Sándor Misik

Forecasting and Correlation Dynamics in Regional and Global Currency Trios

Institute of Finance

Supervisors:
Dr Barbara Dömötör, PhD, Professor
Dr János Száz, Csc, Professor

Corvinus University of Budapest

Doctoral School of Business and Management

Forecasting and Correlation Dynamics in Regional and Global Currency Trios

Ph.D. Thesis

Sándor Misik

Budapest, 2026

Content

List of Figures

List of Tables

Acknowledgements

I would like to express my deepest gratitude to my two supervisors, Barbara Dömötör and János Száz, for their unwavering support, encouragement, and extraordinary patience. Their guidance provided me with the confidence, energy, and belief necessary to complete this decade-long journey. Special thanks are due to the two opponents of my PhD proposal, Ágnes Lubloy and Levente Zsembery, whose valuable comments greatly contributed to improving this thesis. Any remaining errors are entirely my own. I am sincerely grateful to my colleagues at the Institute of Finance of Corvinus University of Budapest for their support throughout my doctoral studies. I would particularly like to thank Edina Berlinger, Zsolt Bihary, Péter Juhász, Gábor Kürthy, Lilla Keresztúri, and Péter Csóka for their valuable academic advice and support in the publication process. I am especially indebted to Dávid Zoltán Szabó, whose contribution as co-author of a manuscript encompassing many of the thesis's key findings significantly enriched this work. I would also like to express my gratitude to Kata Váradi, Ágnes Dancs-Vidovics, Naffa Helena, and Dániel Homolya for their encouragement and occasional but timely pushes towards achieving this goal. I am furthermore thankful to István P. Székely, Jonathan A. Batten, Milán Badics, and Balázs Márkus for generously sharing their time, expertise, and valuable professional insights during our consultations. As this dissertation is largely grounded in empirical research and the practical knowledge of market participants, I owe special thanks to several practitioners whose contributions proved invaluable. I am grateful to Júlia Tóth for writing the very first R code that enabled me to conduct correlation forecasting; to Nóra Hanfeld for connecting me with London-based foreign exchange traders and analysts for the qualitative survey on quotation regimes; and to Gábor Felkai and Mihály Tatár for sharing their expertise on quoting practices in the Hungarian foreign exchange market. I would also like to thank Péter Hidvégi for his assistance in obtaining high-frequency Reuters data. This dissertation would not have been possible without the support of my family. I am especially grateful to my wife, Dóra, who took on an extraordinary share of family responsibilities and cared for our three children while I spent countless evenings and weekends wrestling with Excel spreadsheets and R code, particularly during the pre-AI era. Last but certainly not least, the very idea of pursuing a PhD originated with my father. Throughout my thirties, he never missed an opportunity to ask about the progress of my scientific work, often to my great annoyance at the time. Sadly, he could not keep his promise not to leave this world before I completed my doctorate. Yet, by the time he passed away in 2025, the finish line was already within sight. God bless you, Dad.

1. INTRODUCTION

1.1 Motivation

The foreign exchange market occupies a unique position within global financial markets due to its size, decentralised structure, and the existence of strict no-arbitrage relationships between exchange rates. One of the most fundamental of these relationships is triangular arbitrage, which imposes strong constraints on the joint dynamics of exchange rates within currency trios. Unlike equity or bond markets, where co-movement is largely an empirical phenomenon, correlation in currency markets is to a large extent structurally determined by arbitrage conditions.

My interest in this topic originated during my early exposure to professional currency trading. While working as a university intern at a domestic brokerage firm, I encountered published exchange rate forecasts for the EURUSD, EURJPY, and USDJPY pairs that were mutually inconsistent and violated the triangular arbitrage condition. Although such inconsistencies could be explained by reporting errors or rounding, they raised a more fundamental question: whether expectations, pricing conventions, or market frictions could lead to systematic and persistent differences in correlation structures across currency trios, even in the absence of arbitrage opportunities.

Subsequent academic inquiry led me to the concept of implied correlation derived from option prices, as well as to a rapidly growing literature on the interaction between volatility and correlation in currency markets. These frameworks provide a coherent explanation for how exchange rates can exhibit heterogeneous correlation dynamics without violating no-arbitrage conditions. At the same time, the literature predominantly focuses on major global currencies, while regional currency markets receive considerably less attention. This asymmetry is particularly striking in the context of Central and Eastern European (CEE) and Scandinavian currencies – regional currency markets hereinafter, which are actively traded, economically relevant, and strongly integrated into the European financial system.

1.2 Research gaps and contribution

The analysis and forecasting of correlations play a central role in modern finance. Correlations directly affect portfolio allocation, hedging strategies, derivative pricing, and risk measurement. The global financial crisis of 2008–2009 further highlighted that changes in correlation structures can materially amplify systemic risk. In currency markets, correlations are of particular importance because of the triangular arbitrage relationship, which fundamentally links pairwise correlations to volatility and market structure.

Despite the extensive literature on currency correlation, several important gaps remain. First, there is limited evidence on whether conclusions drawn from major global currency markets also apply to regional currencies. Second, while dynamic properties of exchange rates are widely studied at the pairwise level, much less is known about the dynamic interaction within currency trios. Third, although market microstructure is frequently mentioned as a potential driver of correlation differences across markets, there are few models that explicitly link microstructural assumptions to empirically observed correlation levels.

This dissertation addresses these gaps through three closely related research pillars. Chapter 3 focuses on the forecasting accuracy of regional currency market correlations compared to major currency trios. By comparing implied correlations derived from option prices with time-series-based econometric approaches, the chapter provides the first systematic assessment of correlation forecast accuracy for regional currency trios involving the euro and the US dollar. The results contribute to the broader debate on whether option-implied measures contain informational advantages beyond historical data in regional currency markets. Additionally, the forecasting performance of various approaches is evaluated in both correlation and volatility forecasting, reflecting the close relationship between these two aspects of currency market dynamics. Finally, Chapter 3 introduces a variance-based decomposition of correlation forecasting performance, which, to the best of my knowledge, is the first application of its kind, and uses it to identify the main sources of forecast error.

Correlation in the foreign exchange market is a widely researched/discussed topic in its own right, and recently there has been a number of studies focusing on the relationship between foreign exchange yields and correlation (see, among others: Mueller et al. (2017), Chernov et al. (2024). At the same time, the analysis of regional currency correlations is not prominent and primarily arises in relation to the USD from the perspective of USD investors rather than from the perspective of currency trios. Following the clarification of the impact of currency trios on currency correlations in Chapter 2, Chapter 4 examines the dynamic behaviour of exchange rates within currency trios. Based on vector autoregressive and local projection models estimated on ultra-high-frequency data, and qualitative evidence from market practitioners, the chapter identifies asymmetric adjustment patterns and leader-follower relationships between currencies. These dynamic properties offer an economically meaningful explanation for why correlation structures differ between global and regional currency markets, beyond what can be attributed to differences in aggregate risk or volatility alone. In addition, spillover topologies are presented and contrasted with the predictions of competing microstructure-based explanations and the synthetic quotation hypothesis in order to assess which framework is most consistent with the empirically observed directional spillover structure.

Similarly, there is a lack of research in the literature exploring the long-term level of currency market correlations, with regard to differences between individual markets. The main contribution of this thesis to the literature is to further develop the idea that the levels of correlation in currency trios can be well explained and modelled based on certain market microstructural assumptions (Ciacci, et al., 2020). Chapter 5 introduces a parsimonious market microstructure-based framework that is capable of reproducing the empirically observed correlation levels of different currency trios. Rather than aiming to describe the full complexity of foreign exchange trading, the proposed model isolates those microstructural features that look essential for explaining correlation outcomes. By linking these assumptions to observed market characteristics, the model provides a unifying interpretation of the empirical findings presented in earlier chapters. Although the findings are backed by both empirical and modelling evidence I do not claim observable FX triangle correlation regimes are explained by solely the microstructural properties highlighted in Chapter 4 and Chapter 5. There are certain microstructural agent based behavioural patterns that form the realized correlation structures however these seem to be less successful in reproducing the regional patterns and the related literature do not provide explanation to the different correlation regimes either.

1.3 Research questions and hypotheses

In the foreign exchange market, correlation within a currency trio – which more precisely refers to the linear correlation coefficient measured between the log returns of exchange rates formed by three currencies – is clearly determined by volatility of the three currency pairs formed by the three currencies, both ex post and ex ante. In fact, any correlation between any two currency rates are determined by the volatilities of the related FX rates. This relationship, which is explained in Chapter 2, has many interesting aspects, including the possible value ranges of correlations between currency trios, which are narrower than in the general case of three variables.

Following one of the dominant lines of literature on currency correlation and attempting to exploit the empirical usefulness of theoretical models, I compare the predictive power of several widely used correlation estimates in the regional currency markets and in some of the major currency trios. One key question that has not been addressed in the literature so far is whether the so-called implied correlation forecast – derived from option prices – is more accurate compared to the econometric models based on historical data. **The first research question addressed in the dissertation is therefore which method provides the most accurate forecast for regional currencies (the Hungarian forint, Czech koruna, and Polish zloty, Norwegian krone and Swedish krona) compared to major currencies (Swiss Franc, Japanese yen and British pound) in currency trios involving the euro and the dollar, and how this relates to previous findings in the literature, which examined this issue in relation to major global currencies?**

The empirical investigation of this question requires a careful review of the relevant literature and the formulation of testable hypotheses. For this reason, the corresponding hypotheses are introduced in Chapter 3, following the literature review, where they can be properly motivated and empirically evaluated.

Comparing the correlation structure of regional currency markets with that of major currency markets does not seem to appear in the literature, although the differences are clearly visible, I am not aware of any comprehensive study that explains them. The second research question focuses on the dynamic properties of exchange rates within currency trios measured at ultra-high frequency range. Although a broad spectrum of literature explores the microstructure of currency markets (see: Chapter 4.2), the usual frequency range is above one second, usually on several minute-level. Empirical evidence suggests however that driven by the triangular arbitrage mechanism the misalignments of currency rates are diminished within- or close to one second therefore the dynamic relationship and the order of alignment within the currency trio remains hidden and the currency moves appear concurrent. **The second research question therefore asks: How does the dynamic time-series behaviour of exchange rates differ between regional and global currency markets?** Common sense would suggest that differences in riskiness leads to differences in volatility and – as already mentioned – the close relationship between volatility and correlation could provide the answer. While not disputing that the risk-volatility-correlation channel likely has a significant explanatory power I focus instead on other microstructural sources, namely the quotational differences between regional and major currency markets. This question is addressed empirically in Chapter 4, where testable hypotheses related to dynamic exchange rate behaviour are formulated after the review of the relevant literature and evaluated using both quantitative and qualitative methods.

The third research question is motivated by the need for a coherent theoretical interpretation of observed correlation patterns. While existing studies frequently refer to market microstructure as an explanatory factor, few provide a formal framework that links microstructural assumptions to observable correlation outcomes. **The third research question therefore asks: Does there exist a market microstructure-based model that can reproduce the empirically observed correlation structures of currency markets?**

A corresponding hypothesis is formulated and tested in Chapter 5, where the proposed model is evaluated against empirical benchmarks derived from both global- and regional currency markets.

1.4 Structure of the thesis

Following this introductory chapter, Chapter 2 provides an overview of the institutional structure of currency markets and develops the theoretical link between volatility and correlation arising from triangular arbitrage. The chapter also contains a comprehensive review of the literature relevant to the research questions of the dissertation. A secondary aim of the chapter is to provide an overview of the empirically observed correlation- volatility levels and link them to currency market traded volumes.

Chapter 3 presents the empirical analysis of forecasting accuracy of various time-series- and option-implied methods in the regional and major currency markets. Following the strong connection between correlation and volatility explained in Chapter 2 volatility forecasting and correlation error decomposition is also conducted.

Chapter 4 examines the dynamic behaviour of various currency trios using time-series models and qualitative evidence. This chapter also builds on the concept of synthetic quotation – namely one FX rate within the trio is mechanically quoted based on the other two independently traded legs – which is the main argument of the thesis to be the main source of the differences empirically observed in shock spillover topologies across various currency trios and eventually leads to different correlation regimes. The synthetic quotation hypothesis is also tested against various micro-structure assumptions in reproducing the observed spillover matrices.

Chapter 5 introduces and evaluates the market microstructure-based model of currency correlations, that intends to provide a further support to the quotation argument and closely reproduces the observed correlation regimes in major and regional currency trios. This analysis also extends to the so-called USD-based currencies – like Mexican Peso or Korean Won – representing a "third class" of correlation regimes.

Chapter 6 concludes and outlines directions for future research.

2. DESCRIPTIVE ANALYSIS OF THE FOREIGN EXCHANGE MARKET

The structure and dynamics of foreign exchange markets exhibit notable differences across currency groups, particularly between major global currencies and those of the regional. While the FX market as a whole is characterized by exceptional liquidity, deep integration, and the dominance of the US dollar, these aggregate features mask important heterogeneities in trading patterns, volatility characteristics, and correlation structures across individual currency pairs. Understanding these differences is essential for accurately interpreting exchange rate behaviour and its underlying drivers.

After introducing the FX market specific formalized relationship between volatility and correlation driven by triangular arbitrage this chapter focuses on three closely related dimensions of FX markets: correlation patterns, volatility levels, and market turnover. Using daily exchange rate data over an extended period, the analysis highlights systematic differences between regional currencies (HUF, PLN, CZK, SEK, NOK) and major currencies (JPY, CHF, GBP) in their relationship with the EUR and USD. In particular, the empirical results show that while major currency correlations tend to exhibit stable and moderate relationships, regional currencies display more asymmetric patterns, including near-zero correlations in certain currency pairs.

To explain these findings, the chapter examines whether differences in volatility structures can account for the observed correlation patterns, and demonstrates how relative volatility levels influence correlation outcomes in a mechanical way. In this chapter after demonstrating the long-term differences in correlation and volatility levels the analysis extends to the relative turnover of currency pairs. By linking correlation levels to differences in trading volumes—especially the dominance of EUR versus USD trading in CEE markets—the chapter provides support for the idea that liquidity distribution and trading conventions play a key role in shaping exchange rate dynamics.

2.1 The structure of the foreign exchange market

The foreign exchange market is the largest financial market, according to the daily OTC FX turnover averaged USD 9.6 trillion in April 2025 (BIS, 2025). Volumes in all type of instruments experienced new highs in 2025 most of them after decades of growing trends. The main instruments are FX spot FX forward FX swap Currency swap and options/other derivatives as shown in Figure 1 from BIS (2025).

Figure 1: Foreign exchange market turnover by instrument

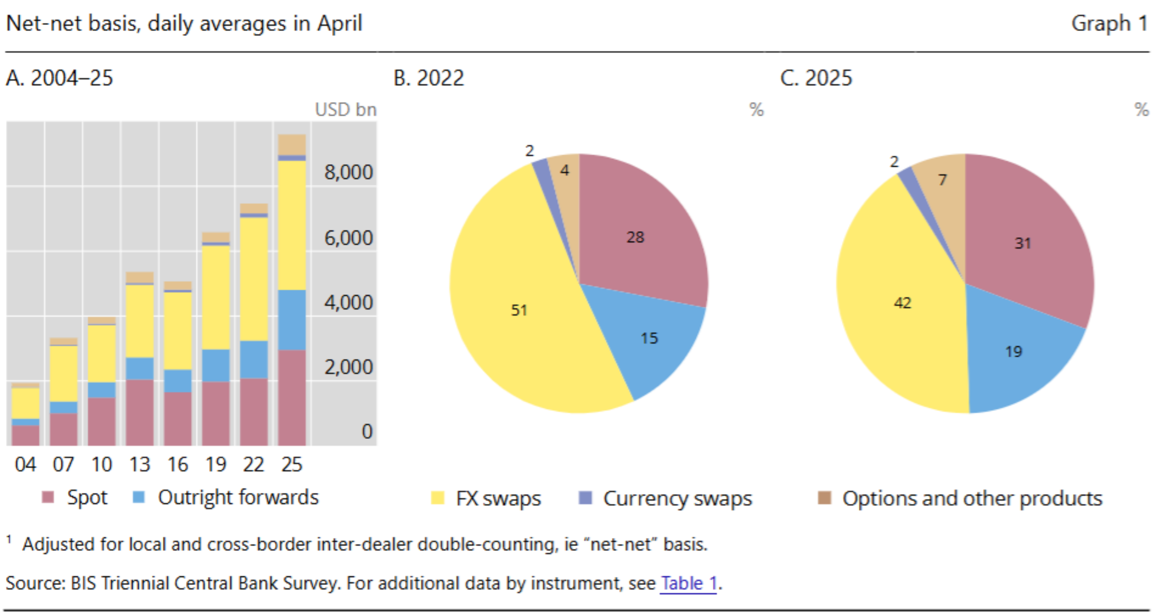

Source: BIS (2025) p. 6

USD remained the main currency of choice being one side of 89.2% of all trades in April 2025. The most heavily traded currency pairs are the EURUSD, USDJPY and GBPUSD.

Figure 2: Foreign exchange market turnover by currency and currency pairs

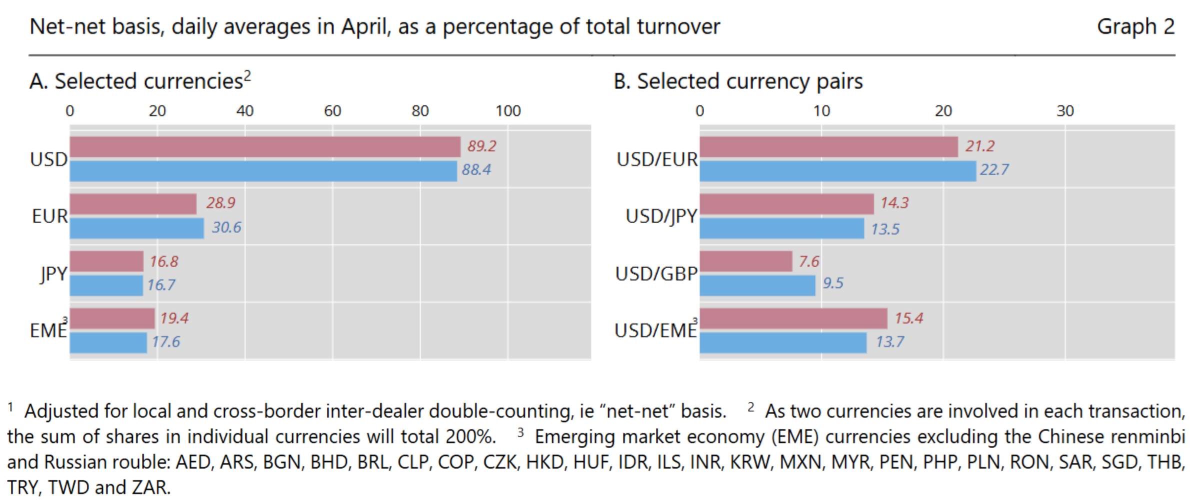

Source: BIS (2025) p. 8

Regarding market participants, the two main group of counterparties are dealers and “other financial institutions” – practically non-reporting banks, institutional investors, hedge funds and proprietary trading firms:

Figure 3: Foreign exchange market turnover by counterparty

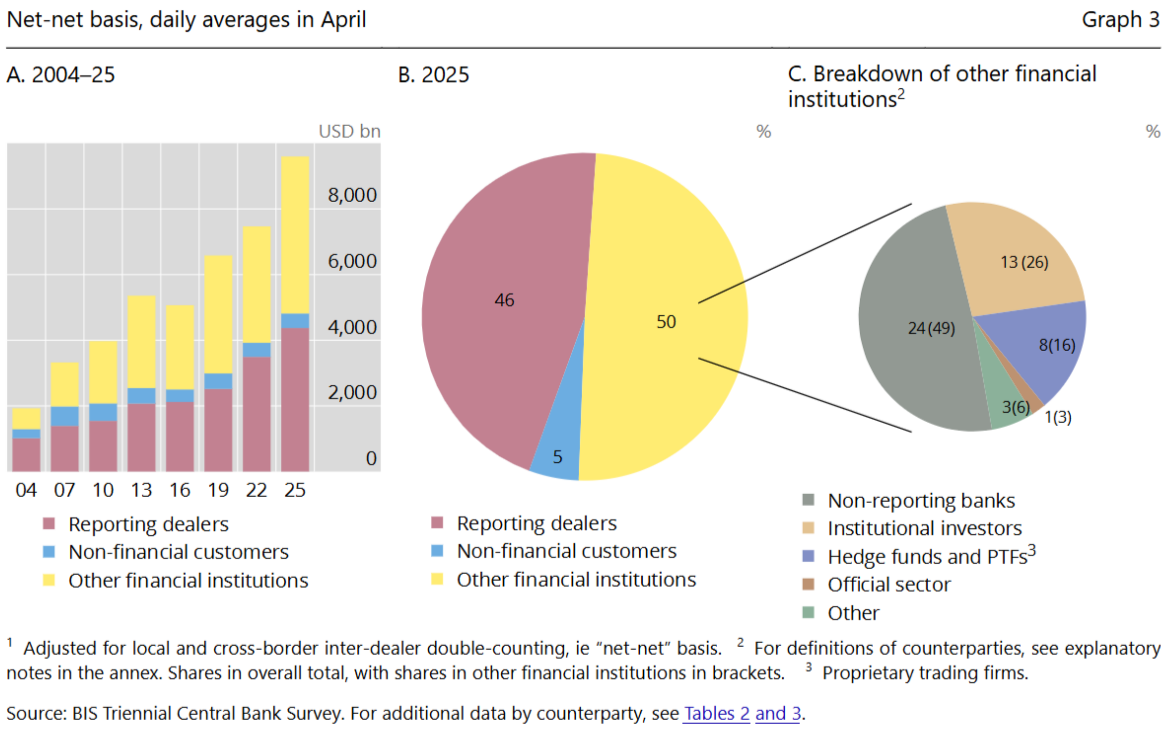

Source: BIS (2025) p. 9

The main trading centres are located in London, New York, Tokyo, and Singapore. The FX market used to be described as a so-called two-tier market meaning the inter-dealer market and the customer-dealer market. This setup meant higher inter-dealer turnovers resulted from trading the inventory imbalances (called “hot-potato trading” based on Lyons (1996)). More recently however in parallel with the rise of new electronic trading venues, high frequency trading and algorithmic trading in general, other participants – mainly proprietary trading firms – have entered the market as market makers and liquidity providers (Drehmann and Suhsko, 2022). This resulted a more even market share of inter-dealer and non-interdealer systems. This said the FX market remained fragmented with the two “primary venues” (Reuters/Refinitiv Matching (replacing Reuters D-3000 since 2013) and Electronic Broking System (EBS)) and more than 30 other venues in the dealer-customer segment (Drehmann and Suhsko, 2022).

2.2 Correlation in the foreign exchange market

Let us assume that between three currency pairs formed from three currencies, there is a product or quotient relationship arising from the absence of arbitrage at every moment. Although based on high-frequency data, "every moment" is not fulfilled, and on average, arbitrage is possible over a time window of 1.5 seconds (Kozhan and Tham, 2012), In the remainder of this section, I consider this time period to be negligible.: $USDHUF = \frac{EURHUF}{EURUSD}$. This means that the correlation between the two currency pairs cannot be arbitrary but will be determined by the volatilities of the three currency pairs. Using the notation of $S_1 = EURHUF$, $S_2 = USDHUF$, $S_3 = EURUSD$ as spot FX rates one can express S_3 as S_1/S_2 meaning $EURUSD = \frac{EURHUF}{USDHUF}$. For the remainder of the thesis I follow the market convention of currency rate quotation where the base currency is put ahead of the numerator currency: EURUSD equals the USD value of 1 EUR.

Let s_1, s_2, s_3 denote the daily log returns of the three exchange rates, which can be calculated for the i-th currency pair as follows: $s_i = \ln\left(\frac{S_t}{S_{t-1}}\right)$ thus s3 = s1 – s2, which can be derived from the logarithms of the exchange rates.[1] Furthermore, the following relationship applies to the variance of the third probability variable formed from the difference between two correlated probability variables, which is also used in the pricing of exchange options (Medvegyev and Száz, 2010):

$$\sigma_{3,t,T}^2 = \sigma_{1,t,T}^2 + \sigma_{2,t,T}^2 - 2\rho_{1,2_{t,T}}\sigma_{1,t,T}\sigma_{2,t,T} \tag{1}$$

If $s_3 = s_1 + s_2$, i.e. the third exchange rate is the product of the other two exchange rates, then

$$\sigma_{3,t,T}^2 = \sigma_{1,t,T}^2 + \sigma_{2,t,T}^2 + 2\rho_{1,2_{t,T}}\sigma_{1,t,T}\sigma_{2,t,T} \tag{2}$$

thus

$$\rho_{1,2_{t,T}} = \frac{\sigma_{1,t,T}^2 + \sigma_{2,t,T}^2 - \sigma_{3,t,T}^2}{2\sigma_{1,t,T}\sigma_{2,t,T}} \tag{3}$$

and

$$\rho_{1,2_{t,T}} = \frac{\sigma_{3,t,T}^2 - \sigma_{1,t,T}^2 - \sigma_{2,t,T}^2}{2\sigma_{1,t,T}\sigma_{2,t,T}} \tag{4}$$

[1] $\ln(S_{3,t}/S_{3,t-1})=\ln(S_{3,t})-\ln(S_{3,t-1})=\ln(S_{1,t}/S_{2,t})-\ln(S_{1,t-1}/S_{2,t-1})=\ln(S_{1,t})-\ln(S_{1,t-1})-[\ln(S_{2,t})-\ln(S_{2,t-1})]=\ln(S_{1,t}/S_{1,t-1})-\ln(S_{2,t}/S_{2,t-1})=s_1-s_2$.

where $\rho_{1,2_{t,T}}$ is the Pearson linear correlation between the log returns of currency pairs 1 and 2 over the period t-T, σ_i is the standard deviation, i.e., the volatility, of the log return of currency pair i. The correlation can also be represented visually, as shown in Figure 4.

Figure 4: Triangle relationship in the FX market

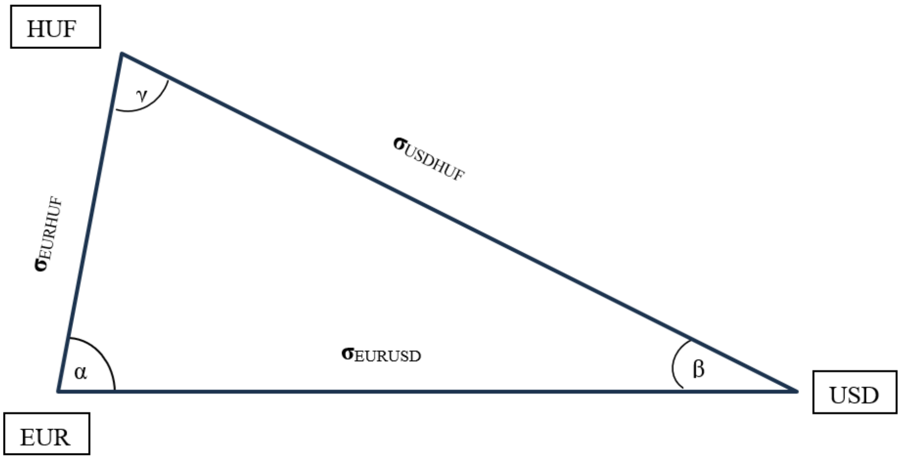

source: Author's own work

The volatility of the three currency pairs that can be formed from the three currencies is directly proportional to the lengths of the sides of the triangle, and the cosine of the enclosed angles is the correlation between the two "adjacent" currency pairs. In Figure 4, for example, $\rho_{EURHUF-EURUSD} = \cos(\alpha)$. Based on relationship (1-4), there is an equation for every currency pair that can be written as the sum of the angles of the triangles[2]:

$$\arccos(\rho_{1,2}) + \arccos(\rho_{1,3}) + \arccos(\rho_{2,3}) = \pi \quad (5)$$

where $\rho_{i,j}$ is the correlation coefficient between the i-th and j-th currency pairs, i.e., the cosine of the respective angle of the triangle. The set of possible correlations in the foreign exchange market is actually a subset of all possible correlation triplets.

The validity condition for correlation matrices is that the matrix is positive definite, which in the case of 3x3 general correlation matrices can be written as the following relationship for the determinant:

[2]It should be noted that the triangle relationship only exists if, in the case of correlation between the three, we take the inverse of one of the exchange rates, i.e., -1 times the log return according to the "normal" market quotation convent.

$$1 + 2\rho_{1,2}\rho_{1,3}\rho_{2,3} - \rho_{1,2}^2 - \rho_{1,3}^2 - \rho_{2,3}^2 \geq 0 \qquad (6)$$

which always equals to zero in the foreign exchange market, i.e.

$$1 + 2\rho_{1,2}\rho_{1,3}\rho_{2,3} - \rho_{1,2}^2 - \rho_{1,3}^2 - \rho_{2,3}^2 = 0 \qquad (7)$$

since all exchange rates can be written as the product or quotient of the other two exchange rates, i.e., the logarithm is the sum or difference of the first two. Since, in the case of currency trios, any exchange rate can be expressed as the quotient or product of the other two exchange rates, the logarithms of the exchange rates are linearly related to each other, so the determinant of the valid correlation matrices in the currency market must be zero.

