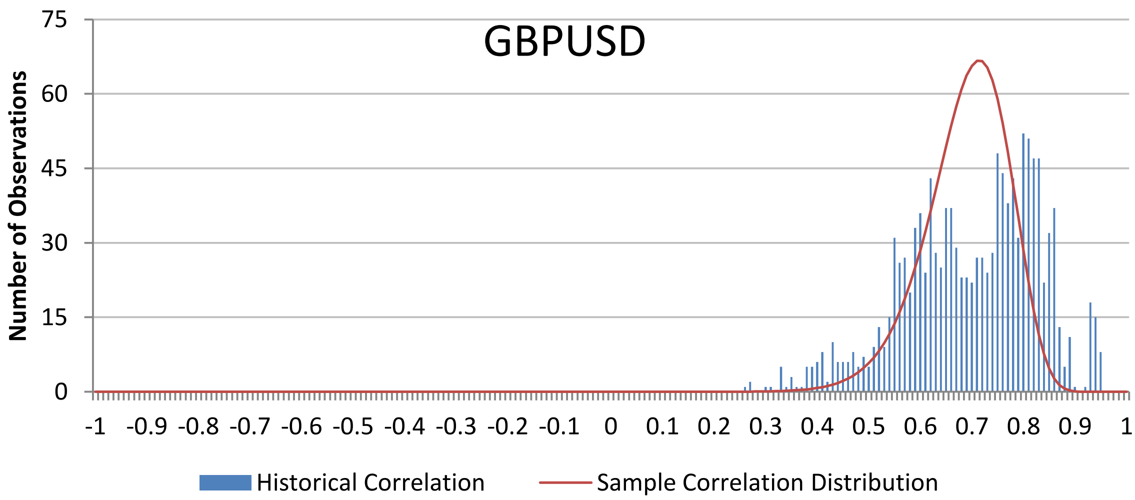

Figure 24a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPJPY. All estimates are based on a 43 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: GBPUSD				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.689	0.698	0.694	0.718
Mode	0.68	0.68	0.80	0.80
Minimum	0.382	0.382	0.253	0.253
Maximum	0.898	0.898	0.942	0.942
Number of observations	1326	1326	1275	1275

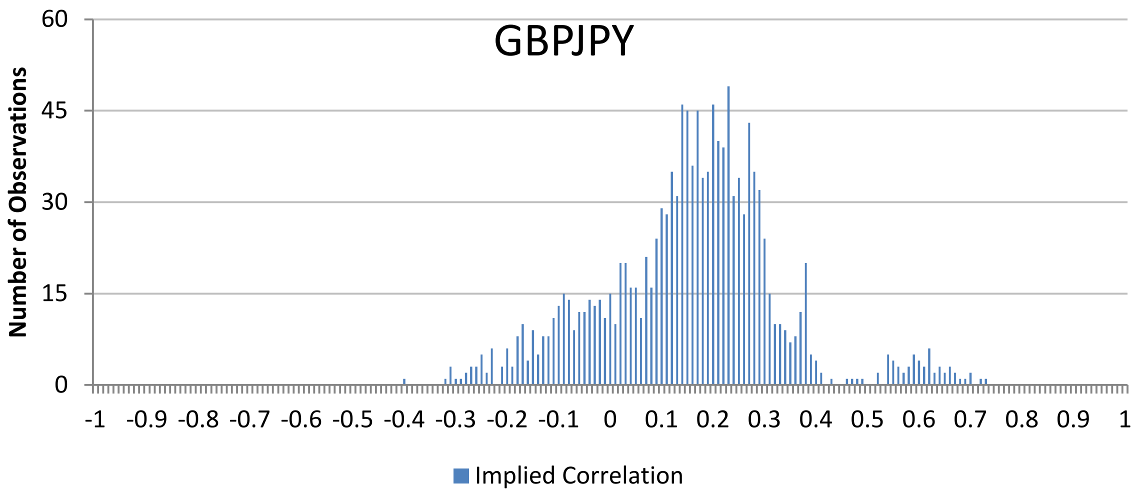

Figure 25a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPUSD and GBPJPY option contracts with 2 months to expiration.

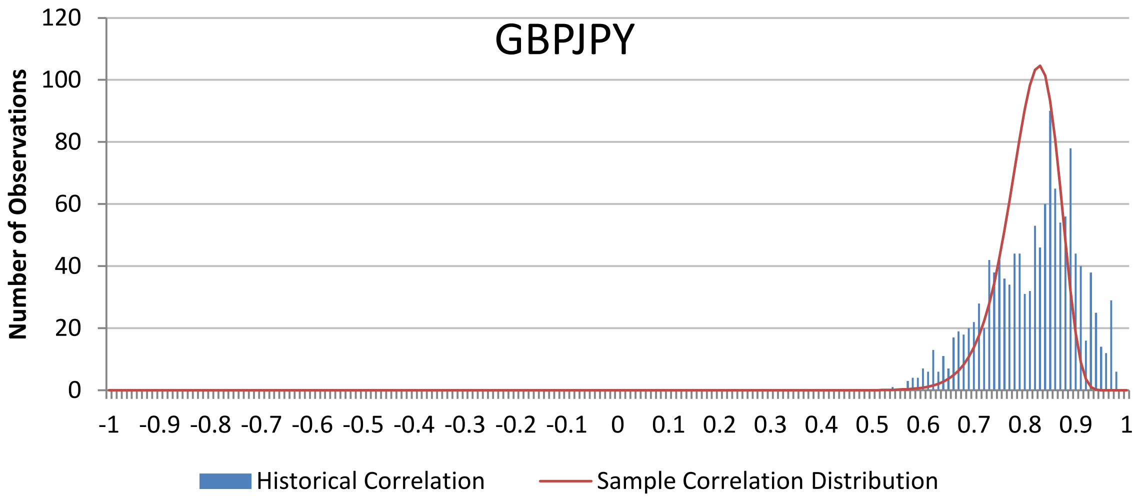

Figure 26a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPUSD. All estimates are based on a 43 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: GBPJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.150	0.156	0.811	0.831
Mode	0.23	0.23	0.85	0.85
Minimum	-0.402	-0.402	0.537	0.537
Maximum	0.725	0.725	0.971	0.971
Number of observations	1326	1326	1275	1275

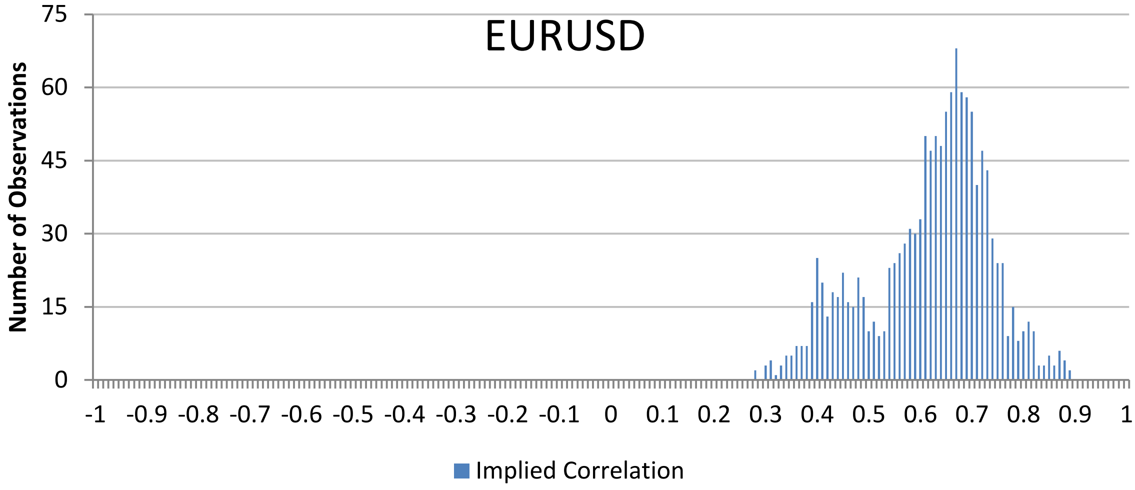

Figure 27a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURJPY and EURUSD option contracts with 2 months to expiration.

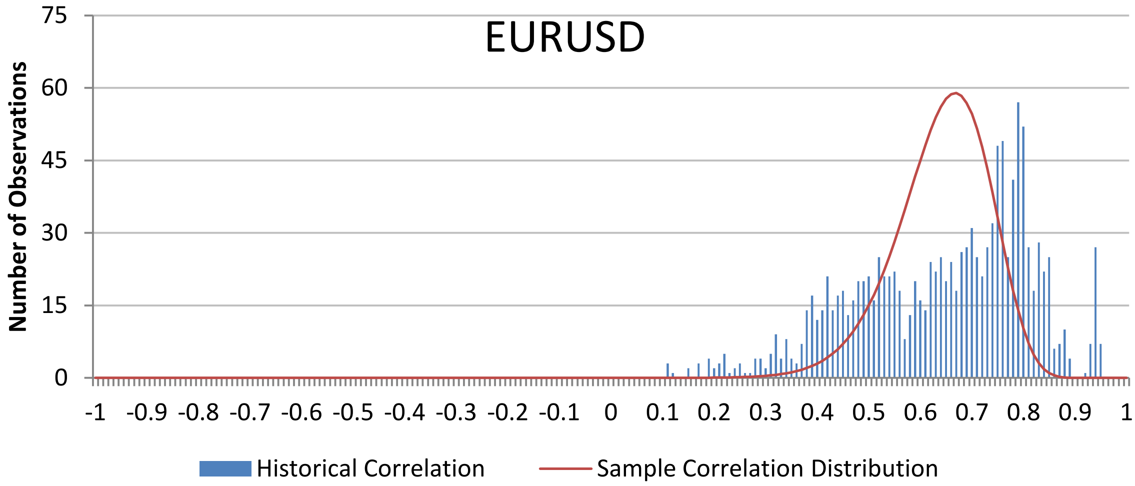

Figure 28a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and EURJPY. All estimates are based on a 43 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURUSD				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.614	0.628	0.638	0.672
Mode	0.67	0.67	0.79	0.79
Minimum	0.275	0.275	0.107	0.107
Maximum	0.885	0.885	0.945	0.945
Number of observations	1326	1326	1275	1275

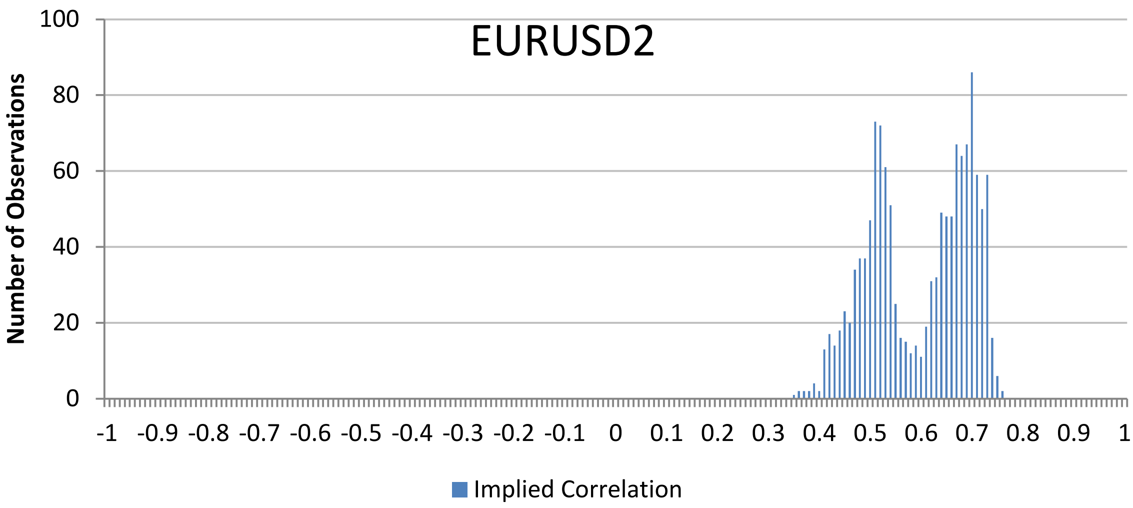

Figure 29a: The empirical distribution of implied correlation based on implied volatility quotes for USDNOK, EURNOK and EURUSD option contracts with 2 months to expiration.

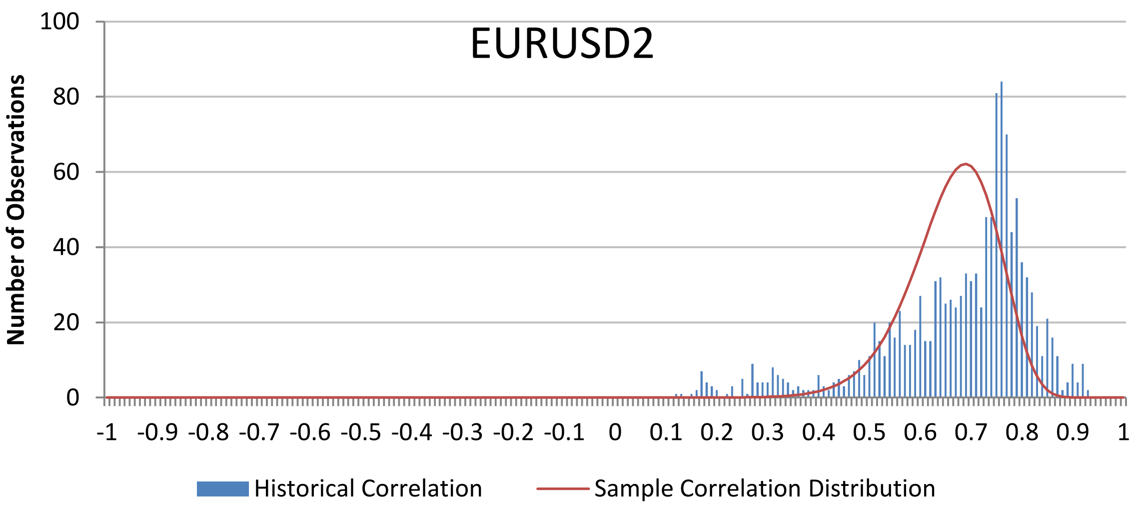

Figure 30a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDNOK and EURNOK. All estimates are based on a 43 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURUSD2				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.593	0.601	0.669	0.693
Mode	0.70	0.70	0.76	0.76
Minimum	0.349	0.349	0.118	0.118
Maximum	0.756	0.756	0.921	0.921
Number of observations	1326	1326	1275	1275

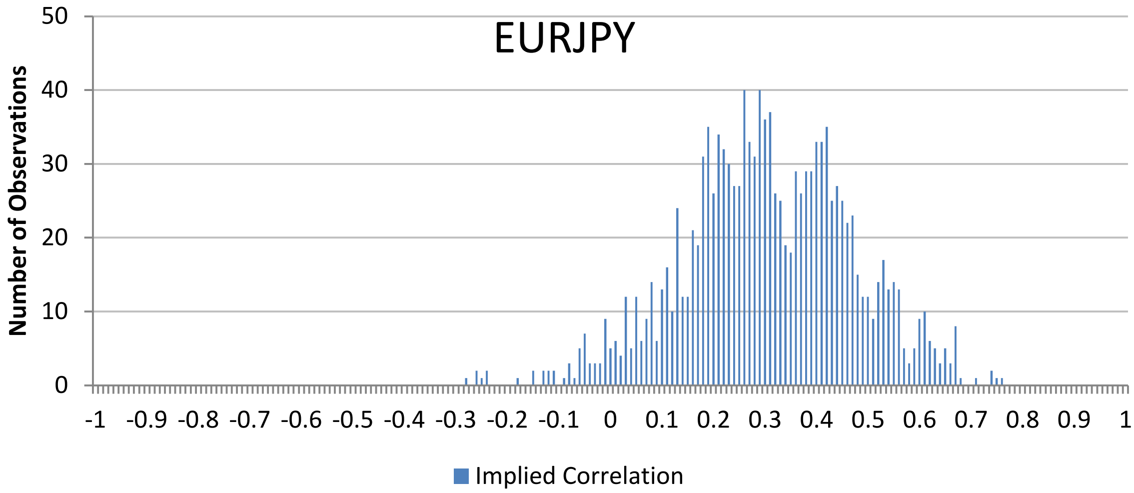

Figure 31a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURUSD and EURJPY option contracts with 2 months to expiration.

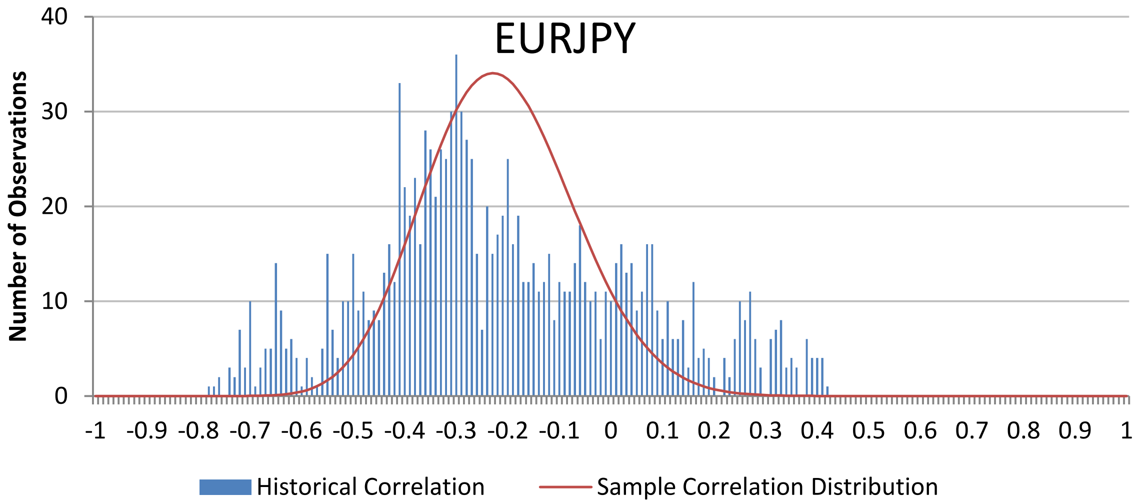

Figure 32a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURUSD and USDJPY. All estimates are based on a 43 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.299	0.309	-0.216	-0.232
Mode	0.26/0.29	0.26/0.29	-0.30	-0.30
Minimum	-0.285	-0.285	-0.783	-0.783
Maximum	0.759	0.759	0.417	0.417
Number of observations	1326	1326	1275	1275

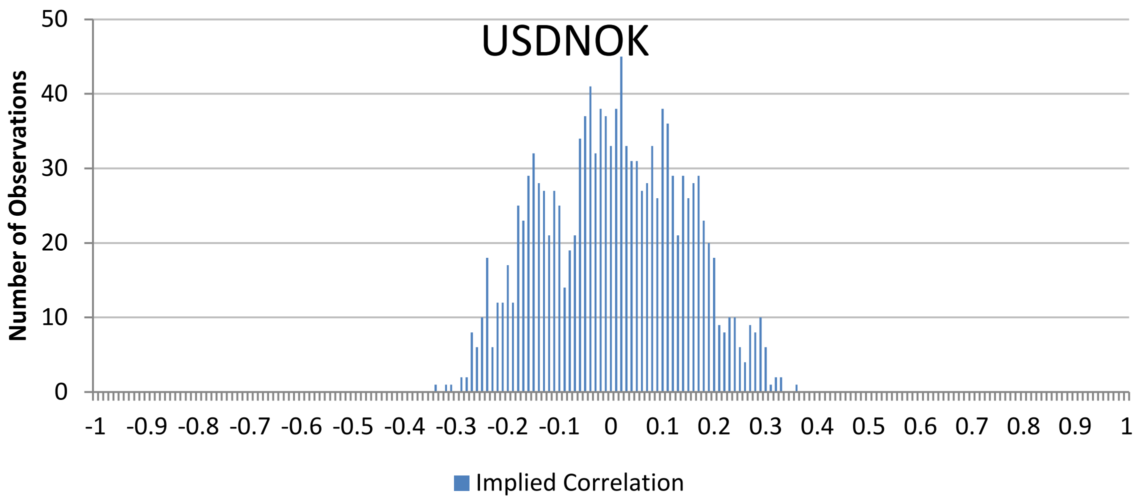

Figure 33a: The empirical distribution of implied correlation based on implied volatility quotes for EURNOK, EURUSD and USDNOK option contracts with 2 months to expiration.

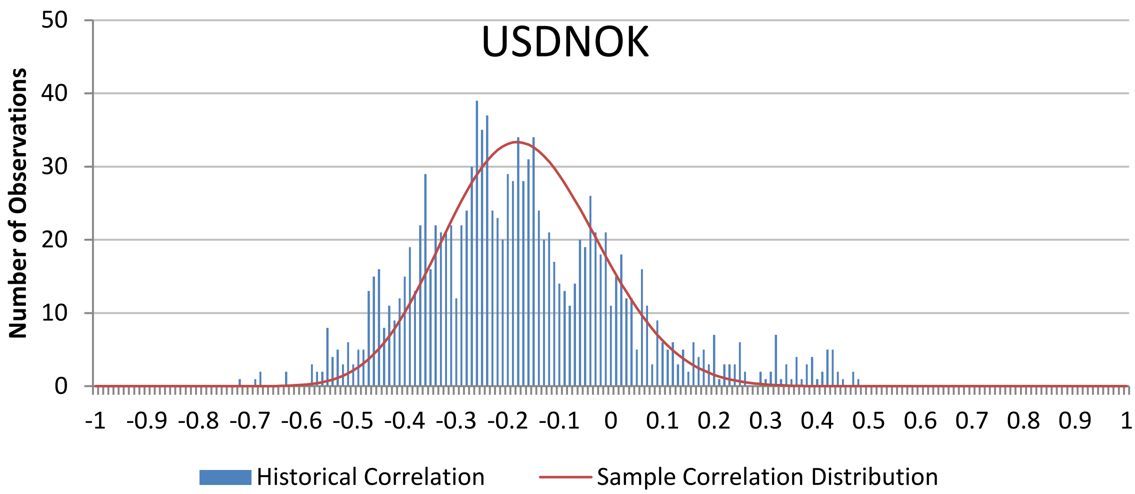

Figure 34a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURNOK and EURUSD. All estimates are based on a 43 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDNOK

	Implied correlation		Historical Correlation	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.002	0.002	-0.177	-0.177
Mode	0.02	0.02	-0.26	-0.26
Minimum	-0.350	-0.350	-0.729	-0.729
Maximum	0.356	0.356	0.478	0.478
Number of observations	1326	1326	1275	1275

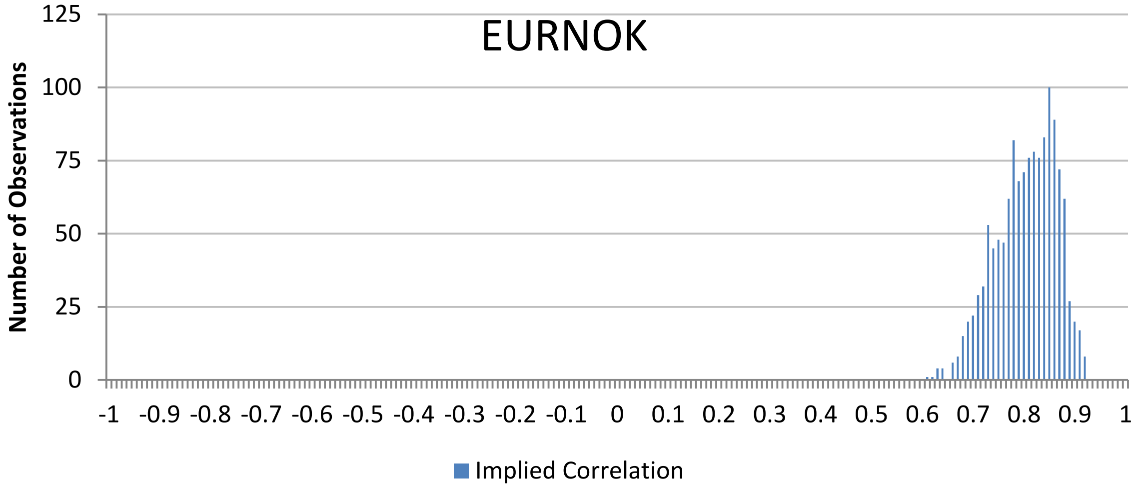

Figure 35a: The empirical distribution of implied correlation based on implied volatility quotes for EURUSD, USDNOK and EURNOK option contracts with 2 months to expiration.