At this point, it is important to draw attention to an apparent contradiction: relationship (5) and (7) appear to contradict each other, i.e., if we take the EUR HUF USD currency trio, the relationship relating to the determinant is valid for EURHUF, USDHUF, and EURUSD exchange rates. At the same time in the case of the triangular relationship of correlations in currency trios, instead of $\rho_{2,3}$, its negative ($-\rho_{2,3}$) makes equation (5) true. This means that the third correlation – between USDHUF and EURUSD – fulfils the triangular relationship if the correlation between USDHUF - USDEUR is applied instead of correlation between USDHUF - EURUSD (the inverse of the later). Since the log of the inverse exchange rate is negative 1 times the log of the original exchange rate, the correlation coefficient only changes sign.

All possible correlation trios in the currency market can be represented in a three-dimensional space (cube) interpreted on an interval of -1 to +1. Figure 5 illustrates the possible currency market correlations according to the triangular relationship that can be applied to currency trios.

Figure 5: Set of possible correlations in the case of currency trios

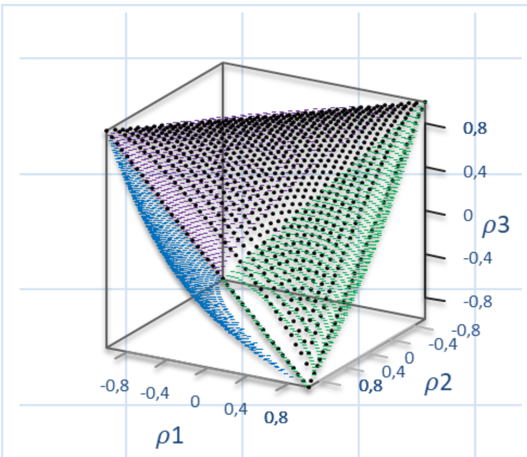

source: Author's own work

In general, valid 3x3 correlation matrices contain all points in the space within the boundary surfaces shown in Figure 5, while in the foreign exchange market, valid correlation triples are indicated by the boundary surfaces themselves.
Based on Rousseeuw and Molenberghs (1994), given a certain level of correlation between variables Y and Z, the possible values of r_{XZ} and r_{XY} must lie within an elliptical range. (If r_{YZ}=0, then within the boundaries of a circle.)

Figure 6: Set of possible correlations in the case of three variables

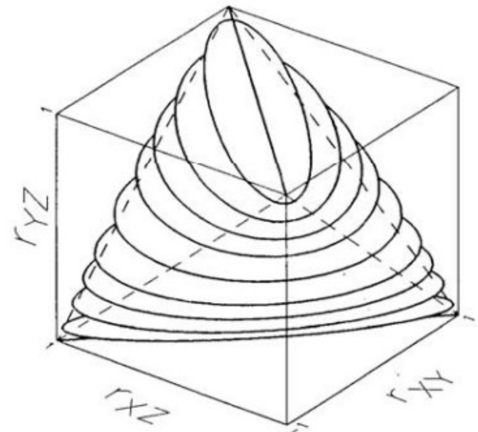

(b)

Figure 1. (a) Set of all Possible Correlations Between X, Y, and Z. (b) Slicing this set at r_{YZ} yields ellipses.

Source: Rousseeuw and Molenberghs (1994) p. 277

The main difference between "general" and foreign exchange market correlation matrices can be clearly illustrated using 3x3 correlation matrices. While for a given correlation value ρ_1, ρ_2 can be located anywhere within the boundaries of the given shape (circle or ellipse) as a function of ρ_3, in the case of valid currency market matrices, only the boundaries of the shape can be considered due to the relationship (7).

Figure 7: Possible values of ρ_3 as a function of ρ_2 when ρ_1=0 and ρ_1=0.9

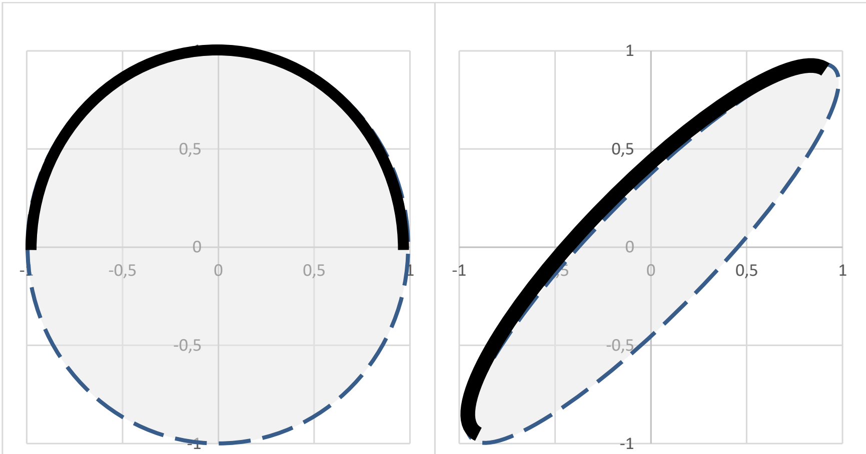

source: Author's own work

The correlation must exist both ex-post and ex-ante due to the absence of arbitrage. Although a correlation matrix based on historical data may be indefinite – for example, in the case of missing data (Szüle, 2012) – great caution is required when the goal is to construct a correlation matrix that reflects individual expectations. This may be the case when a portfolio manager wishes to test scenarios that better reflect his worldview or expectations, or when a risk manager expects extreme correlations in the foreign exchange market. In this case, the goal is to construct a matrix that meets the minimum conditions of valid correlation matrices (symmetry, positive semidefiniteness) and satisfies relation (7) arising from the triple determinacy and arbitrage-free nature of the foreign exchange market. One possibility is to substitute the so-called implied volatilities observed in the options markets into relations (1-4). The resulting so-called implied correlation matrix will satisfy the condition arising from the triple determinacy of currencies, since it was derived by using that condition. However, when assumptions about correlations that differ from those of the market are used, one must keep relations (1-4) in mind! We face this problem when a given currency does not have a (liquid) options market or when there is not a sufficiently long time series available for estimating time series-based volatility/correlation.

This is especially true if we want to determine the correlation matrix between currency pairs formed from more than three currencies: for example, if we include CHF in addition to EUR, HUF, and USD, we already need a 6x6 correlation matrix to satisfy relations (1-4). In general terms, m different currencies give us an $\frac{m(m-1)}{2}$ correlation matrix, and any correlation triplet formed from it must satisfy the no-arbitrage condition.

Another consequence of the linear interconnection of foreign exchange market returns is that partial correlations in the foreign exchange market can take values between -1 and 1 in the presence of spot arbitrage in the foreign exchange market. By definition, partial correlation measures the strength of the linear relationship between two random variables by filtering out a third random variable. This means that if we fix the exchange rate of one currency in a currency trio, the log returns of the other two currencies will be perfectly identical or perfectly opposite to each other, so they will naturally have a perfect positive or perfect negative linear relationship with each other.

In summary, the set of currency market correlations is narrower than in the general case, as currency exchange rates are linked to each other through the arbitrage circle of currency trios. Within currency trios, the correlation between pairs is determined not only by the two currency exchange rates, but also by the third.

2.3 The literature on FX correlation modelling

Mizuno et al. (2004) provides a brief description of the time scale dependence of currency market correlations: examining EUR-USD-JPY and EUR-USD-CHF currency trios, while second-by-second returns are uncorrelated, medium-strong correlations can be measured between most currency pairs at a resolution of a few minutes.

Zheng et al. (2014) used the Dynamic Conditional Correlation (DCC – (Engle (2002)) model to examine the correlation of several currency pairs and the relationship between correlations and certain macroeconomic indicators. Based on their results, the correlation between the Canadian and Australian dollars and the USD index is the most stable, and the relationship between the change in the level of correlations over time and macroeconomic data varies from currency to currency in terms of both statistical significance and sign.

Taking into account the broader scope of research on currency correlation, recent studies have focused on the correlation structure of structures containing multiple currencies using minimum spanning tree and wavelet analysis methods. Keskin et al. (2011) examines the dynamics of the cluster structure of 34 major currencies using the minimum spanning tree method and identifies the main currency clusters that emerge based on geographical regions and economic relationships. The same approach is used by Li and Liao (2020), who compare the pre-crisis (2006-2007) and post-crisis (2011-2012) periods. They find that currency networks became more widespread after the crisis, while relationships within regions strengthened. Barunik and Vacha (2018) propose a new wavelet-based estimation for co-jumps observed in currency correlations, where, according to their results, co-jumps have a significant impact on currency correlations. Wavelet multiple correlation and cross-correlation are also among the methods used by Xu and Kinkyo (2022) to measure the degree of currency market integration in Asian currency markets over a long time horizon (1991-2021). They find that correlations within the region have increased since 2005, when China introduced a managed floating exchange rate system. Kinkyo (2022) finds that the Chinese yuan (CNY) plays a central role in the Asian foreign exchange market.

2.4 Correlation and volatility patterns across currency markets

2.4.1 Correlation patterns

To describe the evolution of correlations over time, I used data on daily closing spot exchange rates for the period from March 31, 2005 to January 22, 2026 – same dataset used for correlation forecasting in Chapter 3. In this chapter, I examined the currencies of the CEE and Scandinavian region, HUF, PLN, CZK, SEK and NOK as well as JPY, CHF, and GBP, as high-volume "major" currencies against the euro and the dollar. Figure 8 show the development of correlations between individual currency pairs over time.

Figure 8: Time series of 1-month correlations between individual currency pairs

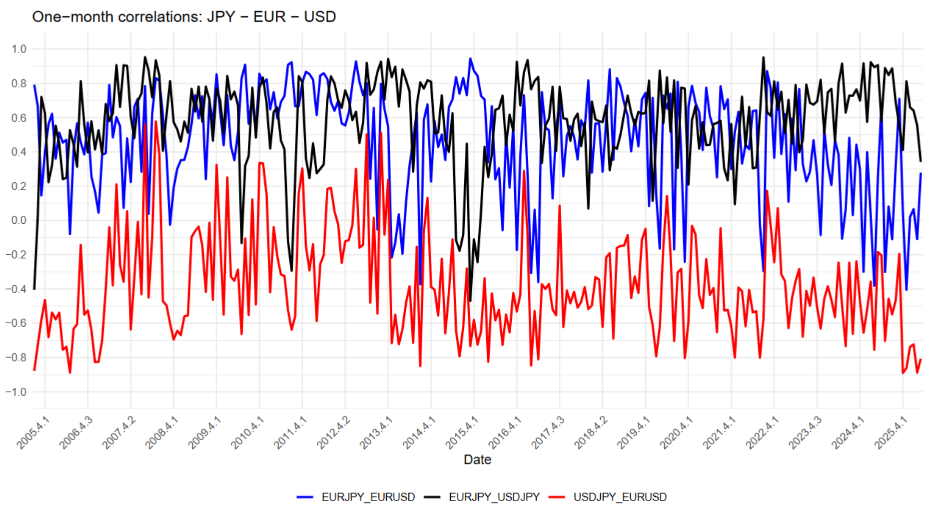

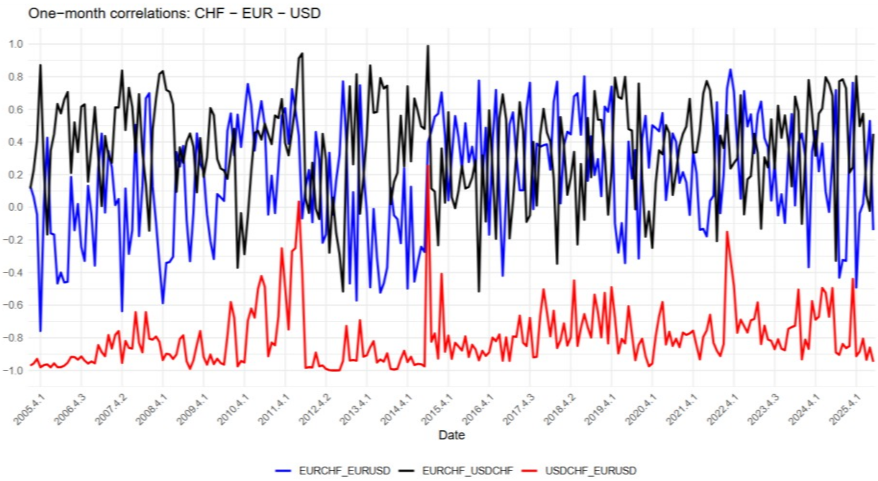
One−month correlations: CHF − EUR − USD
Date
EURCHF_EURUSD
EURCHF_USDCHF
USDCHF_EURUSD

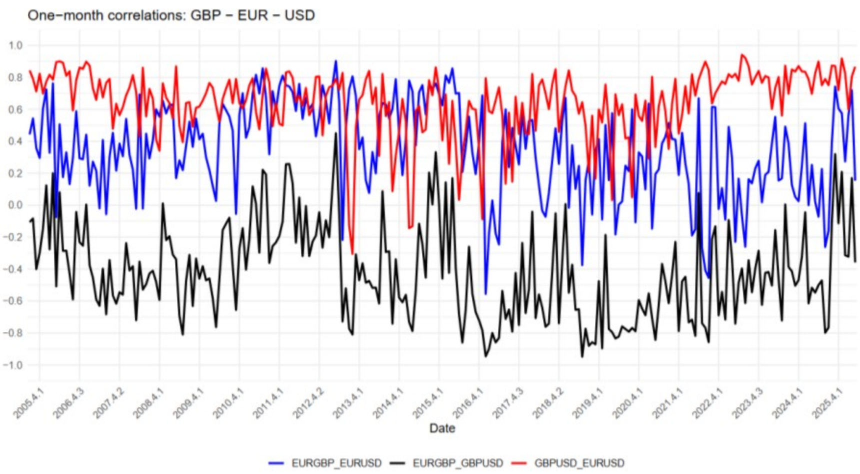
One−month correlations: GBP − EUR − USD
Date
EURGBP_EURUSD
EURGBP_GBPUSD
GBPUSD_EURUSD

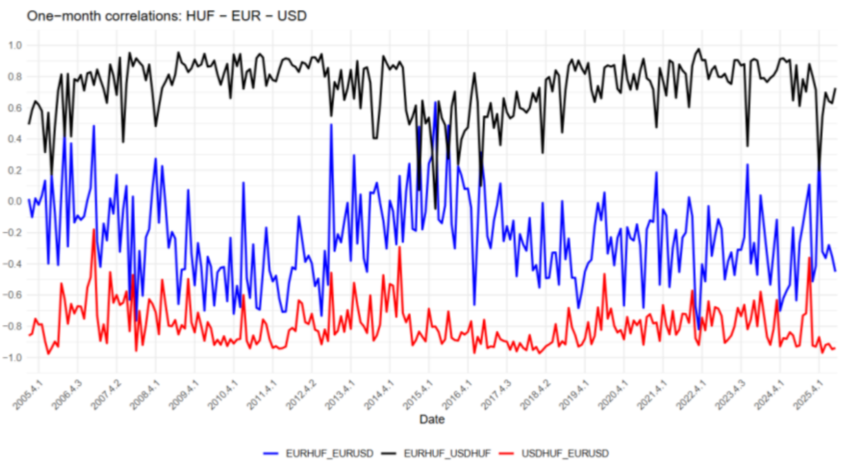
One−month correlations: HUF − EUR − USD
Date
EURHUF_EURUSD
EURHUF_USDHUF
USDHUF_EURUSD

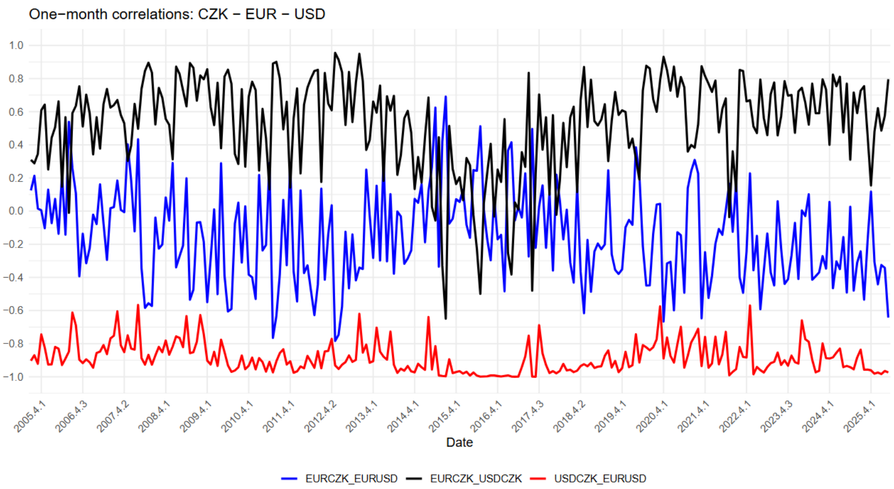
One−month correlations: CZK − EUR − USD
1.0
0.8
0.6
0.4
0.2
0.0
−0.2
−0.4
−0.6
−0.8
−1.0
2005.4.1
2006.4.3
2007.4.2
2008.4.1
2009.4.1
2010.4.1
2011.4.1
2012.4.2
2013.4.1
2014.4.1
2015.4.1
2016.4.1
2017.4.3
2018.4.2
2019.4.1
2020.4.1
2021.4.1
2022.4.1
2023.4.3
2024.4.1
2025.4.1
Date
EURCZK_EURUSD
EURCZK_USDCZK
USDCZK_EURUSD

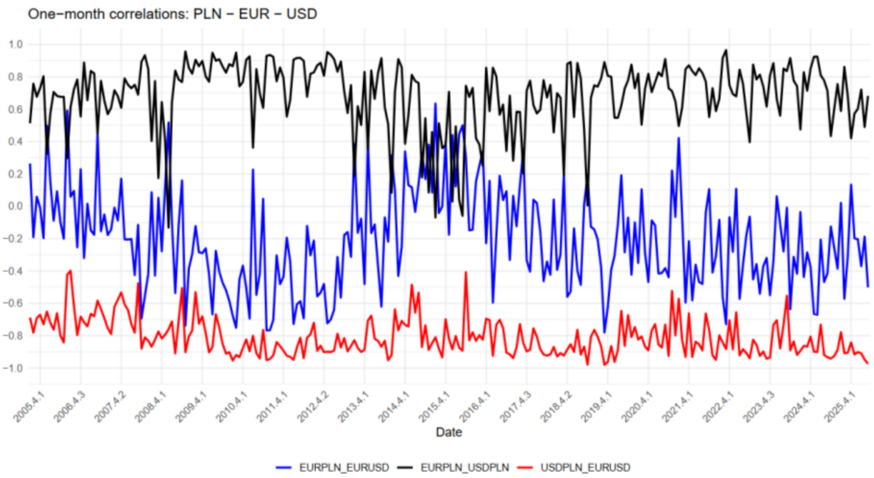
One−month correlations: PLN − EUR − USD
1.0
0.8
0.6
0.4
0.2
0.0
−0.2
−0.4
−0.6
−0.8
−1.0
2005.4.1
2006.4.3
2007.4.2
2008.4.1
2009.4.1
2010.4.1
2011.4.1
2012.4.2
2013.4.1
2014.4.1
2015.4.1
2016.4.1
2017.4.3
2018.4.2
2019.4.1
2020.4.1
2021.4.1
2022.4.1
2023.4.3
2024.4.1
2025.4.1
Date
EURPLN_EURUSD
EURPLN_USDPLN
USDPLN_EURUSD

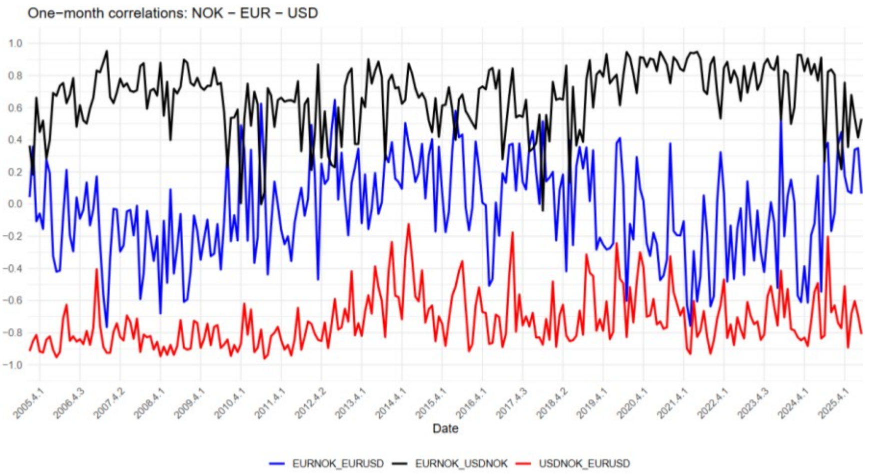
One−month correlations: NOK − EUR − USD
1.0
0.8
0.6
0.4
0.2
0.0
−0.2
−0.4
−0.6
−0.8
−1.0
2005.4.1
2006.4.3
2007.4.2
2008.4.1
2009.4.1
2010.4.1
2011.4.1
2012.4.2
2013.4.1
2014.4.1
2015.4.1
2016.4.1
2017.4.3
2018.4.2
2019.4.1
2020.4.1
2021.4.1
2022.4.1
2023.4.3
2024.4.1
2025.4.1
Date
EURNOK_EURUSD
EURNOK_USDNOK
USDNOK_EURUSD

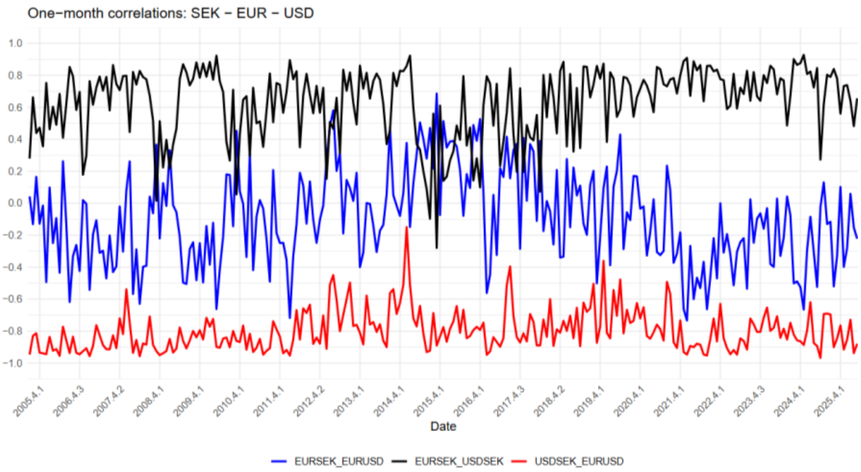

Source: Bloomberg, Author’s own work

At this point, two observations regarding the CHF and GBP are to be made. In the case of CHF, the sudden jump in the period between 2014 and 2015 indicates the Swiss central bank's abolition of the EURCHF 1.2 peg (first appearing in the EURCHF and USDCHF time series on January 15, 2015, with a daily return of ~20%, then disappearing from it). In the case of GBP, the level of correlations differs from that observed for CHF and JPY, simply because in the GBPUSD exchange rate, USD is the counter currency and GBP is the base currency. In the case of inverse pricing, the correlation values of EURUSD-GBPUSD and EURGBP-GBPUSD would change sign and be similar to those of other major currencies.

A very similar pattern can be seen in the HUF and PLN markets, while the correlations in the CZK currency trio have moved in a wider range as far as the EURCZK - EURUSD and USDCZK - EURUSD correlations are concerned. An extremely volatile period can be observed at the beginning of 2017, when the Czech National Bank abandoned the EURCZK exchange rate peg. In general, based on the evolution of the correlations over time, it can be observed that the correlation ρ(1,2) between EURCEE and USDCEE and the correlation ρ(2,4) between USDCEE and USDEUR can be considered much more stable than the relationship between EURCEE and EURUSD ρ(1,3). The evidence are to be presented in Chapter 4 and Chapter 5 but at this point the assumption here is that in the currency trios CEE-USD-EUR there are two independently traded currency pairs and one "derived" or synthetically quoted currency pair: EURUSD and EURCEE can be considered as independent markets, while USDCEE is quoted on the basis of the other two currency pairs. In other words, the value of the CEE national currency is reflected in its exchange rate against the euro (it can be supported by economic integration into the euro area and foreign trade turnover), while the euro-dollar exchange rate is clearly determined independently on the foreign exchange market. In terms of correlations, this means that the USDCEE will typically have an above-average correlation with the other two currency pairs, as the exchange rate can only move in tandem with them. If there is a significant movement in the EURCEE market, it is unlikely to affect the EURUSD market, but it will be immediately visible in the USDCEE exchange rate. Basically, there are two "independent" main currency pairs in the CEE currency trio: EURCEE and EURUSD, and USDCEE as a cross rate.

Looking at the figures, apart from the aforementioned "anomalies," a certain difference between major and regional currencies can be immediately observed: in the case of regional currencies, two of the correlations are strongly negative or positive, while one – the EURUSD-EURCCY correlation – fluctuates around zero within a wide range. The correlation structure across the FX pairs shows a highly heterogeneous but economically intuitive pattern. This dispersion is largely driven by the dominant role of the USD: pairs sharing a common USD leg (e.g. EURCCY vs USDCCY) exhibit strong positive correlations, while EURUSD and USDCCY pairs are strongly negatively correlated due to their mechanical inverse relationship.

Volatility of correlations is moderate to high, indicating that relationships are not stable over time and can shift significantly across market regimes. This is reinforced by the wide min–max ranges, with many series spanning near −1 to +1, suggesting substantial sensitivity to changing macro-financial conditions.

Distributional properties further highlight non-normality – see Table 1. Many series exhibit skewness, implying asymmetric behaviour during market stress, and elevated kurtosis, pointing to fat tails and a higher likelihood of extreme co-movements. Finally, very high autocorrelation across all series could indicate strong persistence, however this is due to overlapping periods.

Overall, the results suggest that FX correlations are structurally driven, time-varying, and prone to tail risks (further statistics including significance test of 12- and 3-month correlations can be found in Appendix A, Tables A1-A2).

In addition to analysing the evolution of volatility and correlation over time, it is worth examining the relative magnitude and position of realised (historical) correlations and those predicted by different methods in the space of possible correlations.

Table 1: Descriptive statistics of realized 1-month correlations between 2005-2026

FX_Pair	Mean	Median	Std_Dev	Min	Max	Skewness	Kurtosis	Autocorrelation
EURCHF_EURUSD	0,1805	0,2069	0,3507	-0,7765	0,8490	-0,2724	-0,7813	0,9762
EURCHF_USDCHF	0,3553	0,3890	0,3019	-0,6094	0,9972	-0,4553	-0,2773	0,9697
EURCZK_EURUSD	-0,1389	-0,1452	0,2772	-0,8088	0,7233	0,1616	-0,3920	0,9624
EURCZK_USDCZK	0,5247	0,5879	0,2771	-0,6716	0,9591	-1,1059	1,3504	0,9764
EURGBP_EURUSD	0,3624	0,3939	0,2996	-0,8171	0,9409	-0,5915	0,0691	0,9751
EURGBP_GBPUSD	-0,4160	-0,4458	0,2919	-0,9848	0,6351	0,5079	-0,1717	0,9742
EURHUF_EURUSD	-0,2283	-0,2403	0,2766	-0,8875	0,7273	0,3744	-0,0382	0,9653
EURHUF_USDHUF	0,7320	0,7868	0,1826	-0,3317	0,9852	-1,5452	3,1911	0,9726
EURJPY_EURUSD	0,4677	0,5306	0,3059	-0,5456	0,9650	-0,8052	0,0915	0,9764
EURJPY_USDJPY	0,5791	0,6285	0,2565	-0,6367	0,9674	-1,3227	2,3725	0,9730
EURNOK_EURUSD	-0,0377	-0,0401	0,3076	-0,8021	0,7426	-0,0226	-0,6672	0,9692
EURNOK_USDNOK	0,6701	0,7108	0,1969	-0,2781	0,9679	-0,9780	0,8769	0,9678
EURPLN_EURUSD	-0,2008	-0,2271	0,3021	-0,8516	0,8212	0,3760	-0,2750	0,9693
EURPLN_USDPLN	0,6833	0,7337	0,2092	-0,5945	0,9793	-1,5702	3,6861	0,9719
EURSEK_EURUSD	-0,0897	-0,1198	0,3097	-0,8580	0,7641	0,2208	-0,6394	0,9713
EURSEK_USDSEK	0,6299	0,6930	0,2237	-0,5933	0,9493	-1,3091	1,8907	0,9723
GBPUSD_EURUSD	0,6423	0,6857	0,2048	-0,6025	0,9566	-1,5962	3,8861	0,9738
USDCHF_EURUSD	-0,8034	-0,8560	0,1834	-0,9998	0,5356	2,2559	7,7743	0,9762
USDCZK_EURUSD	-0,8849	-0,9076	0,0921	-0,9998	-0,4180	1,2423	1,6132	0,9701
USDHUF_EURUSD	-0,7947	-0,8262	0,1388	-0,9821	0,1651	1,7023	4,9853	0,9723
USDJPY_EURUSD	-0,3716	-0,4246	0,3123	-0,9183	0,7367	0,8212	0,3751	0,9765
USDNOK_EURUSD	-0,7236	-0,7688	0,1750	-0,9788	0,2947	1,3179	1,7584	0,9709
USDPLN_EURUSD	-0,8105	-0,8394	0,1172	-0,9836	-0,2603	1,1918	1,3688	0,9659
USDSEK_EURUSD	-0,7959	-0,8253	0,1338	-0,9822	0,0516	1,7585	4,7898	0,9697

Source: Author's own work

Figure 9 shows the one-month implied, EWMA and realised correlations in the space of possible correlation triplets in the case of the HUF currency trio. It can be seen that the implied correlation $\rho_{(EURHUF,USDHUF)}$ and the implied correlation $\rho_{(USDHUF,USDEUR)}$ typically took positive values, while $\rho_{(EURHUF,EURUSD)}$ took both positive and negative values in a wide range.

Figure 9: One-month implied ATM, EWMA and realised correlations in the space of possible correlations

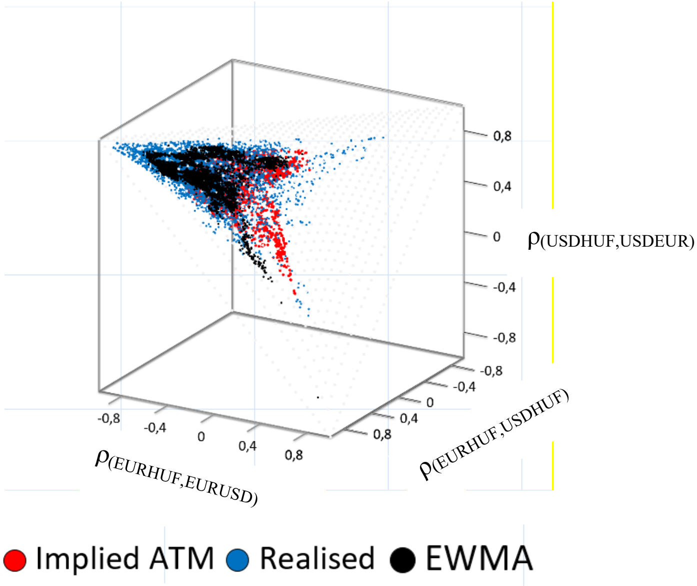

Source: Author's own work

2.4.2 Volatility patterns

In order to compare the volatility of individual exchange rates, I have plotted the daily log returns for the period under review. In Figure 10, I compare the JPY with the HUF.

Figure 10: Log returns of selected individual currency pairs

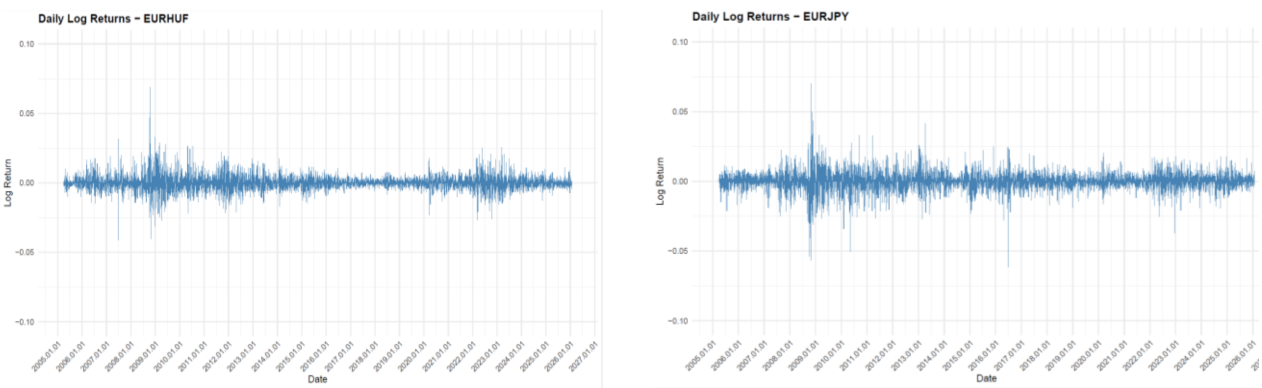

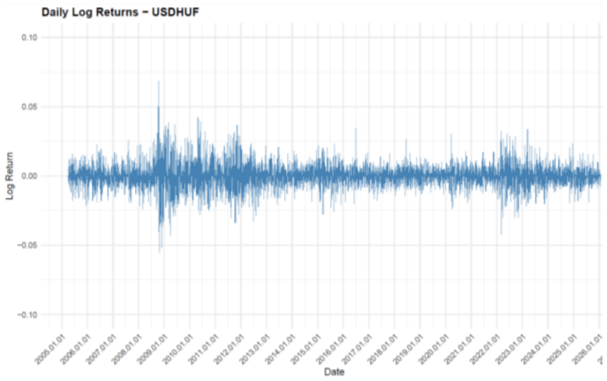

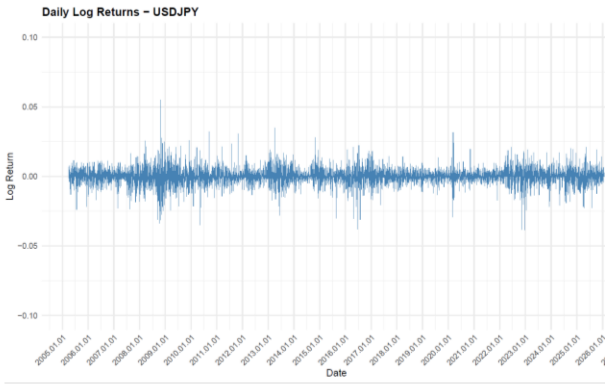

Source: Bloomberg, Author’s own work

The pattern in most major and CEE currencies is very similar to that of the currencies presented in Figure 10. Table 2 shows the volatility of the currency returns examined for the entire period, as well as the minimum and maximum volatility for the 12-month period (the highest values in each row are marked in bold).

Table 2: Historical volatilities (p.a.) of various currency rates

	EUR USD	EUR JPY	USD JPY	EUR HUF	USD HUF	EUR CHF	USD CHF	EUR GBP	USD GBP	EUR CZK	USD CZK	EUR PLN	USD PLN
Entire period	9%	11%	10%	9%	**14%**	8%	11%	8%	9%	6%	12%	9%	**14%**
12-month voatilities minimum	5%	5%	5%	3%	**7%**	2%	5%	4%	5%	0%	5%	3%	6%
12-month volatilities maximum	16%	24%	17%	20%	28%	22%	23%	14%	17%	15%	24%	21%	**29%**

Source: Bloomberg, Author’s own work

For all CEE currencies, the volatility level of USDCEE returns is significantly higher than that of EURCEE currencies, and in the case of HUF and PLN, the volatility of the USD exchange rate is somewhat prominent. Applying the average volatilities in Table 2 and using Equation 1-4 it is easy to model a currency correlation time series that is similar to "typical" major and CEE currency correlations in terms of average correlation levels, as shown in Figure 11.

Figure 11: Simulated correlation structures

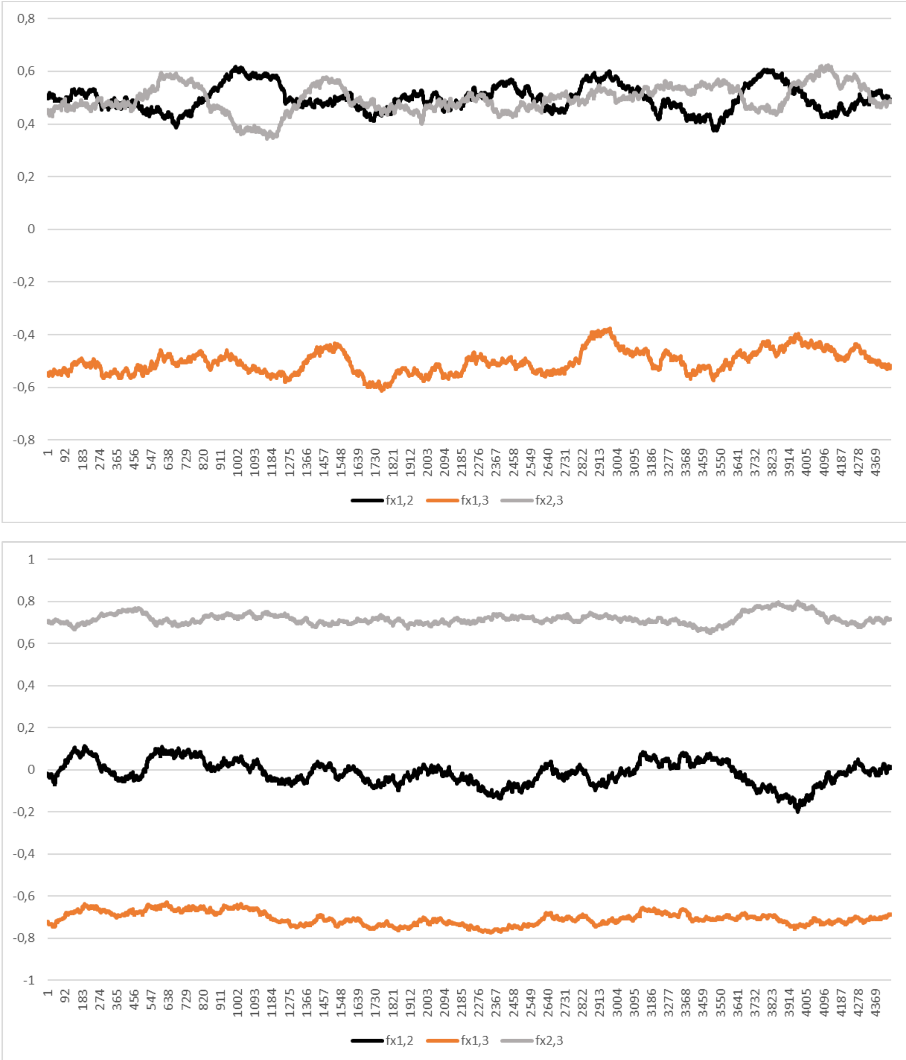

Source: Author's own work

The difference between the two figures is that the correlation between the first and second exchange rates is +0.5 (major) and 0 (regional). In both cases, the volatility of the first and second currencies is 10-10%, while in the major case, the volatility of the third currency is also 10%, and in the regional case, it is 14%.

It is straightforward to see that the level of correlations is determined by the relative position of volatilities, and it is sufficient for the volatility of the exchange rates of two currency trios to differ to get a different picture of correlation. However, a chicken-and-egg problem arises in the causal relationship: does the different volatility of the third currency cause the different levels of correlation, or do the different correlations between the first and second currency pairs cause the higher/lower volatility of the third currency? In other words, the question is why regional currencies exhibit such high volatility against the USD, but not against the EUR.

2.5 Market turnover and trading structures

One obvious difference between the CEE and major currency markets is found in currency market turnover. From the perspective of my research, the key question is how much a given currency's turnover differs from that of the euro and the dollar.

According to data available on the National Bank of Hungary website, since 2001, the average daily spot turnover of EURHUF has been almost six times higher than that of USDHUF – see Figure 12.

Figure 12: Daily average foreign exchange market spot turnover of HUF market

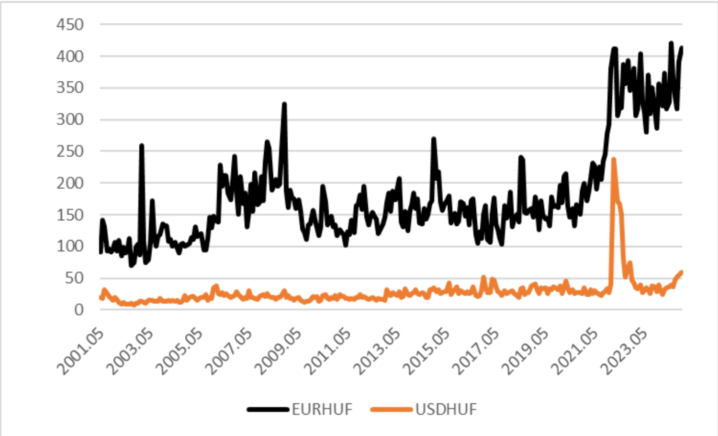

values in HUF bn

Source: NBH (2025)

The picture is similar for the PLN, although not as noticeable: in April 2022, 51% of the daily USD spot PLN turnover of USD 2.455 billion was EURPLN, and 22% was USDPLN.

Figure 13: Currency breakdown of turnover in Polish zloty (2019 vs 2022)

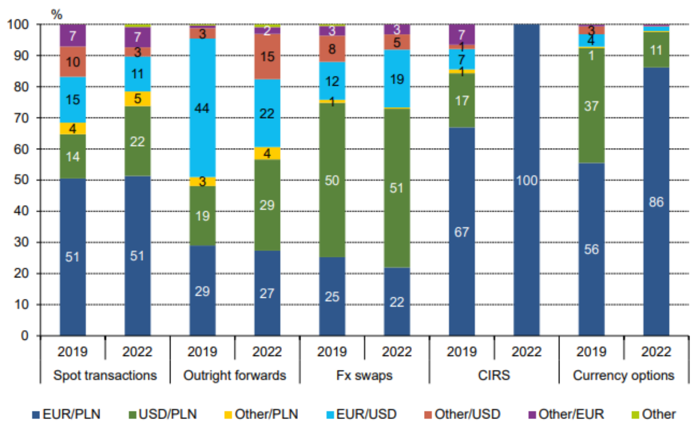

Source: NBP (2022)

Also in April 2022, according to data from the Czech National Bank, daily spot EURCZK turnover amounted to USD 608 million, while USDCZK turnover was USD 90.5 million (CNB, 2025). According to statistics available on the CNB website, EURCZK turnover has been several times higher than USDCZK turnover over the past two decades.

Turnover data for the major currencies covered by the study are based on the BIS's triennial global foreign exchange market report, shown in Table 3.

Table 3: Turnover of some major currency pairs

Spot daily turnover April 2022 (USD bn)			
EURUSD	418.8		
USDJPY	349.7	EURJPY	35.2
USDGBP	158.7	EURGBP	37.1
USDCHF	59.2	EURCHF	20.0

Source: Author's own work (based on BIS (2022), pp. 14–28)

The dominance of the dollar in all currencies is noticeable. If currency market volumes have some explanatory power in terms of correlations, then when comparing the correlations of the HUF and JPY currency trios, one can expect the EURUSD-EURHUF correlation to be replaced by the EURUSD-USDJPY correlation. Looking back at the correlation chart (Figure 8), we can see that this is more or less the case, with both correlations showing mostly weak/moderate relationships. Here, the similarity of the correlation signs points to a difference rather than a similarity, as the negative correlation is justified based on the USDJPY pricing convention, while it seems unjustified at first glance in the EURUSD-EURHUF relationship (to put it bluntly: why does the EUR weaken against the HUF when it strengthens against the USD?).

There is limited literature available on the relative turnover of major currency pairs. Hasbrouck and Levich (2019) presents April 2016 turnover data in Table 4 (similar values were measured between 2010 and 2013). The table, which can be summarized row by row, clearly shows that the foreign exchange market can basically be divided into two camps: currencies traded against the dollar and currencies traded against the euro. For most currencies, the dollar is the dominant currency, while for Scandinavian currencies it is the euro. However, it is important to note that, according to data available on the Swedish central bank's website, USDSEK trading has been more significant than EURSEK on the spot and forward markets since 2019 (Riksbank, 2025).

Table 4: Settlement turnover of the main currencies

Table 3. CLS settlement contra currencies

The sample is all CLS spot settlements in April 2016. A row summarizes all settlements in which the row-currency is involved in the exchange. Percentage entries in the row reflect the total USD equivalent value of the settlement, broken out by the other currency in the exchange. For example, of the total dollar value of all settlements involving the AUD, 1.4% occurred in the AUD/CAD pair.

	AUD	CAD	CHF	DKK	EUR	GBP	HKD	ILS	JPY	KRW	MXN	NOK	NZD	SEK	SGD	USD	ZAR
AUD		1.4	0.2		6.2	1.9			11.4			0.0	6.4	0.0	0.2	72.2	
CAD	1.4		0.2		4.0	1.3	0.0		1.9		0.1	0.0	0.3	0.0		90.7	
CHF	0.4	0.5			30.8	2.8			1.5			0.1	0.1	0.3		63.4	
DKK					76.0	0.7								0.0		23.3	
EUR	1.7	1.1	3.5	1.1		7.5	0.0	0.0	6.4		0.1	3.8	0.2	4.7	0.1	69.7	0.1
GBP	1.3	0.9	0.8	0.0	18.3		0.0		7.8		0.0	0.1	0.3	0.1	0.0	70.2	0.0
HKD		0.1			0.8	0.5			0.2							98.4	
ILS					3.0											97.0	
JPY	4.3	0.7	0.2		8.7	4.4	0.0					0.0	0.5	0.0	0.0	81.1	0.0
KRW																100.0	
MXN		0.3			1.5	0.0										98.2	
NOK	0.0	0.2	0.2		66.9	0.9			0.1				0.0	4.4		27.3	
NZD	20.3	0.9	0.1		2.4	1.6			4.1			0.0		0.0	0.1	70.3	
SEK	0.0	0.0	0.4	0.0	67.0	0.8			0.0			3.5	0.0			28.2	
SGD	1.0				1.2	0.3			0.5				0.2			96.9	
USD	8.2	10.1	2.9	0.1	28.5	11.8	1.3	0.2	24.3	1.8	2.8	0.6	2.5	0.8	2.5		1.4
ZAR					1.7	0.1			0.6							97.5	

Source: Hasbrouck and Levich (2019) p. 41

In the case of the Norwegian krone, central bank statistics broken down by EUR and USD are not available, but according to the Norwegian central bank's 2024 publication on the Norwegian financial system, the spot market was dominated by the EUR, while the forward market was dominated by the USD (Norges Bank, 2024). The study also notes that, in line with market practice, banks always quote USDNOK via EURNOK-EURUSD:

„*In the NOK spot market, turnover has traditionally been highest in EUR, while in the NOK FX swap market, turnover has been highest against USD. If a Norwegian customer wants to buy USD with NOK, the customer's bank will probably first use NOK to buy EUR and sell EUR for USD simultaneously. The transaction is conducted in this manner because these markets are somewhat more liquid than the market for USD purchased with NOK.*".

The British pound and the Swiss franc are located between the two extremes, with the euro also playing a significant role (18% and 31%, respectively).

As empirical observations have shown, different currency pairs have different levels of correlation, and we can rightly assume a relationship between the level of correlation and currency market turnover (ratios). Examining the exchange rates of a few selected currencies against the dollar and the euro and the correlation between the EURUSD, a strong relationship can be seen between the relative turnover in the currency market and the long-term average level of correlation, as shown in Figure 14.

Figure 14: Relationship between average correlation level and relative turnover

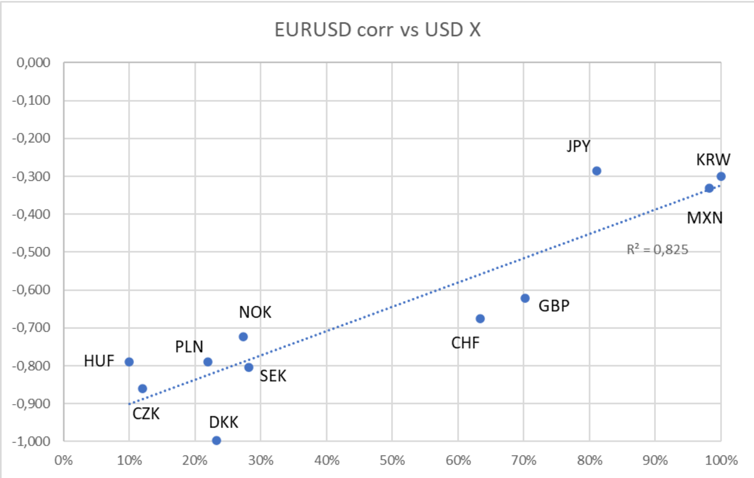

x-axis: the percentage by which the given currency was settled against the USD, based on Hasbrouck-Levich (2019) and central bank data. y-axis: the correlation level of the given currency/USD exchange rate against the EURUSD (calculated based on Reuters/Bloomberg daily closing prices over the past 20 years).

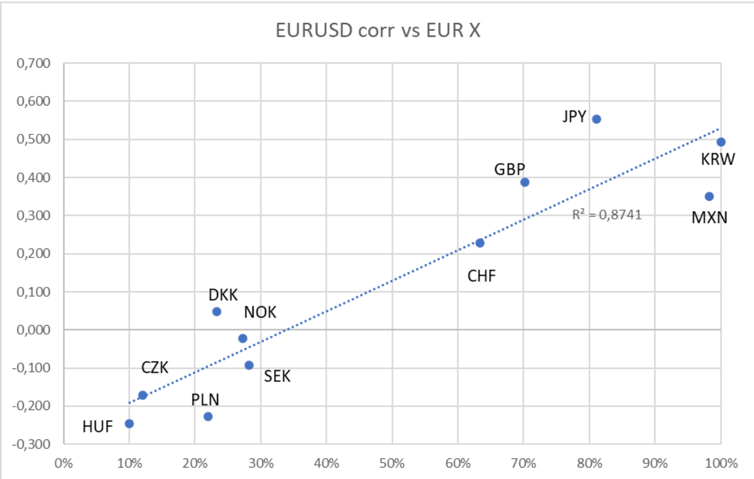

x-axis: the percentage by which the given currency was settled against the USD, based on Hasbrouck-Levich (2019) and central bank data. y-axis: the correlation level of the given currency/EUR exchange rate against the EURUSD (calculated based on Reuters/Bloomberg daily closing prices over the past 20 years).

Source: Author's own work based on Hasbrouck and Levich (2019), Refinitiv, Bloomberg

The relationship appears to be quite strong: there is a particularly strong correlation of 0.91 in case of the EUR rates and 0.935 in case of USD rates between the correlation level of individual currency pairs against the EURUSD and relative currency market turnover.

2.6 Summary and discussion

This chapter has examined the structure and behaviour of the foreign exchange market from three closely related perspectives: correlation patterns, volatility structures, and market turnover. While the global FX market appears highly integrated and liquid—dominated by the US dollar—this analysis demonstrates that important and systematic differences exist across currency groups, particularly between major- and regional currencies.

A key theoretical result underpinning the analysis is the triangular arbitrage relationship, which imposes strict constraints on the joint behaviour of exchange rates. Because any currency pair can be expressed as the product or quotient of two others, the correlation structure of FX returns is not arbitrary but mechanically linked to the relative volatilities of the underlying exchange rates. This implies that, unlike in general financial markets, admissible correlation combinations in the FX market lie on a constrained surface rather than within a broad feasible set.

The empirical analysis confirms that correlations differ markedly across currency groups. Major currencies such as JPY, CHF, and GBP exhibit relatively stable and moderate correlations with EURUSD, reflecting balanced relationships between EUR and USD trading. By contrast, CEE currencies (HUF, PLN, CZK) and Scandinavian currencies (NOK, SEK), display asymmetric correlation structures: two correlations within a currency trio are typically strong (either positive or negative), while the third – most notably the EURUSD–EURCCY correlation – often fluctuates around zero.

The analysis shows that these differences can be largely explained by the relative volatility structure of the exchange rates. In particular, CEE currencies exhibit significantly higher volatility against the USD than against the EUR. Given the triangular relationship between variances and correlations, such asymmetry mechanically leads to lower or near-zero correlations between EURCCY and EURUSD – although the exact mechanism is yet to be clarified in later chapters.

However, volatility differences themselves raise a deeper question regarding their origin. The results are consistent with the interpretation that market structure – and specifically the distribution of trading activity across currency pairs – may contribute to the observed volatility/correlation pattern. In CEE markets, trading is heavily concentrated in EUR crosses, while USD trading plays a secondary role. This is evidenced by turnover data showing that EURHUF, EURPLN, and EURCZK volumes significantly exceed their USD counterparts, often by several multiples. In contrast, major currencies are predominantly traded against the USD, or relative volumes are more balanced.

This uneven distribution of liquidity has important implications. Higher turnover typically implies deeper liquidity, more efficient price discovery, and lower volatility. At the same time the EUR exchange rates of regional currencies tend to be more stable, while USD exchange rates exhibit higher volatility.

The analysis documents a close link between the share of USD trading in a given currency and its correlation with EURUSD, with correlation coefficients exceeding 0.9 in both EUR and USD rate specifications. This finding highlights that correlations in FX markets are not purely statistical artefacts but are closely tied to underlying trading conventions and liquidity conditions.

At the same time, the analysis raises several critical considerations. The direction of causality between turnover, volatility, and correlation remains inherently difficult to establish. While the chapter argues that liquidity drives volatility, which in turn determines correlation, the reverse channel cannot be fully excluded: persistent differences in volatility may themselves influence trading preferences and thus turnover. This endogeneity points to a potential limitation of the framework and suggests that a fully structural model of FX market dynamics would need to jointly determine all three variables. Chapter 4 and Chapter 5 therefore will attempt to provide further evidence of how heterogenous quotational practices may lead to different dynamics and volatility/correlation patterns between regional and global currency trios.

Moreover, the reliance on historical correlations – even when adjusted for volatility considerations – may be problematic in periods of structural change. Episodes such as the Swiss franc peg removal illustrate how sudden regime shifts can invalidate previously stable relationships. This underscores the need for caution when extrapolating past correlation structures into the future, particularly in risk management applications.

In conclusion, the foreign exchange market should be understood not as a single unified system, but as a layered structure in which global and regional dynamics coexist. The dominance of the USD at the global level does not eliminate the importance of regional anchors such as the euro in shaping local market behaviour.

3. FORECASTING CORRELATION IN MAJOR AND REGIONAL CURRENCY TRIANGLES

3.1 Introduction

Modelling and forecasting correlations between different financial instruments has been a key issue for decades, both from a scientific and a practical point of view. Financial globalisation has broadened the toolbox of portfolio managers and investors, adding new assets and asset classes to the possible portfolio elements. Whether we want to construct an optimal portfolio or look for the right hedging instrument to keep market risk at a given level (Dömötör and Havran, 2011), predicting correlation becomes a key focus as it basically determines the choice of the right instrument and its weight in the portfolio. In particular, accurate correlation forecasts are essential inputs to portfolio allocation, hedging, and risk management frameworks, where misestimation can lead to significant under- or overestimation of portfolio risk.

These issues are especially pronounced in foreign exchange (FX) markets, which constitute the largest and most liquid financial market globally, with an average daily turnover of approximately USD 9.6 trillion according to the latest Bank for International Settlements Triennial Central Bank Survey (BIS, 2025). Beyond their aggregate size, FX markets are characterized by intense cross-currency trading activity, where investors actively exploit relative value opportunities across currency pairs and construct positions through combinations of bilateral exchange rates.

While trading activity is dominated by major currencies – accounting for the vast majority of global turnover – the role of regional currencies should not be underestimated. For European investors in particular, smaller currencies can play an important role in diversification strategies, as international portfolios are increasingly exposed to a broader set of currency risks. These currencies often exhibit distinct macroeconomic sensitivities, liquidity conditions, and risk premia compared to major currencies, reflecting differences in monetary policy regimes, external vulnerabilities, and market depth (Mancini et al. (2013)). As a result, they may offer diversification benefits relative to core currencies. At the same time, their correlations with major currencies, as well as among themselves, are highly time-varying and sensitive to both global and regional shocks, particularly during periods of financial stress when correlations tend to increase and diversification benefits diminish (e.g., Engle (2002)). Measuring and predicting the linear correlation coefficient can also help predict periods of market stress Berlinger et al. (2016), which is crucial from a financial stability perspective. This makes accurate modelling and forecasting of FX correlations particularly important in portfolios exposed to multiple currency risks.

As shown in Chapter 2 distinctive feature of FX markets is that exchange rates are not independent assets but are linked through triangular arbitrage relationships. For any three currencies, the corresponding exchange rates must satisfy a no-arbitrage identity, which imposes structural restrictions on their joint distribution. As a consequence, the correlation matrix of FX returns is inherently constrained and of reduced dimensionality, reflecting the linear dependence among exchange rates within a currency trio. This feature differentiates FX markets from other asset classes and provides a natural framework for deriving correlations from variances. Similar to the estimation of volatility, which is a key element in the modelling of market prices, there are currently two widely used methods for predicting the correlation between returns: the time series-based approaches, which are based on historical data, and the so-called implied measure, which can be derived from option market prices. While the former can be calculated from a model fitted to a historical sample, the implied measure can be derived from implied volatility derived from option market prices (see: Chapter 3.4). Although, as Rebonato (1999) famously said, "implied volatility is the wrong number to put in the wrong formulae to get the right price", the correlation estimate based on the implied measure is currently the only one that reflects market expectations.