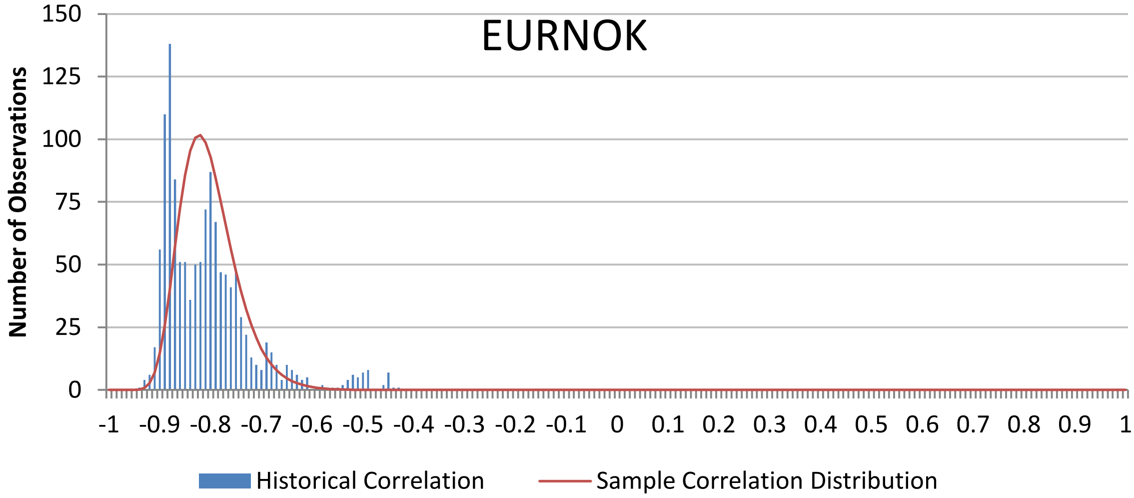

Figure 36a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURUSD and USDNOK. All estimates are based on a 43 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURNOK				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.799	0.807	-0.811	-0.826
Mode	0.85	0.85	-0.88	-0.88
Minimum	0.608	0.608	-0.943	-0.943
Maximum	0.919	0.919	-0.434	0.434
Number of observations	1326	1326	1275	1275

7.3 Time horizon: 3 months

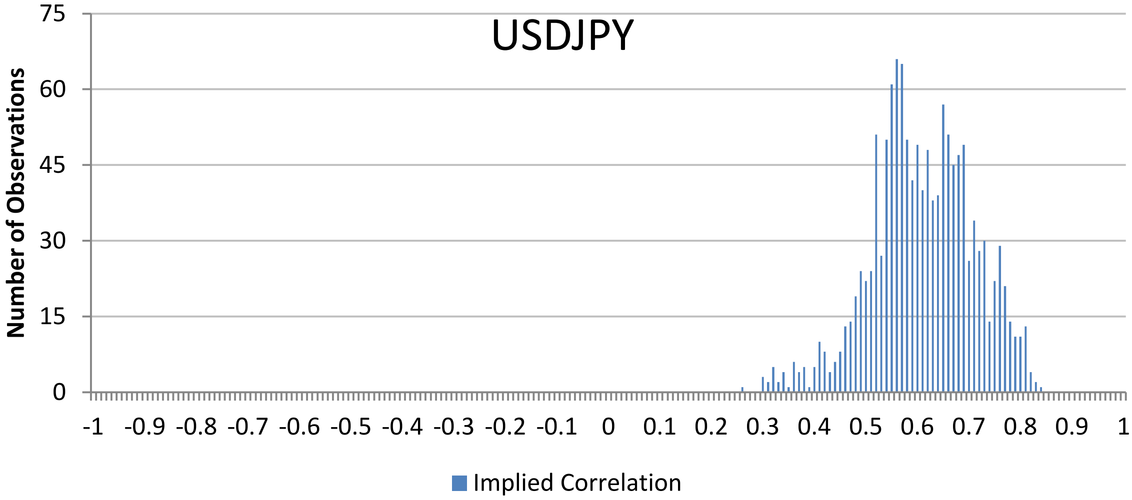

Figure 37a: The empirical distribution of implied correlation based on implied volatility quotes for GBPJPY, GBPUSD and USDJPY option contracts with 3 months to expiration.

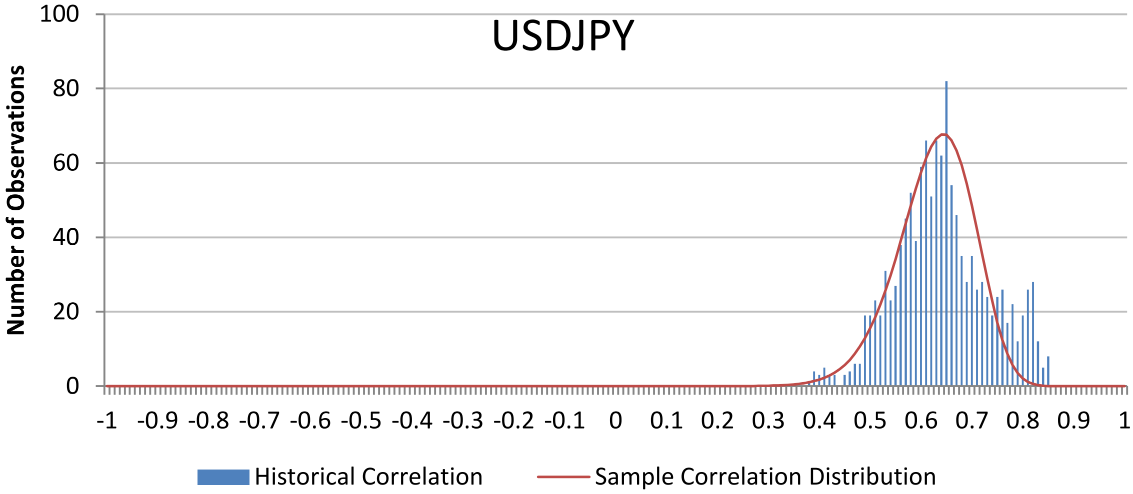

Figure 38a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for GBPJPY and GBPUSD. All estimates are based on a 65 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.605	0.614	0.637	0.647
Mode	0.56	0.56	0.65	0.65
Minimum	0.255	0.255	0.378	0.378
Maximum	0.834	0.834	0.848	0.848
Number of observations	1326	1326	1253	1253

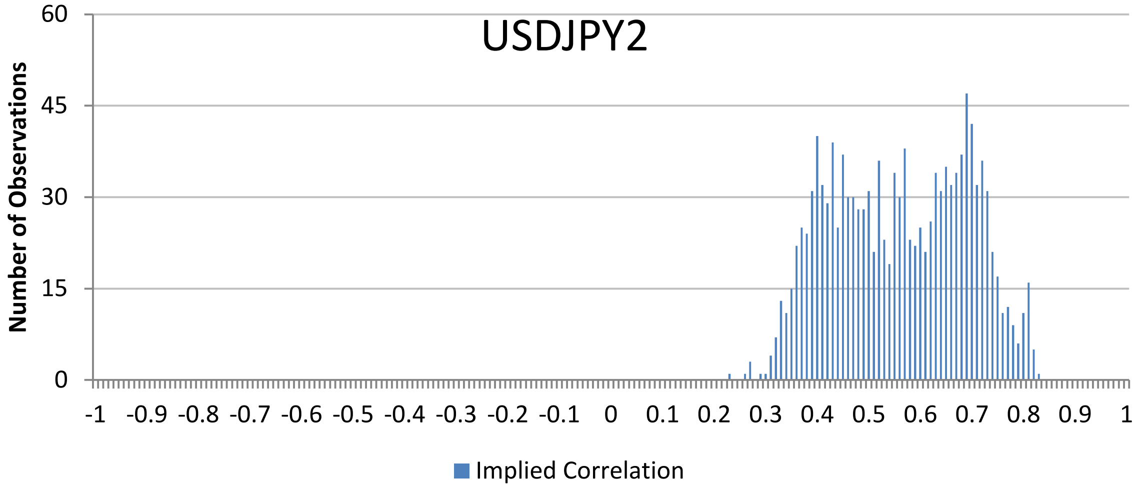

Figure 39a: The empirical distribution of implied correlation based on implied volatility quotes for EURJPY, EURUSD and USDJPY option contracts with 3 months to expiration.

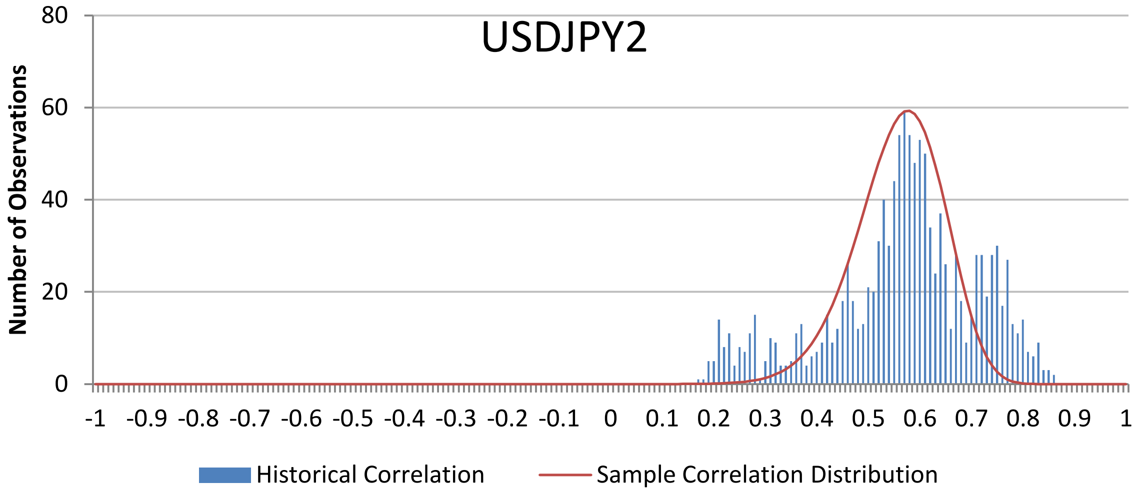

Figure 40a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURJPY and EURUSD. All estimates are based on a 65 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDJPY2				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.555	0.570	0.563	0.580
Mode	0.70	0.70	0.57	0.57
Minimum	0.226	0.226	0.164	0.164
Maximum	0.820	0.820	0.852	0.852
Number of observations	1326	1326	1253	1253

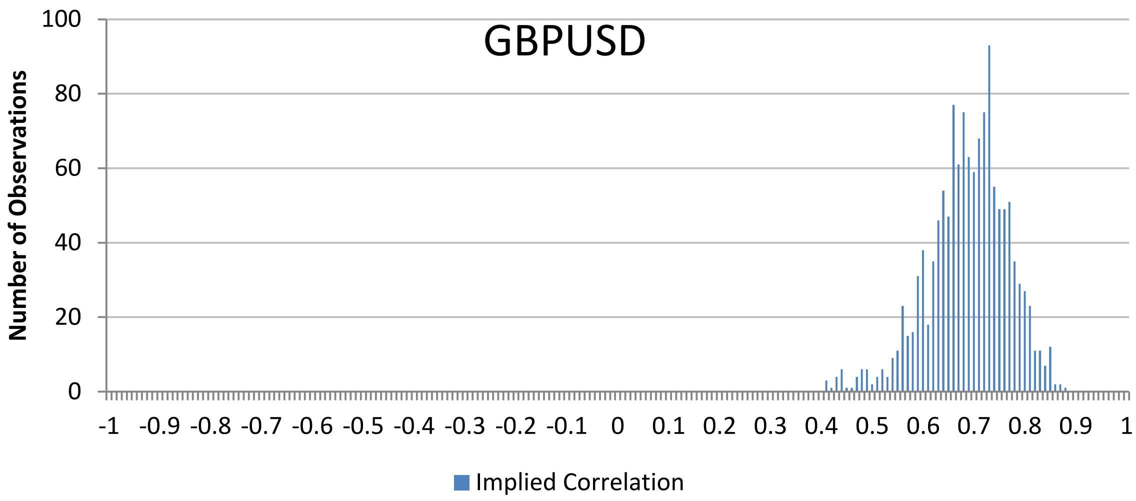

Figure 41a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPJPY and GBPUSD option contracts with 3 months to expiration.

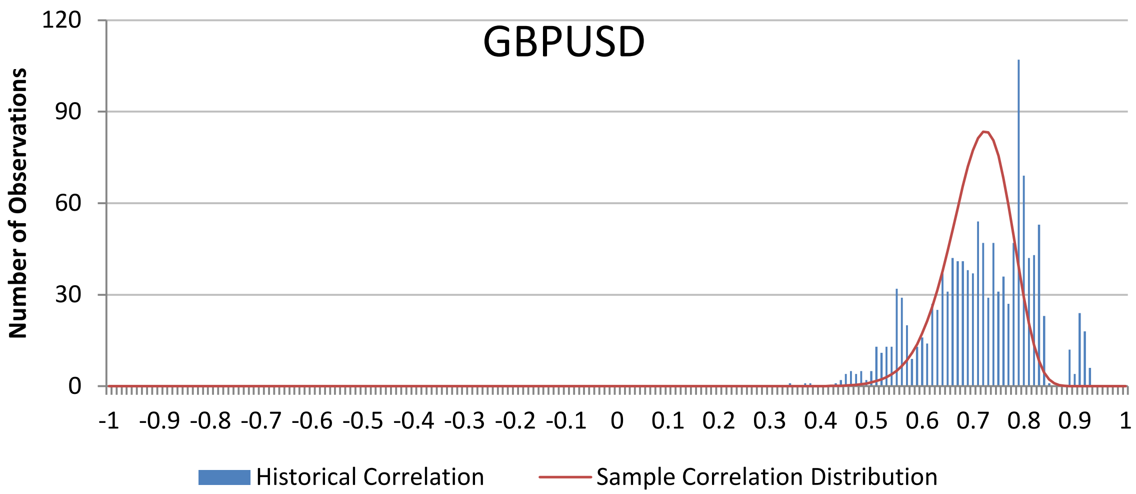

Figure 42a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPJPY. All estimates are based on a 65 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: GBPUSD				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.684	0.692	0.710	0.727
Mode	0.73	0.73	0.79	0.79
Minimum	0.403	0.403	0.340	0.340
Maximum	0.875	0.875	0.927	0.927
Number of observations	1326	1326	1253	1253

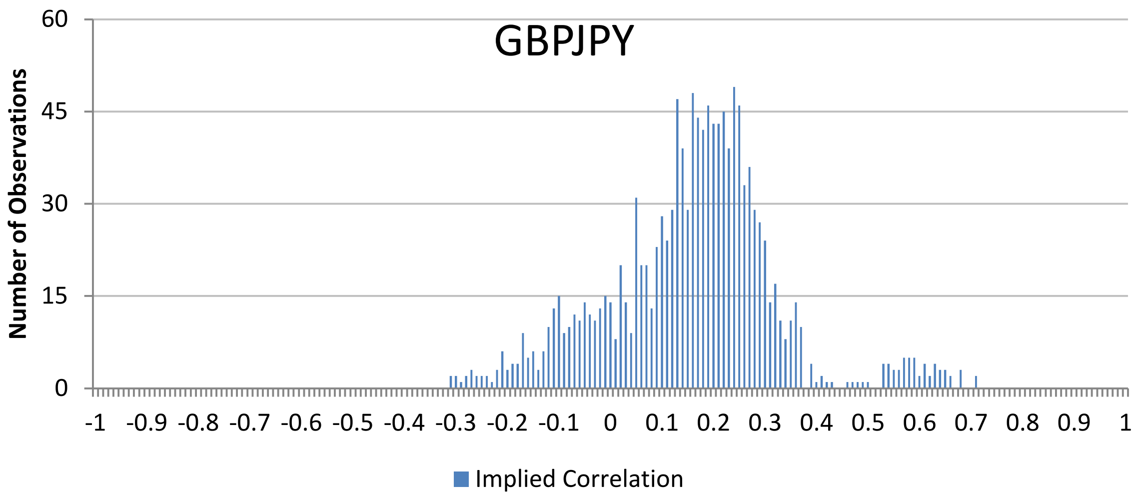

Figure 43a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPUSD and GBPJPY option contracts with 3 months to expiration.

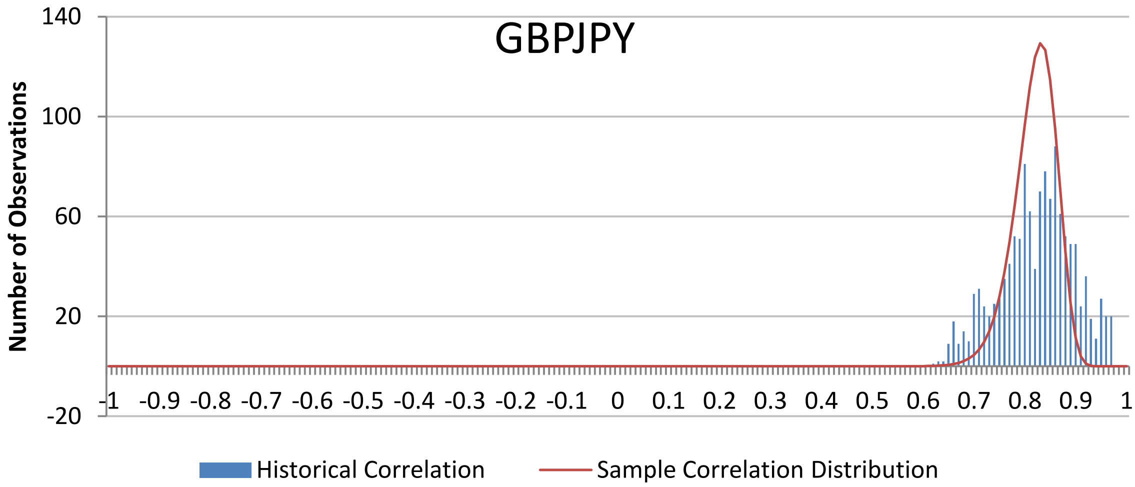

Figure 44a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPUSD. All estimates are based on a 65 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: GBPJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.154	0.159	0.819	0.834
Mode	0.24	0.24	0.86	0.86
Minimum	-0.314	-0.314	0.620	0.620
Maximum	0.708	0.708	0.967	0.967
Number of observations	1326	1326	1253	1253

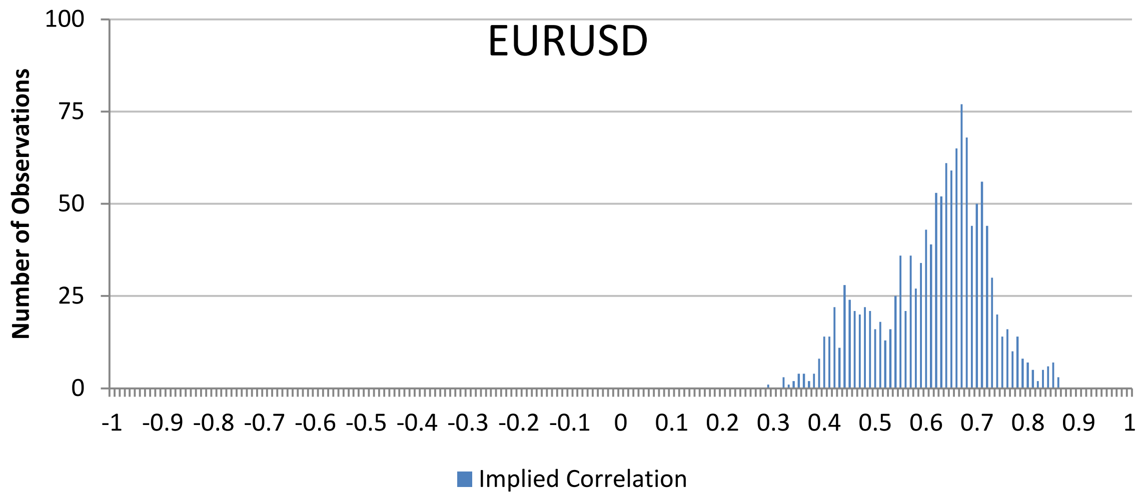

Figure 45a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURJPY and EURUSD option contracts with 3 months to expiration.

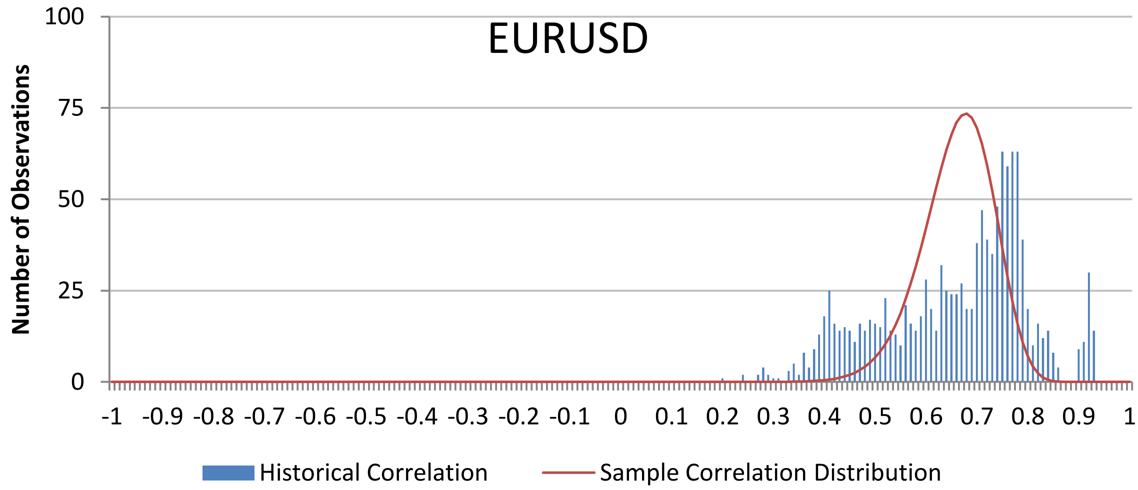

Figure 46a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and EURJPY. All estimates are based on a 65 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURUSD				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.610	0.621	0.654	0.681
Mode	0.67	0.67	0.77/0.78	0.77/0.78
Minimum	0.290	0.290	0.198	0.198
Maximum	0.857	0.857	0.927	0.927
Number of observations	1326	1326	1253	1253

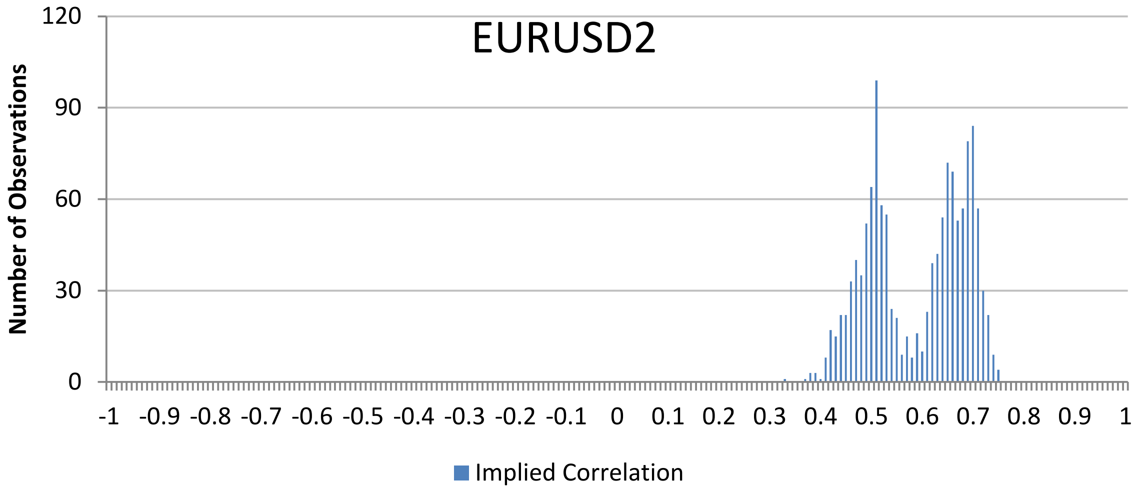

Figure 47a: The empirical distribution of implied correlation based on implied volatility quotes for USDNOK, EURNOK and EURUSD option contracts with 3 months to expiration.

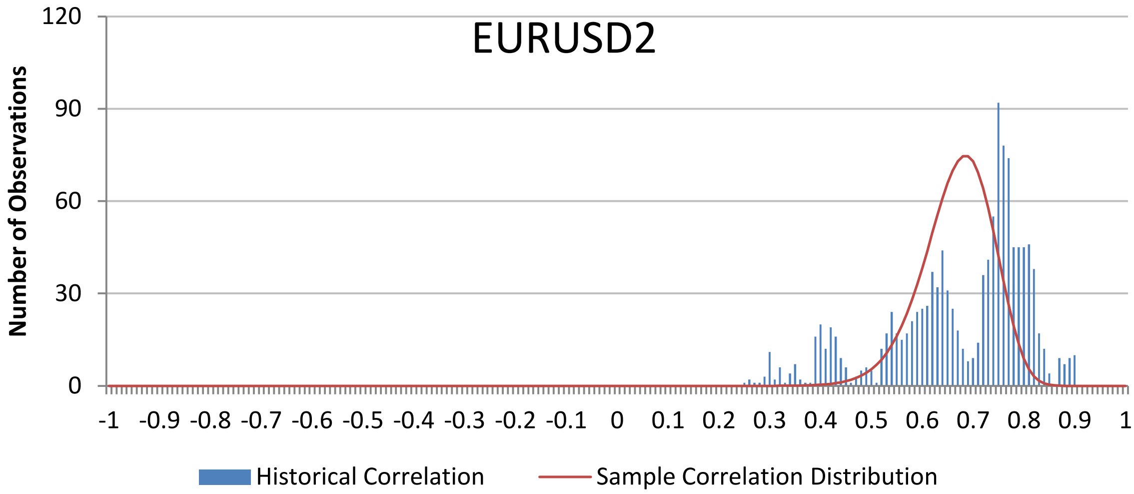

Figure 48a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDNOK and EURNOK. All estimates are based on a 65 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURUSD2				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.585	0.593	0.667	0.688
Mode	0-69	0.69	0.75	0.75
Minimum	0.329	0.329	0.249	0.249
Maximum	0.744	0.744	0.899	0.899
Number of observations	1326	1326	1253	1253

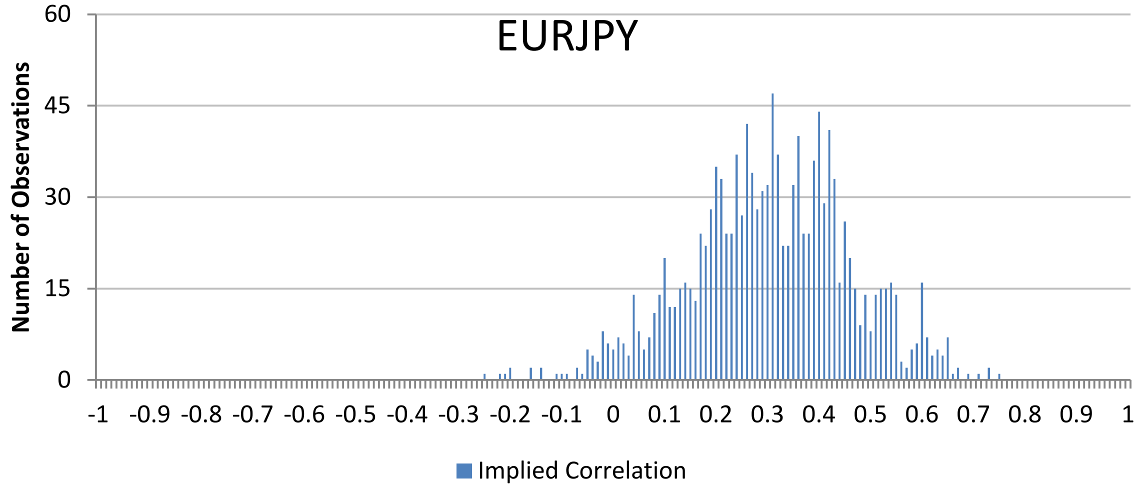

Figure 49a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURUSD and EURJPY option contracts with 3 months to expiration.

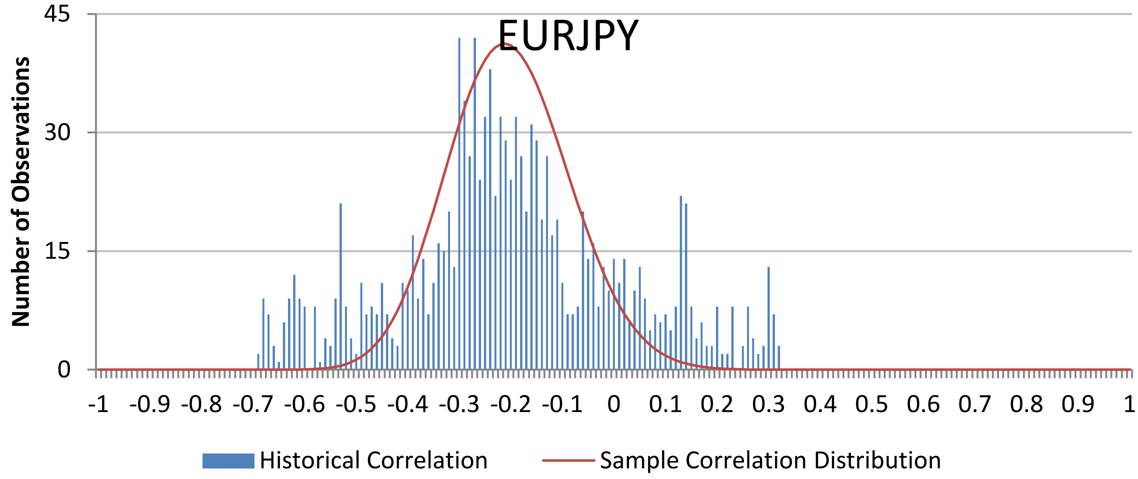

Figure 50a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURUSD and USDJPY. All estimates are based on a 65 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.301	0.309	-0.203	-0.214
Mode	0.31	0.31	-0.27/- 0.30	-0.27/- 0.30
Minimum	-0.257	-0.257	-0.697	-0.697
Maximum	0.745	0.745	0.318	0.318
Number of observations	1326	1326	1253	1253

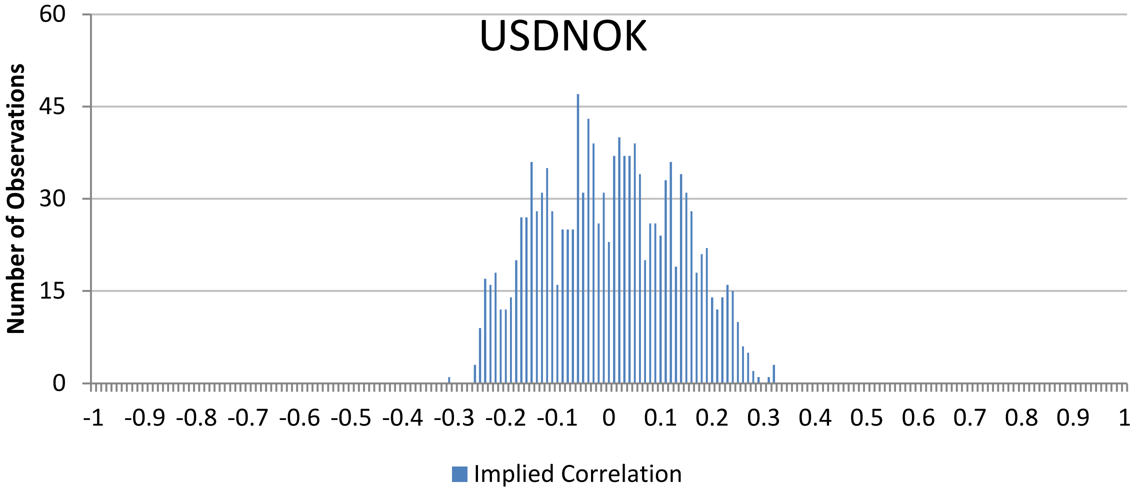

Figure 51a: The empirical distribution of implied correlation based on implied volatility quotes for EURNOK, EURUSD and USDNOK option contracts with 3 months to expiration.

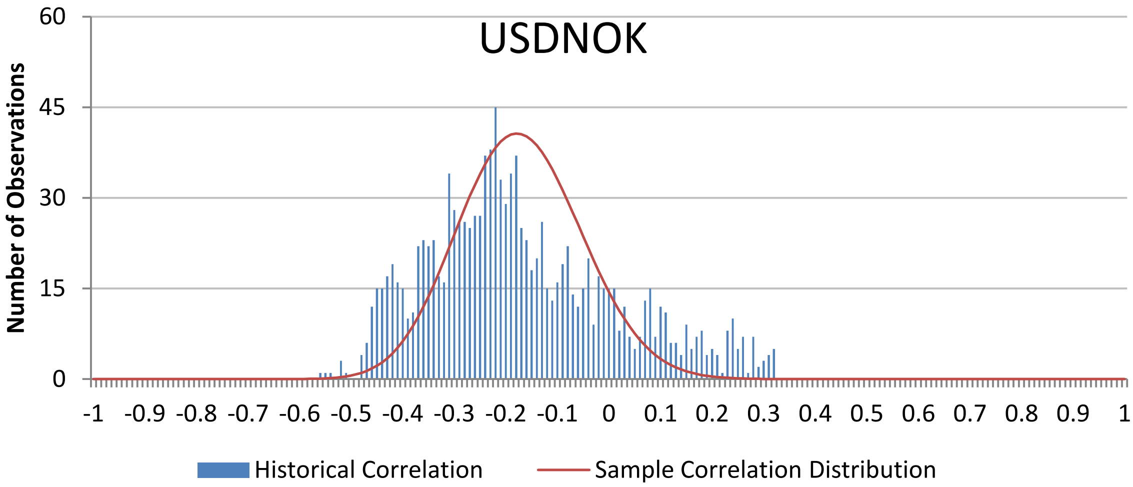

Figure 52a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURUSD and EURNOK. All estimates are based on a 65 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDNOK				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	-0.003	-0.003	-0.174	-0.180
Mode	-0.06	-0.06	-0.22	-0.22
Minimum	-0.316	-0.316	-0.564	-0.564
Maximum	0.313	0.313	0.320	0.320
Number of observations	1326	1326	1253	1253

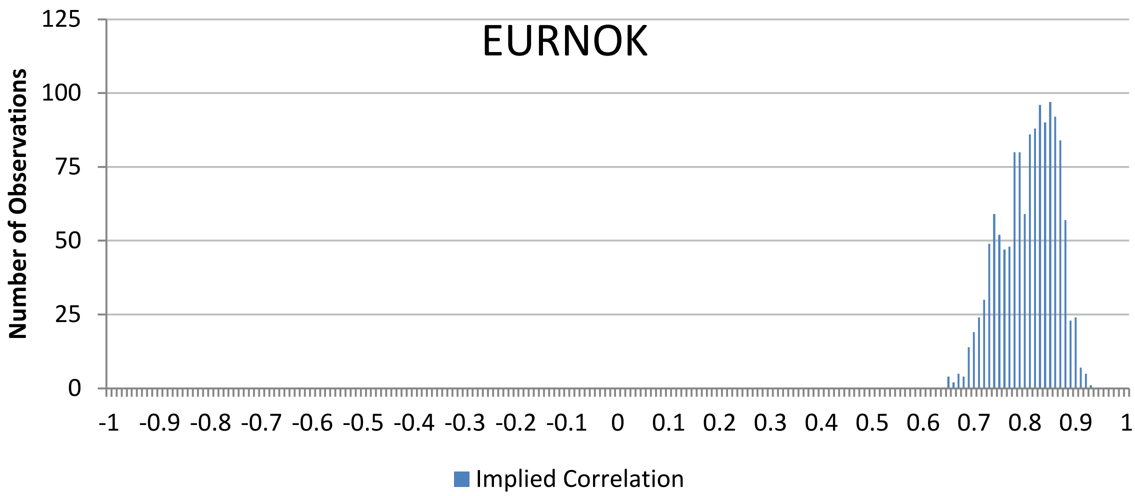

Figure 53a: The empirical distribution of implied correlation based on implied volatility quotes for EURUSD, USDNOK and EURNOK option contracts with 3 months to expiration.

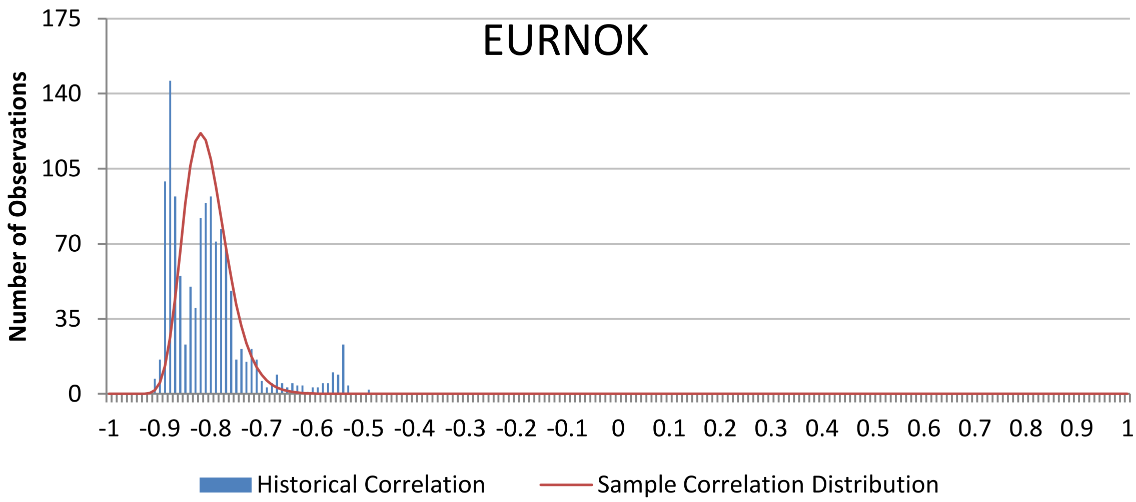

Figure 54a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURUSD and USDNOK. All estimates are based on a 65 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURNOK				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.803	0.810	-0.809	-0.822
Mode	0.85	0.85	-0.88	-0.88
Minimum	0.641	0.641	-0.922	-0.922
Maximum	0.921	0.921	-0.494	-0.494
Number of observations	1326	1326	1253	1253

7.4 Time horizon: 6 months

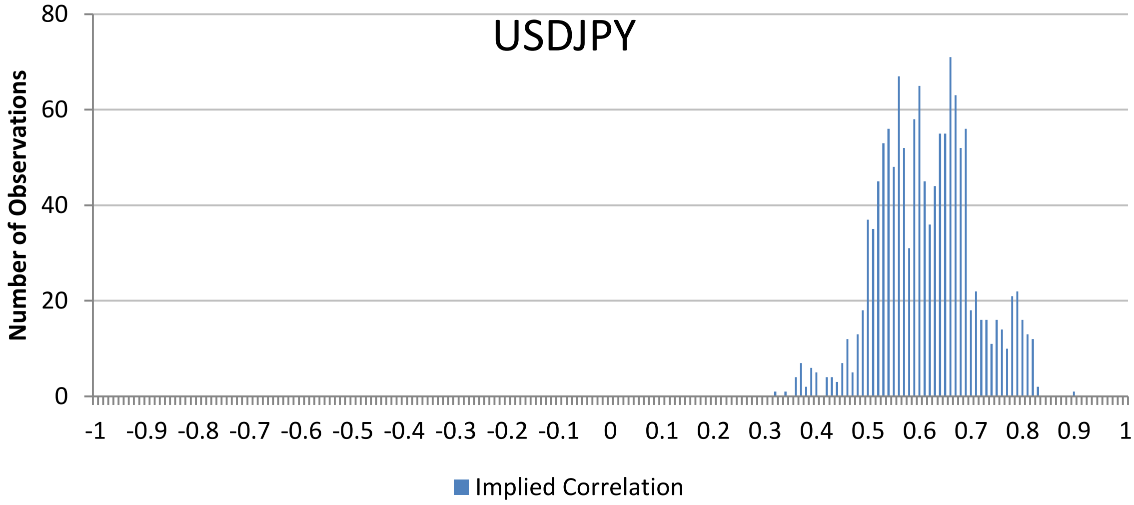

Figure 55a: The empirical distribution of implied correlation based on implied volatility quotes for GBPJPY, GBPUSD and USDJPY option contracts with 6 months to expiration.

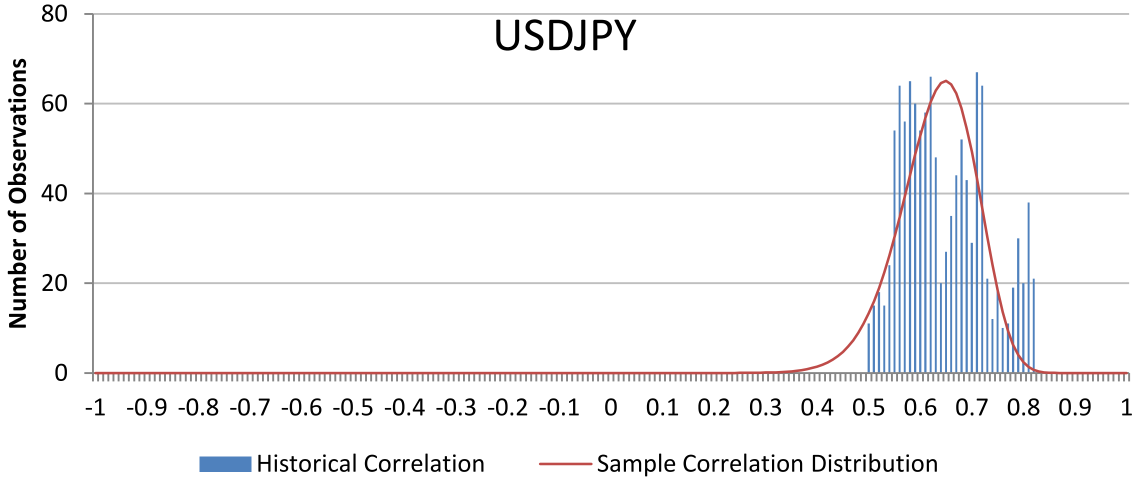

Figure 56a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for GBPJPY and GBPUSD. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.609	0.618	0.644	0.652
Mode	0.66	0.66	0.71	0.71
Minimum	0.317	0.317	0.493	0.493
Maximum	0.900	0.900	0.817	0.817
Number of observations	1326	1326	1190	1190

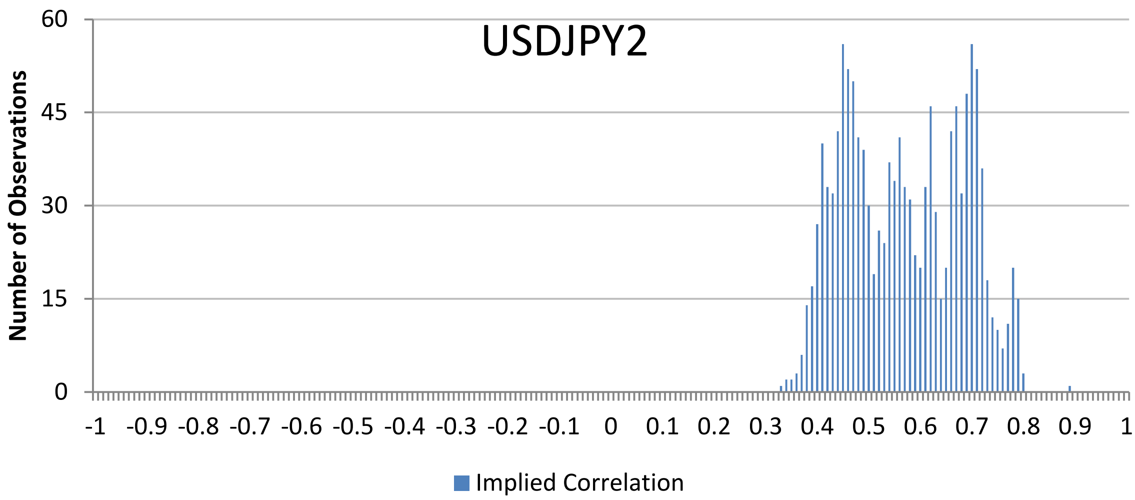

Figure 57a: The empirical distribution of implied correlation based on implied volatility quotes for EURJPY, EURUSD and USDJPY option contracts with 6 months to expiration.

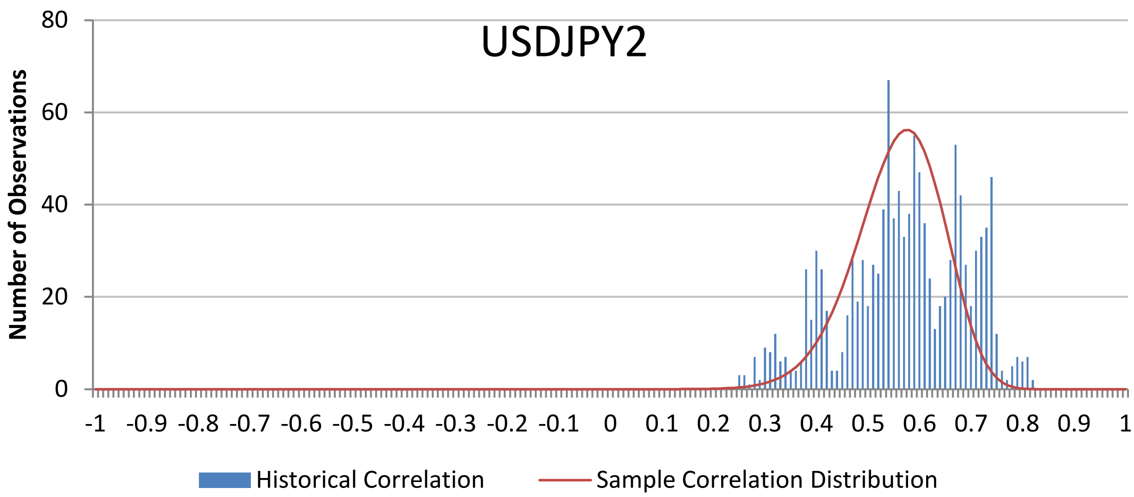

Figure 58a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURJPY and EURUSD. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDJPY2				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.564	0.576	0.567	0.579
Mode	0.45/0.70	0.45/0.70	0.59	0.59
Minimum	0.323	0.323	0.248	0.248
Maximum	0.881	0.881	0.819	0.819
Number of observations	1326	1326	1190	1190

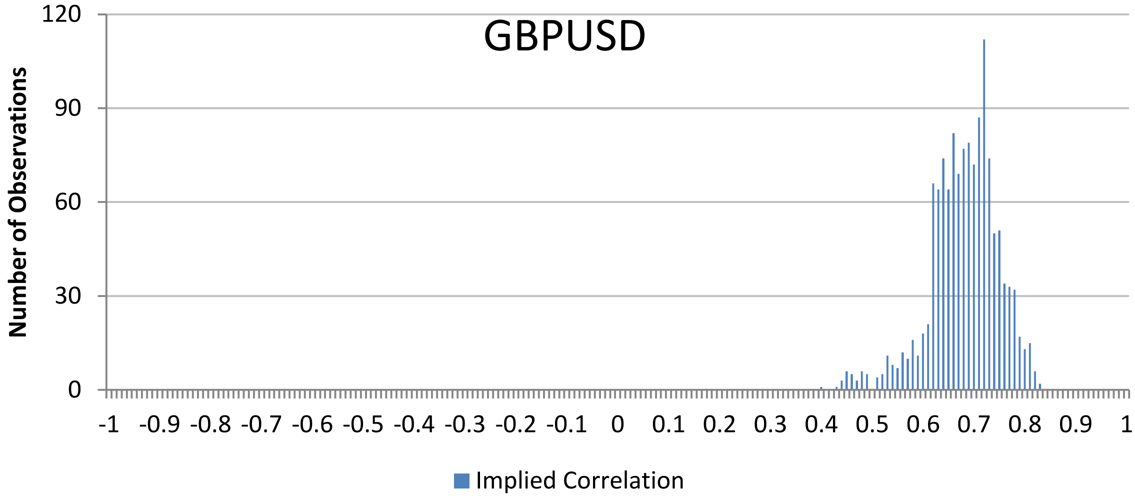

Figure 59a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPJPY and GBPUSD option contracts with 6 months to expiration.

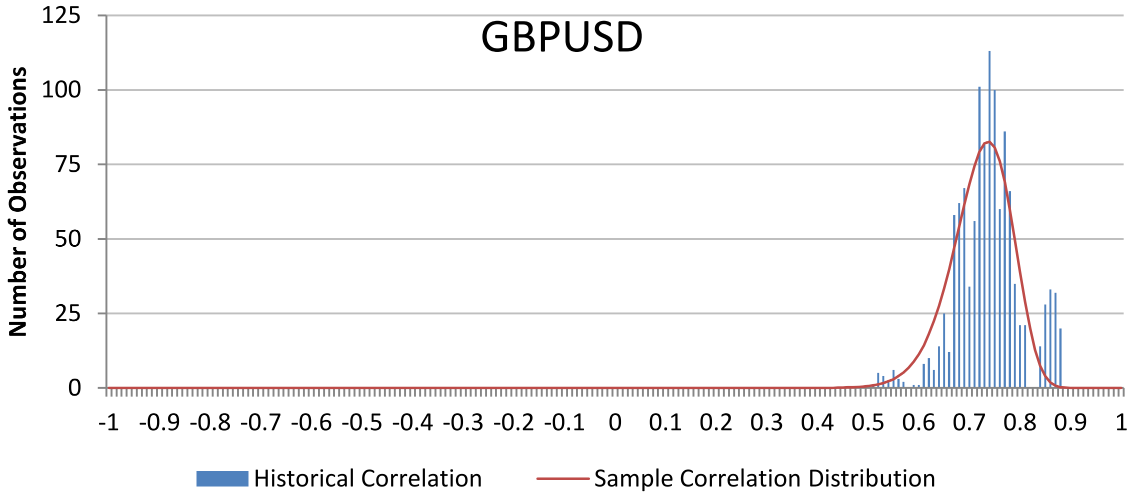

Figure 60a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for GBPJPY and USDJPY. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: GBPUSD				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.676	0.681	0.733	0.740
Mode	0.72	0.72	0.74	0.74
Minimum	0.392	0.392	0.497	0.497
Maximum	0.826	0.826	0.878	0.878
Number of observations	1326	1326	1190	1190

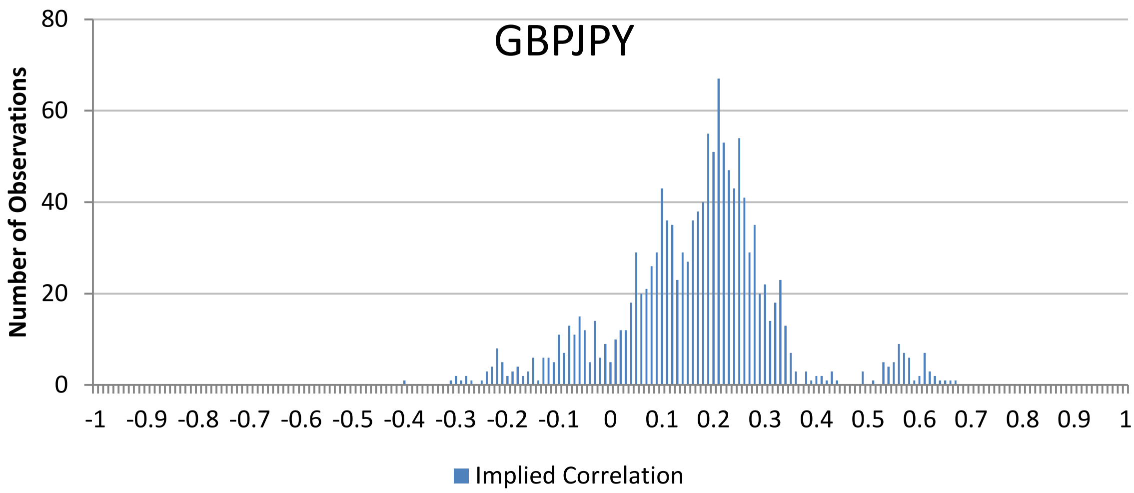

Figure 61a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPUSD and GBPJPY option contracts with 6 months to expiration.