This chapter presents my own research addressing the first research question, which aims to compare the forecasting performance of implied and time series-based models for the regional (CEE, Scandinavian) and some selected global currency markets and to analyse the main reasons for any differences observed. This chapter heavily relies on previous work done regarding the Forint market (Misik (2023)) but involves results with wider scope that is not yet published including data between 2005-2025 and a wider range of currency trios. The results show that the EWMA-based correlation provides the most accurate forecast for regional currency pairs against EUR and USD currencies for the one-month forecast periods. Moreover different methods do not produce the same ranking of forecasts, when comparing exchange rate volatility forecast versus exchange rate correlation forecast.

3.2 Related literature and research hypotheses

It was in the 1970s, following the emergence and popularisation of modern portfolio theory, that research into the correlation between stock markets began. Two decades later, in the 1990s, researchers also began to study foreign exchange markets. The correlation within stock markets and between individual stocks was studied by Longin and Solnik (1995) and Karolyi and Stulz (1996), among others. Their key findings were about the dynamic nature of correlations influenced by business cycles and macroeconomic factors and the importance of increased international stock market correlations in the globalized financial markets. Meanwhile, the first comprehensive studies of the foreign exchange market have been carried out by Bollerslev and Engle (1993) and Campa and Chang (1998). As correlation forecasting based on purely historical data raised obvious questions about its forecasting ability, and as the significant expansion and increasing liquidity of derivatives markets made it possible to examine forecasts based on market expectations, more and more attention was paid to implied measures. Bodurtha and Shen (1995) and Siegel (1997) calculated implied correlation using so-called implied volatility based on option prices.

More recent research using implied measures focuses mainly on equity markets: Driessen et al. (2009) investigated the relationship between expected return and correlation risk in the case of the S&P100 stock index and its components using implied correlations in the stock market.

The results of the literature that focuses on comparing the forecasting ability of time-series and implied correlations are mixed:

Campa and Chang (1998) examined the forecasting ability of the implied correlation between DEM - USD - JPY based on 1600 trading days between 1989 and 1995. According to their results, the implied correlation outperformed the time series-based methods (GARCH (1,1), EWMA and Historical) at both one-month and three-month forecast horizons. Walter and Lopez (2000) studied forecasting accuracy for two currency trios: their results for the DEM - USD - CHF currency trio, in addition to DEM - USD - JPY, do not support the results of Campa and Chang(1998) regarding the superiority of implied correlation. Castrén and Mazzotta (2005) extended his research to other currencies, comparing the forecasting ability of the implied correlation with the GARCH and EWMA methods for currency trios formed by EUR, USD, GBP, JPY. According to the results of the study, based on data from 1992-2004, the predictive ability of implied correlation differs for each currency pair and for each period. More recent work on the foreign exchange market does not support the superiority of the implied measure over GARCH models. Chong (2004) in relation to value-at-risk estimates, concluded that estimates based on implied correlations overestimate VaR in calm periods, but react late to market changes in turbulent periods, i.e. underestimate VaR. Although Chong (2004) studied implied volatility, his conclusions are relevant to the study of implied correlation because, as we have already seen, implied correlation is a function of implied volatility.

The application of implied correlation goes beyond correlation forecasting: Esposito and Laruccia (1999) mapped the underlying structure of dollar exchange rates with a principal component analysis calculated on the basis of the implied correlation matrix. Another major advantage of using implied correlation is that the correlation matrix derived in this way is always positive semidefinite, i.e. in cases where we want to use a correlation/covariance matrix consisting of many currency pairs, we do not need to resort to matrix correction methods.

Comparison of the various correlation forecasts in the foreign exchange market has been somewhat limited recently, and analyses covering the CEE region are more or less absent. This thesis attempts to fill this gap by comparing the forecasting ability of different correlation forecasts in the CEE FX market, covering the currency trio of CZK, HUF and PLN rates against EUR and USD on one - month horizon.

Research based on implied correlation received a new impetus at the turn of the millennium as Demeterfi et al. (1999) and Carr and Madan (2002), who extended the estimation of implied volatility based on ATM option pricing to the entire observable volatility smile, thereby establishing the theoretical and practical background for variance swap pricing. Covariance swaps can be priced in a manner analogous to variance swaps, on the basis of which market estimates of correlation can also be obtained. A widely discussed issue in the literature concerning implied measures is the difference between values calculated based on real and risk-neutral probabilities: while portfolio managers and investors are primarily interested in real-world correlations, risk-neutral implied measures can be calculated from option pricing. In the case of volatilities, it can be stated that the two measures only coincide when the market price of volatility risk is zero (Chernov, 2007). This assumption has been refuted by Carr and Wu (2009) and Driessen et al (2009), among others, making some correction of implied volatilities necessary, for example, based on DeMiguel et al. (2013). The literature supports the view that using information obtained from option prices improves volatility and correlation forecasts, for example Kourtis et al. (2016) and Buss and Vilkov (2012), however considered as a biassed forecast (Gereben and Pintér, 2005).

Despite the extensive literature on FX volatility and correlation modelling—including multivariate GARCH models such as the Dynamic Conditional Correlation (DCC) framework by Engle (2002), realized covariance approaches based on high-frequency data (Andersen et al. (2003); Barndorff-Nielsen and Shephard (2004), as well as option-implied measures of dependence (e.g., Campa and Chang (1998); Carr and Wu (2009))—empirical research has predominantly focused on major currency pairs. Further developments have incorporated realized measures into multivariate volatility and correlation forecasting frameworks

(Noureldin et al. (2012); Hansen et al. (2012); Bauwens and Xu (2025)), improving the modelling of time-varying dependence structures. In parallel, a large empirical literature has compared the predictive performance of historical and implied volatility measures, often documenting that implied information contains forward-looking components that enhance forecast accuracy (e.g., Christensen and Prabhala, (1998); Blair et al. (2001); Benavides and Capistrán (2012)).

Based on the literature outlined before one can find ambiguous results regarding the forecasting ability of different methodologies and very little work done on regional currency markets. Driven by these findings I set up two hypotheses.

H1: In the regional currency markets, implied correlation forecasts do not provide more accurate predictions than time series-based correlation estimates.

Also, following the strong relationship between volatility and correlation one might expect volatility- and correlation forecasting will provide the same result in terms of comparing the predictive power of different methods. My second hypothesis therefore:

H2: Volatility and correlation forecasting in the CEE currency market provides the same ranking when different methodologies are compared.

3.3 Data and sample construction

The time-series based correlation models – to be introduced in Chapter 3.4 – are used with data from Bloomberg BGN spot rates (Px_Last), between 31 March 2005 – 22 January 2026.

For the implied correlation measures the one-month forward ATM, as well as butterfly and daily last prices of implied volatilities of risk reversal strategies are also based on Bloomberg daily close data. Provided implied volatility data are not available for all FX rates until January 2006 and time-series models are using a 500-trading day rolling window the effective correlation forecasting period is between 2 March 2007 – 24 December 2025.

3.4 Methodology

3.4.1 ATM Implied correlation

Foreign exchange option markets provide a rich source of information about investors' expectations regarding future currency movements. One can calculate implied correlations by the use of Formulae (3) that requires respective volatility information. The ATM (At-The-Money) Implied correlation methodology derives market-implied correlations from observed ATM option volatilities, offering a forward-looking measure of co-movement that reflects current market sentiment and pricing of risk. Unlike historical correlations, which are based on past observations, ATM implied correlations incorporate the collective expectations of market participants and are therefore particularly useful for risk management, portfolio construction, and derivative valuation.

3.4.2 Variance swap (VS) Implied correlation

The thesis extends the scope of researching forecasting ability of implied measures by adding the so-called variance swap based methodology (hereafter referred to as "VS"), following the approach of Demeterfi et al. (1999) and Carr and Madan (2002). This approach extends the implied correlation estimate based on the pricing of ATM options to the entire observable volatility smile, i.e. it provides another alternative to the implied correlation estimate when comparing correlation estimates.

Demeterfi et al. (1999) provide an intuitive derivation of how, for a given maturity, an increasing number of option positions with different strike prices can be used to create an exposure that is insensitive to the price change of the underlying asset, but replicates the subsequent realised variance: appropriate weighting of put/call options with strike prices below/above the ATM, proportional to the inverse square of the strike ($1/K^2$ where K^2 is the squared strike price). Based on the fact that the difference between the value at maturity and the current value of a portfolio consisting of put/call options with strike prices below/above the ATM, weighted by $1/K^2$, is captured by half the variance, and that a payment equal to the negative logarithmic return of the exchange rate can be synthetically generated using option and forward positions, the authors derive the following formulae as the current market price of the variance:

$$K_{var} = \frac{2}{T}\left[rT - \left(\frac{S_0}{S_*}e^{rt} - 1\right) - \log\left(\frac{S_*}{S_0}\right) + e^{rt}\int_0^{S_*}\frac{1}{K^2}P(K)dK + e^{rt}\int_{S_*}^{\infty}\frac{1}{K^2}C(K)dK\right] \tag{8}$$

where K_{var} is the annualized variance, T is the respective maturity, r is the risk free log return, S_* is the ATM forward price of the underlying asset – defining the boundary between liquid put and call option strikes, S_0 is the actual spot price of the underlying asset, P(K) and C(K) are the prices of the put and call options with strike price of K. Since in reality there are no quoted option prices for strike prices on a continuous scale from zero to infinity, formulae (8) cannot be applied in practice. Therefore, the authors recommend the following discrete approach:

$$\pi_{CP} = \sum_i w(K_{iP}) P(S, K_{iP}) + \sum_i w(K_{iC}) C(S, K_{iC}) \tag{9}$$

where w denotes the weights to be assigned to the put and call options with the i-th strike, K_{iP} and K_{iC} denote the put and call options with the i-th strike. The weights applied are determined by formulae 10/a and 10/b:

$$w_c(K_0) = \frac{f(K_{1c}) - f(K_0)}{K_{1c} - K_0} \tag{10/a}$$

$$w_c(K_1) = \frac{f(K_{2c}) - f(K_{1c})}{K_{2c} - K_{1c}} - w_c(K_0) \tag{10/b}$$

in general:

$$w_c(K_{n,c}) = \frac{f(K_{n+1,c}) - f(K_{n,c})}{K_{n+1,c} - K_{n,c}} - \sum_{i=0}^{n-1} w_c(K_{i,c}) \tag{11/a}$$

$$w_p(K_{n,p}) = \frac{f(K_{n+1,p}) - f(K_{n,p})}{K_{n,p} - K_{n+1,p}} - \sum_{i=0}^{n-1} w_p(K_{i,p}) \tag{11/b}$$

The same approach is used by, among others, Carr and Madan (2002), who list three alternatives to volatility trading, in which they introduce an alternative variance forecast that takes into account all the "moneyness" options listed for the given maturity (i.e. taking into account the total volatility smile).

The calculations of time-series based correlations are based on daily log returns calculated from spot prices, while the calculation of implied correlations from implied volatilities is based on relationships (12/a-12/b). In most data sources risk-reversal (RR) and butterfly (BF) option strategies are quoted and these can be used to calculate the implied volatility of different moneyness options using a closed formulae (example: options with a delta of 25):

$$\sigma_{25RR} = \sigma_{25\Delta CALL} - \sigma_{25\Delta PUT} \tag{12/a}$$

and

$$\sigma_{25BF} = \frac{\sigma_{25\Delta CALL} + \sigma_{25\Delta PUT}}{2} - \sigma_{ATM} \tag{12/b}$$

where RR is the risk reversal and BF is the butterfly position.

The one-month implied correlation forecasts on a given day are derived from the implied volatilities quoted for the respective maturities.

3.4.3 Time series-based correlation

In line with the related literature, this thesis compares the predictive power of implied correlation with correlation coefficients derived from time series models following the approach of Campa and Chang (1998). These methods are historical correlation, the exponentially weighted moving average (EWMA) and the GARCH model, and additionally a more recent approach – Heterogenous Autoregressive (HAR) model – is applied in this study.
In the case of historical correlation, the prediction for the upcoming period of length T-t is the realised correlation looking back over a similar period:

$$\rho_{1,2_{t-T}}^{HIST} = \frac{\sum_{j=1}^{n}(s_{1,t}-\overline{s_1})(s_{2,t}-\overline{s_2})}{\sqrt{\sum_{j=1}^{n}(s_{1,t}-\overline{s_1})^2}\sqrt{\sum_{j=1}^{n}(s_{2,t}-\overline{s_2})^2}} \tag{13}$$

where $s_{1,t}$ and $s_{2,t}$ are the daily log returns of exchange rates 1 and 2, respectively, while $\bar{s}_1$ and $\bar{s}_2$ are the average daily log returns for the given period.
The EWMA method is a method developed by J.P. Morgan/Reuters (RiskMetrics) in correlation forecasting (RiskMetrics, 1996), the essence of which is that more recent data are given greater weight and these weights decrease exponentially:

$$\rho_{1,2_{t-T}}^{EWMA} = \left[\frac{1}{\sum_{j=1}^{n}\lambda^j}\right]\sum_{j=1}^{n}\lambda^j\rho_{1,2_{t-j-T,T}} \tag{14}$$

which can be rewritten as follows:

$$\rho_{1,2_{t-T}}^{EWMA} = \frac{\sum_{j=1}^{n}\lambda^j(s_{1,t}-\overline{s_1})(s_{2,t}-\overline{s_2})}{\sqrt{\sum_{j=1}^{n}\lambda^j(s_{1,t}-\overline{s_1})^2}\sqrt{\sum_{j=1}^{n}\lambda^j(s_{2,t}-\overline{s_2})^2}} \tag{15}$$

The λ parameter in the formulae determines the rate of decrease of the weights. In line with the literature, I apply a λ parameter of 0.94 and 0.97 and the size of the rolling window is n=500 days.

To model the conditional variance and capture the asymmetric effects of shocks, in this thesis I utilize the Glosten, Jagannathan, and Runkle (GJR-GARCH) specification introduced by Glosten et al. (1993). Standard GARCH models assume that positive and negative shocks of equal magnitude exert a symmetric impact on conditional volatility. However, asymmetric volatility shocks are often observable on CEE currency markets. In these emerging economies, bad news — such as sudden capital flight, geopolitical tensions, or heightened global risk aversion — typically triggers a much sharper increase in exchange rate volatility than positive shocks of a similar scale. Returns are modelled as a zero-mean process with conditionally heteroskedastic errors:

$$r_t = \epsilon_t, \;\; \epsilon_t = \sigma_t z_t \tag{16}$$

where z_t follows a standardized Student-t distribution to accommodate fat tails. The conditional variance evolves as:

$$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 + \gamma\epsilon_{t-1}^2 I(\epsilon_{t-1} > 0) \tag{17}$$

where ε_{t-1} is the t-1 day error term, σ_{t-1} is the square root of the t-1-day conditional variance (t-day GARCH volatility), ω is the long-run equilibrium variance, $\alpha \geq 0$, $\beta \geq 0$ and $\alpha + \gamma \geq 0$ ensures positive variance. The indicator function $I(\epsilon_{t-1} > 0)$captures the asymmetric impact of negative shocks (provided the CEE currency is in the nominator positive returns are considered as potentially negative shocks in the CEE currency trios). In case the asymmetry in Scandinavian or major currencies would not manifest the model collapses to GARCH(1,1) specification – see Bollerslev (1986). The model is estimated using a rolling window of 500 trading days. For each estimation window, 21-day-ahead volatility forecasts are generated by iterating the conditional variance equation and aggregating daily forecasted variances. To ensure robustness, only converged models with persistence below unity are retained. In cases of convergence failure or implausible forecasts, volatility is approximated using scaled historical standard deviation.

A fourth class of models is based on the heterogeneous autoregressive (HAR) framework of Corsi (2009), which captures the persistence of volatility across multiple time horizons. The HAR model is motivated by the heterogeneous market hypothesis, according to which market participants operate on different time scales, generating long-memory behaviour in volatility dynamics.

In this study, realized variance is approximated using squared daily returns. To stabilize variance and improve model fit, a logarithmic transformation is applied. The HAR model is specified for forecasting cumulative multi-period volatility as:

$$\log(RV_{t,t+H}) = \beta_0 + \beta_d \log(RV_t) + \beta_w \log(RV_t^{(w)}) + \beta_m \log(RV_t^{(m)}) + \epsilon_t \quad (18)$$

where $RV_{t,t+H} = \sum_{i=1}^{H} r_{t+i}^2$ denotes the realized cumulative variance over the next $H = 21$ trading days, $RV_t = r_t^2$ is the daily realized variance, and $RV_t^{(w)}$and $RV_t^{(m)}$represent the average of log realized variance over the previous 5 and 22 trading days, respectively. The error term ϵ_t is assumed to be independently distributed.

The model is estimated using a rolling window of 500 observations. For each estimation window, the HAR regression is fitted on historical data, and a direct forecast of the logarithm of cumulative 21-day variance is obtained using realized volatility measures up to time t, ensuring no look-ahead bias.

The resulting forecasts are transformed back to the original scale to obtain cumulative volatility forecasts.

This direct multi-horizon HAR specification provides a simple yet effective framework for capturing volatility persistence and multi-scale effects, while avoiding the error accumulation associated with iterative multi-step forecasting.

3.5 Forecast evaluation

There are various methods applied in related research to quantification and comparison of the predictive power of correlation forecasts. The comparison of implied measures with time series-based forecasts, i.e. the quantification of the accuracy of the forecasts obtained with different methods, was carried out in accordance with the international literature using the following methods: Root mean square error (RMSE), Mean Absolute Error (MAE), Fisher z-Transformed RMSE (z-RMSE), Variance-Weighted RMSE (WRMSE) and Quasi likelihood (QLIKE).

The RMSE indicator calculates the Root Mean Squared Error according to the following formulae – in line with Campa and Chang (1998), Walter and Lopez (2000) and Castrén and Mazzotta (2005):

$$RMSE = \sqrt{\frac{\sum_{t=1}^{n}(\hat{\rho}_t - \rho_t)^2}{n}} \quad (19)$$

where $\hat{\rho}$ denotes the estimated correlation, while ρ denotes the realised correlation at time t for the next one month.

The mean absolute error (MAE) – that is often used in FX forecasting since Meese and Rogoff (1983) – measures the average absolute deviation between forecasted and realized correlations and is defined as:

$$MAE = \frac{1}{n}\sum_{t=1}^{n} \mid \hat{\rho}_t - \rho_t \mid \tag{20}$$

where $\hat{\rho}_t$denotes the forecasted correlation and $\rho_t$denotes the realized correlation for the subsequent period.

Compared to RMSE, MAE is less sensitive to large forecast errors and therefore provides a more robust measure of predictive accuracy in the presence of extreme observations.

Since correlation coefficients are bounded between $-1$and 1, their variance depends on the level of correlation. To address this issue, forecast accuracy is also evaluated using the Fisher z-transformation:

$$z_t = \frac{1}{2}\ln\left(\frac{1+\rho_t}{1-\rho_t}\right) \tag{21}$$

The corresponding z-transformed RMSE is calculated as:

$$z\text{-}RMSE = \sqrt{\frac{\sum_{t=1}^{n}(\hat{z}_t - z_t)^2}{n}} \tag{22}$$

where $\hat{z}_t$and $z_t$denote the transformed forecasted and realized correlations, respectively.

The Fisher transformation removes the bounded nature of correlation coefficients and stabilizes their variance, allowing for more meaningful comparisons of forecast accuracy across different correlation regimes. Highlighting these benefits the approach has been applied in various fields including meteorology or psychology: Branstator (1986), Silver and Dunlap (1987).

Following Campa and Chang (1998), we also report variance-weighted RMSE, where correlation forecast errors receive greater weight in periods when the corresponding exchange rates are more volatile and correlation errors are therefore more relevant for portfolio risk.

$$WRMSE = \sqrt{\frac{\sum_{t=1}^{n} w_t(\hat{\rho}_t - \rho_t)^2}{\sum_{t=1}^{n} w_t}} \tag{23}$$

where w_t denotes the weighting variable. Following the triangular variance relationship underlying FX correlations, the weights are defined as the product of the realized volatilities of the two corresponding exchange rates:

$$w_t = \sigma_{1,t}\sigma_{2,t} \tag{24}$$

In addition to error-based measures, forecast accuracy is evaluated using a likelihood-based loss function analogous to the quasi-likelihood (QLIKE) measure commonly employed in volatility forecasting (Patton and Sheppard, 2009). While the original QLIKE measure is defined for variance forecasts, a correlation-based analogue can be derived from the Gaussian likelihood of a 2×2 correlation matrix. Specifically, Equation (25) corresponds to the negative Gaussian log-likelihood obtained when the forecasted and realized dependence structures are represented by 2×2 correlation matrices.

For a given forecast-realization pair, the QLIKE loss is defined as:

$$L_t^{QLIKE} = \log(1 - \hat{\rho}_t^2) + \frac{2(1 - \rho_t \hat{\rho}_t)}{1 - \hat{\rho}_t^2} - 2 \tag{25}$$

where $\hat{\rho}_t$ denotes the forecasted correlation and ρ_t denotes the realized correlation for the subsequent period.

Overall forecasting performance is then evaluated by averaging the loss across all forecast periods:

$$QLIKE = \frac{1}{n}\sum_{t=1}^{n} L_t^{QLIKE} \tag{26}$$

Unlike RMSE and MAE, which evaluate forecast accuracy solely on the basis of forecast errors, the QLIKE measure assesses the adequacy of the entire dependence structure implied by the forecast. In particular, the loss function penalizes overconfident forecasts, i.e., forecasts with correlations close to ± 1, especially when realized correlations differ substantially in magnitude or sign. Consequently, QLIKE is particularly sensitive to forecasting failures during periods of heightened market dependence and stress. Lower QLIKE values indicate superior forecasting performance.

3.6 Results

3.6.1 Overall forecasting performance

At each trading day t, a one-month-ahead correlation forecast is generated. The realized correlation is subsequently computed from daily returns over the following 21 trading days. Forecast accuracy measures are calculated across the entire out-of-sample period, yielding approximately N forecast-realization pairs for each correlation series.

Table 5 reports the RMSE values obtained by the different correlation forecasting methods for each currency-pair correlation. The best-performing method (the one with the lowest RMSE) is highlighted in green. The results confirm the relevance of distinguishing between major and regional currency trios. For major currencies, option-implied correlations generally provide the most accurate forecasts and the differences between the two best forecasts in general (EWMA_97 and ATM implied) are statistically significant at least at 5% significance level for almost all FX pairs. This finding is consistent with the existence of deep and liquid FX option markets, where option prices aggregate forward-looking information about future exchange rate dynamics and dependence structures.

For regional currencies, however, historical methods, particularly EWMA, frequently outperform implied correlations, however statistical significance can usually be stated at the EURCEE – EURUSD and at USD_ Scandinavian – EURUSD FX pairs – where ATM outperforms. This suggests that the informational advantage of option-implied measures diminishes in less liquid option markets, where quoted volatilities may contain larger pricing noise and variance premia.

The related literature confirms the valuable information content of the option markets regarding the volatility forecast (Jorion (1995) but very little attention is paid to connection between the liquidity of option markets and the quality of information. The investigation of this relationship is beyond the scope of the thesis but the inverse relationship between forecasting accuracy and option market liquidity is a feasible assumption.

Table 5: RMSE of one-month FX correlation forecasts, 2007–2025

FX – Pair	Historical	EWMA_94	EWMA_97	GARCH	HAR	ATM_imp	VS_imp
EURCHF_EURUSD	0,348	0,316	0,302	0,462	0,456	0,301	0,304
EURCHF_USDCHF	0,335	0,312	0,313	0,399	0,397	0,274***	0,278
USDCHF_EURUSD	0,183	0,181	0,188	0,225	0,191	0,152***	0,160
EURGBP_EURUSD	0,317	0,285	0,271	0,329	0,294	0,245	0,244***
EURGBP_GBPUSD	0,301	0,266	0,250	0,281	0,285	0,229***	0,230
GBPUSD_EURUSD	0,237	0,208	0,197	0,221	0,206	0,178***	0,178
EURJPY_EURUSD	0,333	0,300	0,285	0,314	0,299	0,242***	0,243
EURJPY_USDJPY	0,255	0,229	0,221	0,223	0,252	0,205**	0,209
USDJPY_EURUSD	0,360	0,321	0,303	0,324	0,327	0,268***	0,270
EURCZK_EURUSD	0,346	0,306	0,282***	0,431	0,423	0,313	0,310
EURCZK_USDCZK	0,278	0,251	0,235	0,338	0,317	0,230	0,245
USDCZK_EURUSD	0,099	0,090	0,087	0,100	0,103	0,093	0,100
EURHUF_EURUSD	0,312	0,294	0,276***	0,374	0,364	0,316	0,291
EURHUF_USDHUF	0,184	0,161	0,149**	0,170	0,171	0,167	0,149
USDHUF_EURUSD	0,137	0,169	0,166	0,172	0,155	0,118	0,126
EURPLN_EURUSD	0,324	0,298	0,279***	0,370	0,346	0,306	0,336
EURPLN_USDPLN	0,237	0,208	0,194	0,232	0,218	0,187	0,205
USDPLN_EURUSD	0,124	0,115	0,108	0,136	0,128	0,110	0,116
EURNOK_EURUSD	0,316	0,287	0,274	0,378	0,369	0,282	0,292
EURNOK_USDNOK	0,199	0,177	0,168	0,199	0,202	0,165	0,169
USDNOK_EURUSD	0,185	0,170	0,163	0,195	0,205	0,148***	0,149
EURSEK_EURUSD	0,351	0,311	0,288	0,372	0,350	0,296	0,302
EURSEK_USDSEK	0,247	0,215	0,200	0,245	0,242	0,195	0,196
USDSEK_EURUSD	0,156	0,137	0,129	0,148	0,137	0,120***	0,122

Note**:** The table reports the RMSE values of the alternative correlation forecasting methods for the examined currency trios over the 2007–2025 sample period. The lowest value, indicating the most accurate forecast, is highlighted in green. Diebold–Mariano tests compare the predictive accuracy of EWMA(0.97) and ATM-implied correlation forecasts. Long-run variance is estimated using Newey–West HAC standard errors with 21 lags (Newey and West, 1987). ***, ** and * denote significance at the 1%, 5% and 10% levels, respectively.

Source: Author's own work

An even more interesting pattern emerges when distinguishing between independently quoted and derived exchange rates within regional currency trios (the distinction between "independent" and "derived" or "synthetic" exchange rate regimes are explained in detail in Chapter 4 and Chapter 5). For correlations involving independently traded exchange rates (i.e., EURCEE–EURUSD and EUR_Scandinavian–EURUSD), the EWMA approach consistently provides the most accurate forecasts. In contrast, for correlations involving derived exchange rates, such as EURCEE–USDCEE or USDCEE–EURUSD, option-implied correlations often perform comparatively better.

A possible explanation is that independently traded exchange rates exhibit relatively stable correlation patterns, making them well suited to historical forecasting methods. For derived exchange rates, however, the correlation is mechanically linked to the expected future volatilities of the underlying exchange rates through the triangular arbitrage relationship. As option-implied correlations directly reflect these volatility expectations, they may contain information that is not captured by historical correlations alone.

The results based on the MAE criterion largely reinforce the conclusions obtained from the RMSE analysis. For major currency trios, ATM implied correlations provide the most accurate forecasts in all cases. For regional currencies, the EWMA(0.97) specification is the dominant forecasting method, here as well, with only a few exceptions, primarily involving correlations between USD/regional exchange rates and the other members of the currency trio. Overall, the MAE results confirm that implied information is particularly valuable in major currency markets, while historical dependence structures appear more informative for regional currencies.

The zRMSE results lead to very similar conclusions. Since the Fisher transformation places greater emphasis on forecasting accuracy when correlations approach their boundaries, these findings indicate that the superiority of ATM implied correlations in major currency markets is not limited to average forecasting performance but also extends to periods characterized by strong dependence. Likewise, the dominance of EWMA for regional currencies remains largely unchanged, suggesting that the observed differences between major and regional currency regimes are robust across different forecast evaluation measures.

A somewhat different picture emerges from the variance-weighted RMSE (WRMSE) results. Here, historical methods become considerably more competitive and outperform implied measures in a larger number of cases, including several major-currency correlations involving CHF. Since WRMSE assigns greater weight to periods of elevated volatility, this finding suggests that the advantage of implied correlations diminishes when forecast performance is evaluated primarily during turbulent market conditions. In particular, periods of high volatility are often associated with rapid changes in correlation structures, which may be difficult to capture using option-implied measures alone.