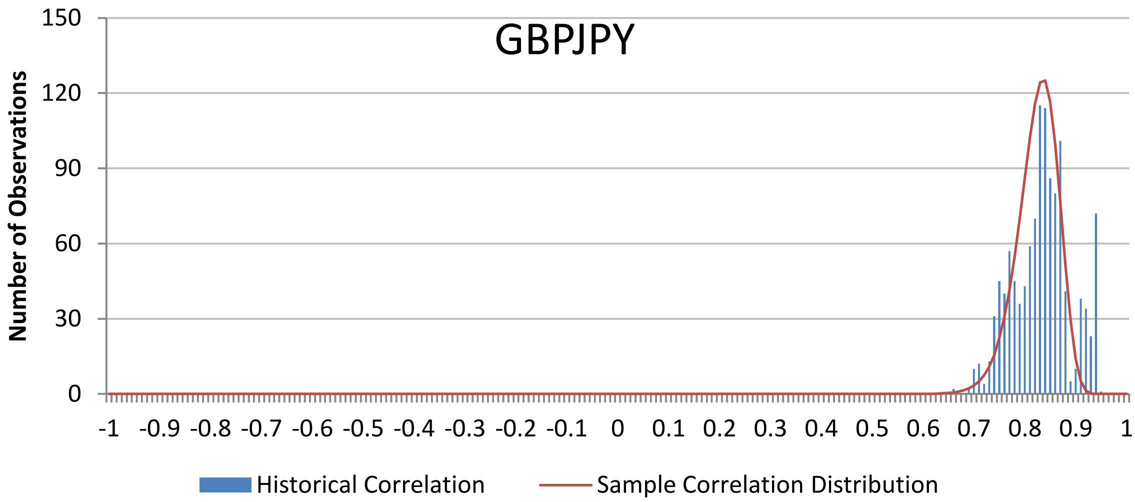

Figure 62a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPUSD. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: GBPJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.162	0.166	0.828	0.838
Mode	0.21	0.21	0.83	0.83
Minimum	-0.406	-0.406	0.658	0.658
Maximum	0.661	0.661	0.943	0.943
Number of observations	1326	1326	1190	1190

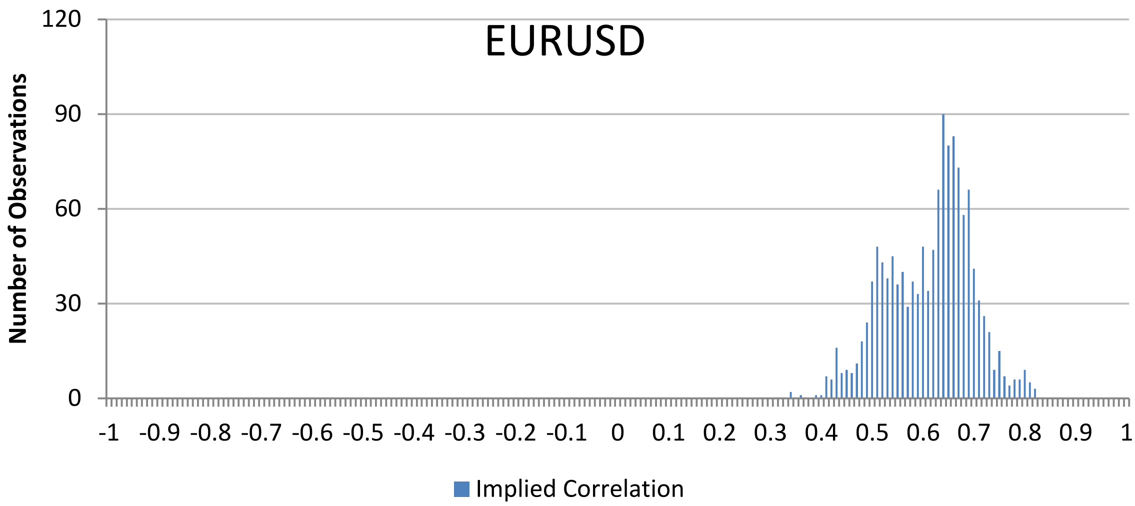

Figure 63a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURJPY and EURUSD option contracts with 6 months to expiration.

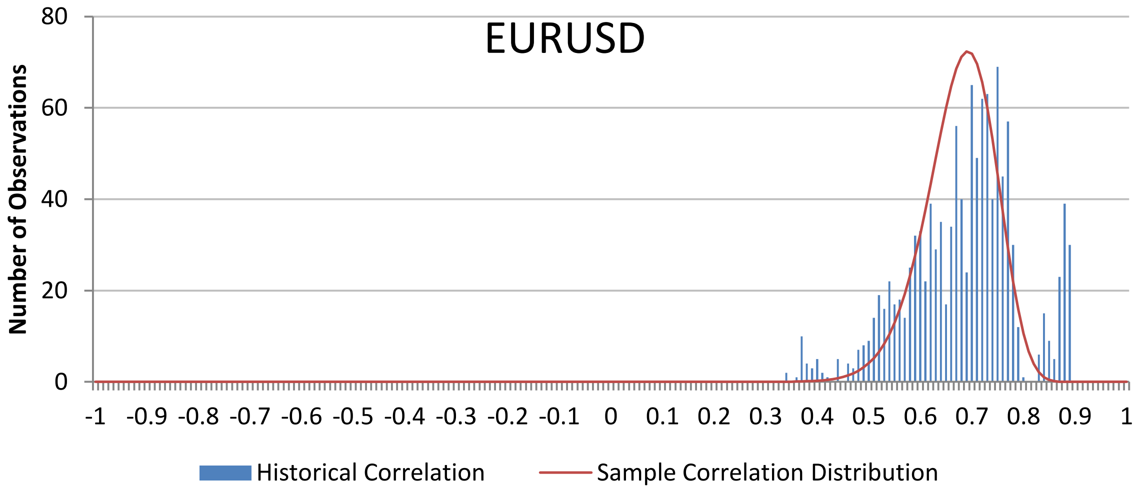

Figure 64a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and EURJPY. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURUSD				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.608	0.615	0.679	0.695
Mode	0.64	0.64	0.75	0.75
Minimum	0.332	0.332	0.333	0.333
Maximum	0.812	0.812	0.888	0.888
Number of observations	1326	1326	1190	1190

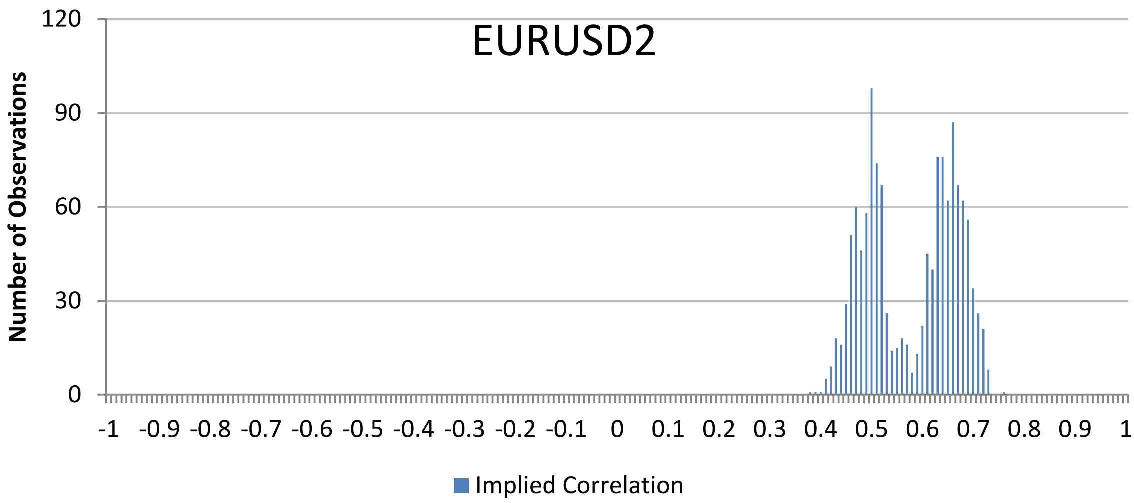

Figure 65a: The empirical distribution of implied correlation based on implied volatility quotes for USDNOK, EURNOK and EURUSD option contracts with 6 months to expiration.

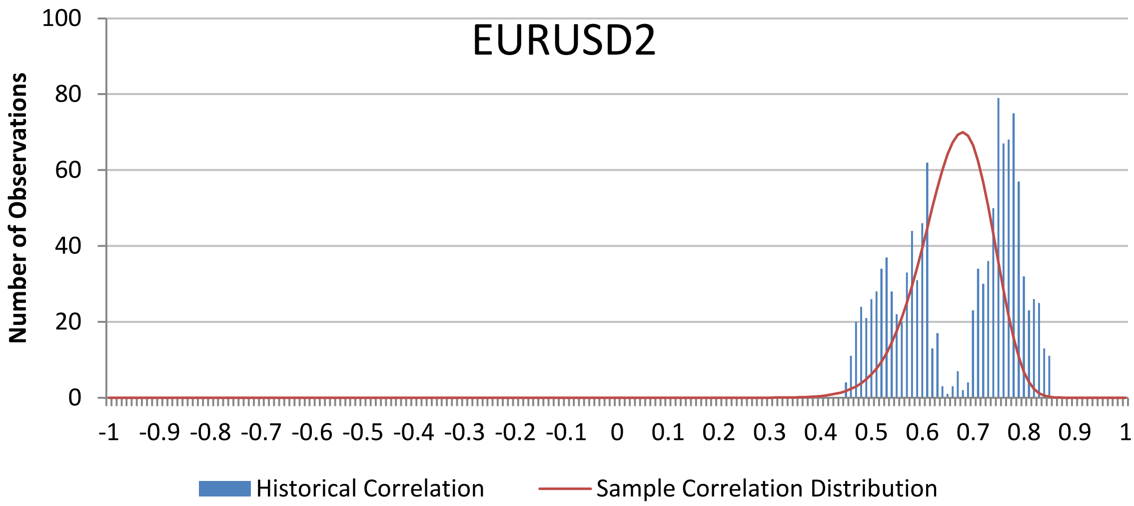

Figure 66a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDNOK and EURNOK. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURUSD2				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.575	0.581	0.666	0.682
Mode	0.50	0.50	0.75	0.75
Minimum	0.379	0.379	0.448	0.448
Maximum	0.753	0.753	0.847	0.847
Number of observations	1326	1326	1190	1190

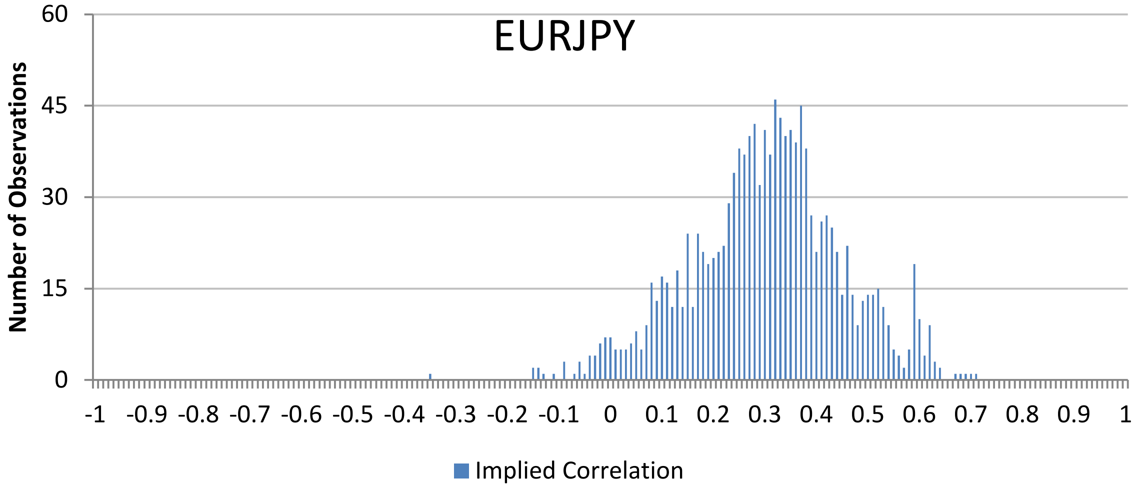

Figure 67a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURUSD and EURJPY option contracts with 6 months to expiration.

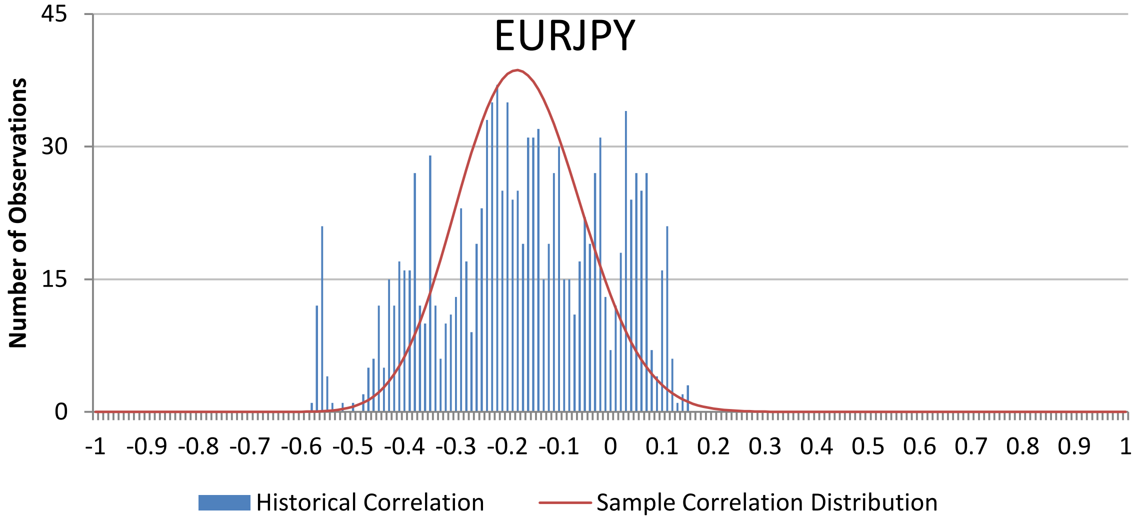

Figure 68a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and EURUSD. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.298	0.305	-0.177	-0.183
Mode	0.32	0.32	-0.22	-0.22
Minimum	-0.356	-0.356	-0.586	-0.586
Maximum	0.709	0.709	0.143	0.143
Number of observations	1326	1326	1190	1190

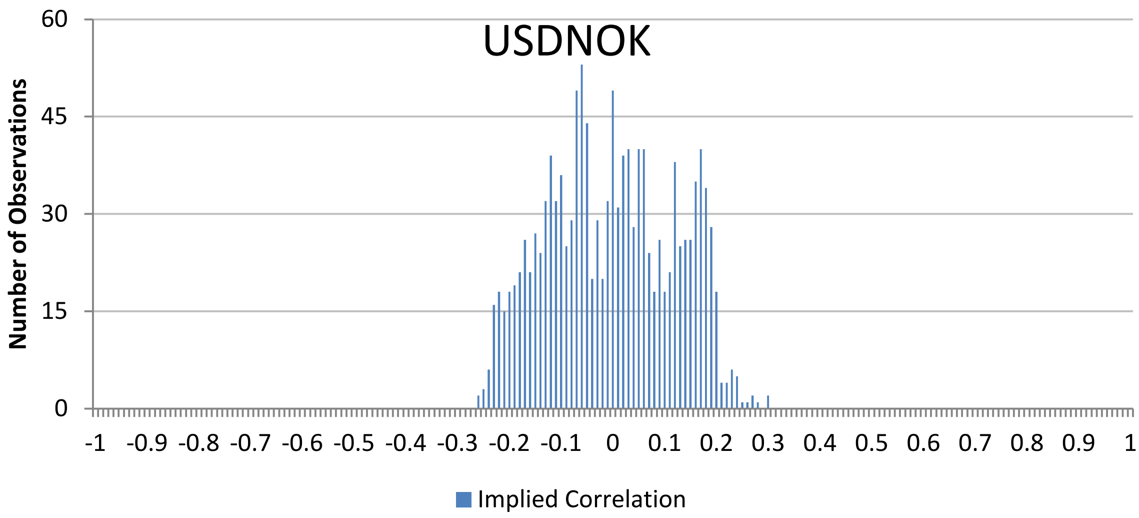

Figure 69a: The empirical distribution of implied correlation based on implied volatility quotes for EURNOK, EURUSD and USDNOK option contracts with 6 months to expiration.

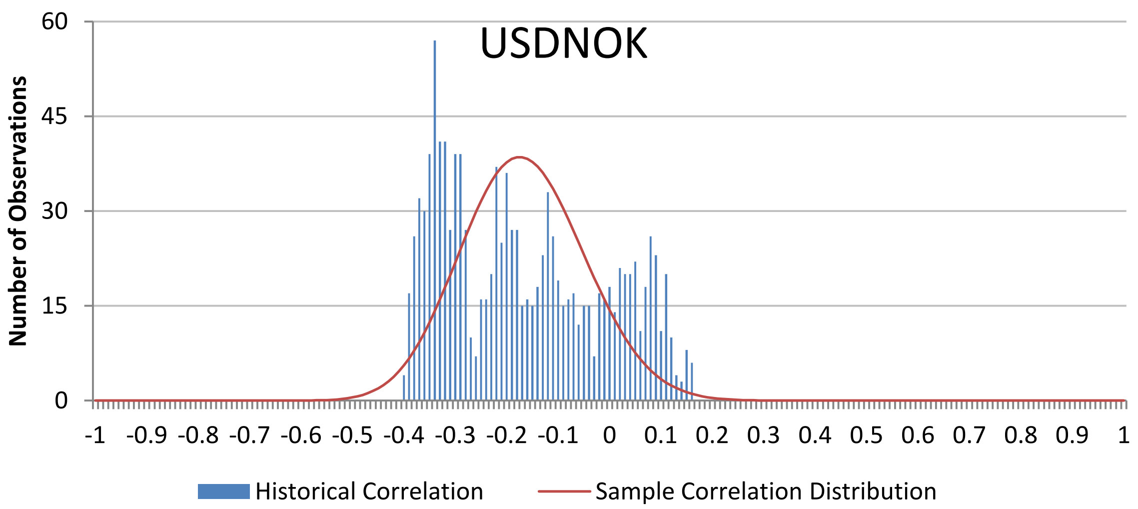

Figure 70a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURNOK and EURUSD. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDNOK				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	-0.009	-0.009	-0.172	-0.176
Mode	-0.06	-0.06	-0.34	-0.34
Minimum	-0.268	-0.268	-0.407	-0.407
Maximum	0.293	0.293	0.153	0.153
Number of observations	1326	1326	1190	1190

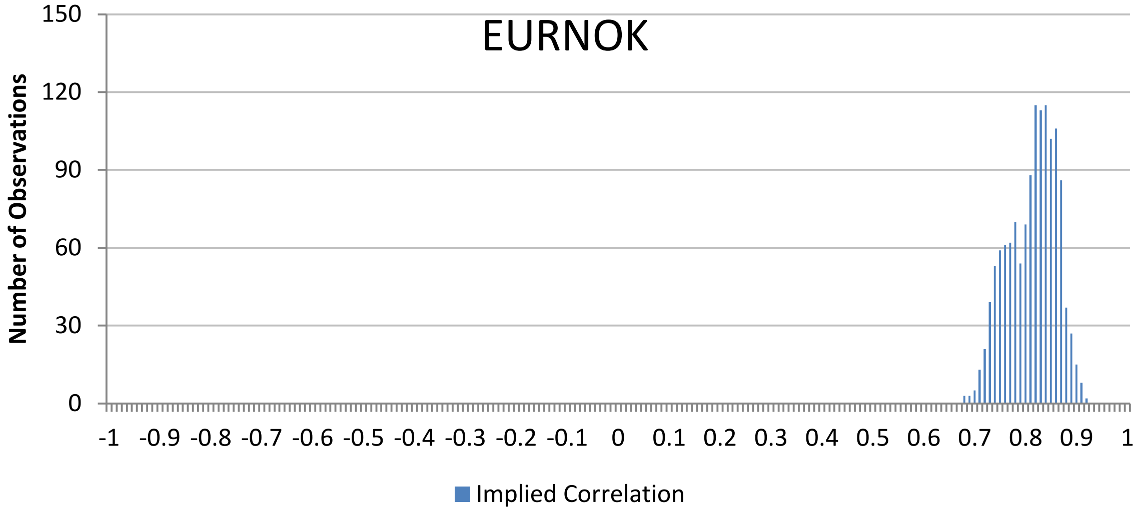

Figure 71a: The empirical distribution of implied correlation based on implied volatility quotes for EURUSD, USDNOK and EURNOK option contracts with 6 months to expiration.

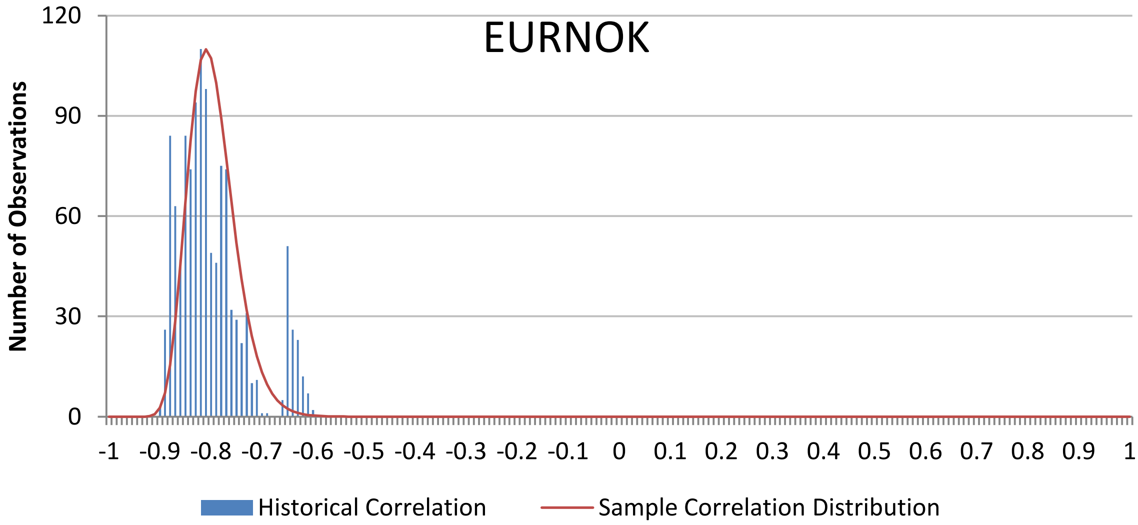

Figure 72a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURUSD and EURNOK. All estimates are based on a 128 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURNOK				
	Implied correlation		Historical Correlation	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.808	0.813	-0.803	-0.812
Mode	0.82/0.84	0.82/0.84	-0.82	-0.82
Minimum	0.676	0.676	-0.904	-0.904
Maximum	0.911	0.911	-0.609	-0.609
Number of observations	1326	1326	1190	1190

7.5 Time horizon: 9 months

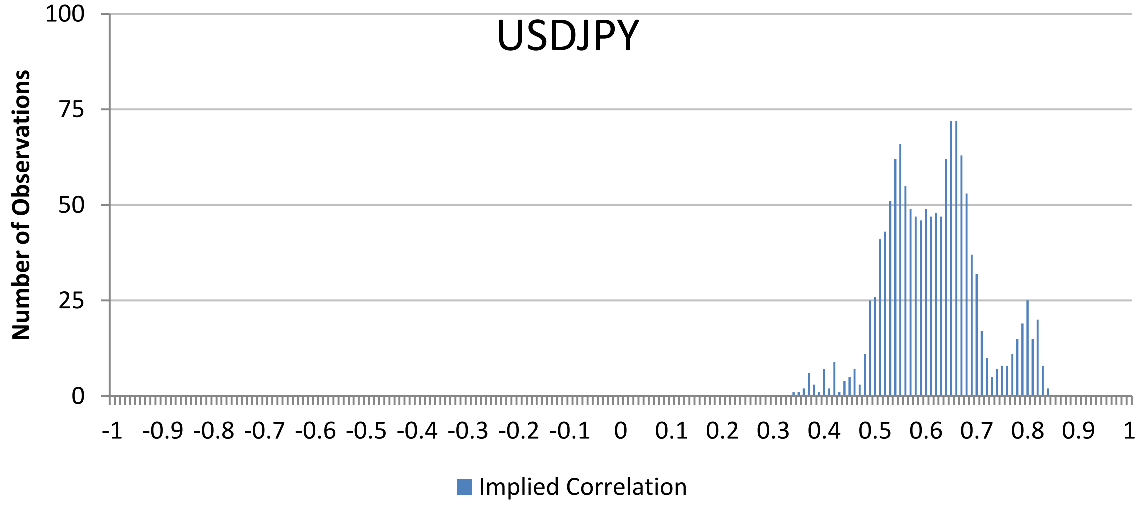

Figure 73a: The empirical distribution of implied correlation based on implied volatility quotes for GBPJPY, GBPUSD and USDJPY option contracts with 9 months to expiration.

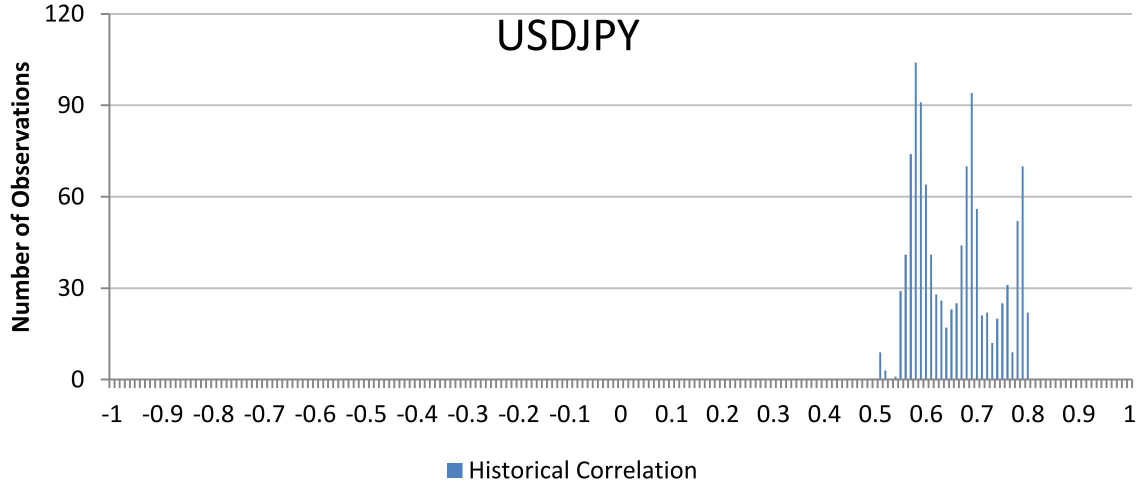

Figure 74a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for GBPJPY and GBPUSD. All estimates are based on a 193 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.610	0.619	0.653	0.660
Mode	0.65/0.66	0.65/0.66	0.58	0.58
Minimum	0.337	0.337	0.503	0.503
Maximum	0.833	0.833	0.794	0.794
Number of observations	1326	1326	1124	1124

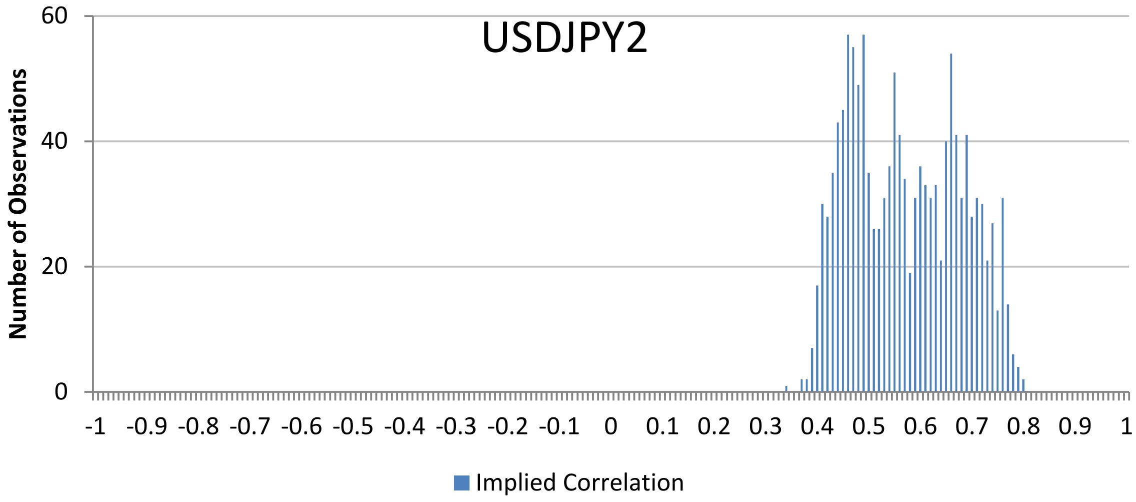

Figure 75a: The empirical distribution of implied correlation based on implied volatility quotes for EURJPY, EURUSD and USDJPY option contracts with 9 months to expiration.

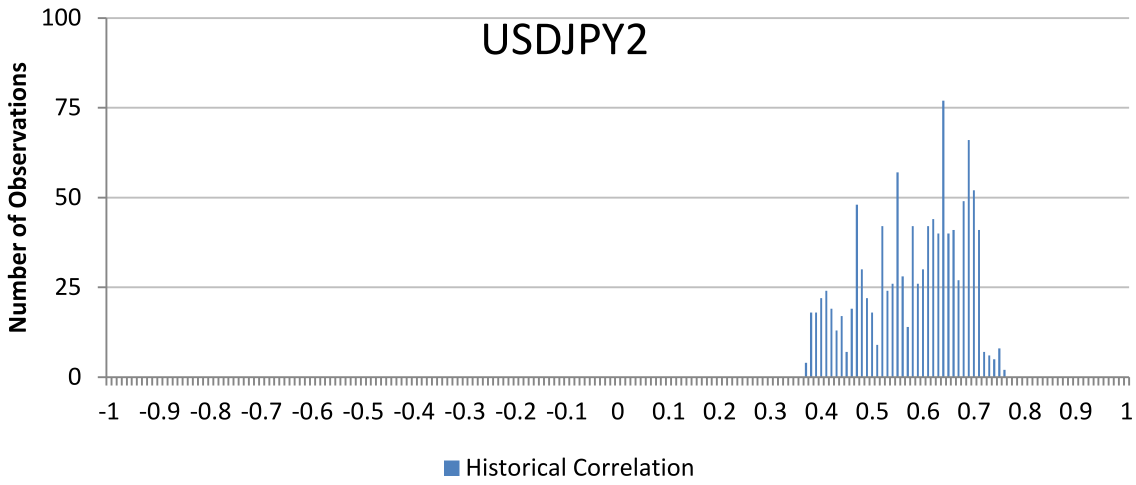

Figure 76a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURJPY and EURUSD. All estimates are based on a 193 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDJPY2				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.567	0.577	0.576	0.584
Mode	0.46/0.49	0.46/0.49	0.64	0.64
Minimum	0.340	0.340	0.368	0.368
Maximum	0.795	0.795	0.755	0.755
Number of observations	1326	1326	1124	1124

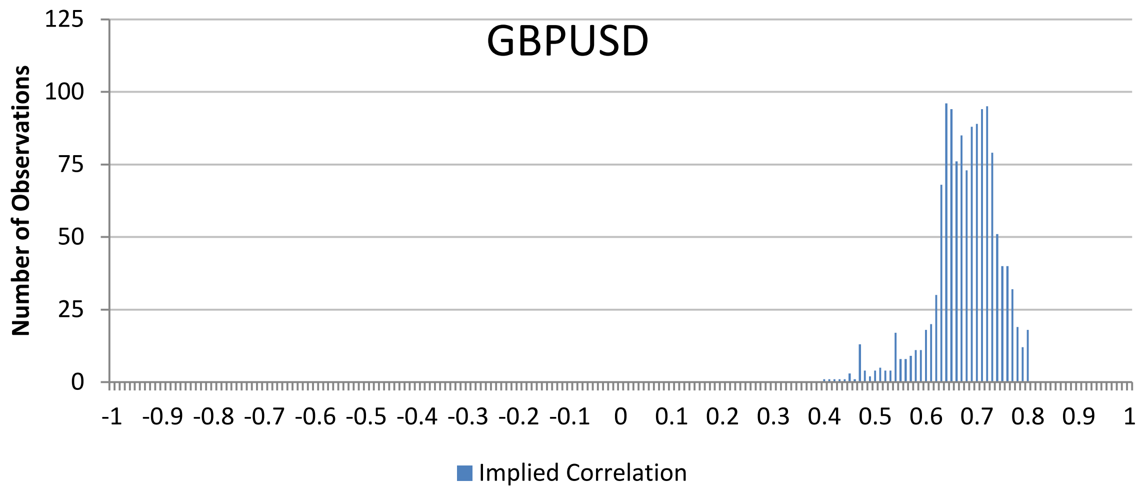

Figure 77a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPJPY and GBPUSD option contracts with 9 months to expiration.

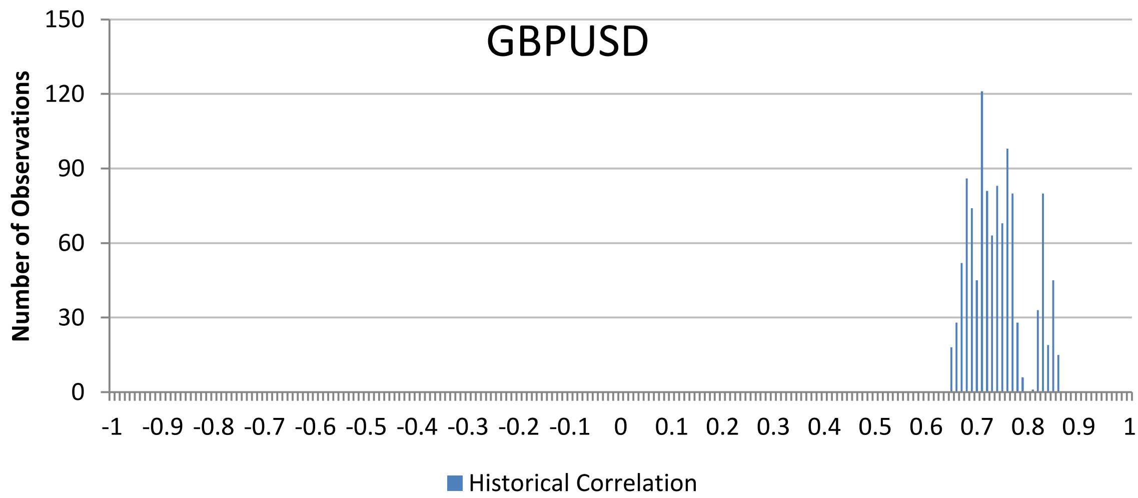

Figure 78a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for GBPJPY and USDJPY. All estimates are based on a 193 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: GBPUSD				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.673	0.678	0.736	0.741
Mode	0.64	0.64	0.71	0.71
Minimum	0.394	0.394	0.640	0.640
Maximum	0.800	0.800	0.859	0.859
Number of observations	1326	1326	1124	1124

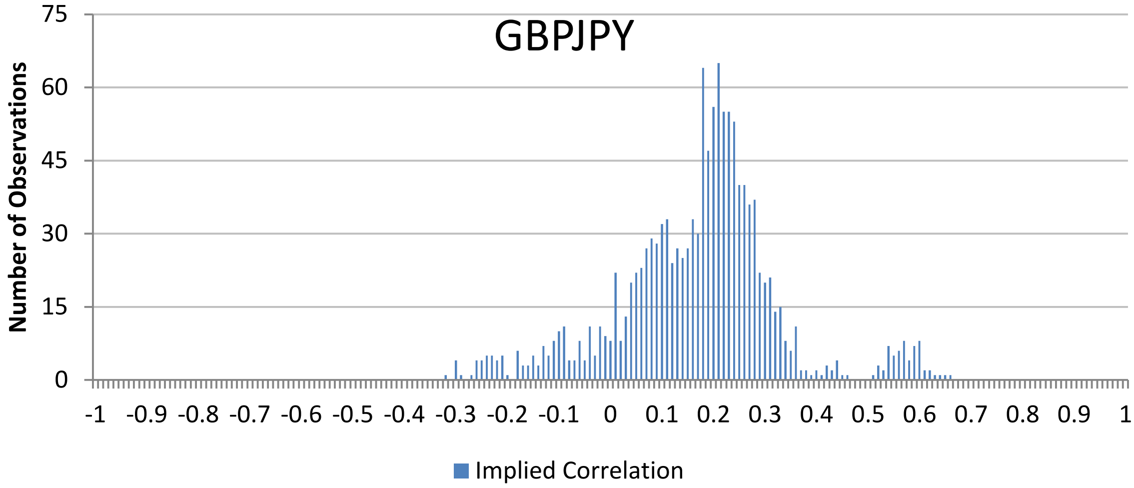

Figure 79a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPUSD and GBPJPY option contracts with 9 months to expiration.

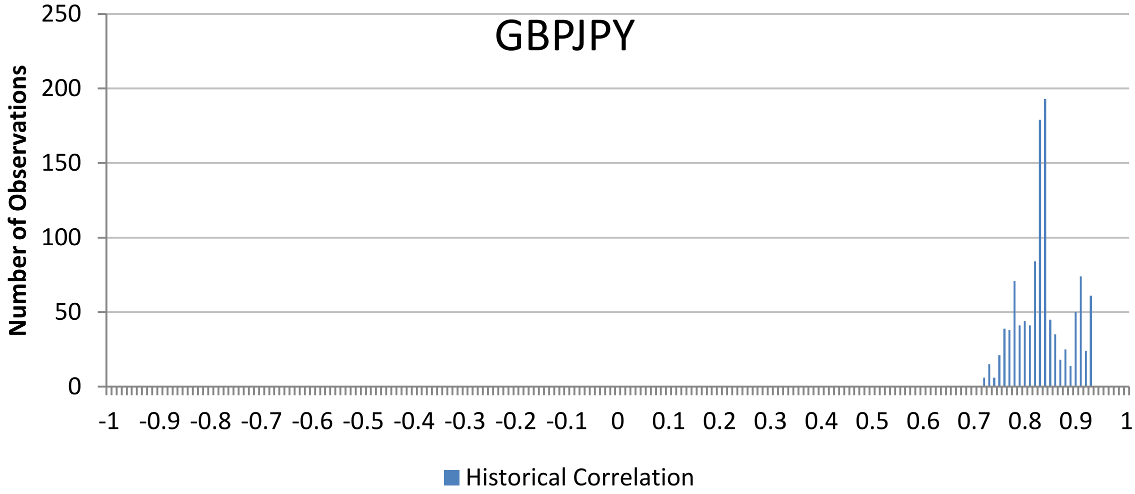

Figure 80a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPUSD. All estimates are based on a 193 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: GBPJPY				
	Implied correlation		Historical Correlation	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.165	0.170	0.831	0.838
Mode	0.21	0.21	0.83	0.83
Minimum	-0.328	-0.328	0.713	0.713
Maximum	0.651	0.651	0.927	0.927
Number of observations	1326	1326	1124	1124

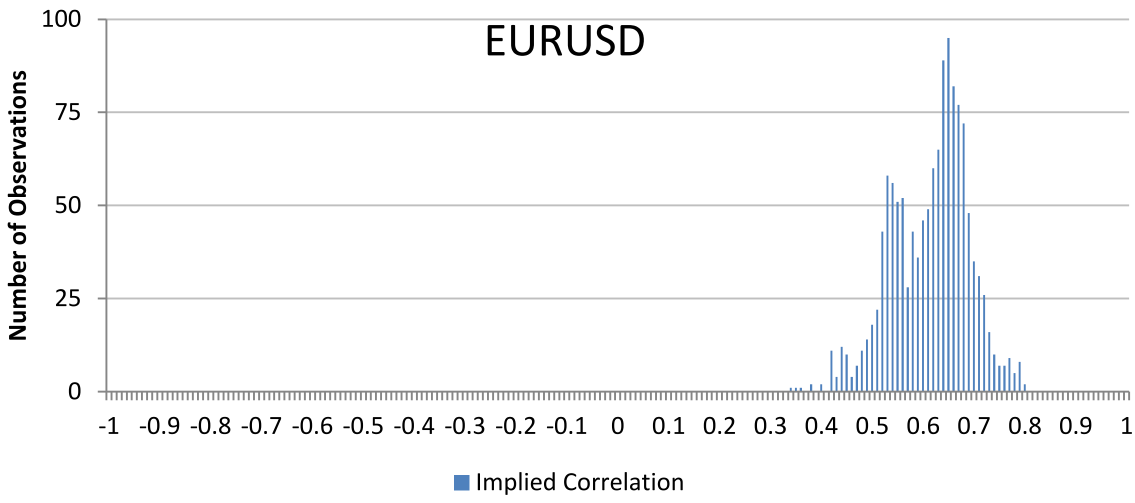

Figure 81a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURJPY and EURUSD option contracts with 9 months to expiration.

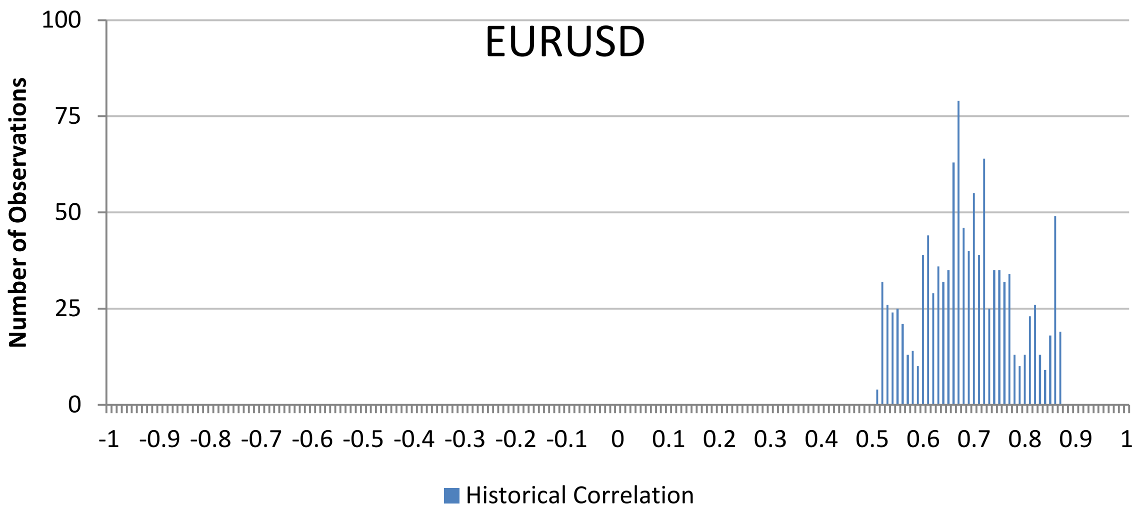

Figure 82a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and EURJPY. All estimates are based on a 193 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURUSD				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.609	0.615	0.683	0.696
Mode	0.65	0.65	0.67	0.67
Minimum	0.332	0.332	0.504	0.504
Maximum	0.798	0.798	0.866	0.866
Number of observations	1326	1326	1124	1124

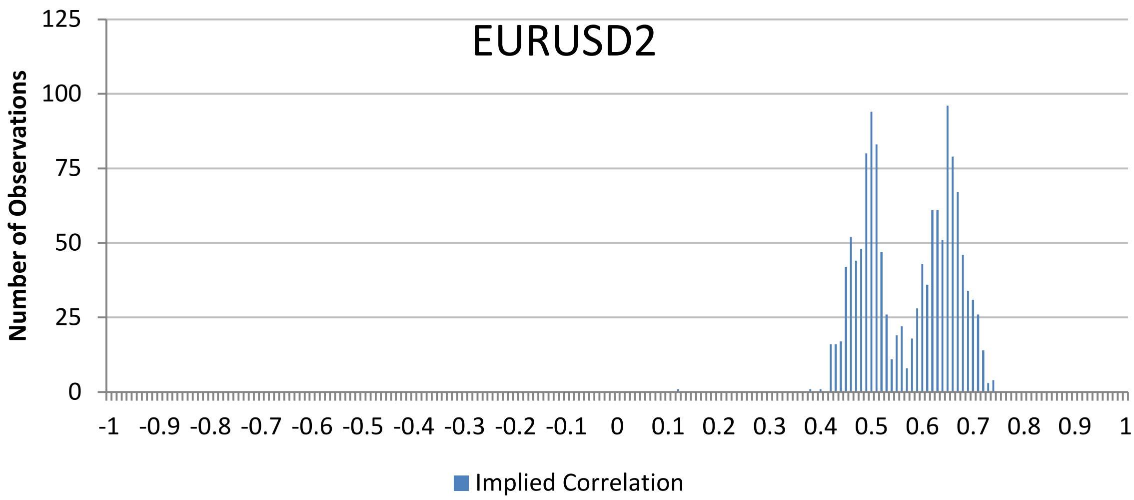

Figure 83a: The empirical distribution of implied correlation based on implied volatility quotes for USDNOK, EURNOK and EURUSD option contracts with 9 months to expiration.

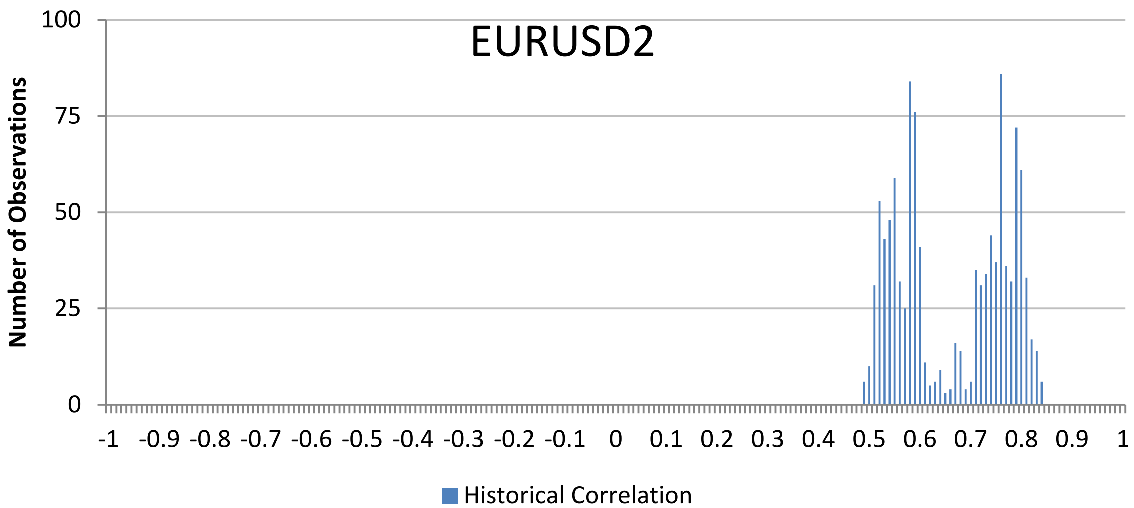

Figure 84a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURNOK and USDNOK. All estimates are based on a 193 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURUSD2				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.569	0.576	0.660	0.674
Mode	0.65	0.65	0.76	0.76
Minimum	0.116	0.116	0.488	0.488
Maximum	0.738	0.738	0.836	0.836
Number of observations	1326	1326	1124	1124

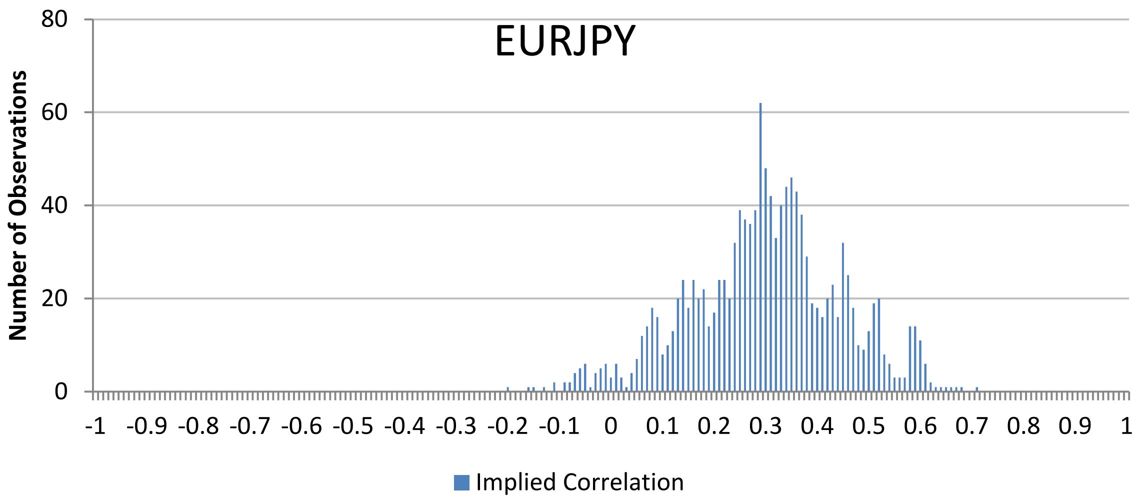

Figure 85a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURUSD and EURJPY option contracts with 9 months to expiration.