Table 6: Best forecasting methods according to alternative accuracy measures for one-month FX correlation forecasts, 2007–2025

FX Pair	RMSE	MAE	zRMSE	WRMSE	QLIKE
EURCHF_EURUSD	ATM Imp.	ATM Imp.	ATM Imp.	Hist Vol	Hist Vol
EURCHF_USDCHF	ATM Imp.	ATM Imp.	ATM Imp.	HAR	Hist Vol
USDCHF_EURUSD	ATM Imp.	ATM Imp.	ATM Imp.	Hist Vol	GJR-GARCH
EURGBP_EURUSD	VS Imp.	ATM Imp.	VS Imp.	VS Imp.	Hist Vol
EURGBP_GBPUSD	ATM Imp.	ATM Imp.	ATM Imp.	VS Imp.	GJR-GARCH
GBPUSD_EURUSD	ATM Imp.	ATM Imp.	ATM Imp.	VS Imp.	GJR-GARCH
EURJPY_EURUSD	ATM Imp.	ATM Imp.	ATM Imp.	ATM Imp.	Hist Vol
EURJPY_USDJPY	ATM Imp.	ATM Imp.	ATM Imp.	VS Imp.	GJR-GARCH
USDJPY_EURUSD	ATM Imp.	ATM Imp.	ATM Imp.	ATM Imp.	GJR-GARCH
EURCZK_EURUSD	EWMA97	EWMA97	EWMA97	Hist Vol	Hist Vol
EURCZK_USDCZK	ATM Imp.	EWMA97	ATM Imp.	Hist Vol	Hist Vol
USDCZK_EURUSD	EWMA97	EWMA97	EWMA97	Hist Vol	GJR-GARCH
EURHUF_EURUSD	EWMA97	EWMA97	VS Imp.	Hist Vol	Hist Vol
EURHUF_USDHUF	EWMA97	EWMA97	VS Imp.	EWMA97	GJR-GARCH
USDHUF_EURUSD	ATM Imp.	ATM Imp.	ATM Imp.	Hist Vol	GJR-GARCH
EURPLN_EURUSD	EWMA97	EWMA97	EWMA97	Hist Vol	Hist Vol
EURPLN_USDPLN	ATM Imp.	EWMA97	EWMA97	EWMA97	GJR-GARCH
USDPLN_EURUSD	EWMA97	EWMA97	EWMA97	EWMA97	GJR-GARCH
EURNOK_EURUSD	EWMA97	EWMA97	EWMA97	Hist Vol	Hist Vol
EURNOK_USDNOK	ATM Imp.	EWMA97	EWMA97	EWMA97	GJR-GARCH
USDNOK_EURUSD	ATM Imp.	ATM Imp.	ATM Imp.	ATM Imp.	GJR-GARCH
EURSEK_EURUSD	EWMA97	EWMA97	EWMA97	EWMA97	EWMA94
EURSEK_USDSEK	ATM Imp.	EWMA97	ATM Imp.	EWMA97	GJR-GARCH
USDSEK_EURUSD	ATM Imp	ATM Imp.	ATM Imp.	ATM Imp.	GJR-GARCH

Note: The table reports the best-performing forecasting method according to each forecast accuracy measure over the 2007–2025 sample period. Historical methods are highlighted in blue, while option-implied methods are highlighted in red.

Source: Author's own work

This conclusion is further reinforced by the QLIKE results. Under this likelihood-based measure, which penalizes forecasting failures more heavily during periods of strong dependence and market stress, historical approaches clearly dominate. In particular, forecasts based on historical volatility and GARCH-type models provide the most accurate predictions for the majority of correlation series. The WRMSE and QLIKE results indicate that although option-implied correlations are highly informative under normal market conditions, historical methods appear more robust during periods of heightened uncertainty and market stress.

The consistency of the RMSE, MAE, and zRMSE results suggests that the observed differences between major and regional currency trios are not driven by a particular forecast evaluation metric but reflect genuine differences in the underlying correlation-generating processes. It is also noteworthy that the effectiveness of correlation forecasting methods depends not only on the currency group but also on the underlying market structure. While option-implied correlations perform best for major currencies and for several derived exchange-rate relationships within regional currency trios, historical methods dominate for independently traded exchange rates within the same currency group (EURCCY-EURUSD).

Table 7 shows the forecasting improvements of the ATM implied and EWMA methods relative to the naïve benchmark forecast based on the previous one-month realized correlation (historical correlation).

Table 7: Average RMSE and Forecast Improvements Across Currency Regimes

Regime	No. of correlations	Actual RMSE	EWMA(0.97) RMSE	ATM implied RMSE	EWMA(0.97) improvement	ATM implied improvement
Major currencies	9	0.297	0.259	0.233	−12.2%	−21.5%
Regional currencies	15	0.233	0.200	0.203	−14.0%	−12.8%
Independent regional correlations	5	0.330	0.279	0.303	−14.9%	−8.2%
Derived regional correlations	10	0.185	0.160	0.153	−13.2%	−17.0%

Note: The table shows average RMSE values across the specified correlation groups. Major currencies comprise the CHF, GBP, and JPY currency trios, while regional currencies comprise the CZK, HUF, PLN, NOK, and SEK trios. Independent correlations refer to correlations between independently traded exchange rates within regional currency trios (e.g., EUR/regional–EURUSD), whereas derived correlations involve at least one exchange rate obtained as a cross rate through the triangular arbitrage relationship (e.g., EUR/regional–USD/regional or USD/regional–EURUSD). Improvements are calculated relative to the naïve benchmark based on the previous one-month realized correlation. Negative values indicate lower RMSE and therefore superior forecasting performance.

Source: Author’s own work

For major currency trios, ATM implied correlations reduce RMSE by 21.5% relative to the naïve benchmark, compared to 12.2% for EWMA(0.97). For regional currencies the advantage of implied correlations disappears, and EWMA(0.97) slightly outperforms ATM implied correlations. An even more pronounced difference emerges when distinguishing between independently traded and derived exchange-rate correlations within the regional currency group. For independently traded exchange rates, EWMA achieves substantially larger forecast improvements than ATM implied correlations (14.9% versus 8.2%). For derived exchange-rate correlations, ATM implied correlations become superior, reducing RMSE by 17.0% compared with 13.2% for EWMA.

3.6.2 Forecast performance across market regimes

To examine the stability of the results over time, the full sample is divided into three subperiods: 2007–2013, covering the Global Financial Crisis and the European sovereign debt crisis; 2014–2019, representing a relatively tranquil period characterized by low volatility and accommodative monetary policies; and 2020–2025, which includes the COVID-19 pandemic, the subsequent inflationary period, and heightened geopolitical uncertainty. For each subperiod, I calculate the average RMSE improvement of the alternative forecasting methods relative to the naïve benchmark based on the previous one-month realized correlation. Table 8 reports the average forecasting improvements for the different methods.

The results reported in Table 8 provide strong support for the robustness of the main findings. While the absolute forecasting performance of the individual methods varies across subperiods, the relative ranking of the forecasting approaches remains remarkably stable within the different correlation regimes. For major currency trios, ATM implied correlations consistently outperform all historical forecasting methods throughout the entire sample period. The improvement relative to the naïve benchmark ranges from 20.0% during 2007–2013 to more than 22% in the subsequent subperiods. A very different pattern emerges for independently traded regional exchange rates. Here, EWMA(0.97) is the best-performing forecasting method in all three subperiods, reducing RMSE by approximately 13–18%, while ATM implied correlations provide substantially smaller improvements. During the first subperiod, ATM implied correlations do not improve upon the naïve benchmark at all.

Table 8: Average RMSE improvement by subperiods

Period / Regime	EWMA(0.97)	GARCH	HAR	ATM Imp.	Variance-Swap Imp.
2007–2013 Major	-13.0%	+13.1%	+9.1%	-20.0%	-19.8%
2007–2013 Independent regional	-12.9%	+18.6%	+17.4%	+0.8%	+2.9%
2007–2013 Derived regional	-8.1%	+7.8%	+7.6%	-11.0%	-9.3%
2014–2019 Major	-11.0%	+2.2%	+1.6%	-22.7%	-21.5%
2014–2019 Independent regional	-18.3%	+12.1%	+3.9%	-15.7%	-14.5%

2014–2019 Derived regional	-17.1%	0.0%	-4.6%	-20.7%	-17.2%
2020–2025 Major	-13.3%	-3.8%	-6.7%	-22.0%	-21.1%
2020–2025 Independent regional	-13.0%	+19.1%	+15.5%	-10.1%	-11.3%
2020–2025 Derived regional	-14.2%	+7.9%	+3.9%	-16.6%	-15.6%

Note: Improvements are calculated relative to the naïve benchmark based on the previous one-month realized correlation. Negative values indicate lower RMSE and therefore superior forecasting performance, while positive values indicate inferior performance relative to the benchmark.

Source: Author's own work

For derived exchange-rate correlations, however, ATM implied correlations again dominate historical methods across all subperiods. The forecasting advantage is particularly pronounced during the 2014–2019 period but remains economically significant before and after that period as well.

The subperiod analysis suggests that market structure is more important than the prevailing market environment in determining the relative performance of correlation forecasting methods. The distinction between major currencies, independently traded regional exchange rates, and derived exchange-rate correlations remains visible throughout the entire sample period, despite substantial differences in global financial conditions confirming that the forecasting performance of correlation models depends fundamentally on the underlying FX correlation regime.

3.6.3 Volatility forecasting performance

As one could have already seen, correlation forecasting in the currency market is closely linked to volatility forecasting, given the triangulation relationship explained in Chapter 2.2. The word "almost" has a special meaning here: in theory, the best volatility estimates should be the best correlation estimates for a given currency trio, but it is even more important to forecast the ratios of volatilities that determine the angles of the currency triangle. Provided the correlation forecasts required to perform volatility forecasts I compared the methods by the same evaluations as performed in Chapter 3.6.1. As shown in Table 9 according to RMSE implied measures overperform in almost all cases confirming the results of the related literature (worth noting other measures provide the same overall picture).

Table 9: RMSE of one-month FX volatility forecasts, 2007–2025

FX	VS_imp	ATM_imp	EWMA_94	EWMA_97	GARCH	HAR
EURCHF	0,0150	0,0150**	0,0193	0,0194	0,0163	0,0181
EURCZK	0,0080	0,0075	0,0076	0,0077	0,0076	0,0108
EURGBP	0,0058**	0,0060	0,0072	0,0071	0,0069	0,0124
EURHUF	0,0096	0,0083	0,0090	0,0088	0,0095	0,0157
EURJPY	0,0105	0,0105**	0,0115	0,0115	0,0110	0,0192
EURNOK	0,0093***	0,0094	0,0105	0,0104	0,0102	0,0156
EURPLN	0,0097	0,0094	0,0087	0,0087	0,0092	0,0142
EURSEK	0,0058**	0,0058	0,0062	0,0062	0,0063	0,0116
EURUSD	0,0131	0,0062	0,0064	0,0064	0,0063	0,0138
GBPUSD	0,0073**	0,0074	0,0096	0,0094	0,0090	0,0156
USDCHF	0,0135**	0,0135	0,0175	0,0170	0,0154	0,0200
USDCZK	0,0089	0,0089***	0,0099	0,0100	0,0101	0,0182
USDHUF	0,0107	0,0095***	0,0108	0,0109	0,0112	0,0232
USDJPY	0,0090***	0,0090	0,0100	0,0098	0,0097	0,0173
USDNOK	0,0097***	0,0098	0,0113	0,0112	0,0113	0,0203
USDPLN	0,0119	0,0116	0,0119	0,0120	0,0119	0,0219
USDSEK	0,0076***	0,0078	0,0086	0,0086	0,0088	0,0182

Note**:** The table reports the RMSE values of the alternative volatility forecasting methods for the examined currency trios over the 2007–2025 sample period. The lowest value, indicating the most accurate forecast, is highlighted in green. Diebold–Mariano tests compare the predictive accuracy of EWMA(0.97) and the most accurate volatility forecasts. Long-run variance is estimated using Newey–West HAC standard errors with 21 lags. ***, ** and * denote significance at the 1%, 5% and 10% levels, respectively.

Source: Author's own work

3.6.4 Decomposition of correlation forecast errors

To better understand the sources of correlation forecast errors, I decompose the squared error of the correlation estimate into components associated with the volatility ratios underlying the triangular arbitrage relationship. As shown in Chapter 2.2, the correlation between two exchange rates within a currency trio can be expressed as a function of the volatilities of the three exchange rates forming the triangle.

Rewriting Equation (3), the correlation between exchange rates 1 and 3 can be expressed as:

$$\rho_{1,3,t,T} = \frac{1}{2}\left(\frac{\sigma_{1,t,T}}{\sigma_{3,t,T}} + \frac{\sigma_{3,t,T}}{\sigma_{1,t,T}} - \frac{\sigma_{2,t,T}^2}{\sigma_{1,t,T}\sigma_{3,t,T}}\right), \tag{27}$$

which shows that correlation is determined by a combination of three volatility-ratio components. For notational convenience, define:

- $a = \frac{\sigma_{1,t,T}}{\sigma_{3,t,T}}$
- $b = \frac{\sigma_{3,t,T}}{\sigma_{1,t,T}}$

- $c = \frac{\sigma_{2,t,T}^2}{\sigma_{1,t,T}\sigma_{3,t,T}}$

so that the correlation can be written compactly as $\rho = \frac{1}{2}(a + b - c)$.

The squared forecasting error can then be decomposed as:

$$(\hat{\rho}_t - \rho_t)^2 = \frac{1}{4}\left[(\hat{a} - a) + (\hat{b} - b) - (\hat{c} - c)\right]^2 \tag{28}$$

which, using the identity $(x + y - z)^2 = x^2 + y^2 + z^2 + 2xy - 2xz - 2yz$, yields a decomposition into three main components and their interaction terms. This representation allows us to attribute total correlation forecast error to errors in forecasting the underlying volatility ratios.

A crucial feature of this decomposition is that the three components have distinct economic interpretations. The terms a and b depend exclusively on the volatilities of the two independently traded exchange rates in the currency trio (EURCCY and EURUSD in the case of CEE and Scandinavian currencies). By contrast, component c is fundamentally different: it is the only term that directly depends on the volatility of the USDCCY exchange rate. In CEE and Scandinavian FX markets, this exchange rate is typically not independently traded but is instead constructed as a cross rate from EURCCY and EURUSD. As a result, c captures the contribution of the synthetically derived leg of the currency triangle.

The decomposition results reported in Table 10 indicate that the dominant contribution to total forecasting error arises from the c^2component across all currency groups. This finding suggests that inaccuracies in forecasting the USDCCY volatility—and, more broadly, in capturing the volatility structure associated with the derived exchange rate—are the primary driver of correlation forecast errors. In other words, while both EWMA and implied approaches perform reasonably well in forecasting the volatilities of the independently traded exchange rates, their relative performance in correlation forecasting is largely determined by how accurately they capture the volatility of the cross rate and, crucially, the associated volatility ratios.

Table 10: Decomposition of total squared forecast error – one-month correlation estimates between EURCCY-EURUSD

		a^2	b^2	c^2	2ab	-2ac	-2bc	sum
HUF	EWMA_97	0,1334	0,1493	0,3565	-0,2114	0,0531	-0,1272	0,3536
	ATM implied	0,1028	0,1495	0,5201	-0,1794	-0,0104	-0,1504	0,4322
PLN	EWMA_97	0,0858	0,2462	0,3851	-0,2433	0,0052	-0,1597	0,3193
	ATM implied	0,0794	0,2551	0,4989	-0,2362	0,0359	-0,2552	0,3779
NOK	EWMA_97	0,1232	0,1309	0,3388	-0,2025	-0,0192	-0,0811	0,2901
	ATM implied	0,1087	0,1125	0,3700	-0,1737	-0,0536	-0,0628	0,3009
SEK	EWMA_97	0,0725	0,1861	0,3694	-0,1957	0,0536	-0,1631	0,3229
	ATM implied	0,0598	0,1541	0,3962	-0,1554	0,0262	-0,1386	0,3422
JPY	EWMA_97	0,1558	0,0643	0,3570	-0,1783	-0,1147	0,0447	0,3288
	ATM implied	0,1145	0,0539	0,2649	-0,1354	-0,0846	0,0196	0,2330
GBP	EWMA_97	0,0800	0,1005	0,3401	-0,1493	-0,0709	-0,0094	0,2910
	ATM implied	0,0605	0,0859	0,2687	-0,1198	-0,0462	-0,0177	0,2315

Note: In theory, the "sum" column should be consistent with the RMSE numbers shown in Table 5 satisfying the following relationship: $RMSE^2$=1/4 x (sum). Minor inconsistencies occur as correlation forecast are compared against realized correlations meanwhile volatility decomposition is applied against realized volatilities. Based on Equation (1)-(4) realized correlations should be fully consistent with realized volatilities however this is not perfectly satisfied on empirical data.

Source: Author's own work

This result provides a natural explanation for the contrasting performance of forecasting methods across currency regimes. For major currency trios, where all exchange rates are actively traded and supported by deep and liquid derivative markets, implied volatility measures yield more accurate forecasts for all components, including c, resulting in superior correlation forecasts. In contrast, for regional currencies, the USDCCY exchange rate is synthetically determined and associated option markets are relatively illiquid. Consequently, implied volatilities for this leg may contain additional noise or volatility premia, which disproportionately affect the c component and, in turn, the overall correlation forecast. By comparison, the EWMA approach – despite being backward-looking – appears to capture the relative stability of the volatility ratios more effectively in these markets.

Taken together, the decomposition highlights that accurate correlation forecasting in FX markets depends not only on forecasting individual volatility levels but, more importantly, on accurately capturing the relative structure of volatilities within the currency triangle. This insight helps reconcile the apparent discrepancy between volatility and correlation forecasting performance and underscores the role of market structure in shaping the effectiveness of alternative forecasting approaches.

3.7 Discussion and conclusions

The results of this study provide strong evidence that the forecasting performance of FX correlation models depends not only on the forecasting methodology itself, but also on the underlying market structure and the liquidity of the currencies involved. The analysis reveals a clear distinction between major currency markets (CHF, GBP, JPY) and regional currency markets (CZK, HUF, PLN, NOK, SEK), as well as between independently traded and derived exchange-rate relationships within currency trios. These findings contribute to the relatively limited literature on correlation forecasting in regional FX markets and offer a more nuanced view than the existing empirical evidence.

The first hypothesis (H1) stated that implied correlations would not provide more accurate forecasts than time-series-based methods in the CEE/Scandinavian region, similarly to findings reported for some global currency markets. The results largely support this hypothesis.

The literature reviewed in Table 11 presents mixed evidence regarding the forecasting superiority of implied correlations. Campa and Chang (1998) found that implied correlations outperformed historical approaches for the DEM–USD–JPY currency trio, whereas Walter and Lopez (2000) and Castrén and Mazzotta (2005) reported considerably less consistent advantages. More recent studies have also questioned the forecasting superiority of implied measures, particularly in less liquid markets.

The findings of this thesis are largely consistent with this latter strand of the literature. For the major currency trios analysed, ATM-implied correlations clearly outperform historical alternatives. On average, ATM-implied forecasts reduced RMSE by 21.5% relative to the naïve benchmark, substantially exceeding the improvement achieved by EWMA forecasts. This result confirms that option markets for major currencies contain valuable forward-looking information about future dependence structures and supports the conclusions of Campa and Chang (1998).

Table 11: Currency Correlation Forecasting Methods Comparison

Source	FX Rates & Period	Correlation Forecast Methods Compared	Key Findings on Method Comparison
Campa, J. M., Chang, P. H. K, (1998).	USD/DEM/JPY: Jan 3, 1989 – May 23, 1995 (1,600 obs.)	• Implied correlation • Historical correlation • EWMA correlation • bivariate GARCH(1,1)	Implied correlations clearly outperforms over 1- and 3-month forecasting period based on all evaluation approaches
Walter, C., and Lopez, J. A. (2000).	USD/DEM/JPY: Oct 2, 1990 – Apr 2, 1997 (1,679 obs.) USD/DEM/CHF: Sep 13, 1993 – Apr 2, 1997 (910 obs.)	• Implied correlation • Historical correlation • EWMA correlation • bivariate GARCH(1,1)	For USD/DEM/JPY, implied and GARCH-based correlations generate lower RMSE than simple time-series forecasts and provide useful information not present in other forecasts. For USD/DEM/CHF, implied correlation is useful for p(DEM/USD, CHF/USD) but economic benefits are small
Castrén, O., and Mazzotta, S. (2005).	Major currencies (DEM/EUR, USD, JPY, GBP): 1992 – 2004	• Implied correlation • Historical correlation • EWMA correlation • bivariate GARCH(1,1)	Implied correlations are not uniformly superior to return-based measures but provide the most consistent forecasting performance across currencies. Combining implied and historical correlation measures often yields the best forecasts, with predictive regressions explaining up to approximately 42% of realized correlation variation.

Source: Author’s own work

However, a very different pattern emerges for regional and regional currencies. For CZK, HUF, PLN, NOK and SEK, the EWMA(0.97) specification generally delivers the most accurate forecasts, particularly for correlations between independently traded exchange rates. This finding is consistent with the view that the informational advantage of implied measures diminishes when option markets are less liquid and option prices contain larger pricing noise. In this respect, the results are closer to those of Walter and Lopez (2000) and Castrén and Mazzotta (2005), who found no universal superiority of implied correlations.

The most important contribution of this chapter is the identification of a structural distinction that has received little attention in the previous literature. Existing studies have primarily focused on the distinction between major and regional currencies, whereas the present analysis suggests that the role of the exchange rate within the currency triangle is equally important.

For independently traded exchange rates, such as EURHUF–EURUSD or EURPLN–EURUSD, historical methods consistently outperform implied approaches. By contrast, for correlations involving derived exchange rates, ATM-implied correlations frequently become the best-performing forecasting method, even within regional currency groups.

This result is economically intuitive. Independently traded exchange rates are determined in separate markets and exhibit comparatively stable dependence structures. Consequently, historical methods such as EWMA can efficiently extract relevant information from past observations. Derived exchange rates, however, are mechanically linked to the expected future volatilities of the underlying exchange rates through triangular arbitrage relationships. Since implied correlations are themselves derived from option-implied volatility expectations, they naturally incorporate information about future dependence structures that may not yet be visible in historical data.

To my knowledge, this distinction between independent and derived FX correlation regimes has not been explicitly documented in the correlation forecasting literature and therefore represents an important extension to previous research.

An additional insight emerges from the comparison of forecast evaluation metrics. RMSE, MAE and zRMSE generally favour ATM-implied correlations for major currencies and EWMA for independently traded regional currencies. However, the results based on WRMSE and QLIKE indicate that historical methods become relatively more competitive during periods of elevated volatility and market stress.

This finding partially contrasts with the common argument that option-implied measures should perform best precisely when uncertainty is high because they immediately incorporate market expectations. Instead, the evidence suggests that historical models, particularly volatility-based specifications such as GARCH and realized-volatility approaches, are often more robust during turbulent periods. One possible explanation is that option-implied measures may contain substantial risk-premium and liquidity components during crisis periods, reducing their ability to forecast future realized correlations accurately.

The subperiod analysis further strengthens this conclusion. Although the sample includes fundamentally different market environments—from the Global Financial Crisis through the low-volatility period of the 2010s to the COVID-19 and post-pandemic period—the relative ranking of forecasting methods remains remarkably stable. This suggests that market structure is more important than the prevailing macro-financial environment in determining forecasting performance.

The second hypothesis (H2) proposed that the ranking of forecasting methodologies would be similar for volatility and correlation forecasting because of the close theoretical relationship between volatility and correlation. The results do not fully support this hypothesis. Therefore, H2 is rejected. The empirical evidence shows that methodologies that perform best for volatility forecasting are not necessarily those that provide the most accurate correlation forecasts.

While implied information often improves volatility forecasting in the literature, the same superiority does not consistently carry over to correlation forecasting. The results show that the best-performing volatility forecasting models are not necessarily the best-performing correlation forecasting models. This discrepancy is particularly evident in the CEE markets, where EWMA emerges as the dominant correlation forecasting method despite the availability of forward-looking implied information.

Therefore, the relationship between volatility forecasting and correlation forecasting appears weaker than theoretical considerations alone would suggest. Correlation forecasting involves additional structural features related to the joint dynamics of exchange rates that cannot be inferred solely from volatility forecasting performance.

From a practical perspective, the results imply that there is no universally superior correlation forecasting method. Portfolio managers, risk managers and financial institutions should tailor their forecasting approach to the characteristics of the currencies and exchange rates involved. For major currencies, option-implied correlations provide the most reliable forecasts and should be preferred whenever liquid option market data are available. For regional currencies and independently traded exchange rates, however, simple historical approaches such as EWMA remain difficult to outperform. Finally, for derived FX relationships generated through triangular arbitrage, implied measures appear to contain economically meaningful forward-looking information and therefore deserve greater attention in practical forecasting applications. Overall, the findings reconcile previously conflicting results in the literature by demonstrating that forecasting performance is fundamentally regime-dependent. The apparent disagreement among earlier studies may therefore reflect differences in currency groups, option market liquidity and exchange-rate structures rather than methodological inconsistencies alone.

4. QUOTATION STRUCTURES, PRICE DISCOVERY AND SHOCK TRANSMISSION IN CURRENCY TRIANGLES

4.1 Introduction and research hypotheses

Why do some currency triangles behave as tightly integrated systems while others display persistent directional asymmetries in shock transmission? At its core, this is a question of whether observed FX correlations reflect genuine bilateral information transmission, or whether they instead arise mechanically from the quotation architecture of currency triangles. Foreign exchange markets are tightly linked through triangular arbitrage relations, which mechanically connect exchange rates across currency pairs. At high frequencies, deviations from no-arbitrage conditions are typically eliminated very rapidly, indicating a high degree of integration across FX markets. Considering the similarity of quotational architectures among various currency trios and adding the triangular arbitrage relationship, the different correlation and volatility patterns raise demand for deeper insight.

Understanding the origin of these asymmetries is important for both empirical and practical reasons. Correlation structures are a key input into currency risk measurement, portfolio allocation, and stress testing. Moreover, asymmetric spillovers affect how shocks propagate across currency markets and how exchange rate movements are interpreted in empirical models. If some currency pairs do not engage in independent price discovery but are instead priced mechanically from others, observed correlations may reflect market architecture rather than information transmission. This matters in practice: correlation-based risk measures, contagion metrics, and indicators of financial integration may then misread mechanically induced co-movement as genuine information transmission, distorting cross-country inferences. Differences in quotation practices between global and regional markets may therefore contribute not only to heterogeneous spillover patterns but also to observed differences in volatility regimes across FX systems. The primary focus of this chapter, however, is to explore the existence of distinct impulse-response dynamics and spillover topologies. The analytical relation between volatility and correlation implied by synthetic construction is derived separately in Appendix D.

The existing literature offers several explanations for exchange rate co-movement. From a microstructure perspective, volatility is often treated as an outcome of liquidity provision, trading intensity, or information asymmetry, while correlation patterns are typically attributed to common information or cross-market order flow. Relatively little attention, however, has been paid to triangular arbitrage and quotational differences as a structural mechanism that may itself shape the long-run geometry of correlations within currency trios. In particular, the possibility that one leg of a currency triangle may be priced synthetically from the other two has received limited empirical scrutiny. When one leg of a currency triangle is mechanically constructed from the other two, it may absorb shocks without transmitting them back symmetrically. This mechanism predicts directional spillovers, weak reverse feedback, and impulse responses that differ systematically from those implied by explanations based solely on volatility, liquidity, or cross-market information transmission. While existing literature documents links between these phenomena and FX volatility or co-movement, it does not explicitly model the quotation architecture generating such patterns. This chapter aims to fill that gap. Using high-frequency bid–ask data sampled at 100 and 200 milliseconds, I compare impulse-response and spillover patterns across major and regional currency triangles.

The following hypothesis is formulated based on the introduction above – the third hypothesis of the thesis:

H3: Regional currency triangles exhibit stronger directional spillovers, and weaker feedback effects than major currency triangles at ultra-high-frequencies, as predicted by synthetic quotation mechanisms.

This Chapter heavily relies on the results of Misik and Szabó (2026).

4.2 Related literature

VAR-type systems are standard tools for analysing multivariate dynamics and transmission mechanisms (see among others: Sims (1980), Engle and Granger (1987), Johansen (1988), Lütkepohl (1991) Pesaran and Shin (1998)). In FX markets they are widely used to study order-flow interactions and price discovery (Hasbrouck (1991), Evans and Lyons (2002), Love and Payne (2008), Chaboud et al. (2014), Chaboud et al. (2021)). A key result is that information flow is strongly directional, especially at high frequencies, and often dominated by USD-centric trading activity. A large body of microstructure research shows that quoting behaviour is highly heterogeneous across currencies and platforms. Payne (2003) demonstrates that some currency pairs update more slowly and largely follow information generated elsewhere. Berger et al. (2008) and Mancini et al. (2013) document substantial cross-platform differences in liquidity, order flow, and price-updating intensity. These findings complement institutional evidence summarised in King et al. (2011), who emphasise that in regional and emerging-market FX, one or more legs of a triangular system are often synthetically priced and updated less frequently. More recent evidence reinforces this picture of a structured, hub-based market: Hasbrouck and Levich (2019) document that the FX settlement network is organised around vehicle currencies, with most flows routed through a small number of hub pairs; Ranaldo and Somogyi (2021) show that the information content of order flow is strongly heterogeneous across currency pairs; and Huang et al. (2025) show that dealers' constrained liquidity provision concentrates in specific pairs and times. Routing through hubs, heterogeneous information content, and concentrated liquidity provision are precisely the ingredients under which one leg of a regional triangle can end up contributing little independent price discovery. Information-leadership studies provide a second complementary mechanism. Baillie et al. (2002) and Rime et al. (2010) show that price discovery is systematically asymmetric across currency pairs, with USD-denominated quotes typically dominating non-USD crosses. Building on this evidence, Menkhoff et al. (2016) show how information flows across different customer segments shape the directionality of FX adjustment. Taken together, these results establish that heterogeneous quoting intensity and asymmetric information transmission are first-order features of FX dynamics. The most developed structural accounts of these frictions sharpen, rather than displace, this question. Vayanos and Wang (2012) show how asymmetric information and imperfect competition shape liquidity and price impact, while Babus and Kondor (2018) model how private information diffuses through a dealer network. Both are single-asset frameworks: they describe how information and liquidity govern the transmission of price-relevant information, but they do not, by construction, eliminate a market's capacity to contribute price discovery. This distinction matters for currency triangles, where the relevant question is not which leg leads but whether all three legs contribute independent channels at all. Despite these advances, the existing microstructure literature has not systematically linked such heterogeneities to long-run correlation structures within currency trios. The empirical analysis shows that exactly these features generate sharp and persistent differences between major and regional FX systems.

4.3 Data and sample construction

For the analysis, I use two publicly available sources of tick-by-tick exchange rate data: TrueFX.com (https://www.truefx.com/truefx-historical-downloads/) and Dukascopy Bank SA (https://www.dukascopy.com/swiss/english/marketwatch/historical/). Dukascopy is a well-known, long-established trading platform provider in the world of online currency trading, while TrueFX is a subsidiary of Integral Group, which primarily operates as a technology provider. In order to compare the exchange rate dynamics of different currency markets, I use

Dukascopy data for the main EUR-USD-JPY currency triangle and for the EUR-USD-HUF trio. Since HUF data is not available in TrueFX, I used this data primarily for robustness testing of the results obtained on the JPY, CHF, and GBP trios based on Dukascopy. The Dukascopy Bank SA website provides tick-by-tick data for a wide range of currency exchange rates going back several years, meaning that each price quote update represents a new bid and ask data point with a millisecond identifier. This can mean couple of 100,000s bid/ask quotes per day for each currency pair.

High-frequency data and non-equal intervals

A long-standing problem in the literature concerning ultra-high-frequency (also known as tick-by-tick, i.e., transaction-level) data is that transactions – or, in our case, price quote updates – typically do not occur at regular intervals, which can make it difficult to apply traditional time series analysis methods (see Engle (2000)). When examining the data, a considerable misalignment can be observed when comparing actual- and theoretical equal-interval between consecutive quotes. We can compare two counters: the actual time and the time elapsed in milliseconds between transactions occurring at regular intervals. Figure 15 shows the difference between the cumulative times of the uniform and actual elapsed times, measured in millions of milliseconds, for the USDJPY (based on a sample from January 20, 2025).

Figure 15: Actual vs. regular time elapsed - USDJPY

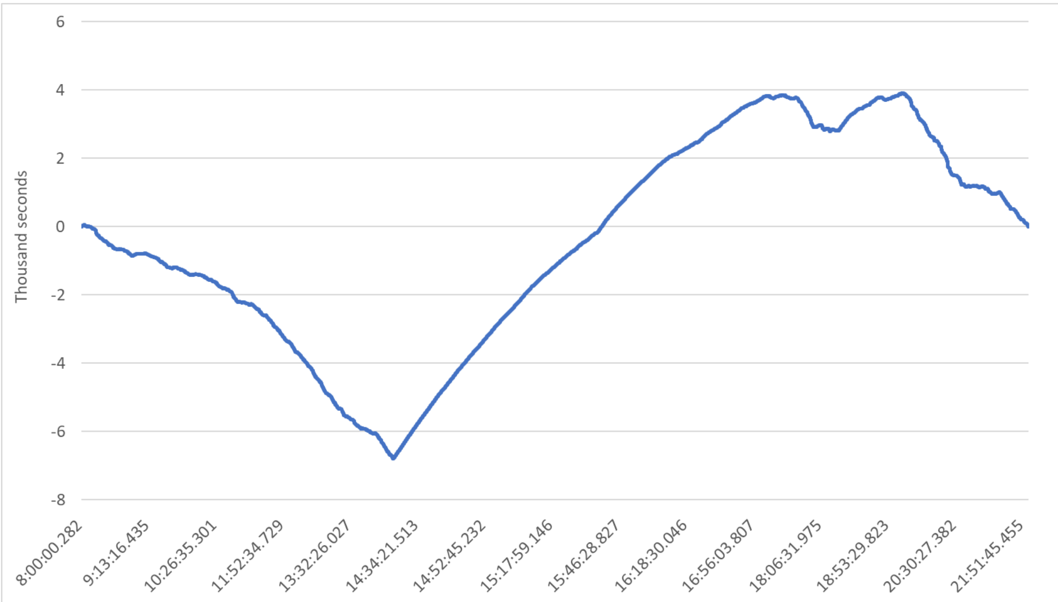

Source: Author’s own work

It can be seen that until around 2:30 p.m. (U.S. stock market opening), the spread is on an increasing negative trend, which means that until 2:30 p.m., price updates are more frequent than usual, and then the trend reverses until 5:00 p.m. At its minimum, the shift is 1.9 hours.

The literature basically uses two approaches to handle time series patterns containing irregular intervals: smoothing the time series or correcting the methodology used. An example of the first solution is converting an intraday time series into a minute-by-minute time series: in this case, the price information for a given minute will either be the same as the last price preceding it, or the interpolated value of the two price information points closest to that time point, preceding and following it (Zivot and Wang (2006) page 348).

From the previous example, it is apparent that information loss occurs, as in some markets, several hundred or even thousands of transactions/quotes can occur within one minute. In the samples examined, I found that the VAR model calculated based on transaction-level data leads to almost identical results as when, based on the samples, I convert the quotes into a time series with 100-millisecond intervals, i.e., with equal time steps. Thus, by keeping information loss sufficiently low, we can exclude the problem of non-equal intervals, which is often identified as a limitation in time series analysis. During data filtering, all quotes with a millisecond timestamp are assigned to the nearest 100-millisecond time stamp, and then I take the first data point with the appropriate time stamp as the basis (in case of missing data, the exchange rate takes the last available value), as Figure 16 shows.

Figure 16: From millisecond to 100-millisecond aggregation

Filtered data	
	EURJPY ask
14:42:16.900	174,409
14:42:17.000	174,456
14:42:17.100	174,449

	original data	
timestamp rounded to the nearest 100 millisec.	timestamp	EURJPY ask
14:42:16.900	14:42:16.895	**174,409**
14:42:16.900	14:42:16.899	174,409
14:42:16.900	14:42:16.946	174,456
14:42:17.000	14:42:16.950	**174,456**
14:42:17.000	14:42:16.951	174,456
14:42:17.000	14:42:17.001	174,449
14:42:17.000	14:42:17.003	174,449
14:42:17.100	14:42:17.055	**174,449**
14:42:17.100	14:42:17.104	174,459

Source: Author’s own work

The same approach was applied at 200 millisecond frequency for Local Projection method. As an interesting note, Figure B1 of Appendix B presents the results of an analysis performed on aggregated data from a given sample, where the transaction-level data are aggregated to seconds. In the example shown (the Bank of Japan's intervention on July 11, 2024), I find almost no significant impulse-response relationship within the currency triangle, which indicates that the impulse-response relationship occurs within a timeframe shorter than one second. In another case, however (sample from January 20, 2025), the VAR model signalled an error on the per-second data, as it found a significant linear relationship between the FX log returns, meaning that at a resolution of one second, the order of adjustment within the currency trio can no longer be detected.

4.4 Methodology

4.4.1 Vector autoregressive model

One of the main hypotheses of the thesis is that the observable differences in correlation regimes between regional major currency trios can be explained by differences in market quotations and the fingerprint can be spotted by examining the sub-second triangular dynamics.

A widely used method for examining and explaining time series variables and their relationships is the vector autoregressive model (hereinafter VAR) based on Lütkepohl (2005). In the most general form of VAR, each endogenous variable is explained by its own lagged values and the lagged values of other endogenous variables. The general form of a VAR model with p-order lags written for k variables is:

$$y_t = \mu + A_1 y_{t-1} + \cdots + + A_p y_{t-p} + u_t \tag{29}$$
$$u_t \sim N(0, \sum)$$

where y_t is the vector of endogenous variables kx1, μ is the vector containing constants kx1, the matrix A kxk contains the coefficients corresponding to the given delay, while u_t is an error term with zero expected value, $\sum$ covariance matrix. Written in matrix form for k=3 variables:

$$\begin{pmatrix} y_t^1 \\ y_t^2 \\ y_t^3 \end{pmatrix} = \begin{pmatrix} \mu^1 \\ \mu^2 \\ \mu^3 \end{pmatrix} + \begin{pmatrix} \alpha_{t-1}^{1,1} & \alpha_{t-1}^{1,2} & \alpha_{t-1}^{1,3} \\ \alpha_{t-1}^{2,1} & \alpha_{t-1}^{2,2} & \alpha_{t-1}^{2,3} \\ \alpha_{t-1}^{3,1} & \alpha_{t-1}^{3,2} & \alpha_{t-1}^{3,3} \end{pmatrix} \begin{pmatrix} y_{t-1}^1 \\ y_{t-1}^2 \\ y_{t-1}^3 \end{pmatrix} + \cdots + \begin{pmatrix} \alpha_{t-p}^{1,1} & \alpha_{t-p}^{1,2} & \alpha_{t-p}^{1,3} \\ \alpha_{t-p}^{2,1} & \alpha_{t-p}^{2,2} & \alpha_{t-p}^{2,3} \\ \alpha_{t-p}^{3,1} & \alpha_{t-p}^{3,2} & \alpha_{t-p}^{3,3} \end{pmatrix} \begin{pmatrix} y_{t-p}^1 \\ y_{t-p}^2 \\ y_{t-p}^3 \end{pmatrix} + \begin{pmatrix} u_t^1 \\ u_t^2 \\ u_t^3 \end{pmatrix} \tag{30}$$

where the superscript of the α coefficient indicates the direction of interaction between the two variables.

4.4.2 Impulse-response function

The VAR model's so-called moving average representation allows for the quantification of the spillover effects of shocks affecting individual variables.:

$$y_t = \hat{y}_t + \sum_{p=0}^{\infty} \Psi_p u_{t-p} \tag{31}$$

where $\hat{y}_t$ is the deterministic component of y_t, and the elements of the Ψ_0 unit matrix Ψ_p quantify the effect of the p-lag shock on all variables. The same logic can be used to interpret the effect of the shock arriving at time t after a period of p.

Performing the Cholesky decomposition of the covariance matrix $\sum$ of the errors u_t, we obtain the following result:

$$\sum = PP' \tag{32}$$

and introducing independent and autocorrelation-free ε_t error terms

$$\varepsilon_t = P^{-1} u_t \tag{33}$$

$$E(\varepsilon_t {\varepsilon_t}') = I \ ,$$

We can write the moving average representation of the VAR model as follows:

$$y_t = \hat{y}_t + \sum_{p=0}^{\infty} \Psi_p^* \varepsilon_{t-p} \tag{34}$$

where Ψ_p^*=$\Psi_p P$ matrix contains the orthogonal impulse response coefficients.

For the purposes of the study, I use non-orthogonal impulse response analysis, avoiding the uncertainty arising from the order of variables in the orthogonal impulse response function (see Lütkepohl (2005)). In general, non-orthogonalized impulse responses should be interpreted as the dynamic effects of a one-unit innovation in the disturbance term of the selected VAR equation rather than as the effects of an isolated shock to a particular variable. The reason is that contemporaneous correlation among the reduced-form innovations may prevent a clean separation of shocks across variables. In the present application, however, the underlying ultra-high-frequency exchange-rate returns exhibit negligible contemporaneous cross-correlations prior to aggregation to the 100-millisecond frequency. Consistent with this finding, orthogonalized and non-orthogonalized VAR specifications produce virtually identical impulse response functions. This suggests that the estimated reduced-form innovations are already close to orthogonal, and therefore the distinction between innovations and isolated shocks is of limited practical importance in the present setting. Consequently, while the impulse responses are formally interpreted as responses to one-unit innovations in the selected currency return equation, they may also be viewed as close approximations to responses generated by isolated shocks to the corresponding exchange-rate return series. In Chapter 4.5.6, Local Projection estimates are reported as a robustness analysis. The close correspondence between the VAR- and LP-based impulse response functions provides additional evidence that the estimated dynamic relationships are not driven by the particular econometric framework employed.

The VAR model looks to be suitable for examining foreign exchange market effects in cases where a sequence of individual exchange rate shocks can be identified, for example in high-frequency data. In such cases, in addition to simultaneity, a certain sequence pattern can also be identified in the movements of individual exchange rates, allowing the direction and strength of the effect to be examined. In contrast, in the case of low-frequency data – which in the foreign exchange market presumably means a frequency of less than a few seconds – many cases will appear to be simultaneous effects, even though the movement of one exchange rate actually preceded the movement of the other (two) exchange rates.

This problem of simultaneity has long been recognized in the literature. The study by Daníelsson and Love (2006) eliminates the shortcoming of the original order flow theory, according to which the simultaneous relationship between price movement and order flow (traded volume with a sign) is unidirectional: price movement is a consequence of order flow. This directionality of the relationship is clear in non-aggregated (ultra-high frequency) data, but with a certain degree of aggregation, the effect can be considered both simultaneous and bidirectional, i.e., price movements can trigger trading.

Hasbrouck (1991) examined the unidirectional relationship between order flow and price movement using ultra-high-frequency (tick-by-tick) data. This methodology became the basis for subsequent order flow studies: Dufour and Engle (2000) or Engle and Patton (2004) in relation to stock markets, Evans and Lyons (2002), Love and Payne (2008) for exchange rates, and Cohen and Shin (2003) or Green (2004) for bond markets. In these models, the VAR model explains price movements with delayed price movements and simultaneous and delayed values of order flow, while in the case of order flow, only delayed values of order flow and price movements play a role. These models therefore logically exclude the possibility of simultaneous feedback trading, which would cause significant information loss when examining dynamic relationships based on aggregated data.

Based on VAR models run on different currency triples, it is possible to examine the interaction between exchange rates and its temporal progression. The latter is the impulse response function.

4.4.3 Local Projection

Based on the seminal paper of Jordá (2005) an alternative methodology is available to conduct research on multivariate dynamic systems called local projections (LP henceforth). Major advantage of LP over VAR methodology is the improved robustness – meaning less sensitive to misspecification. The major tradeoff between the two approaches is the bias-variance tradeoff, namely VAR results higher bias/lower variance and LP results higher variance/lower bias. In our case the assumption is the data generating process is a VAR process so in expectation I assume the IRFs based on the LP will be consistent with the IRFs based on the VAR model.

The core idea of LP is to separately model the relationship to all horizons (s=0…h) by estimating the coefficients of the following regressions in a p-lag specification:

$$r_{i,t+h} = \alpha_h + \beta_h r_{j,t} + \Gamma_h' Z_t + \varepsilon_{t+h}, \tag{35}$$

where α_h is a vector of constants, β_h measures the effect of a shock at time t in $r_{j,t}$ on r_i at horizon $t + h$, Γ_h' contains the coefficients on the control variables Z_t, which in our case are lagged returns of all FX rates, and ε_{t+h} is a vector of error terms. By this specification a sizeable magnitude of past autocorrelation effect is filtered however not perfectly.

LP-IRF then calculated as follows:

$$IRF^{(h)} = \hat{B}^{(h)} d_i \tag{36}$$

where $d_i = B_0^{-1}$. For more details see Jordá (2005) and Adammer (2019).

The LP impulse-response functions calculated based on the same dataset as for VAR IRFs are shown in the Appendices. Both LP and VAR based charts are representing the impulse-response functions assuming a unit shock from the selected currency logreturn. The results consistent with the findings of the VAR modelling: the differences in dynamics between currency trios are shown noting also the variability between periods and market makers. It is worth emphasizing the symmetry of impulse-responses within the JPY trio versus the asymmetry within the HUF trio based on the sample from 20th January 2025 (Figure B7 – B9 in Appendix B).

4.4.4 Spillover Topology

To compare the empirical implications of synthetic quotation with those of the competing explanations based on volatility, liquidity, or information diffusion, I examine the full 3×3 cumulative spillover topology of each currency triangle. Following Jordà (2005), horizon-specific local projection coefficients are interpreted as impulse response functions. The resulting coefficient matrices are subsequently viewed as weighted directed networks, where each element represents the magnitude and direction of shock transmission between currency returns. This interpretation is closely related to the connectedness and network literature of Diebold and Yilmaz (2014), although the present framework is based on cumulative local-projection impulse responses rather than forecast-error variance decompositions. Aggregating the horizon-specific responses yields a cumulative spillover matrix that summarizes the overall transmission of shocks throughout the currency system. Considering formulae 36, for each horizon s, each coefficient matrix $\beta_h r_{j,t}$ measures the response of variable i at horizon h to a contemporaneous shock in variable j at time t, conditional on the control vector Z_t.
Collect the coefficients into the horizon-h response matrix:

$$\mathrm{B}^{(h)} = \begin{bmatrix} \beta_{11,h} & \beta_{12,h} & \cdots & \beta_{1N,h} \\ \beta_{21,h} & \beta_{22,h} & \cdots & \beta_{2N,h} \\ \vdots & \vdots & \ddots & \vdots \\ \beta_{N1,h} & \beta_{N2,h} & \cdots & \beta_{NN,h} \end{bmatrix}. \tag{37}$$

Each $\mathrm{B}^{(h)}$is an $N \times N$ horizon-specific spillover matrix.
Let's define

$$\Theta^{(h)} = \mathrm{B}^{(h)}. \tag{38}$$

thus

$$\Theta^{(h)} = \begin{bmatrix} \theta_{11}^{(h)} & \theta_{12}^{(h)} & \cdots & \theta_{1N}^{(h)} \\ \theta_{21}^{(h)} & \theta_{22}^{(h)} & \cdots & \theta_{2N}^{(h)} \\ \vdots & \vdots & \ddots & \vdots \\ \theta_{N1}^{(h)} & \theta_{N2}^{(h)} & \cdots & \theta_{NN}^{(h)} \end{bmatrix}, \tag{39}$$

where

$$\theta_{ij}^{(h)} = \beta_{ij,h}. \tag{40}$$

Each $\Theta^{(h)}$can be interpreted as the spillover network at horizon h.

The cumulative spillover from variable j to variable i is:

$$S_{ij} = \sum_{h=0}^{H} \theta_{ij}^{(h)} = \sum_{h=0}^{H} \beta_{ij,h}. \tag{41}$$

Collecting all entries gives the cumulative spillover matrix:

$$\mathbf{S} = \sum_{h=0}^{H} \Theta^{(h)} = \begin{bmatrix} S_{11} & S_{12} & \cdots & S_{1N} \\ S_{21} & S_{22} & \cdots & S_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ S_{N1} & S_{N2} & \cdots & S_{NN} \end{bmatrix}. \tag{42}$$

where

$$S_{ij} = \sum_{h=0}^{H} \beta_{ij,h}. \tag{43}$$

Interpretation is the following:

- positive S_{ij}: variable j pushes variable $i$upward overall,
- negative S_{ij}: variable j depresses variable i.

The directional spillovers for node i is calculated :

Spillovers TO others:

$$TO_i = \sum_{j \neq i} S_{ji} \tag{44}$$

Spillovers FROM others:

$$FROM_i = \sum_{j \neq i} S_{ij} \tag{45}$$

Net spillover:

$$NET_i = TO_i - FROM_i \tag{46}$$

Interpretation:

- $NET_i > 0$: net transmitter,
- $NET_i < 0$: net receiver.

4.5 Results

4.5.1 The impulse-response analysis of JPY currency trio

First, I examined the dynamic relationships within the EUR-USD-JPY currency trio, using the “vars package” in R. The reference date selected was 20 January 2025, and data from Dukascopy were used to fit VAR models for both the yen and forint currency trios. Due to computational capacity limits, the maximum sample size used for impulse-response analysis was 200,000 price quotes for each currency pair and for both bid and ask sides. Sample size means number of consecutive millisecond-stamped quotes for any of the three currencies within the triangle.

The augmented Dickey-Fuller test indicated stationarity in all cases for the log returns. The appropriate lag length was determined based on the Schwarz criterion, which suggested lags between 3 and 30, depending on the sample.

Although bid and ask information was available for all three currency pairs, separate analysis of these generally did not provide additional insights – except in certain cases – so the analysis was primarily based on log returns calculated from the mid-rate, which is derived from the average of the bid and ask quotes. Based on the first 200,000 observations the IRFs of the EURUSD are shown in Figure 17.

Figure 17: IRF of EURUSD in JPY Trio - 20/01/2025 – 1st sample

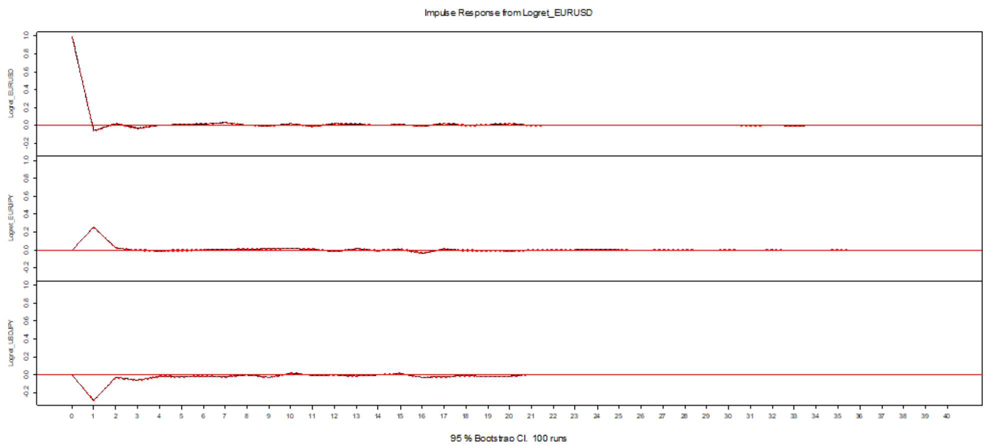

Source: Author’s own work

The response functions clearly show that a change in the EURUSD moves both the EURJPY and USDJPY rates one period later. Similarly, the EURJPY has an effect on both of the other exchange rates, with the magnitude of the impact being roughly similar, as shown in Figure 18.

Figure 18: IRF of EURJPY in JPY Trio - 20/01/2025 – 1st sample

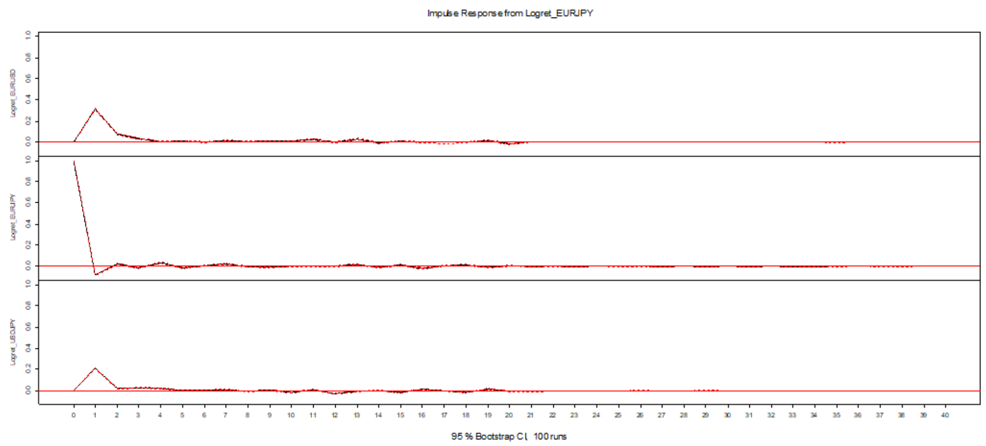

Source: Author’s own work

The EURJPY has a positive effect on all exchange rates, while in the case of USDJPY, in accordance with quotation conventions, the impact on EURUSD is negative and somewhat smaller, see Figure 19.

Figure 19: IRF of USDJPY in JPY Trio - 20/01/2025 – 1st sample

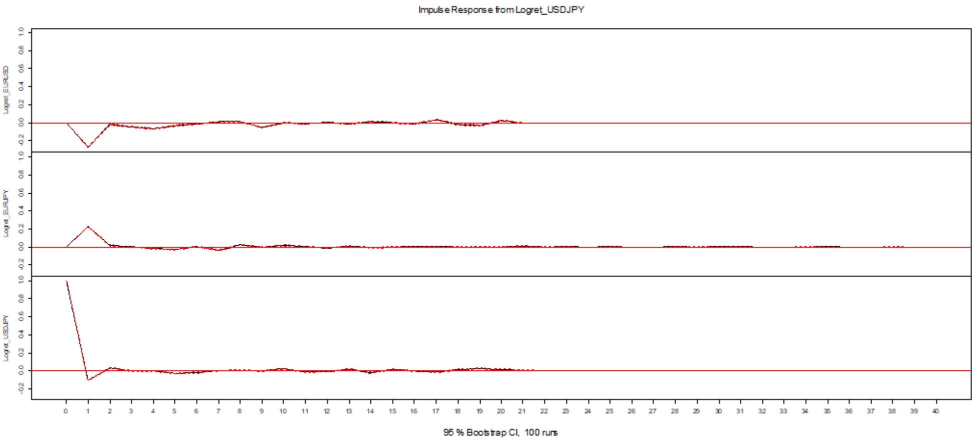

Source: Author’s own work

In the JPY currency trio, the magnitude of the effect can vary depending on the sample. For instance, in the first 200,000-sample presented earlier, the USDJPY had a greater influence on the EURJPY, whereas, based on the third 100,000-sample taken on the same day, the effect on the EURUSD was more pronounced (here, broken down separately for bid and ask quotations), shown in Figure 20.

Figure 20: IRF of USDJPY_bid in JPY Trio - 20/01/2025 – 3rd sample

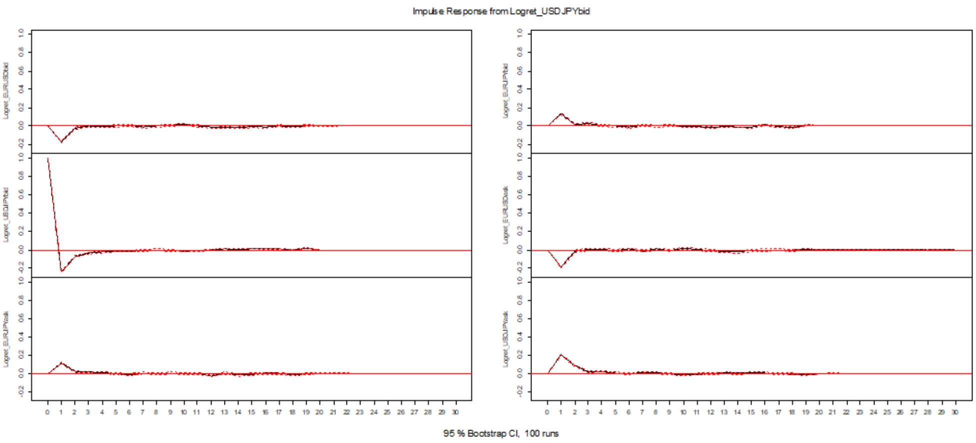

Source: Author’s own work

Based on these findings, we can draw the conclusion that within the EUR-USD-JPY currency trio, innovations arising in the currency trio produces largely symmetric responses.

4.5.2 The impulse-response analysis of HUF currency trio

For the HUF, I also examined samples maximum 200 000 observations, and in every case, the augmented Dickey-Fuller test indicated stationarity. The appropriate number of lags was determined using the Schwarz criterion, which suggested a lag order of 11.

In both samples of the HUF currency trio, the first noticeable difference compared to the JPY trio, shown in Figure 21, is that the effect of the EURUSD is essentially only reflected in the USDHUF; its impact on the EURHUF is very low and, for the most part, not significant.