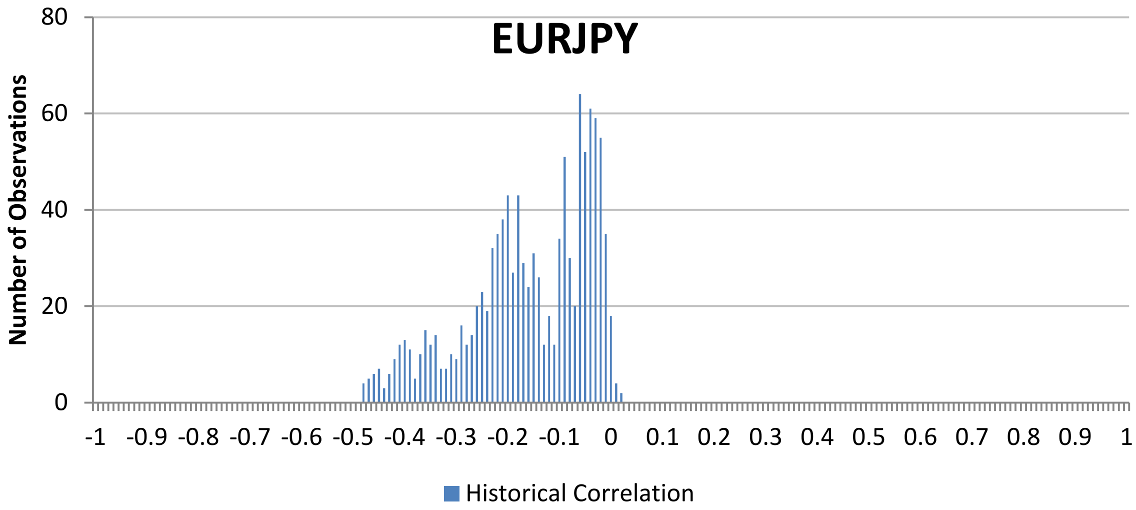

Figure 86a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and EURUSD. All estimates are based on a 193 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.295	0.302	-0.166	-0.169
Mode	0.29	0.29	-0.06	-0.06
Minimum	-0.200	-0.200	-0.485	-0.485
Maximum	0.704	0.704	0.018	0.018
Number of observations	1326	1326	1124	1124

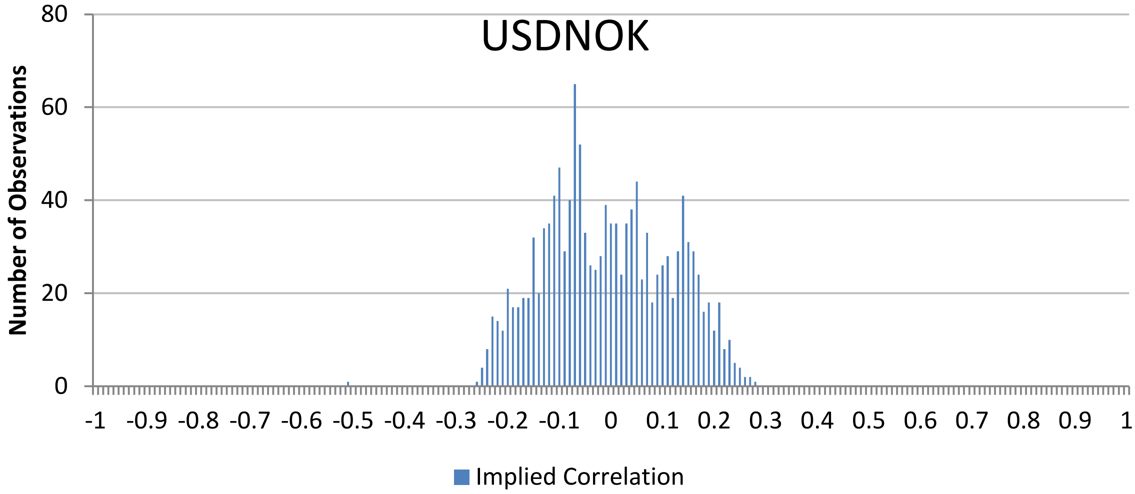

Figure 87a: The empirical distribution of implied correlation based on implied volatility quotes for EURNOK, EURUSD and USDNOK option contracts with 9 months to expiration.

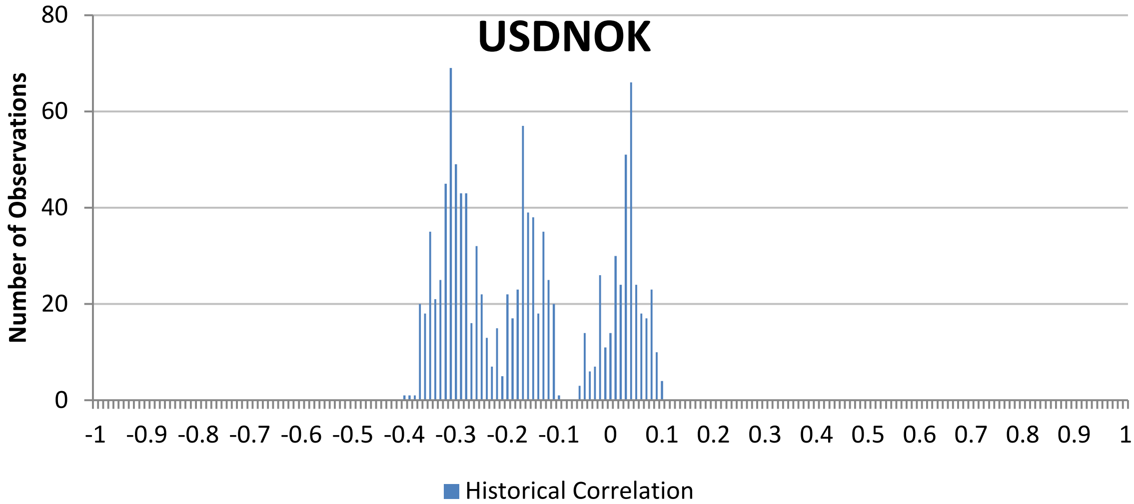

Figure 88a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURNOK and EURUSD. All estimates are based on a 193 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDNOK				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	-0.012	-0.012	-0.166	-0.169
Mode	-0.07	-0.07	-0.31	-0.31
Minimum	-0.511	-0.511	-0.409	-0.409
Maximum	0.271	0.271	0.097	0.097
Number of observations	1326	1326	1124	1124

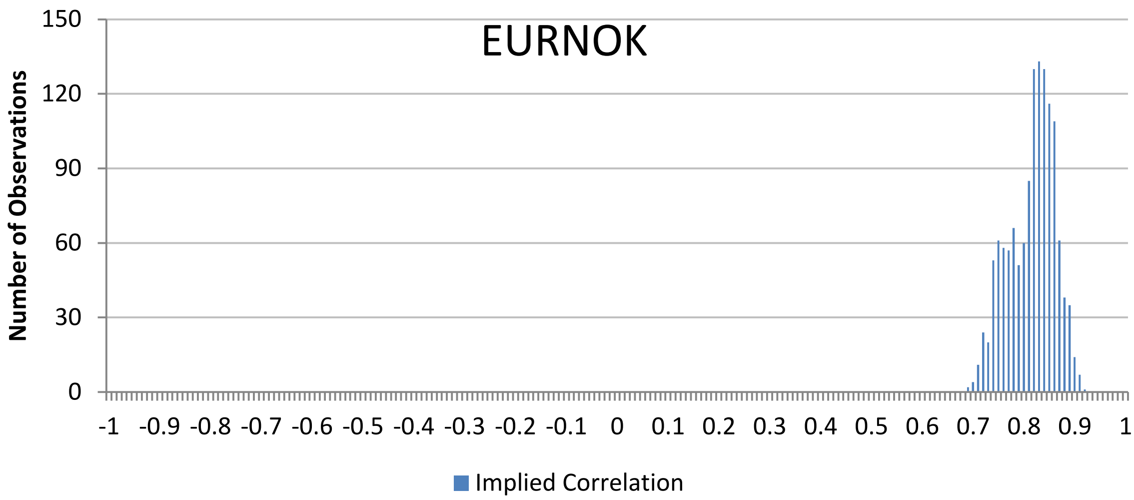

Figure 89a: The empirical distribution of implied correlation based on implied volatility quotes for EURUSD, USDNOK and EURNOK option contracts with 9 months to expiration.

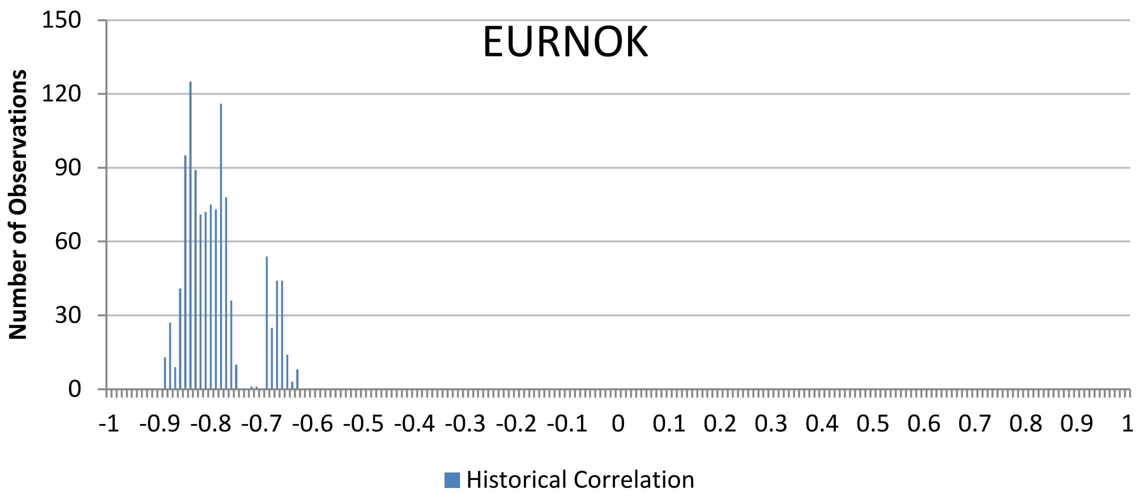

Figure 90a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDNOK and EURUSD. All estimates are based on a 193 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURNOK				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.811	0.815	-0.795	-0.803
Mode	0.65	0.65	-0.84	-0.84
Minimum	0.687	0.687	-0.898	-0.898
Maximum	0.911	0.911	-0.634	-0.634
Number of observations	1326	1326	1124	1124

7.6 Time horizon: 12 months

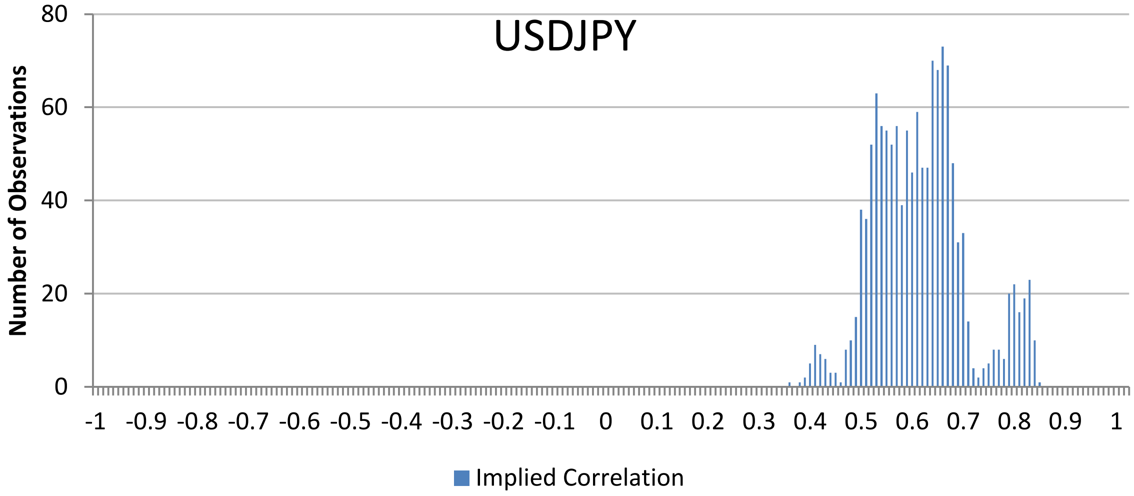

Figure 91a: The empirical distribution of implied correlation based on implied volatility quotes for GBPJPY, GBPUSD and USDJPY option contracts with 12 months to expiration.

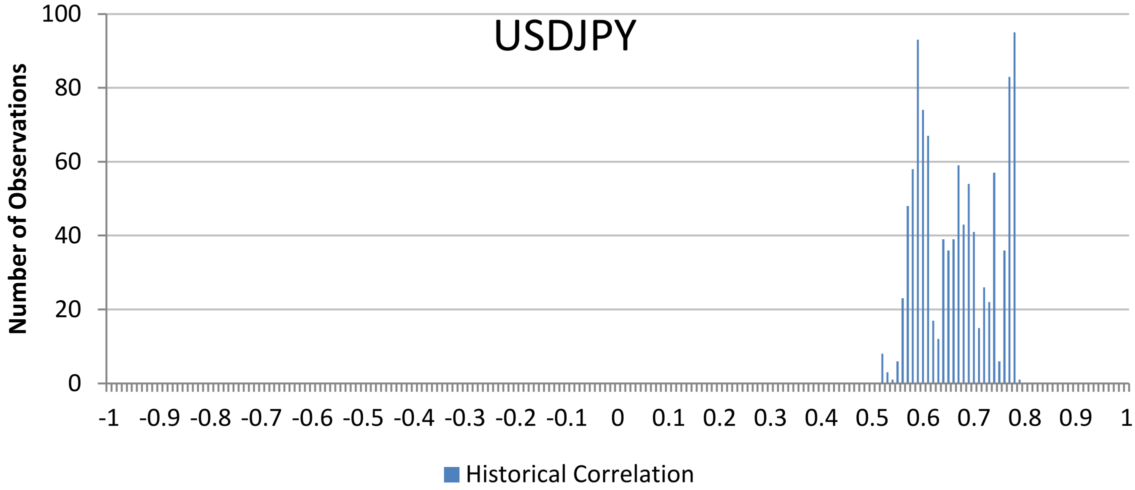

Figure 92a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for GBPJPY and GBPUSD. All estimates are based on a 256 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.611	0.621	0.662	0.669
Mode	0.66	0.66	0.78	0.78
Minimum	0.350	0.350	0.512	0.512
Maximum	0.846	0.846	0.780	0.780
Number of observations	1326	1326	1062	1062

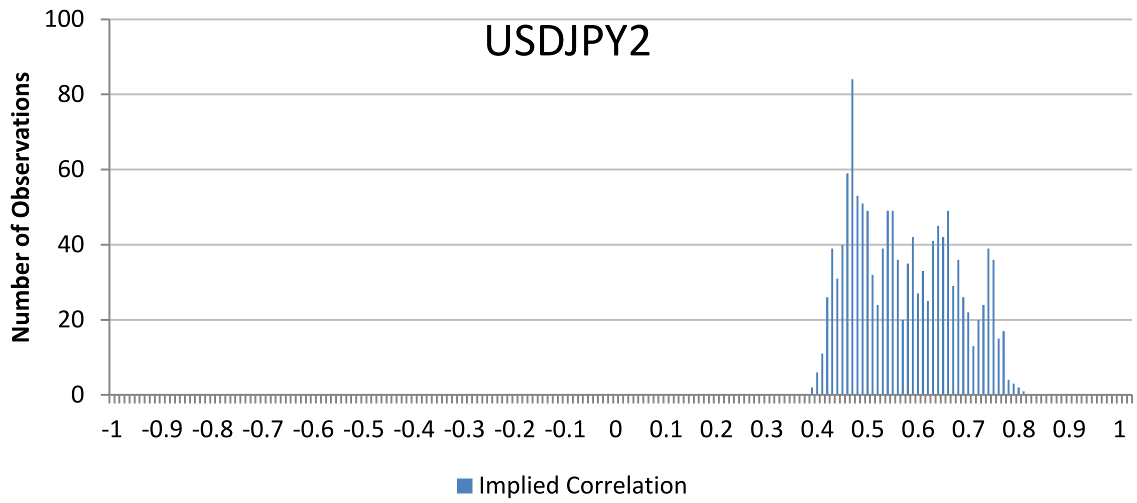

Figure 93a: The empirical distribution of implied correlation based on implied volatility quotes for EURJPY, EURUSD and USDJPY option contracts with 12 months to expiration.

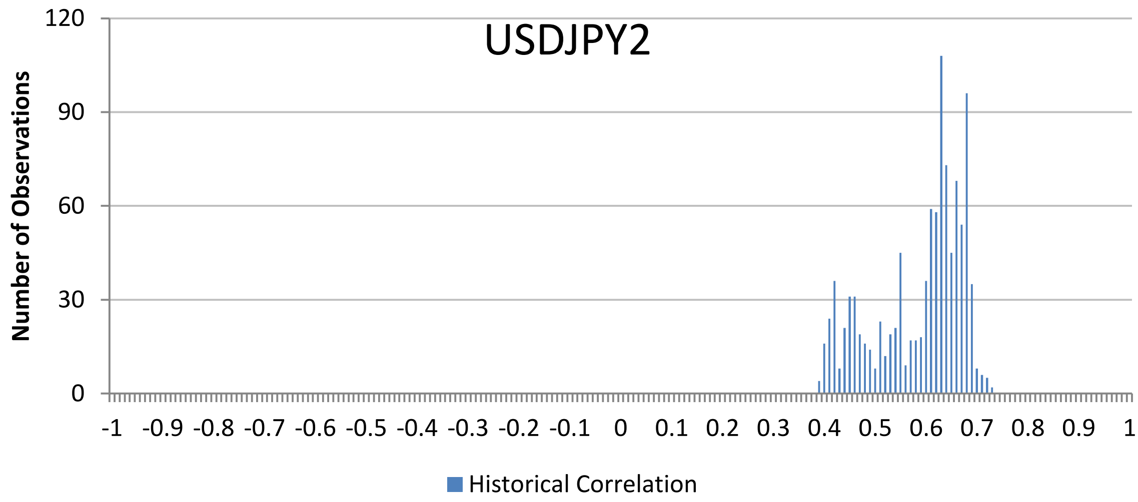

Figure 94a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURJPY and EURUSD. All estimates are based on a 256 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDJPY2				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.569	0.578	0.582	0.588
Mode	0.47	0.47	0.63	0.63
Minimum	0.385	0.385	0.381	0.381
Maximum	0.804	0.804	0.724	0.724
Number of observations	1326	1326	1062	1062

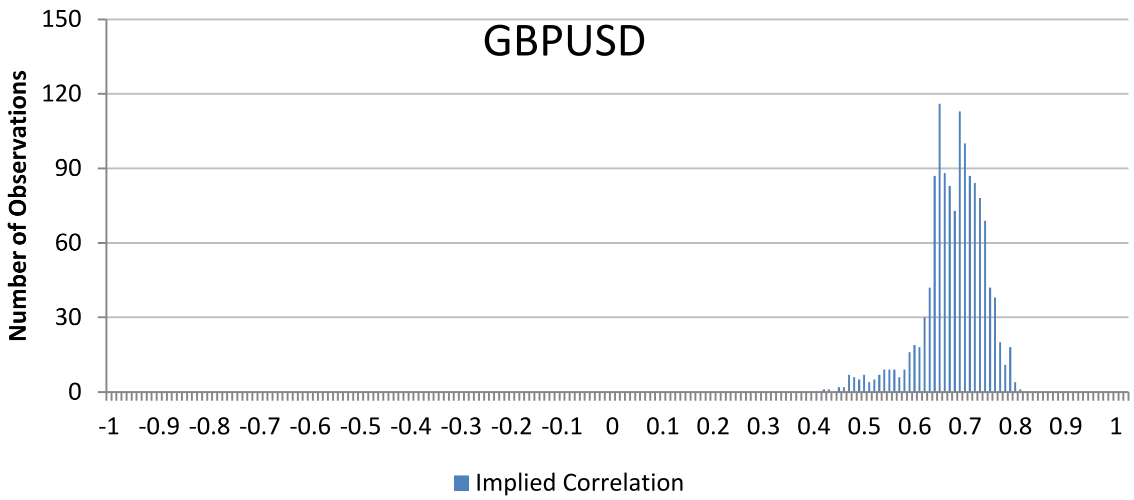

Figure 95a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPJPY and GBPUSD option contracts with 12 months to expiration.

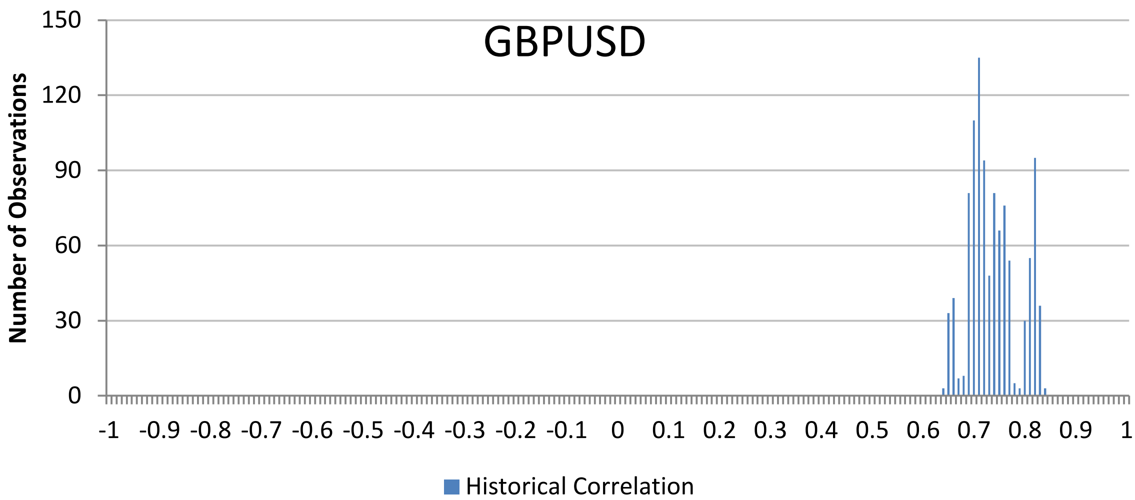

Figure 96a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPJPY. All estimates are based on a 256 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: GBPUSD				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.673	0.677	0.734	0.738
Mode	0.65	0.65	0.71	0.71
Minimum	0.411	0.411	0.638	0.638
Maximum	0.808	0.808	0.832	0.832
Number of observations	1326	1326	1062	1062

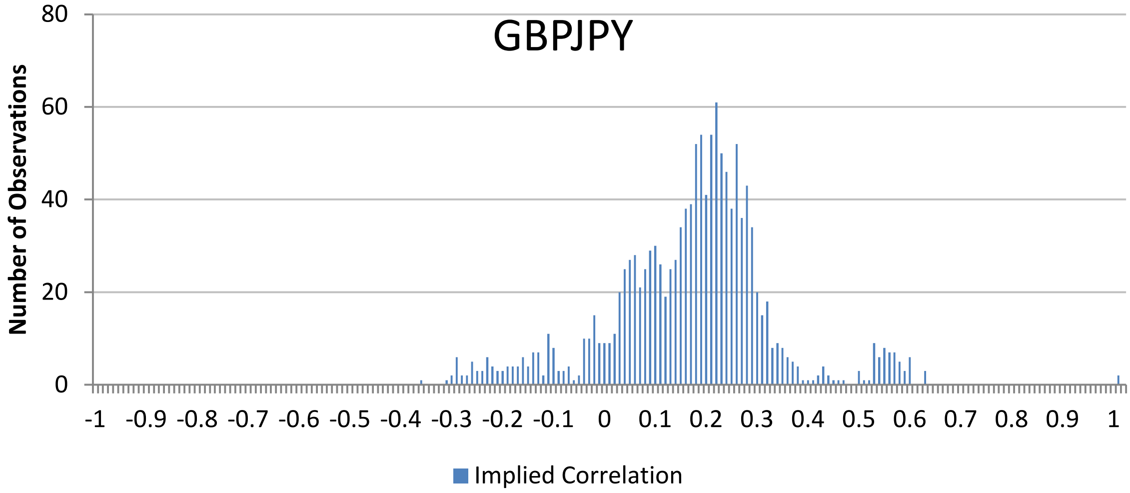

Figure 97a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, GBPUSD and GBPJPY option contracts with 12 months to expiration.