Figure 21: IRF of EURUSD in HUF Trio - 20/01/2025 – 1st sample

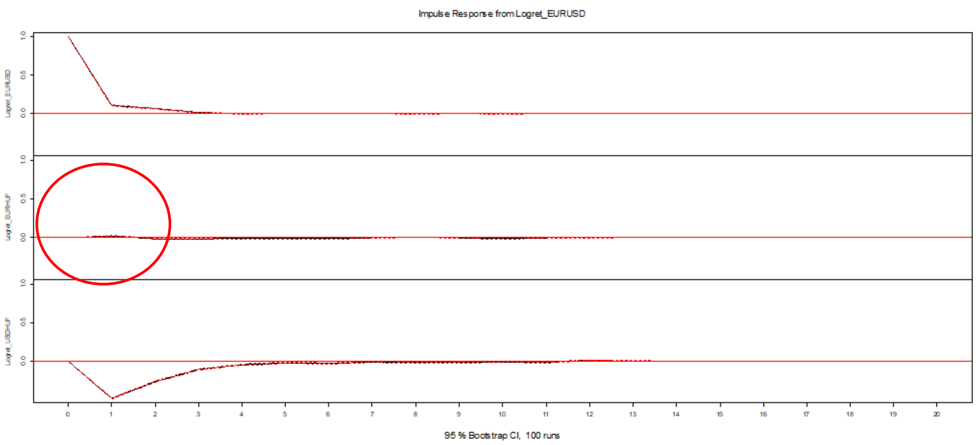

Source: Author’s own work

We can see in Figure 22, that regarding the EURHUF, the empirical results confirm the expectations: innovations in the EURHUF have no effect on the EURUSD, but only influence the USDHUF.

Figure 22: IRF of EURHUF in HUF Trio - 20/01/2025 – 1st sample

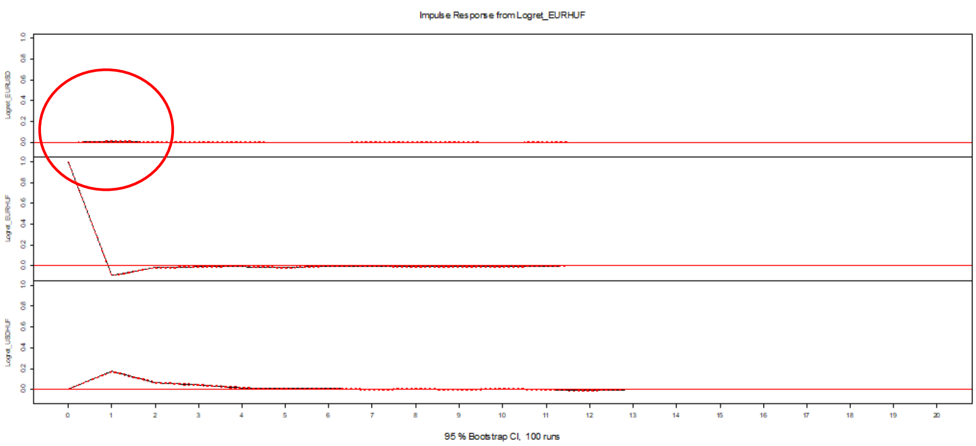

Source: Author’s own work

Figure 23 presents that the USDHUF displays a very similar pattern to the EURHUF, with the distinction that, in relation to the EURUSD, it shows a clearly low but significant negative effect.

Figure 23: IRF of USDHUF in HUF Trio - 20/01/2025 – 1st sample

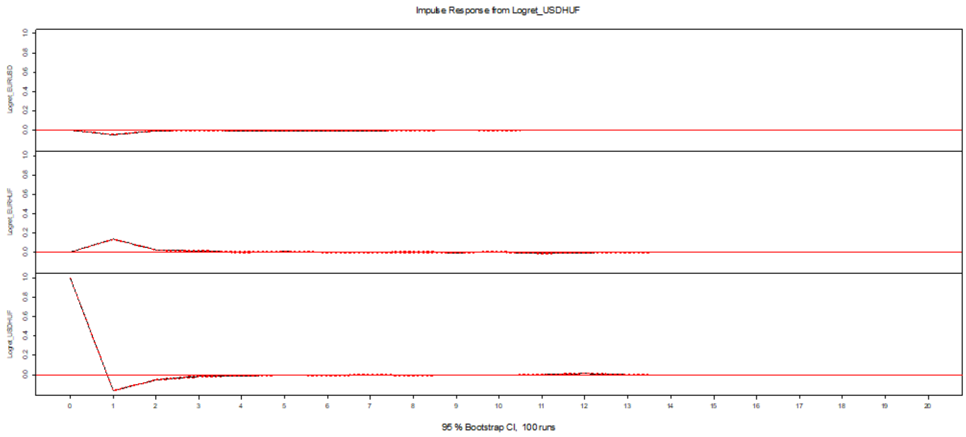

Source: Author's own work

In addition, based on the qualitative survey of quoting practices (see: Appendix C), there seem to be further evidence why empirical dataset provide the above results: In the CEE currency trio, according to four banks, the USDHUF rate is always determined by the EURUSD and EURHUF rates, which is mainly explained by the low volume of USDHUF trading (estimated at 10–20%). However, according to Bank 4, the USDCEE rate is also quoted independently, but due to low liquidity and arbitrage, it more closely tracks the other two currency rates.

In summary, based on the impulse response functions, the main finding is that in the case of the yen currency trio, although to varying degrees, all exchange rates influence the other two. In contrast, for the forint, only the EURUSD rate primarily affects the USDHUF rate, while the two forint exchange rates influence only each other.

4.5.3 Robustness analysis: The results of impulse-response analysis under special market conditions

It is noteworthy that the previously presented empirical results are derived from a randomly selected period January 20th 2025, that essentially describes a calm market environment. However, it is important to examine how robust and generalisable the findings regarding foreign exchange market dynamics actually are. Robustness should be assessed firstly by considering whether results obtained from different data sources (market makers) for the same period show similar patterns, and secondly by examining whether behavioural patterns change across different time periods. The Japanese yen market offers an excellent opportunity for the latter, as currency market events – specifically Bank of Japan interventions – are well documented and can be clearly observed across multiple data sources (CLS, 2023):

On 21/10/2022 and 24/10/2022, the BOJ sold 96.1 and 12.6 billion USD against the JPY, respectively. The first intervention was a 1.5-hour operation (15:30-17:00), while the second lasted just a quarter of an hour (0:30-0:45). Comparing the two events is also interesting because the first took place during US/European trading hours, and the second during Japanese market hours, although in this respect the differences are mainly seen in the effect on the ask-side quotes (rather than the direction of the intervention itself, which fundamentally should have influenced USDJPY bid quotes). As an important note: the exact time of the interventions are not published officially rather deduced from the respective price actions during the intervention days.

A much more recent series of interventions took place in July 2024. For this period, only a monthly aggregated figure is available: 5.5 trillion yen, roughly 32 billion USD (MOF, 2024). Both contemporary financial press reports and the observed exchange rate movements indicate that the first intervention occurred on 11/07/2024 between 12:30 and 13:00 – in fact, during a little over three minutes between 12:42 and 12:45, the USDJPY fell from 161.50 to below 157 (this information does not appear in the half-hourly Refinitiv chart, as shown in Figure 24).

Figure 24: USDJPY spot rate at BOJ intervention - 11/07/2024

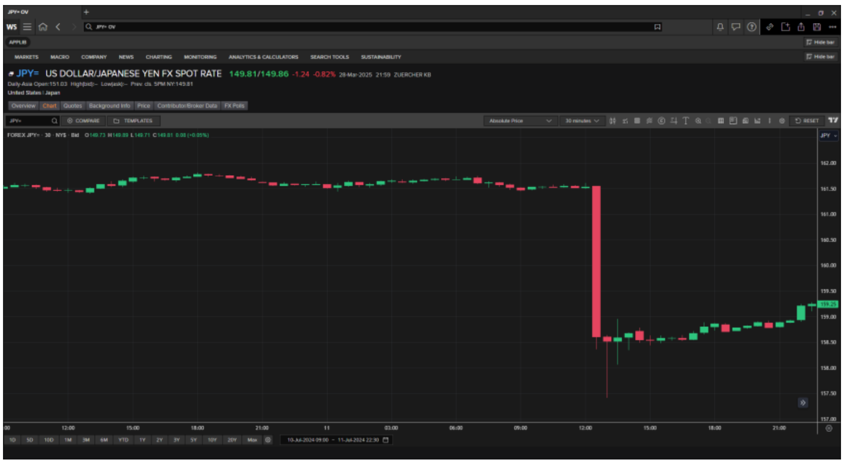

Source: Refinitiv

The "advantage" of the 11 July 2024 intervention is that tick-by-tick data is available from two online sources, meaning that the robustness of the results can at least be verified from this perspective.

Before comparing the BOJ interventions, I selected a period of just over 1.6 hours before 00:36:40 on 24 October 2022 as the pre-intervention period, based on Dukascopy data. The results reveal dynamics very similar to those found in the 2025 sample, with perhaps the most notable difference being that the yen exchange rates have less impact on the EURUSD quotations (Figure B2 in Appendix B illustrates the impulse response functions for the bid quotations).

Based on the actual intervention that occurred a few minutes after 00:37, the following changes compared to the previous period were observed: the effect of the EURUSD does not appear to be significant, and where it is, it has a negative sign and also affects the EURJPY. The impact of the JPY exchange rates in the direction of the EURUSD also decreases – and essentially disappears (Figure B3 in Appendix B shows the impulse response functions of the bid-side quotations).

Interestingly, the intervention on 21 October 2022 lasted about an hour and a half and took place during US trading hours. The IRFs modelling the event represents a blend between a "normal" period and the very short intervention that occurred during Japanese trading hours, as described above, as shown in Figure B4 in Appendix B.

While the EURUSD clearly and significantly influences the USDJPY, only the EURUSD ask-side impacts the EURJPY. The channel between the yen pairs remains strong, but the EURJPY does not affect the EURUSD. The intervention on 11 July 2024 (also during US trading hours), based on the Dukascopy database, supports the previously observed negative or non-significant effect of the EURUSD on both JPY exchange rates, as well as the finding that the JPY exchange rates have no impact on the EURUSD, as shown in Figure B5 in Appendix B.

It is also noticeable for both JPY exchange rates that the impact of the bid-side quotations is greater (significant), meaning the market reacts more strongly to JPY buying pressure during intervention than to JPY sales.

Figure B6 in the Appendix B illustrates that based on the TrueFX database, very similar conclusions can be drawn regarding the EURUSD and EURJPY. At the same time, the dominant role of USDJPY is evident, which aligns with the fact of the intervention – namely, that the Bank of Japan intervenes through USDJPY sales.

We can summarize the above analysis that compared to the symmetrical situation observed during "normal" periods (when each exchange rate affects the other two), different patterns can be seen during interventions. Although the degree varies depending on the period and the sample, the impact of the EURUSD exchange rate on the JPY exchange rates is either negligible or negative during interventions. This is only surprising in the case of EURJPY, where, due to quoting conventions, the effect would normally be expected to be positive. Similarly, in some samples, the impact of the JPY bid-side quotations becomes stronger – which is not surprising – and the effect of the yen exchange rates on the EURUSD either disappears or diminishes significantly. In a certain sense, during interventions, the JPY currency trio shows similarities to the HUF trio, in that the JPY (HUF) exchange rates have a strong influence on each other. It can also be stated that impulse response functions calculated from different data sources sometimes yield different results even for the same periods, which aligns with the findings of the bank questionnaire to be presented in Appendices, namely that quoting behaviour within currency trios is not completely homogeneous.

4.5.4 Impulse-response analysis of GBP and CHF currency trios

Below, I examine the markets for the British pound and the Swiss franc in order to compare them with the currency trios based on the yen and the forint. The samples are based on TrueFX data, covering the period between 10:00 and 19:00 on 4 September 2024. For the Swiss franc, the effect of the EURUSD is very similar to what is observed in the JPY trio during normal periods, significantly influencing both EURCHF and USDCHF, but much more strongly the former, shown in Figure 25.

Figure 25: IRF of EURUSD in CHF Trio - 04/09/2024

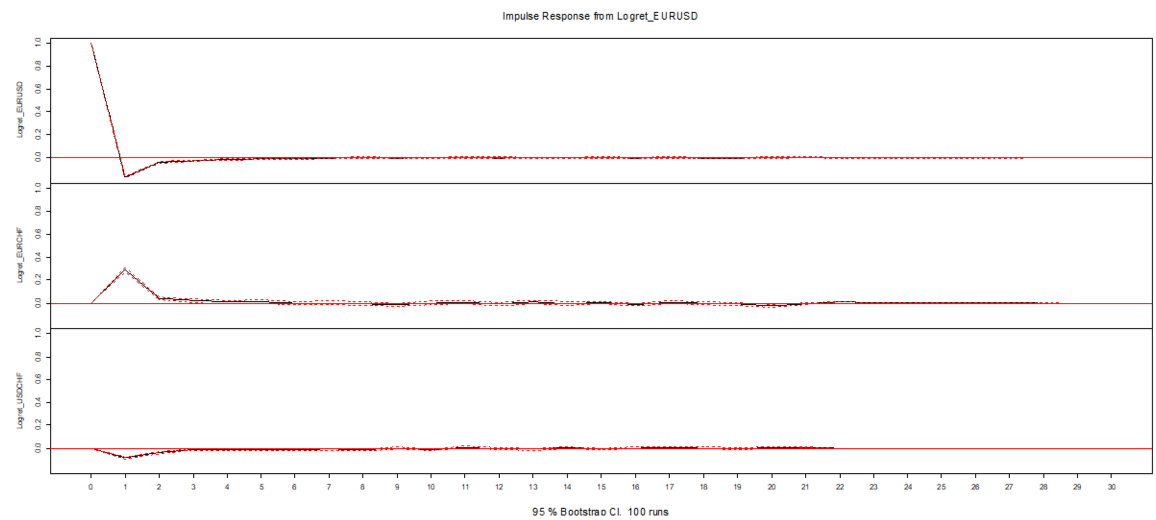

Source: Author’s own work

Figure 26 presents that in contrast to what we can find at the JPY trio, the effect of the EURCHF is barely detectable, while the USDCHF primarily moves the EURCHF.

Figure 26: IRF of EURCHF and USDCHF in CHF Trio - 04/09/2024

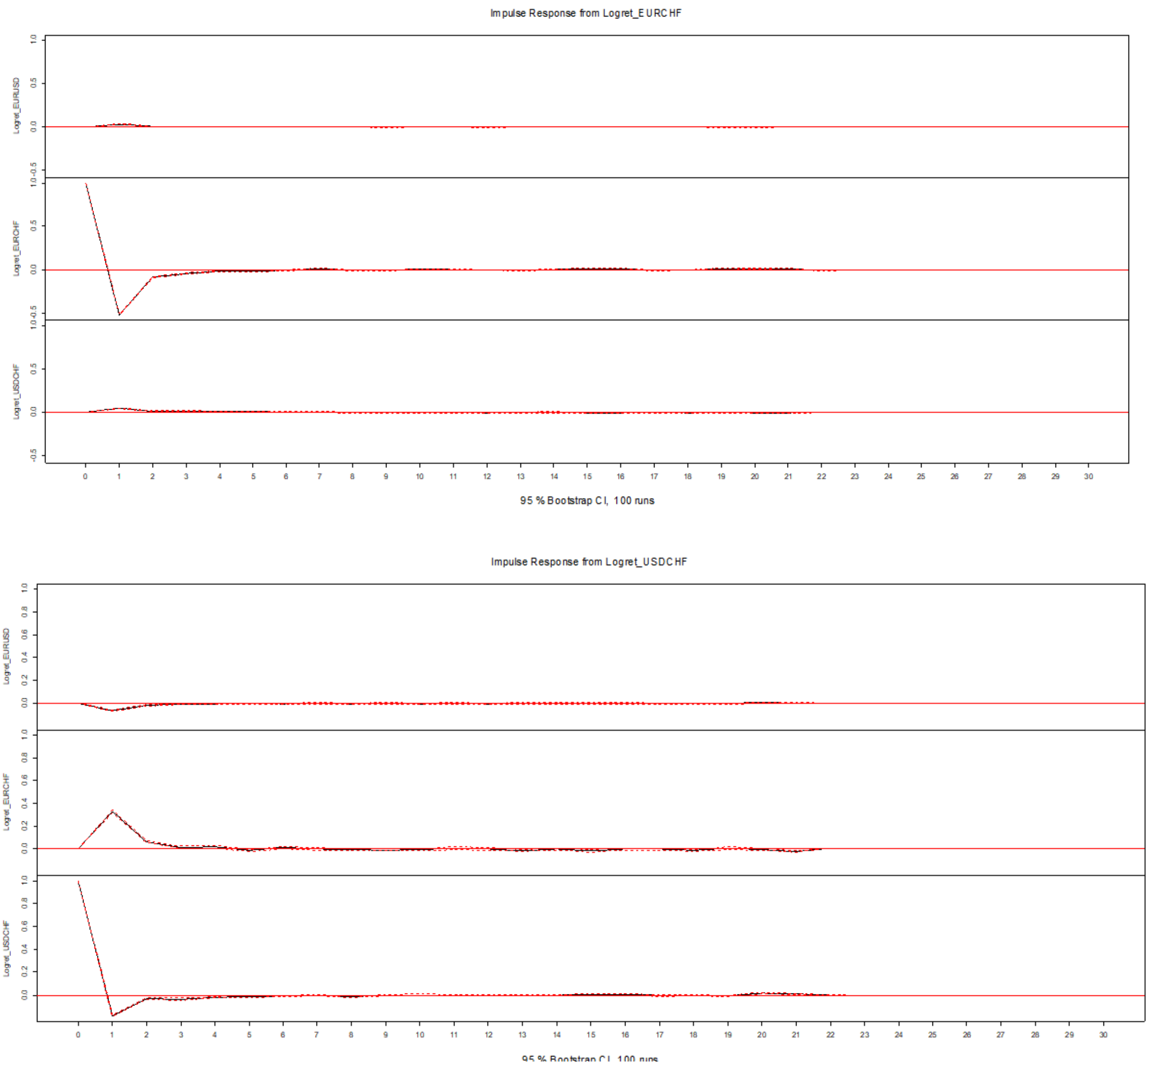

Source: Author's own work

Of course, the limited impulse from the EURCHF in this sample is apparent, but it is easy to imagine a period when there is heavy trading in EURCHF and adjustment takes place via USDCHF.

The British pound shows a similar pattern to the Swiss franc market (due to quoting conventions, the effects on GBPUSD and USDCHF are reversed in sign). We can see in Figure 27, that the EURUSD moves both exchange rates to a similar extent, the impact of EURCHF and EURGBP is rather low, while GBPUSD and USDCHF have a slight effect on the EURUSD and a greater effect on the EURGBP and EURCHF.

Figure 27: IRFs in GBP Trio - 04/09/2024

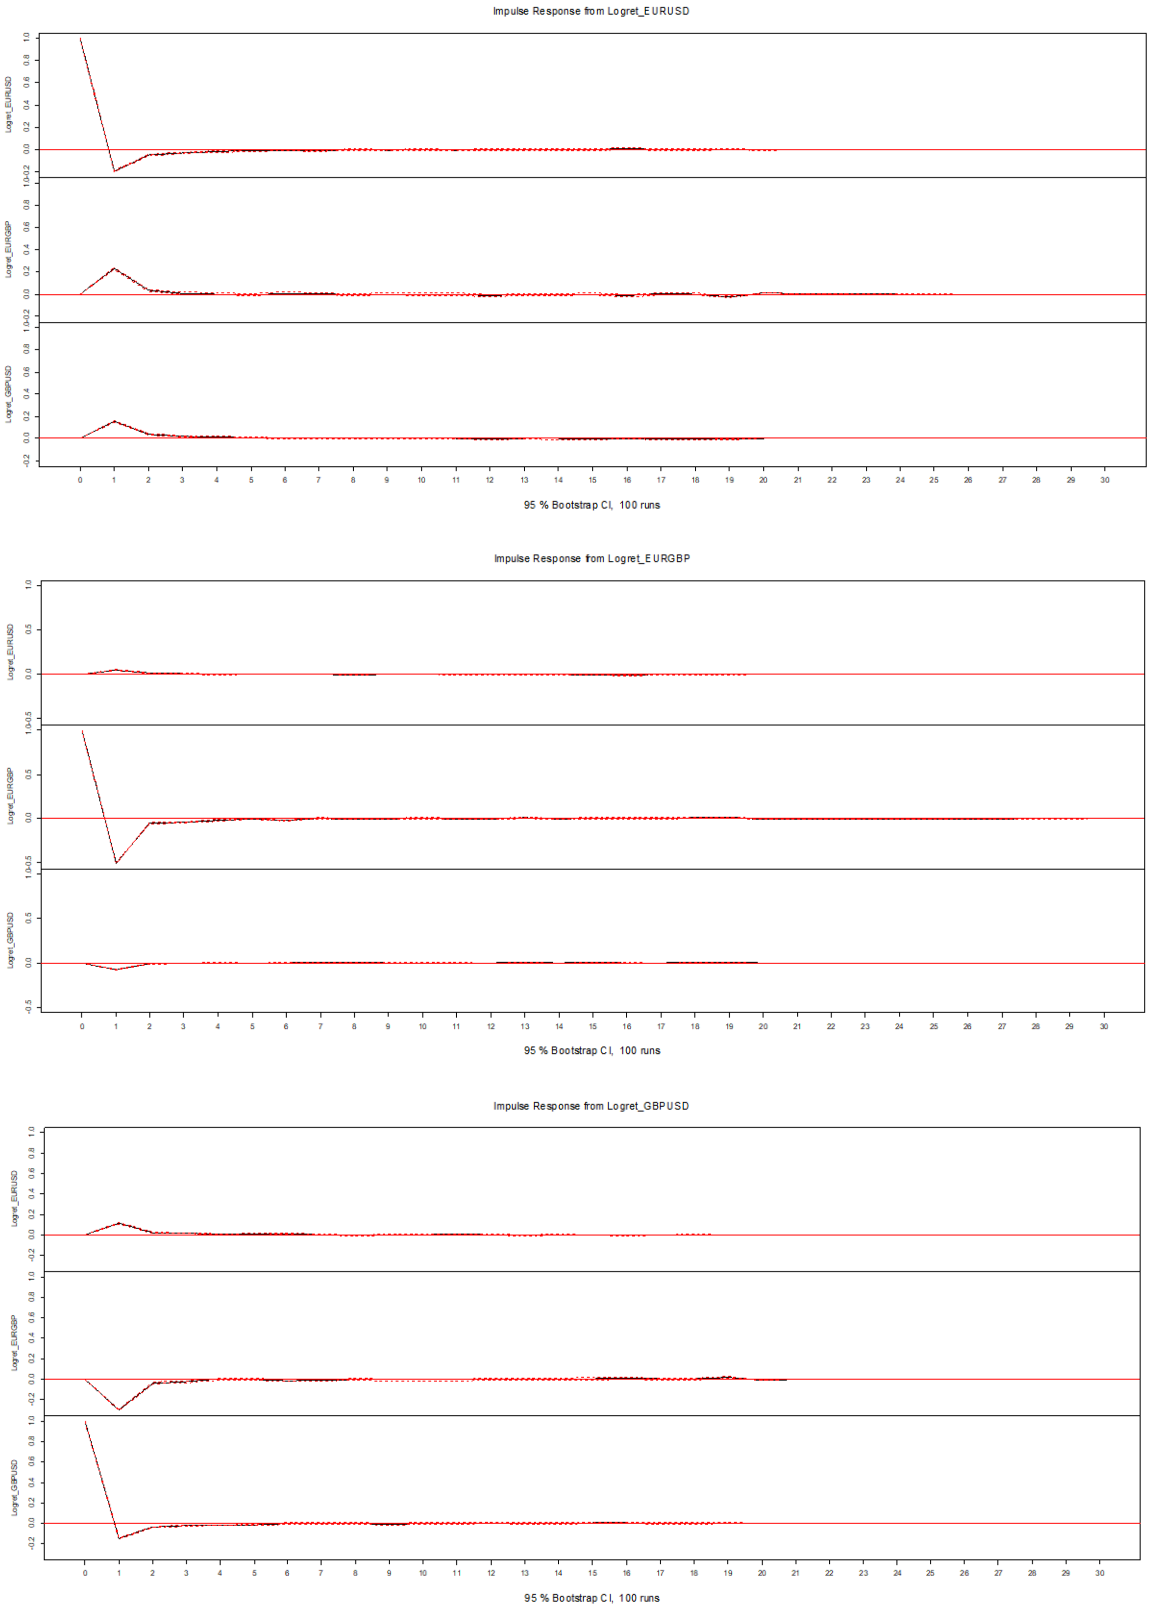

Source: Author's own work

One can conclude from the above analysis that, based on the sample used, the GBP and CHF show a similar pattern to the JPY, with the difference that the significance of the EURCHF and EURGBP appears to be less than the impact of the EURJPY on the other two exchange rates. These findings are further supported by the institutional evidence (see: Appendix C). According to four banks, all three currency rates are quoted independently in the JPY and GBP markets, while according to Bank 2, the EURJPY and either the EURGBP or GBPUSD are quoted based on the other two currency rates (the latter meaning that some market makers quote GBPUSD independently, while others quote EURGBP independently) and Bank 6 considers direct vs indirect EURJPY trading is dependent on market circumstances. On this basis, we can assume that price quoting in the JPY and GBP markets is not necessarily uniform, but the dominant practice appears to be the independent quoting of all three currencies, which are kept in line by arbitrage.

4.5.5 Impulse-response analysis of SEK and NOK currency trios

Given the fact that the trading volume of the Scandinavian currencies – similarly to the CEE region – is more significant against the euro, it seems worthwhile to carry out impulse response analysis for these currencies as well. In line with expectations, based on the sample from 20 January 2025, I find a pattern in the Swedish and Norwegian krona-based currency trios that is very similar to that observed in the forint market. Using the first 100,000 observations taken from more than 500,000 requotes between 08:00 and 22:59 UTC, Figure 28 shows the results regarding the effect of the EURUSD.

Figure 28: IRF of EURUSD in SEK and NOK Trio - 20/01/2025

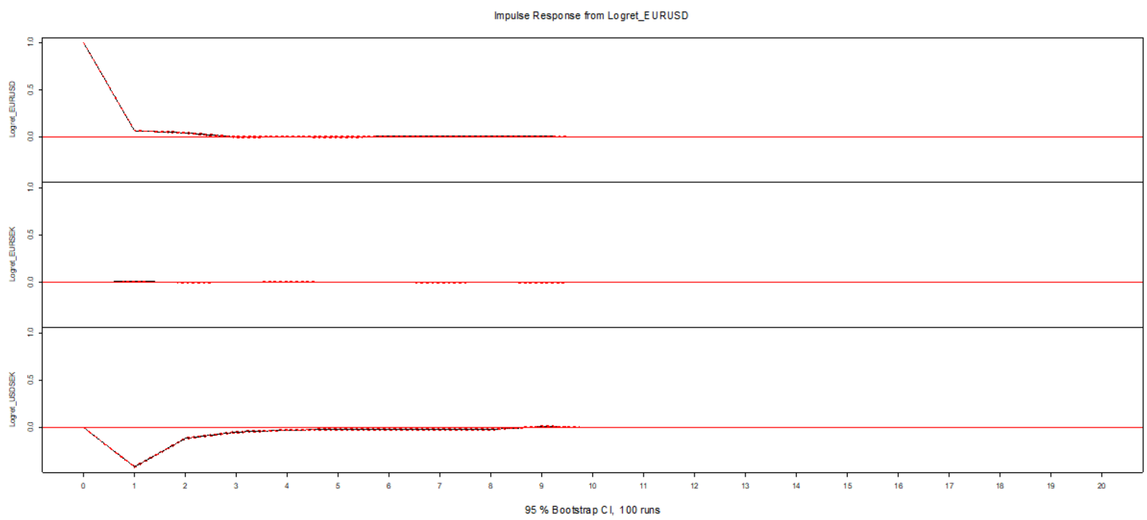

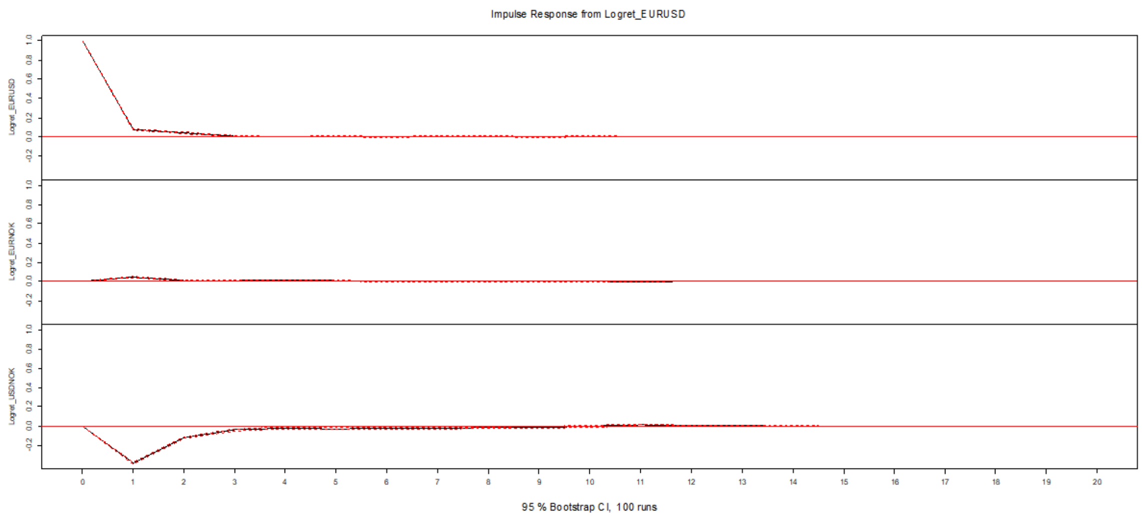

Source: Author’s own work

Similarly to the forint market, the EURUSD essentially only moves the USDSEK and USDNOK exchange rate. The EURSEK and EURNOK exchange rate influences also only the USDSEK and USDNOK rate; thus, the statistical relationship between EURUSD and EURSEK and EURNOK is either non-existent or extremely negligible, an observation which appears to support the near-zero correlation coefficient between EURUSD and EURSEK or EURNOK, see Figure 29.