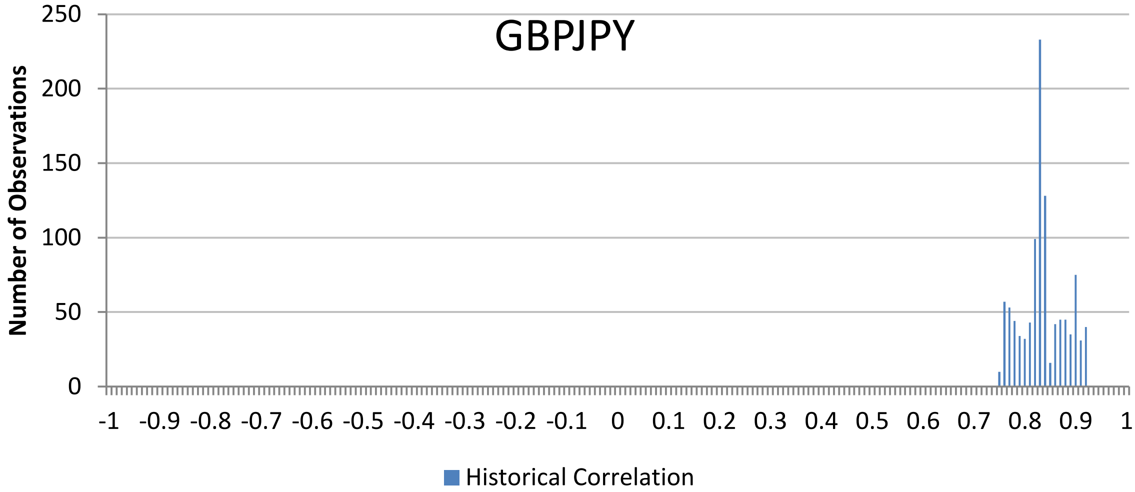

Figure 98a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and GBPUSD. All estimates are based on a 256 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: GBPJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.165	0.168	0.831	0.836
Mode	0.22	0.22	0.83	0.83
Minimum	-0.369	-0.369	0.747	0.747
Maximum	1.008	1.008	0.918	0.918
Number of observations	1326	1326	1062	1062

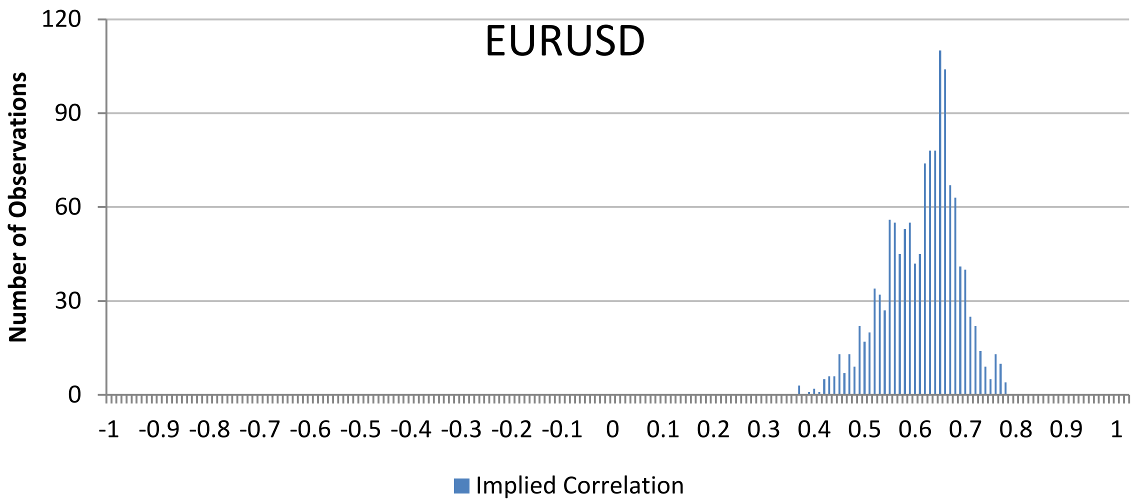

Figure 99a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURJPY and EURUSD option contracts with 12 months to expiration.

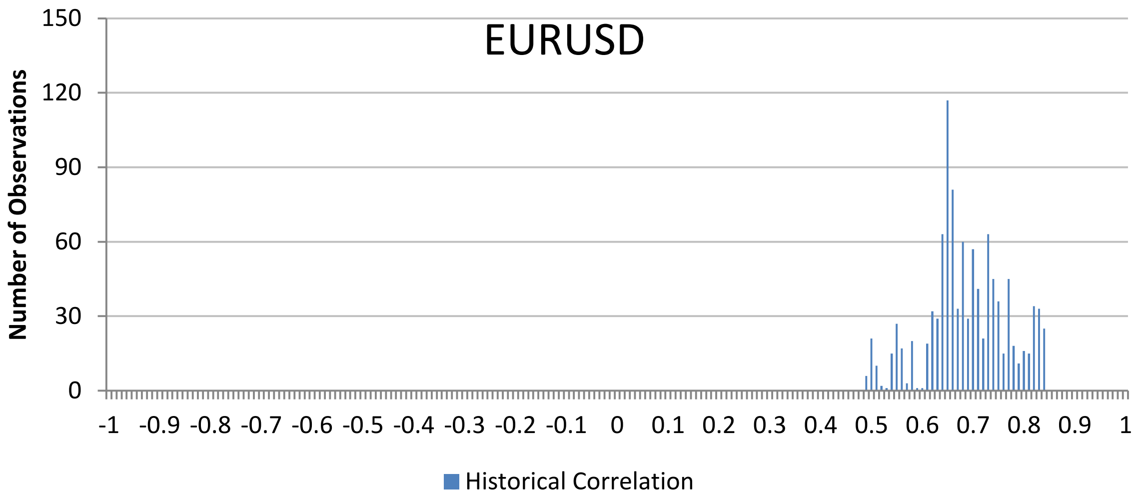

Figure 100a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and EURJPY. All estimates are based on a 256 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURUSD				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.610	0.615	0.682	0.691
Mode	0.65	0.65	0.65	0.65
Minimum	0.367	0.367	0.481	0.481
Maximum	0.777	0.777	0.837	0.837
Number of observations	1326	1326	1062	1062

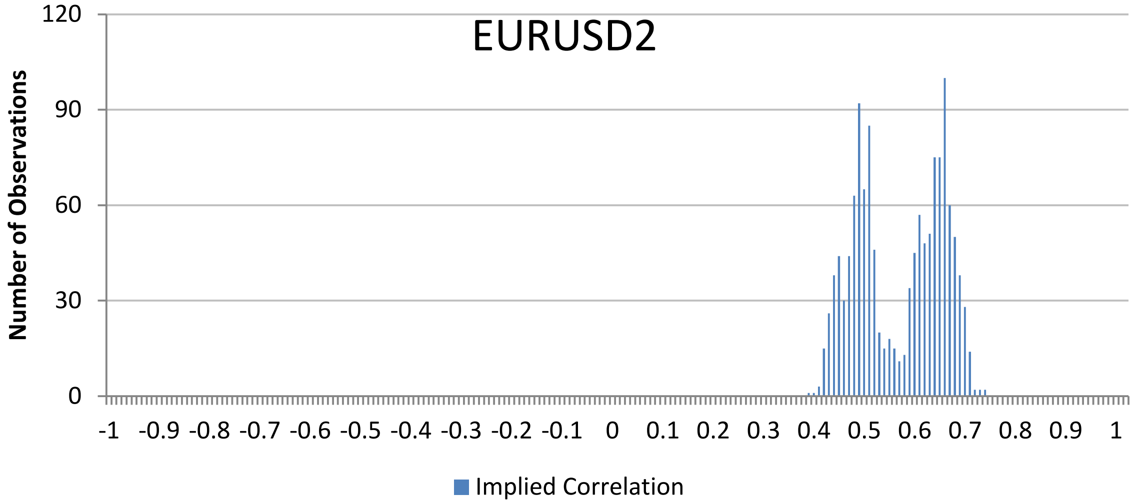

Figure 101a: The empirical distribution of implied correlation based on implied volatility quotes for USDNOK, EURNOK and EURUSD option contracts with 12 months to expiration.

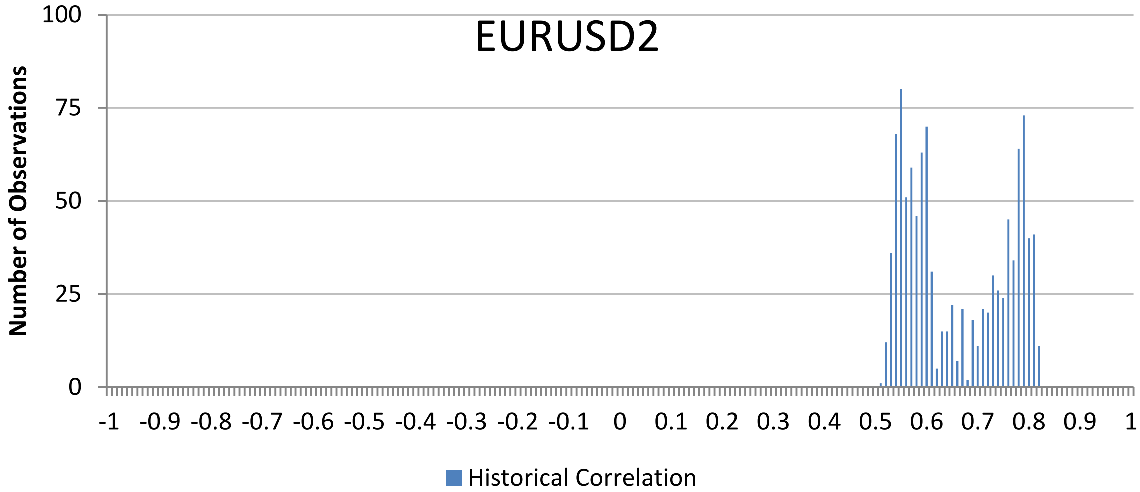

Figure 102a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURNOK and USDNOK. All estimates are based on a 256 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURUSD2				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.566	0.572	0.655	0.667
Mode	0.66	0.66	0.55	0.55
Minimum	0.385	0.385	0.508	0.508
Maximum	0.738	0.738	0.816	0.816
Number of observations	1326	1326	1062	1062

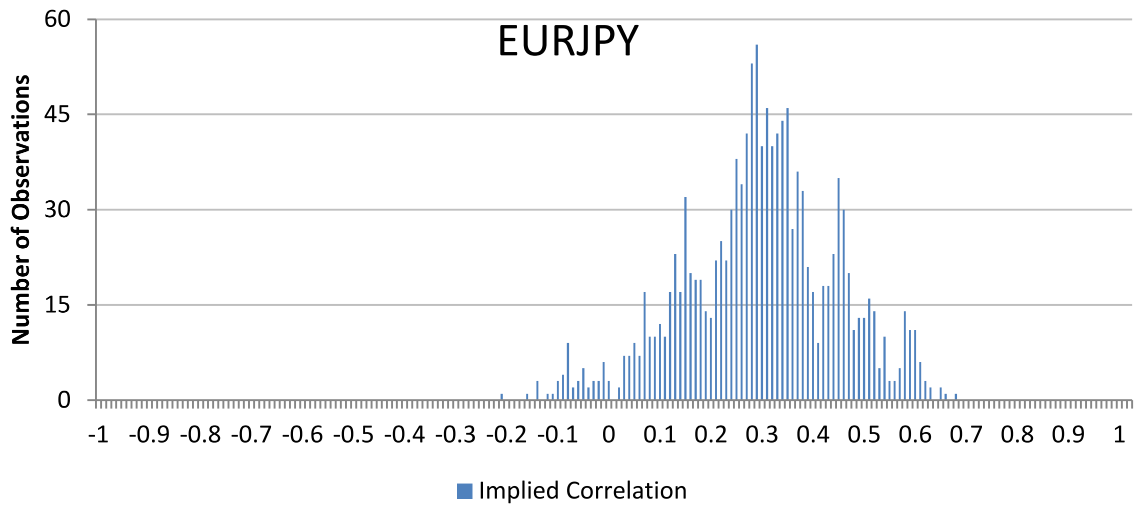

Figure 103a: The empirical distribution of implied correlation based on implied volatility quotes for USDJPY, EURUSD and EURJPY option contracts with 12 months to expiration.

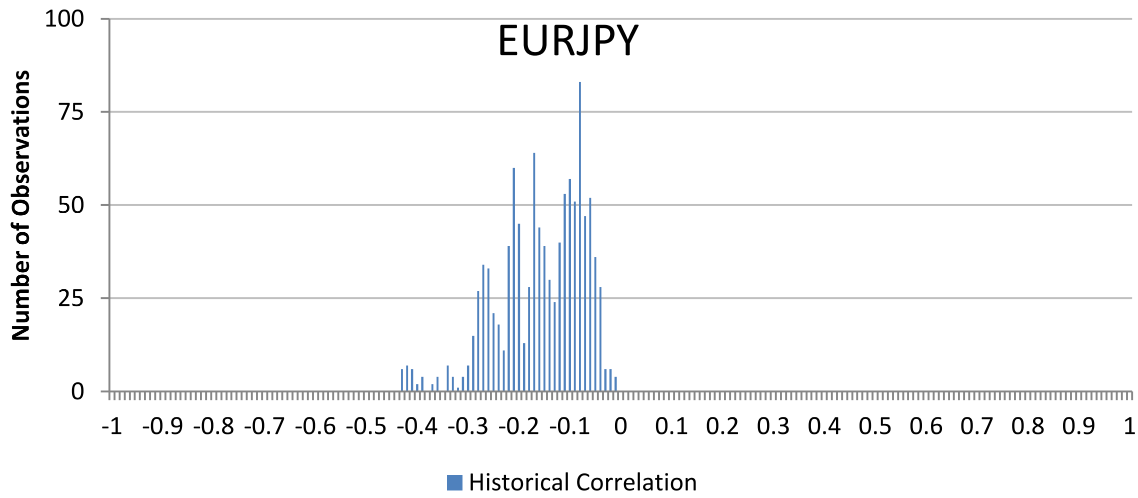

Figure 104a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDJPY and EURUSD. All estimates are based on a 256 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURJPY				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.293	0.300	-0.163	-0.164
Mode	0.29	0.29	-0.08	- 0.08
Minimum	-0.216	-0.216	-0.435	-0.435
Maximum	0.678	0.678	-0.018	-0.018
Number of observations	1326	1326	1062	1062

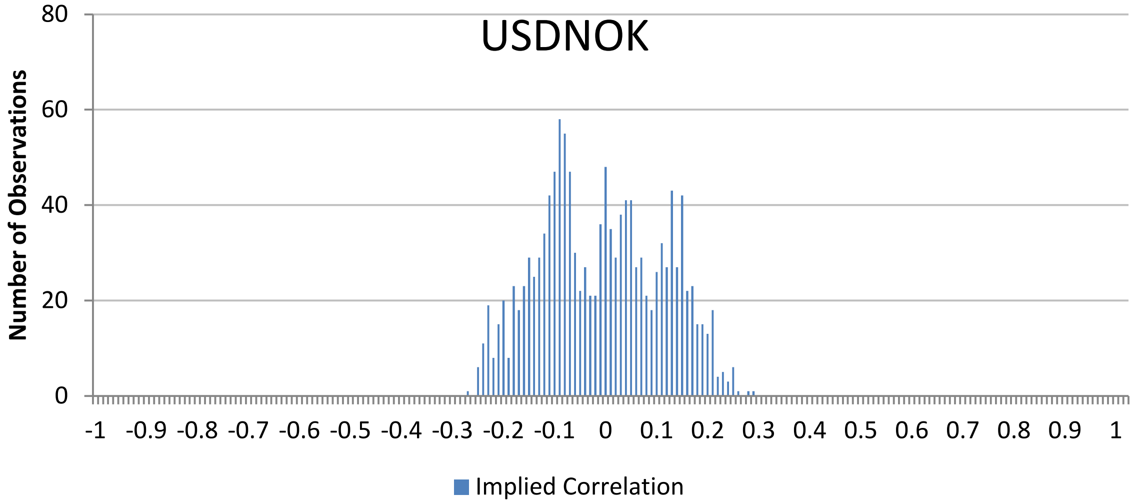

Figure 105a: The empirical distribution of implied correlation based on implied volatility quotes for EURNOK, EURUSD and USDNOK option contracts with 12 months to expiration.

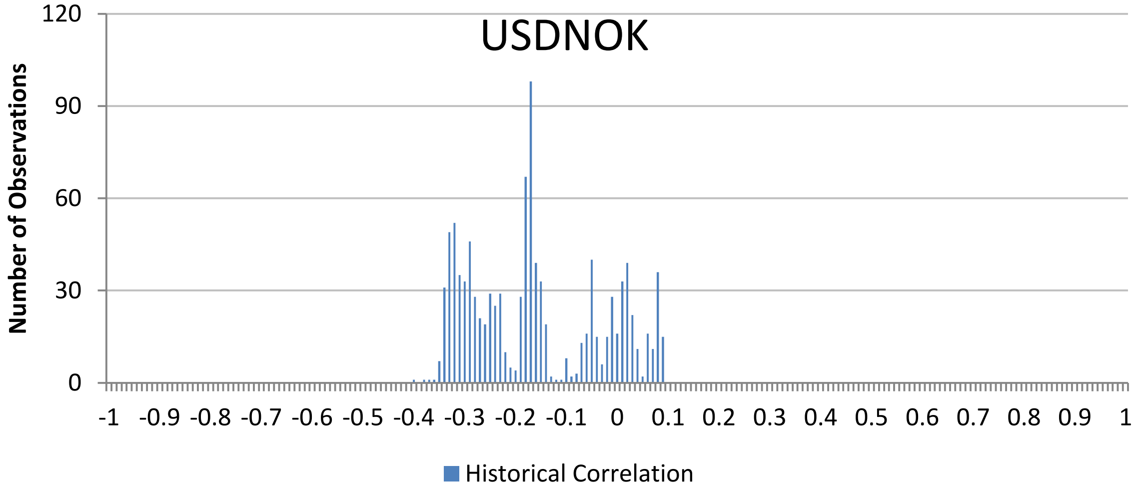

Figure 106a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for EURNOK and EURUSD. All estimates are based on a 256 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: USDNOK				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	-0.013	-0.013	-0.165	-0.168
Mode	-0.09	-0.09	-0.17	-0.17
Minimum	-0.279	-0.279	-0.403	-0.403
Maximum	0.280	0.280	0.086	0.086
Number of observations	1326	1326	1062	1062

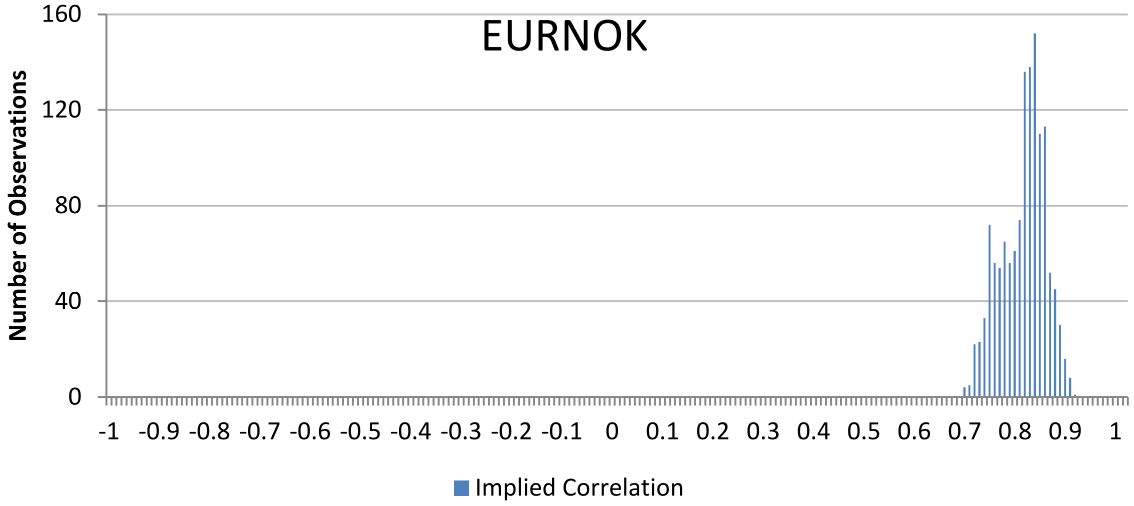

Figure 107a: The empirical distribution of implied correlation based on implied volatility quotes for EURUSD, USDNOK and EURNOK option contracts with 12 months to expiration.

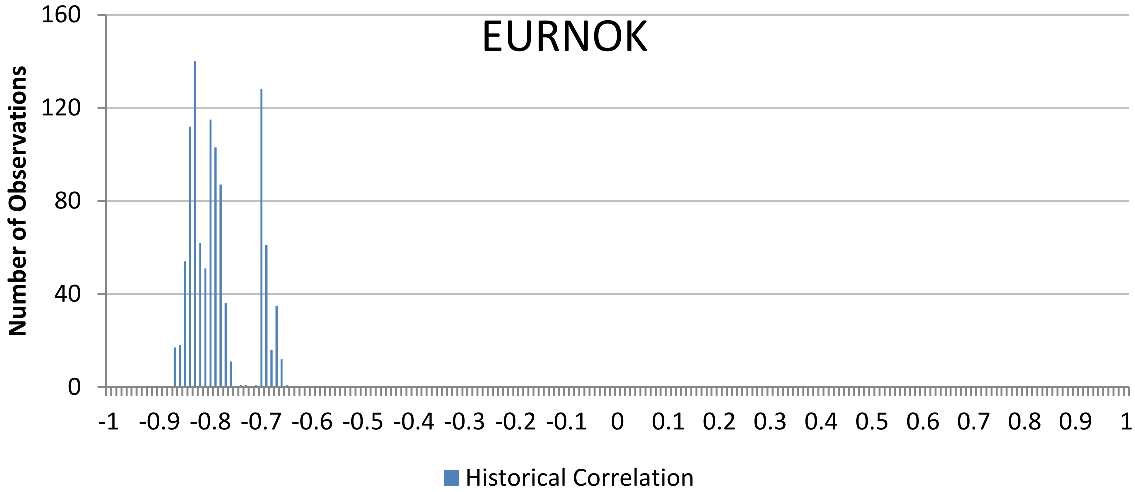

Figure 108a: The empirical distribution of historical correlation, calculated from daily log-returns of exchange rate quotes for USDNOK and EURUSD. All estimates are based on a 256 day rolling window. The red line represents the theoretical sample correlation distribution, and it is based on the mean of the actual data set.

Descriptive statistics: EURNOK				
	Implied correlation		**Historical Correlation**	
	r_{raw}	$r_{transform}$	r_{raw}	$r_{transform}$
Average	0.812	0.817	-0.789	-0.795
Mode	0.84	0.84	-0.83	-0.83
Minimum	0.692	0.692	-0.878	-0.878
Maximum	0.910	0.910	-0.659	-0.659
Number of observations	1326	1326	1062	1062

8 Appendix B

Table 1b: Error statistics for the 1 month correlation forecasts

	USDJPY		USDJPY2		GBPUSD	
2006-2011	**MFE**	**RMSFE**	**MFE**	**RMSFE**	**MFE**	**RMSFE**
Implied Correlation	-0.084 (-0,37)	0.295 (1.04)	-0.059 (-0,26)	**0.321**	-0.006 (-0,03)	0.347 (1.02)
Historical (n=22)	-0.018 (-0,06)	0.383* (1.38)	-0.022 (-0,07)	0.434* (1.40)	-0.002 (-0,01)	0.430* (1.29)
Historical (n=65)	-0.026 (-0,10)	**0.285**	-0.033 (-0,13)	0.342 (1.07)	0.014 (0,06)	0.416* (1.25)
Historical (n=128)	-0.028 (-0,12)	0.320* (1.13)	-0.075 (-0,31)	0.380* (1.20)	0.046 (0,19)	**0.320**
2006-2007						
Implied correlation	-0.096 (-0,42)	**0.306**	-0.066 (-0,29)	**0.349**	-0.020 (-0,09)	**0.387**
Historical (n=22)	-0.032 (-0,10)	0.451* (1.54)	-0.042 (-0,13)	0.528* (1.61)	-0.041 (-0,13)	0.507* (1.37)
Historical (n=65)	-0.033 (-0,13)	0.317 (1.04)	-0.014 (-0,05)	0.415 (1.21)	-0.004 (-0,01)	0.572* (1.59)
Historical (n=128)	-0.125 (-0,52)	0.333 (1.09)	-0.098 (-0,40)	0.434* (1.28)	-0.142 (-0,59)	0.490* (1.31)
2008-2009						
Implied correlation	-0.106 (-0,47)	0.331* (1.16)	-0.062 (-0,27)	**0.329**	0.030 (0,13)	0.302 (1.10)
Historical (n=22)	-0.002 (-0,01)	0.409* (1.46)	0.015 (0,05)	0.424* (1.33)	-0.001 (0,00)	0.369* (1.36)
Historical (n=65)	-0.003 (-0,01)	**0.288**	-0.001 (0,00)	0.333 (1.01)	-0.017 (-0,07)	0.304 (1.11)
Historical (n=128)	-0.012 (-0,05)	0.358* (1.26)	0.005 (0,02)	0.401* (1.24)	0.063 (0,26)	**0.276**
2010-2011						
Implied correlation	-0.051 (-0,22)	**0.235**	-0.053 (-0,23)	**0.289**	-0.040 (-0,18)	0.364 (1.04)
Historical (n=22)	-0.027 (-0,08)	0.285* (1.23)	-0.051 (-0,16)	0.363* (1.28)	0.023 (0,07)	0.435* (1.27)
Historical (n=65)	-0.048 (-0,19)	0.260 (1.11)	-0.080 (-0,31)	0.303 (1.05)	0.061 (0,23)	0.402* (1.16)
Historical (n=128)	-0.019 (-0,08)	0.263 (1.13)	-0.159 (-0,66)	0.334* (1.17)	0.078 (0,32)	**0.351**

* Indicates statistical significance at the 5 % level, F-values in brackets

Number of observations: 2006-2011: 1228, 2006-2007: 252, 2008-2009: 522, 2010-2011: 454

H_0: The mean forecast error (MFE) is zero

H_0: The forecasts are identical in the ability to predict future realized correlation, i.e. $\frac{RMSFE}{RMSFE_{lowest}} = 1$

Table 2b: Error statistics for the 3 month correlation forecasts

	USDJPY		**USDJPY2**		**GBPUSD**	
2006-2011	**MFE**	**RMSFE**	**MFE**	**RMSFE**	**MFE**	**RMSFE**
Implied Correlation	-0.057 (-0,45)	**0.190**	-0.034 (-0,27)	**0.210**	-0.073 (-0,58)	0.276* (1.25)
Historical (n=22)	0.016 (0,06)	0.293* (1.57)	0.004 (0,02)	0.339* (1.65)	-0.019 (-0,07)	0.419* (1.97)
Historical (n=65)	-0.002 (-0,01)	0.201 (1.06)	-0.027 (-0,15)	0.251* (1.20)	0.010 (0,06)	0.339* (1.55)
Historical (n=128)	-0.002 (-0,01)	0.227* (1.20)	-0.060 (-0,39)	0.278* (1.34)	0.053 (0,34)	**0.223**
2006-2007						
Implied correlation	-0.099 (-0,78)	0.157* (1.52)	-0.052 (-0,41)	0.167 (1.16)	-0.175 (-1,40)	0.353 (1.15)
Historical (n=22)	0.027 (0,10)	0.292* (2.89)	0.008 (0,03)	0.382* (2.77)	-0.143 (-0,56)	0.576* (2.05)
Historical (n=65)	-0.019 (-0,11)	0.130* (1.26)	-0.036 (-0,20)	0.255* (1.79)	-0.058 (-0,32)	0.507* (1.75)
Historical (n=128)	-0.091 (-0,59)	**0.103**	-0.112 (-0,72)	**0.145**	-0.057 (-0,37)	**0.309**
2008-2009						
Implied correlation	-0.073 (-0,58)	**0.223**	-0.013 (-0,10)	**0.223**	-0.001 (-0,01)	**0.205**
Historical (n=22)	0.002 (0,01)	0.338* (1.55)	0.022 (0,08)	0.362* (1.68)	0.010 (0,04)	0.306* (1.52)
Historical (n=65)	0.001 (0,00)	0.243 (1.09)	0.006 (0,03)	0.272* (1.23)	-0.006 (-0,03)	0.221 (1.08)
Historical (n=128)	-0.009 (-0,06)	0.278* (1.26)	0.012 (0,08)	0.324* (1.48)	0.075 (0,48)	0.206 (1.01)
2010-2011						
Implied correlation	-0.008 (-0,07)	**0.160**	-0.053 (-0,42)	**0.215**	-0.062 (-0,49)	0.278* (1.33)
Historical (n=22)	0.026 (0,10)	0.219* (1.38)	-0.021 (-0,08)	0.266* (1.24)	0.028 (0,11)	0.393* (1.94)
Historical (n=65)	0.004 (0,02)	0.175 (1.09)	-0.064 (-0,36)	0.218 (1.01)	0.072 (0,40)	0.310* (1.49)
Historical (n=128)	0.035 (0,22)	0.174 (1.09)	-0.135 (-0,88)	0.240 (1.12)	0.058 (0,37)	**0.211**

* Indicates statistical significance at the 5 % level, F-values in brackets

Number of observations: 2006-2011: 1187, 2006-2007: 211, 2008-2009: 522, 2010-2011: 454

H_0: The mean forecast error (MFE) is zero

H_0: The forecasts are identical in the ability to predict future realized correlation, i.e. $\frac{RMSFE}{RMSFE_{lowest}} = 1$

Table 3b: Error statistics for the 6 month correlation forecasts

	USDJPY		USDJPY2		GBPUSD	
2006-2011	**MFE**	**RMSFE**	**MFE**	**RMSFE**	**MFE**	**RMSFE**
Implied Correlation	-0.056 (-0,63)	**0.195**	-0.034 (-0,39)	**0.208**	-0.115 (-1,30)	0.230* (1.24)
Historical (n=22)	-0.004 (-0,02)	0.315* (1.65)	-0.036 (-0,15)	0.354* (1.75)	-0.025 (-0,10)	0.372* (2.07)
Historical (n=65)	-0.014 (-0,09)	0.213 (1.09)	-0.052 (-0,33)	0.257* (1.24)	-0.012 (-0,08)	0.262* (1.42)
Historical (n=128)	-0.009 (-0,07)	0.201 (1.03)	-0.056 (-0,44)	0.245* (1.18)	0.025 (0,20)	**0.187**
2006-2007						
Implied correlation	-0.075 (-0,84)	0.119* (1.54)	0.003 0,03	**0.119**	-0.289* (-3,33)	0.332 (1.13)
Historical (n=22)	0.005 (0,02)	0.301* (4.00)	0.000 0,00	0.361* (3.18)	-0.196 (-0,82)	0.478* (1.70)
Historical (n=65)	-0.014 (-0,06)	0.210 (1.18)	-0.052 -0,08	0.252* (1.50)	-0.012 (-0,79)	0.256* (1.38)
Historical (n=128)	-0.046 (-0,37)	**0.077**	-0.035 -0,28	0.144 (1.21)	-0.015 (-0,12)	**0.296**
2008-2009						
Implied correlation	-0.088 (-0,99)	0.250 (1.01)	-0.029 (-0,33)	**0.248**	-0.028 (-0,31)	0.145 (1.01)
Historical (n=22)	-0.039 (-0,16)	0.366* (1.52)	-0.028 (-0,11)	0.393* (1.64)	0.007 (0,03)	0.267* (1.89)
Historical (n=65)	-0.040 (-0,26)	0.267 (1.08)	-0.043 (-0,28)	0.309* (1.26)	-0.009 (-0,06)	0.167* (1.16)
Historical (n=128)	-0.029 (-0,23)	**0.247**	-0.031 (-0,24)	0.290* (1.18)	0.024 (0,19)	**0.143**
2010-2011						
Implied correlation	0.009 (0,10)	**0.149**	-0.076 (-0,85)	**0.205**	-0.083 (-0,93)	0.211* (1.31)
Historical (n=22)	0.042 (0,17)	0.227* (1.54)	-0.080 (-0,33)	0.270* (1.33)	0.074 (0,31)	0.390* (2.52)
Historical (n=65)	0.021 (0,14)	0.178* (1.20)	-0.094 (-0,61)	0.213 (1.04)	0.065 (0,429	0.241* (1.50)
Historical (n=128)	0.041 (0,33)	0.166 (1.11)	-0.103 (-0,82)	0.209 (1.02)	0.048 (0,38)	**0.162**

* Indicates statistical significance at the 5 % level, F-values in brackets

Number of observations: 2006-2011: 1125, 2006-2007: 149, 2008-2009: 522, 2010-2011: 454

H_0: The mean forecast error (MFE) is zero

H_0: The forecasts are identical in the ability to predict future realized correlation, i.e. $\frac{RMSFE}{RMSFE_{lowest}} = 1$

Table 4b: Error statistics for the 1 month correlation forecasts

	GBPJPY		EURUSD		EURUSD2	
2006-2011	**MFE**	**RMSFE**	**MFE**	**RMSFE**	**MFE**	**RMSFE**
Implied Correlation	-0.774* (-4,58)	0.803* (2.71)	-0.044 (-0,20)	**0.343**	-0.153 (-0,68)	**0.314**
Historical (n=22)	0.000 (0,00)	0.429* (1.13)	0.003 (0,01)	0.435* (1.30)	0.003 (0,01)	0.347 (1.11)
Historical (n=65)	0.026 (0,10)	0.414 (1.08)	0.013 (0,05)	0.408* (1.21)	-0.018 (-0,07)	0.336 (1.07)
Historical (n=128)	0.054 (0,22)	**0.386**	0.068 (0,28)	0.398* (1.18)	-0.006 (-0,02)	0.325 (1.04)
2006-2007						
Implied correlation	-0.799* (-4,88)	0.833* (2.39)	-0.076 (-0,34)	**0.341**	-0.184 (-0,82)	0.314 (1.02)
Historical (n=22)	-0.005 (-0,02)	**0.464**	-0.005 (-0,02)	0.497* (1.53)	-0.003 (-0,01)	0.356 (1.17)
Historical (n=65)	0.088 (0,34)	0.488 (1.06)	0.001 (0,00)	0.534* (1.68)	-0.018 (-0,07)	**0.309**
Historical (n=128)	-0.042 (-0,17)	0.499 (1.09)	-0.121 (-0,50)	0.543* (1.71)	-0.172 (-0,72)	0.357 (1.17)
2008-2009						
Implied correlation	-0.797* (-4,85)	0.823* (2.72)	-0.023 (-0,10)	**0.330**	-0.201 (-0,90)	0.341 (1.06)
Historical (n=22)	-0.002 (-0,01)	0.445 (1.12)	-0.010 (-0,03)	0.402* (1.25)	0.001 (0,00)	0.372* (1.16)
Historical (n=65)	0.008 (0,03)	0.427 (1.07)	-0.027 (-0,10)	0.335 (1.02)	-0.039 (-0,15)	0.366 (1.14)
Historical (n=128)	0.096 (0,40)	**0.404**	0.042 (0,17)	0.359 (1.10)	-0.043 (-0,18)	**0.323**
2010-2011						
Implied correlation	-0.724* (-4,07)	0.745* (2.91)	-0.047 (-0,21)	**0.359**	-0.077 (-0,34)	**0.280**
Historical (n=22)	0.006 (0,02)	0.382* (1.22)	0.024 (0,08)	0.424* (1.21)	0.010 (0,03)	0.308 (1.11)
Historical (n=65)	0.012 (0,05)	0.341 (1.08)	0.065 (0,25)	0.391 (1.10)	0.007 (0,03)	0.312 (1.12)
Historical (n=128)	0.033 (0,14)	**0.318**	0.150 (0,62)	0.385 (1.08)	0.084 (0,35)	0.316 (1.14)

* Indicates statistical significance at the 5 % level, F-values in brackets

Number of observations: 2006-2011: 1228, 2006-2007: 252, 2008-2009: 522, 2010-2011: 454

H_0: The mean forecast error (MFE) is zero

H_0: The forecasts are identical in the ability to predict future realized correlation, i.e. $\frac{RMSFE}{RMSFE_{lowest}} = 1$

Table 5b: Error statistics for the 3 month correlation forecasts

	GBPJPY		EURUSD		EURUSD2	
2006-2011	**MFE**	**RMSFE**	**MFE**	**RMSFE**	**MFE**	**RMSFE**
Implied Correlation	-0.782* (-8,31)	0.795* (3.65)	-0.099 (-0,78)	**0.287**	-0.153 (-1,21)	**0.246**
Historical (n=22)	-0.013 (-0,05)	0.427* (1.53)	-0.007 (-0,03)	0.422* (1.52)	0.038 (0,15)	0.337* (1.40)
Historical (n=65)	0.013 (0,07)	0.368* (1.30)	0.018 (0,10)	0.350* (1.24)	0.017 (0,09)	0.294* (1.21)
Historical (n=128)	0.069 (0,45)	**0.289**	0.089 (0,57)	0.309 (1.08)	0.037 (0,24)	0.282* (1.16)
2006-2007						
Implied correlation	-0.814* (-9,03)	0.832* (3.65)	-0.203 (-1,62)	**0.363**	-0.214 (-1,72)	0.258 (1.10)
Historical (n=22)	-0.060 (-0,23)	0.491* (1.65)	-0.092 (-0,36)	0.563* (1.67)	0.001 (0,01)	0.312* (1.34)
Historical (n=65)	0.021 (0,12)	0.466* (1.55)	-0.027 (-0,15)	0.504* (1.56)	-0.022 (-0,12)	**0.236**
Historical (n=128)	0.028 (0,18)	**0.316**	0.020 (0,13)	0.402 (1.12)	-0.175 (-1,14)	0.241 (1.02)
2008-2009						
Implied correlation	-0.797* (-8,63)	0.807* (3.27)	-0.069 (-0,55)	**0.237**	-0.179 (-1,43)	**0.263**
Historical (n=22)	-0.005 (-0,02)	0.453* (1.43)	0.001 (0,01)	0.345* (1.49)	0.064 (0,25)	0.364* (1.42)
Historical (n=65)	0.005 (0,03)	0.394* (1.22)	-0.016 (-0,09)	0.261 (1.11)	0.025 (0,14)	0.332* (1.28)
Historical (n=128)	0.094 (0,61)	**0.329**	0.053 (0,35)	0.307* (1.31)	0.020 (0,13)	0.316* (1.22)
2010-2011						
Implied correlation	-0.720* (-7,19)	0.729* (4.27)	-0.034 (-0,27)	**0.262**	-0.071 (-0,56)	**0.200**
Historical (n=22)	0.009 (0,03)	0.329* (1.57)	0.040 (0,15)	0.380* (1.49)	0.029 (0,11)	0.316* (1.61)
Historical (n=65)	0.017 (0,10)	0.227 (1.07)	0.088 (0,49)	0.315* (1.22)	0.031 (0,17)	0.271* (1.37)
Historical (n=128)	0.051 (0,33)	**0.213**	0.153 (0,99)	0.275 (1.05)	0.121 (0,79)	0.244* (1.22)

* Indicates statistical significance at the 5 % level, F-values in brackets

Number of observations: 2006-2011: 1187, 2006-2007: 211, 2008-2009: 522, 2010-2011: 454

H_0: The mean forecast error (MFE) is zero

H_0: The forecasts are identical in the ability to predict future realized correlation, i.e. $\frac{RMSFE}{RMSFE_{lowest}} = 1$

Table 6b: Error statistics for the 6 month correlation forecasts						
	GBPJPY		**EURUSD**		**EURUSD2**	
2006-2011	**MFE**	**RMSFE**	**MFE**	**RMSFE**	**MFE**	**RMSFE**
Implied Correlation	-0.782* (-11,77)	0.790* (3.91)	-0.126 (-1,41)	**0.265**	-0.152 (-1,71)	**0.215**
Historical (n=22)	-0.019 (-0,08)	0.394* (1.52)	0.000 (0,00)	0.392* (1.52)	0.062 (0,25)	0.309* (1.46)
Historical (n=65)	0.007 (0,04)	0.316* (1.19)	0.011 (0,07)	0.304* (1.16)	0.039 (0,25)	0.253* (1.19)
Historical (n=128)	0.045 (0,36)	**0.267**	0.056 (0,44)	0.265 (1.00)	0.042 (0,33)	0.245* (1.15)
2006-2007						
Implied correlation	-0.810* (-12,65)	0.821* (3.08)	-0.286* (-3,30)	**0.379**	-0.242* (-2,76)	0.259* (1.58)
Historical (n=22)	-0.082 (-0,34)	0.436 (1.24)	-0.105 (-0,43)	0.487 (1.33)	-0.005 (-0,02)	0.277* (1.70)
Historical (n=65)	0.007 (-0,01)	0.306 (1.10)	0.011 (-0,26)	0.295 (1.21)	0.039 (-0,12)	0.248 (1.14)
Historical (n=128)	0.089 (0,71)	**0.360**	0.093 (0,74)	0.410 (1.09)	-0.067 (-0,53)	**0.166**
2008-2009						
Implied correlation	-0.804* (-12,45)	0.809* (4.08)	-0.118 (-1,33)	0.211 (1.03)	-0.148 (-1,66)	**0.226**
Historical (n=22)	-0.022 (-0,09)	0.420* (1.62)	-0.020 (-0,08)	0.317* (1.57)	0.114 (0,47)	0.330* (1.49)
Historical (n=65)	-0.012 (-0,08)	0.329* (1.24)	-0.036 (-0,23)	0.211 (1.03)	0.074 (0,48)	0.294* (1.31)
Historical (n=128)	0.027 (0,21)	**0.269**	-0.010 (-0,08)	**0.206**	0.055 (0,43)	0.297* (1.33)
2010-2011						
Implied correlation	-0.710* (-9,94)	0.712* (4.62)	0.015 (0,17)	**0.190**	-0.075 (-0,84)	**0.136**
Historical (n=22)	0.041 (0,17)	0.299* (1.60)	0.117 (0,48)	0.392* (2.15)	0.040 (0,16)	0.300* (2.26)
Historical (n=65)	0.042 (0,27)	0.208 (1.10)	0.117 (0,76)	0.276* (1.47)	0.030 (0,19)	0.224* (1.67)
Historical (n=128)	0.049 (0,39)	**0.191**	0.133 (1,06)	0.235* (1.24)	0.081 (0,64)	0.183* (1.36)

* Indicates statistical significance at the 5 % level, F-values in brackets

Number of observations: 2006-2011: 1125, 2006-2007: 149, 2008-2009: 522, 2010-2011: 454

H_0: The mean forecast error (MFE) is zero

H_0: The forecasts are identical in the ability to predict future realized correlation, i.e. $\frac{RMSFE}{RMSFE_{lowest}} = 1$

Table 7b: Error statistics for the 1 month correlation forecasts

	EURJPY		USDNOK		EURNOK	
2006-2011	**MFE**	**RMSFE**	**MFE**	**RMSFE**	**MFE**	**RMSFE**
Implied Correlation	0.520* (2,53)	0.633* (1.84)	0.196 (0,87)	0.392* (1.39)	0.980* (10,24)	0.982* (7.36)
Historical (n=22)	-0.016 (-0,05)	0.469* (1.25)	0.001 (0,00)	0.334* (1.17)	0.005 (0,02)	0.343 (1.12)
Historical (n=65)	0.021 (0,08)	**0.385**	0.008 (0,03)	**0.288**	0.040 (0,16)	**0.307**
Historical (n=128)	0.043 (0,18)	0.445* (1.18)	0.009 (0,04)	0.296 (1.03)	0.099 (0,41)	0.367* (1.21)
2006-2007						
Implied correlation	0.600* (3,04)	0.696* (1.96)	0.271 (1,22)	0.381* (1.50)	0.967* (9,11)	0.969* (8.61)
Historical (n=22)	-0.051 (-0,16)	0.496 (1.24)	0.019 (0,06)	0.345* (1.35)	0.026 (0,08)	0.364* (1.59)
Historical (n=65)	0.021 (0,08)	**0.413**	0.006 (0,02)	**0.261**	0.011 (0,04)	**0.236**
Historical (n=128)	-0.099 (-0,41)	0.530* (1.34)	0.136 (0,56)	0.307 (1.19)	0.067 (0,28)	0.237 (1.00)
2008-2009						
Implied correlation	0.521* (2,53)	0.656* (1.76)	0.359 (1,65)	0.463* (2.01)	0.983* (10,63)	0.985* (7.62)
Historical (n=22)	0.006 (0,02)	0.519* (1.29)	0.011 (0,03)	0.307* (1.27)	0.013 (0,04)	0.312 (1.01)
Historical (n=65)	0.020 (0,08)	**0.419**	0.037 (0,14)	0.256 (1.05)	0.079 (0,31)	**0.308**
Historical (n=128)	0.083 (0,34)	0.479* (1.17)	0.053 (0,21)	**0.244**	0.156 (0,65)	0.396* (1.32)
2010-2011						
Implied correlation	0.455* (2,15)	0.536* (1.81)	-0.060 (-0,26)	**0.287**	0.983* (10,59)	0.984* (6.82)
Historical (n=22)	-0.020 (-0,06)	0.373* (1.19)	-0.023 (-0,07)	0.356* (1.26)	-0.016 (-0,05)	0.360 (1.07)
Historical (n=65)	0.022 (0,09)	**0.319**	-0.023 (-0,09)	0.333* (1.17)	0.011 (0,04)	**0.339**
Historical (n=128)	0.036 (0,15)	0.367* (1.16)	-0.076 (-0,31)	0.342* (1.21)	0.041 (0,17)	0.359 (1.06)

* Indicates statistical significance at the 5 % level, F-values in brackets

Number of observations: 2006-2011: 1228, 2006-2007: 252, 2008-2009: 522, 2010-2011: 454

H_0: The mean forecast error (MFE) is zero

H_0: The forecasts are identical in the ability to predict future realized correlation, i.e. $\frac{RMSFE}{RMSFE_{lowest}} = 1$

Table 8b: Error statistics for the 3 month correlation forecasts						
	EURJPY		**USDNOK**		**EURNOK**	
2006-2011	**MFE**	**RMSFE**	**MFE**	**RMSFE**	**MFE**	**RMSFE**
Implied Correlation	0.496* (4,28)	0.562* (1.81)	0.192 (1,53)	0.323* (1.50)	0.980* (18,11)	0.980* (7.63)
Historical (n=22)	-0.039 (-0,15)	0.404* (1.22)	-0.027 (-0,10)	0.294* (1.35)	-0.048 (-0,19)	0.327 (1.12)
Historical (n=65)	-0.008 (-0,05)	**0.338**	-0.011 (-0,06)	0.230 (1.04)	0.002 (0,01)	**0.293**
Historical (n=128)	0.027 (0,17)	0.349 (1.04)	-0.017 (-0,11)	**0.220**	0.058 (0,37)	0.332* (1.14)
2006-2007						
Implied correlation	0.538* (4,74)	0.610* (2.22)	0.294* (2,39)	0.325* (2.43)	0.969* (16,49)	0.970* (16.82)
Historical (n=22)	-0.110 (-0,42)	0.390* (1.29)	0.018 (0,07)	0.255* (1.88)	0.005 (0,02)	0.253* (2.09)
Historical (n=65)	-0.063 (-0,35)	0.340 (1.11)	0.034 (0,19)	**0.138**	0.028 (0,16)	**0.124**
Historical (n=128)	-0.091 (-0,59)	**0.308**	0.166 (1,08)	0.177* (1.29)	0.091 (0,59)	0.148 (1.20)
2008-2009						
Implied correlation	0.510* (4,43)	0.581* (1.69)	0.318* (2,60)	0.389* (2.23)	0.983* (18,73)	0.983* (7.65)
Historical (n=22)	-0.019 (-0,07)	0.460* (1.27)	-0.026 (-0,10)	0.265* (1.47)	-0.041 (-0,16)	0.311 (1.03)
Historical (n=65)	-0.005 (-0,03)	**0.374**	-0.001 (0,00)	0.192 (1.05)	0.026 (0,14)	**0.303**
Historical (n=128)	0.059 (0,38)	0.401 (1.08)	0.015 (0,10)	**0.182**	0.104 (0,67)	0.397* (1.34)
2010-2011						
Implied correlation	0.445* (3,77)	0.481* (1.84)	-0.086 (-0,68)	**0.175**	0.982* (18,58)	0.982* (8.50)
Historical (n=22)	-0.017 (-0,07)	0.321* (1.17)	-0.058 (-0,22)	0.347* (2.05)	-0.092 (-0,36)	0.382* (1.45)
Historical (n=65)	0.021 (0,11)	0.281 (1.02)	-0.051 (-0,29)	0.303* (1.77)	-0.044 (-0,25)	0.345* (1.30)
Historical (n=128)	0.022 (0,14)	**0.277**	-0.113 (-0,73)	0.269* (1.56)	-0.010 (-0,07)	**0.270**

* Indicates statistical significance at the 5 % level, F-values in brackets

Number of observations: 2006-2011: 1187, 2006-2007: 211, 2008-2009: 522, 2010-2011: 454

H_0: The mean forecast error (MFE) is zero

H_0: The forecasts are identical in the ability to predict future realized correlation, i.e. $\frac{RMSFE}{RMSFE_{lowest}} = 1$

Table 9b: Error statistics for the 6 month correlation forecasts

	EURJPY		USDNOK		EURNOK	
2006-2011	**MFE**	**RMSFE**	**MFE**	**RMSFE**	**MFE**	**RMSFE**
Implied Correlation	0.472* (5,73)	0.505* (1.69)	0.177* (2,00)	0.286* (1.54)	0.979* (25,42)	0.979* (7.22)
Historical (n=22)	-0.078 (-0,32)	0.437* (1.42)	-0.032 (-0,13)	0.288* (1.55)	-0.061 (-0,25)	0.349* (1.16)
Historical (n=65)	-0.045 (-0,29)	0.348 (1.10)	-0.019 (-0,12)	0.208 (1.11)	-0.018 (-0,12)	0.309 (1.01)
Historical (n=128)	-0.006 (-0,04)	**0.318**	-0.021 (-0,16)	**0.189**	0.019 (0,15)	**0.305**
2006-2007						
Implied correlation	0.537* (6,71)	0.579* (1.77)	0.311* (3,59)	0.337* (2.94)	0.973* (23,96)	0.973* (13.52)
Historical (n=22)	-0.144 (-0,59)	0.419 (1.20)	0.057 (0,23)	0.236* (2.01)	0.058 (0,24)	0.260* (1.68)
Historical (n=65)	-0.045 (-0,40)	**0.335**	-0.019 (0,39)	**0.205**	-0.018 (0,46)	**0.300**
Historical (n=128)	0.021 (0,16)	0.365 (1.03)	0.115 (0,91)	0.129 (1.08)	0.114 (0,91)	0.173 (1.10)
2008-2009						
Implied correlation	0.449* (5,40)	0.483* (1.52)	0.271* (3,11)	0.324* (1.97)	0.980* (25,83)	0.981* (6.34)
Historical (n=22)	-0.090 (-0,37)	0.487* (1.54)	-0.068 (-0,28)	0.262* (1.57)	-0.113 (-0,47)	0.359 (1.03)
Historical (n=65)	-0.075 (-0,48)	0.381* (1.16)	-0.042 (-0,27)	0.191 (1.14)	-0.046 (-0,30)	**0.350**
Historical (n=128)	-0.034 (-0,27)	**0.333**	-0.029 (-0,23)	**0.169**	0.014 (0,11)	0.363 (1.04)
2010-2011						
Implied correlation	0.444* (5,34)	0.452* (1.83)	-0.098 (-1,10)	**0.116**	0.981* (26,12)	0.981* (8.86)
Historical (n=22)	-0.005 (-0,02)	0.356* (1.40)	-0.052 (-0,21)	0.354* (3.17)	-0.084 (-0,35)	0.393* (1.58)
Historical (n=65)	0.012 (0,07)	0.281 (1.08)	-0.041 (-0,26)	0.271* (2.39)	-0.042 (-0,27)	0.321* (1.26)
Historical (n=128)	0.022 (0,18)	**0.261**	-0.082 (-0,65)	0.237* (2.07)	-0.025 (-0,20)	**0.257**

* Indicates statistical significance at the 5 % level, F-values in brackets

Number of observations: 2006-2011: 1125, 2006-2007: 149, 2008-2009: 522, 2010-2011: 454

H_0: The mean forecast error (MFE) is zero

H_0: The forecasts are identical in the ability to predict future realized correlation, i.e. $\frac{RMSFE}{RMSFE_{lowest}} = 1$