Figure 29: IRF of EURSEK and EURNOK in SEK and NOK Trio - 20/01/2025

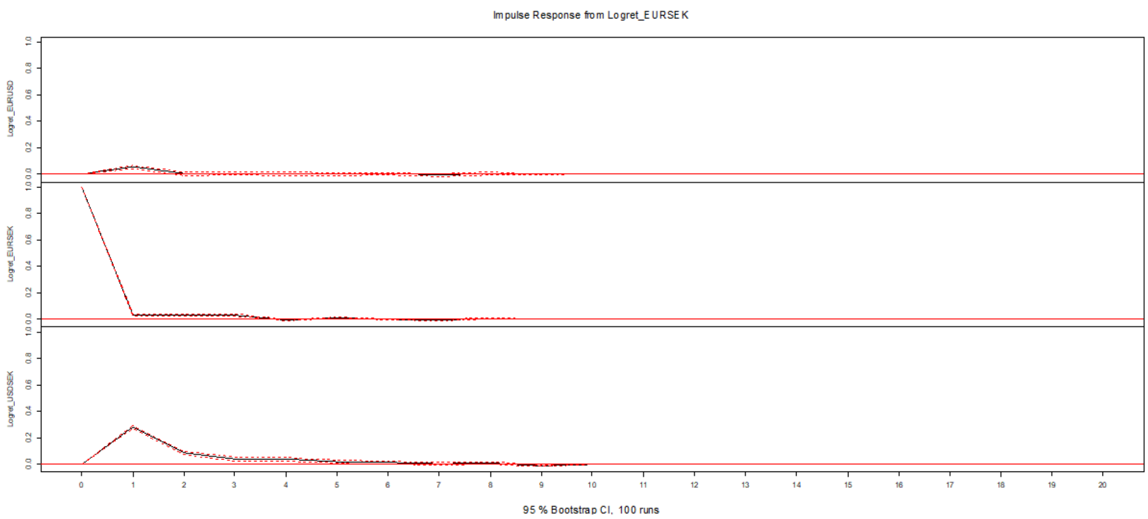

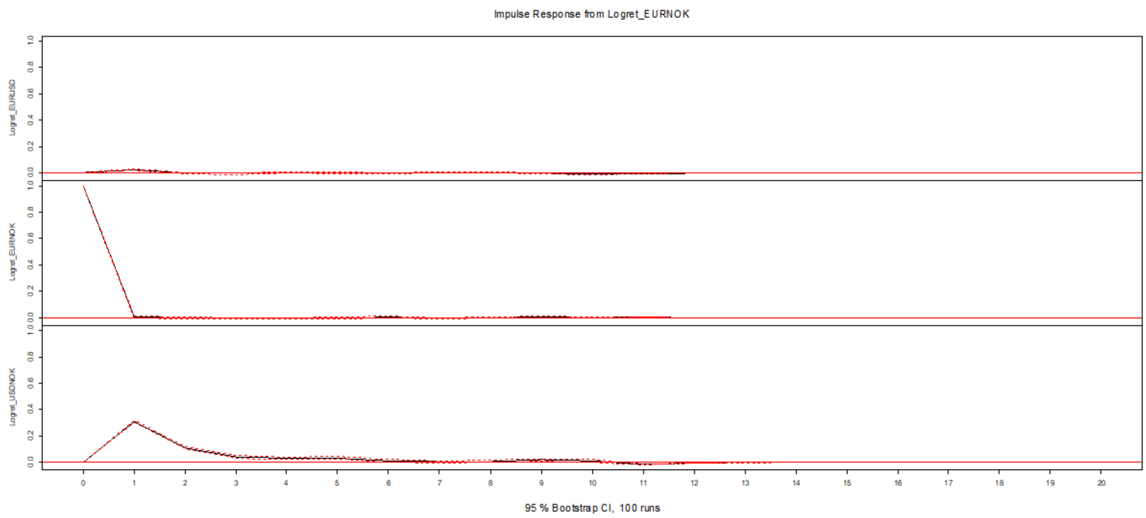

Source: Author's own work

The difference compared to the forint market, however, is that for the USDSEK exchange rate, the influence of the EURUSD appears to be stronger, rather than in the direction of the EURSEK (while for the NOK, the effect is roughly of the same magnitude), as shown in Figure 30.

Figure 30: IRF of USDSEK and USDNOK in SEK and NOK Trio - 20/01/2025

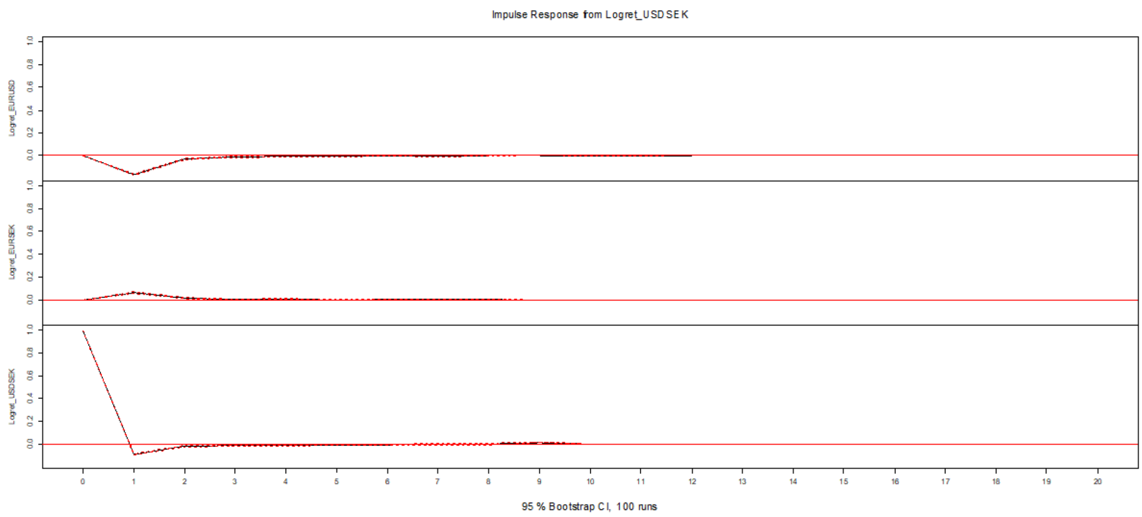

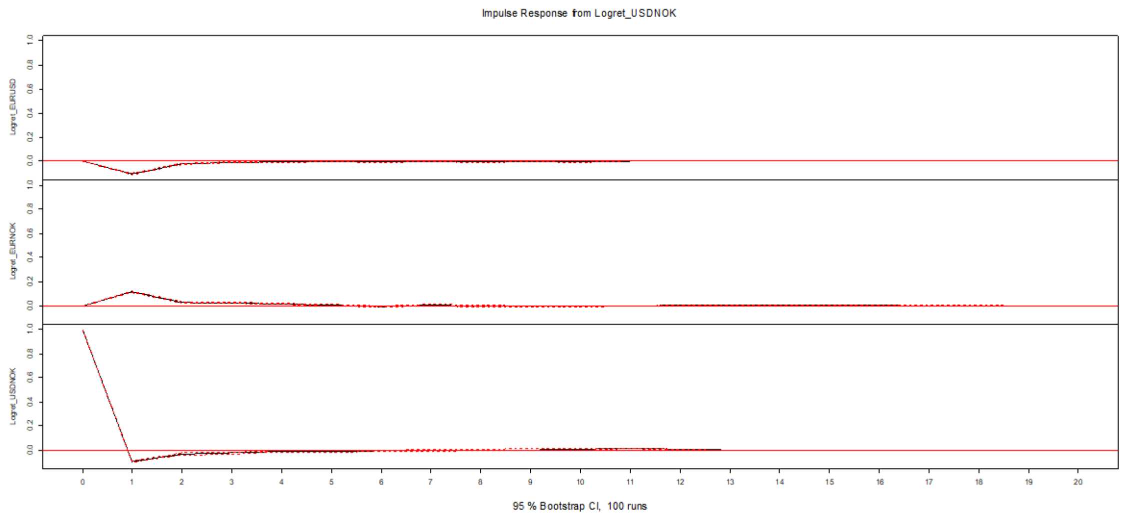

Source: Author's own work

This result is somewhat contradictory in the sense that, based on the relative weight of the two currency markets, it is unlikely that the EURUSD would react to movements in USDSEK or even USDNOK. In contrast, a possible explanation for the anomaly is that if transactional-level movements in the EURUSD are autocorrelated, and the EURUSD influences the USDSEK and USDNOK quotations, then the model captures this as a reciprocal effect; that is, it creates the appearance as if the USDSEK also influences the EURUSD.

Figure 31 presents the autocorrelation function of the EURUSD. It clearly shows a positive correlation, especially for the first 25 lags. The autocorrelation-based explanation is further elaborated in Chapter 4.5.6.

Figure 31: Autocorrelation of EURUSD log returns - 20/01/2025

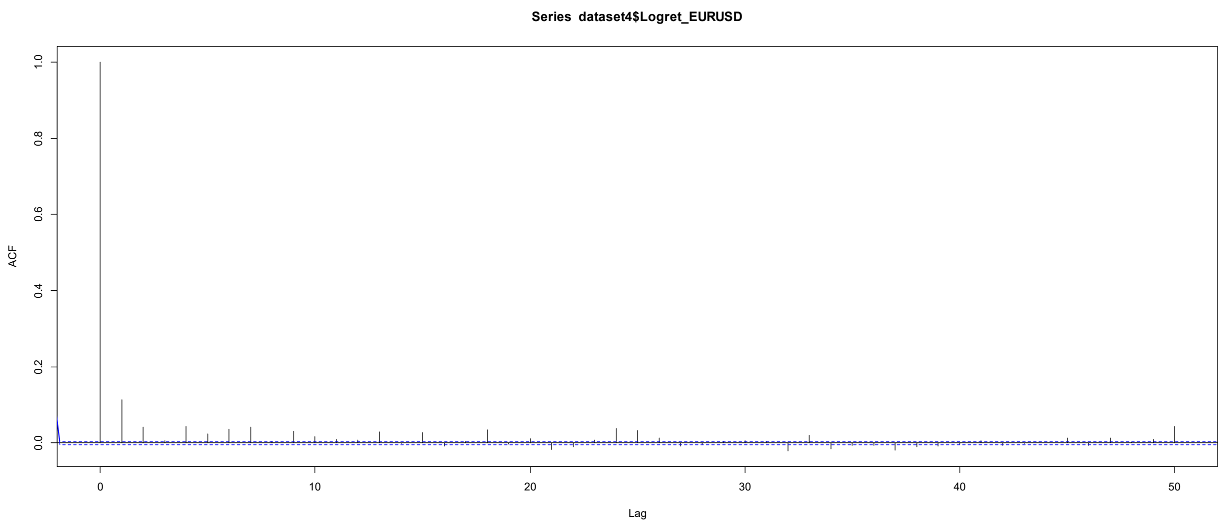

Source: Author's own work

Based on the responses on questionnaire on quoting practices involving six major international banks (see: Appendix C) we can find further evidence why empirical data results the above IRFs for Scandinavian currencies. Half of the responding banks support the dominance of EUR-based quoting, Bank 2 determines whether there is greater liquidity against the EUR or USD depending on the currency, while Bank 4 and Bank 6 basically classifies these currency markets in the same category as JPY and GBP. As seen in Chapter 2.5, the trading volumes of Scandinavian currencies split between EUR and USD vary by instrument and over time; however, the quoting practice revealed by the Norwegian central bank (Norges Bank, 2024) with NOK seems to support quoting USDNOK typically via EURNOK-EURUSD.

Based on the impulse response analysis of currency trios, we can identify certain patterns that distinguish the individual trios. Although, depending on the sample used, there may be differences within a given currency trio, we can still observe characteristic distinctions. One category includes the JPY, GBP, CHF – that is, the major currency trios, in which, under normal market conditions, all exchange rates mutually influence each other. The other main category comprises the HUF, CZK, PLN, NOK, SEK trios, where the key difference compared to the previous category is that the EURUSD does not, or only to a very small extent, affect the EUR-CEE/Scandinavian exchange rate.

Further significant considerations were revealed by the qualitative survey on quoting practices that are relevant to the more nuanced interpretation of the above results. In the foreign exchange market, quoting against the USD is generally the standard, while quoting against the EUR typically represents the cross rate – that is, the rate not quoted independently (according to Bank 2). The method of quoting is significantly determined by the relative weight of each exchange rate in terms of traded volumes; in other words, the more liquid, higher-volume market becomes independent, while the less liquid market is quoted indirectly. On the other hand, the relative weight of volumes can change over time (see: Appendix C: GBP according to Bank 2, CHF according to Bank 3 and G7 currencies according to Bank 6), which can in turn influence quoting practices. Due to increased forex trading volumes, certain exchange rates that were previously quoted synthetically are now quoted independently (Bank 4). If an exchange rate is quoted on one of the major forex trading platforms – such as EBS or Reuters – then independent quoting of that currency is more likely. Bank dealers often do not hedge themselves in the quoted cross rate, but, regardless of liquidity, if they find it advantageous, they hedge via the two other legs (thereby influencing the correlations present in the currency trio). Moreover, this practice is also related to liquidity (see "vehicle currency" in Lyons and Moore (2009)).

4.5.6 Robustness check: Analysis of FX triangle dynamics by local projection method

Table 12 reports the cumulative local-projection coefficients over H = 20 horizons (two seconds at 100 ms frequency). Each entry $\Gamma_{i \leftarrow j}$ measures the total response of return i to a shock in return j. Diagonal elements are normalised to unity. The table therefore provides a complete empirical mapping of directional price discovery within each triad. Each regional triangle contains two independently quoted, deep legs, the global pair EURUSD and the EUR cross (e.g. EURHUF), together with one leg that is synthetically constructed from them, the USD cross (USDHUF = EURHUF/EURUSD). Two features follow. First, because the synthetic leg is by construction a deterministic combination of the other two, its reduced-form innovations are mechanically correlated with theirs; under the non-orthogonal generalised impulse responses used here, a shock to the synthetic USD cross therefore carries a mechanical, non-zero co-movement with EURUSD that is not independent transmission. The diagnostically informative object is instead the feedback from the independently quoted EUR cross into the global leg, $\Gamma_{\text{EURUSD} \leftarrow \text{EUR}}$ cross, with the analogous quantities for the NOK and SEK triangles. Second, as I report point estimates rather than a formal test, I classify reverse channels by an ex-ante magnitude threshold applied uniformly to all six regional off-diagonal entries: with the diagonal own-response normalised to unity, a channel is negligible if $|\Gamma| < 0.05$, attenuated if $0.05 \leq |\Gamma| \leq 0.30$, and strong if $|\Gamma| > 0.30$. A negligible EUR-cross channel, together with economically large forward transmission, is consistent with the synthetic-quotation interpretation and with a reduced number of independent price-discovery channels in regional triangles. This is treated as a reduced-form diagnostic comparison across the empirical implications not as a structural or causal test of the mechanism.

In the benchmark EUR–USD–JPY system, reverse transmission is clearly present. A shock to EURJPY generates a cumulative feedback of $\Gamma_{\text{EURUSD} \leftarrow \text{EURJPY}} = 1.0136$. Similarly, USDJPY exerts economically significant influence on EURUSD. The topology is bidirectional and dense: no leg is informationally isolated. This structure is consistent with independent quotation and integrated price discovery. In contrast, the EUR–USD–HUF system exhibits a structured asymmetry. The global leg transmits strongly into the synthetic USD cross, $\Gamma_{\text{USDHUF} \leftarrow \text{EURUSD}} = -1.3005$, and the synthetic leg in turn shows a sizeable reverse coefficient, $\Gamma_{\text{EURUSD} \leftarrow \text{USDHUF}} = -0.2061$. As noted above, this reverse coefficient is the mechanical image of synthetic construction (USDHUF ≡ EURHUF − EURUSD in logs embeds −EURUSD) rather than independent feedback. The diagnostic channel, the feedback from the independently quoted

EUR cross, is by contrast negligible, $\Gamma_{\text{EURUSD}\leftarrow\text{EURHUF}} = -0.0184$ ($|\Gamma| < 0.05$). The same ordering holds across the regional systems: the EUR-cross channel is negligible

Table 12: Cumulative 3 × 3 Spillover Matrices: Major vs. Regional Trios

Major Trio: EUR–USD–JPY				Regional Trio: EUR–USD–HUF			
Shock →	EURUSD	USDJPY	EURJPY	Shock →	EURUSD	USDHUF	EURHUF
→ EURUSD	1.0000	-0.8901	**1.0136**	→ EURUSD	1.0000	-0.2061	**-0.0184**
→ USDJPY	-0.7499	1.0000	0.1136	→ USDHUF	-1.3005	1.0000	0.4874
→ EURJPY	0.5408	0.2588	1.0000	→ EURHUF	-0.1640	0.3091	1.0000
Satellite: EUR–USD–NOK				**Satellite: EUR–USD–SEK**			
Shock →	EURUSD	USDNOK	EURNOK	Shock →	EURUSD	USDSEK	EURSEK
→ EURUSD	1.0000	-0.4100	**0.0314**	→ EURUSD	1.0000	-0.5734	**-0.0163**
→ USDNOK	-1.1008	1.0000	0.9218	→ USDSEK	-1.1517	1.0000	1.0433
→ EURNOK	0.0314	0.4043	1.0000	→ EURSEK	-0.0163	0.2331	1.0000

Notes: The table reports cumulative coefficients from local projections over 20 horizons (H = 20 at 100 ms frequency). Rows represent the responding variable; columns represent the variable receiving the shock. Diagonal elements are normalized to unity. Because the impulse responses are generalised (non-orthogonal) and cumulated over twenty horizons, a cumulative cross-response can exceed the normalised own-response in absolute value; the coefficients are therefore relative measures of channel strength rather than bounded shares.
Source: Misik and Szabó (2026)

in every case (−0.0184 for HUF, +0.0314 for NOK, −0.0163 for SEK), whereas in the major benchmark both non-global legs feed back strongly into EURUSD (−0.8901 and +1.0136). The regularity is therefore not that a single regional leg is silent in isolation, but that the regional system carries one fewer independent feedback channel than the major system: its two independently quoted legs, EURUSD and the EUR cross, are only weakly coupled, and the third leg is a mechanical combination of them. This is precisely the geometry to be derived in Appendix D, in which the weak correlation falls between the two independent legs while the synthetic leg co-moves strongly with both.

Why reverse spillover is still there: They do not remove dependence generated by common, latent factors, delayed information transmission, asynchronous quote updates, short-lived trends within the current interval.

Suppose there is a short-lived trend: $r_{\text{EURUSD},t-2}, r_{\text{EURUSD},t-1}, r_{\text{EURUSD},t}, r_{\text{EURUSD},t+1}$ are all positively related. Now suppose USDHUF is correlated with the same trend. Even after controlling for p-lags, finite-order autoregressions never perfectly remove persistence. Then $r_{\text{USDHUF},t}$ can still contain information about $r_{\text{EURUSD},t+h}$. The effect is not causal it's simply that both variables are loading on the same short-lived trend. This becomes especially important at very high frequencies where market microstructure generates correlated sequences of quote revisions.

4.6 Alternative microstructural explanations

A natural concern, raised by the structural microstructure literature, is whether mechanisms operating within a fully and independently quoted triangle can reproduce the regional terminal-node topology without invoking synthetic quotation. The two most developed structural representatives are Vayanos and Wang (2012), in which asymmetric information and imperfect competition shape liquidity and price impact, and Babus and Kondor (2018), in which private information diffuses through a dealer network. My point is structural rather than quantitative: information diffusion and liquidity frictions govern when and how much information is transmitted across legs, but they operate within a fixed set of independent price-discovery channels and do not remove a leg's capacity to transmit. In my benchmark, a terminal-node topology arises only when one independent price-discovery channel is mechanically removed, as under synthetic construction. I simulate four data-generating regimes for a triangle in log returns (r1, r2, r3) = (EURUSD, USD cross, EUR cross), linked by the triangular no-arbitrage condition r3 = r1 + r2, with deviations partially or fully corrected according to the regime-specific arbitrage-adjustment speed: (A) symmetric, independent quotation; (B) information diffusion, in which a common shock reaches the global leg first and the regional cross with a lag, all legs independently quoted; (C) liquidity-driven delayed adjustment, in which the USD leg adjusts slowly but remains independently quoted; and (D) synthetic quotation, in which the USD leg is built mechanically as r2 = r3 − r1.

Let $p_{(EURUSD,t)}$, $p_{(USDCCY,t)}$, and $p_{(EURCCY,t)}$ denote the log prices of EURUSD, the USD cross, and the corresponding EUR cross, respectively. Let's define the triangular pricing error as:

$$q_t = p_{EURCCY,t} - p_{EURUSD,t} - p_{USDCCY,t} \tag{47}$$

Alternative (B) is represented by the following information-diffusion model:

$$\Delta p_{i,t} = \beta_i f_{(t-d_i)} + \alpha q_{t-1} + \varepsilon_{i,t} \quad (i = EURUSD, USDCCY, EURCCY) \tag{48}$$

where β_i measures pair-specific exposure to common information f_t, d_i allows information to arrive at different times – representing a pair-specific delay, α is arbitrage adjustment coefficient and $\varepsilon_{i,t}$ is a pair-specific innovation. For example, EURUSD may respond contemporaneously ($d_{EURUSD} = 0$) while information reaches the regional cross with a positive lag. Such a structure readily generates asymmetric leadership, lead-lag relations, and unequal

impulse responses even though all three markets remain potential contributors to price discovery (Evans and Lyons (2002); Rime et al. (2010); Menkhoff et al. (2016); Babus and Kondor (2018)).
The actual equations of the data generating process are as follows:

$$p_{EURUSD,t} = f_t + u_{EURUSD,t} \tag{49}$$

meaning the EURUSD instantaneously to common information. The USD cross simply driven by its own innovations, do not react to common information:

$$p_{USDCCY,t} = u_{USDCCY,t} \tag{50}$$

The EUR cross however reacts by one period delay:

$$p_{EURCCY,t} = \beta_{EURCCY} f_{t-1} + u_{EURCCY,t} \tag{51}$$

The arbitrage error is distributed equally among the three legs:

$$q_{t,i} = \frac{q_t}{3} \quad (i = EURUSD, USDCCY, EURCCY) \tag{52}$$

Considering the modelling range of α -coefficient during modelling (see Table 14), the arbitrage error eliminated at a given step is a function of α.
The distinguishing point is that asymmetric information arrival is not the same as asymmetric capacity to transmit information. Because the regional cross retains its own innovation and remains linked to the other two legs, shocks originating there can still feed back into EURUSD and the USD cross. Information diffusion can therefore generate temporary dominance and asymmetric leadership without mechanically creating a terminal node in the 3 x 3 spillover network
The data generating process in liquidity-driven delayed adjustment model is essentially an arbitrage adjustment process allowing various adjustment speed for the FX rates. This is a two-step approach. In the first step the EURUSD and the EUR cross are modelled as independent innovations, meanwhile the USD cross follows a partial adjustment process. This consist of a target price, a state price equations and an actual quote (denoted as p):

$$Target\, p_{USDCCY,t} \;=\; p_{EURCCY,t} - p_{EURUSD,t} + 0.5u_{USDCCY,t} \tag{53}$$

$$State\, p_{USDCCY,t} \;=\; (1-\lambda)\text{State}\, p_{USDCCY,t-1} + \lambda\text{Target}\, p_{USDCCY,t} \tag{54}$$

$$p_{USDCCY,t} \;=\; State\, p_{USDCCY,t} + 0.4u_{USDCCY,t} \tag{55}$$

$$u_{USDCCY,t} \sim iid(0, \sigma^2)$$

the Target price can considered as a theoretical value driven by the triangular arbitrage and the own innovations of the USD cross. The State price can be considered as a partially adjusted price "moving towards" the Target price and this is adjusted by the noise component when the actual observed quotation price is modelled. The Target- and the actual quotation equations consider 50% and 40% of USD cross related information are considered as fundamentally relevant (Target) and immediately priced in (actual quotation) respectively. The coefficients 0.5 and 0.4 are fixed calibration parameters chosen to represent a moderate transmission of USDCCY-specific innovations through both the target-adjustment and direct-impact channels. Their exact values are not central to the qualitative mechanism of the model. Alternative positive values within a reasonable range would typically preserve the qualitative behaviour of the simulation, although the magnitude of the resulting spillovers and feedback effects would differ. The "State" equation basically states that the USD cross converges to the "Target" price at a rate of λ and driven by its previous quote price by (1- λ).

In the second step all three FX rates corrects the arbitrage mispricing with the same speed (α).

$$\Delta p_t \;=\; \alpha q_{t-1} + \varepsilon_t \tag{56}$$

where $p_t = (p_{(EURUSD,t)}, p_{(USDCCY,t)}, p_{(EURCCY,t)})'$, and ε_t is a vector of FX rate-specific i.i.d error terms.

The key restriction defining the liquidity-delay hypothesis is that every leg remains independently quoted: each component of ε_t has positive innovation variance and can therefore contribute new price information. Thinner liquidity can be represented by a smaller partial adjustment coefficient, for the less liquid leg. This produces slower convergence, weaker contemporaneous responses, and potentially more persistent triangular deviations. This mechanism changes the speed and magnitude of adjustment, but it does not by itself remove an independent price-discovery channel. A shock to one leg creates a triangular pricing error and the subsequent error-correction dynamics transmit part of that disturbance through the other independently quoted legs. Reverse spillovers can therefore be attenuated or shifted to later horizons, but should generally become visible cumulatively as the system returns toward triangular consistency.

For the modelling purposes regime (A) is set up based on Ciacci et al (2020) and its modified version for regime (D). I consider three exchange rates as for the previous models, each represented by a mid-price process. Two types of agents operate in the market:

Market makers update quotes using a simple trend-following-plus-noise rule,

$$p_{i,t+1} = p_{i,t} + \theta_i \, \Delta p_{i,t} + \epsilon_{i,t} \tag{57}$$

where θ_i reflects the degree of trend responsiveness and $\varepsilon_{i,t}$ is an i.i.d. noise term.

Arbitrageurs enforce the triangular identity by monitoring the misalignment measure:

$$q_t = p_{EURCCY,t} - p_{EURUSD,t} - p_{USDCCY,t} \tag{58}$$

They apply corrective trades whenever the implied cross-rate deviation $|q_t|$ exceeds a small threshold that proxies for bid–ask frictions. The adjustment affects all three legs proportionally to q_t, moving the system back toward arbitrage consistency. This structure captures only the two essential forces relevant for the paper: (i) independent price discovery in actively traded pairs, and (ii) triangular arbitrage alignment across all three legs. Table 13 presents the modelling parameters for the synthetic model (D) where additionally a synthetic pricing noise is included: $p_{USDCCY} = p_{EURCCY} - p_{EURUSD} + \eta_t$.

Table 13: Simulation parameters of base model

Parameter	Value
Trend-following coefficient θ_i	0.15
Noise variance $\mathrm{Var}(\varepsilon_{i,t})$	10^{-5}
Arbitrage threshold	10^{-4}
Synthetic pricing noise η_t	$N(0, 5 \times 10^{-6})$
Simulation length	200,000 iterations

Source: Misik and Szabó (2026)

Table 14 summarizes the parameter settings of the competing alternatives.

Table 14: Simulation parameters of competing models

Alternative	Varied parameters	Parameter values	number of scenarios	Resulting reverse spillover range min	max
Symmetric independent	None	–	1	0.5059	0.5059
Information diffusion	Common information loading of EUR cross – β(EURCCY); arbitrage adjustment speed applied to the whole triangle – α	β ∈ {0.2, 0.4, 0.6, 0.8, 1.0}; α ∈ {0.5, 0.7, 0.9, 1.0}	20	0.4736	0.9221
Liquidity delay	Partial adjustment speed of USD cross – λ (USDCCY); arbitrage speed fixed at α =0.5	λ ∈ {0.10, 0.20, 0.35, 0.50, 0.70, 0.90}	6	0.0098	0.2117
Synthetic quotation	standard deviation of synthetic pricing noise - η	sd(η) ∈ {1e−4, 5e−4, 1e−3, 5e−3, 1e−2}	5	0.0048	0.0148

Source: Author's own work

For each regime, we estimate a VAR(5) by OLS and compute generalized impulse responses following Pesaran and Shin (1998). We report the cumulative reverse spillover Γ(r1⟵r3) over 20 horizons, with the own-response normalised to unity, exactly as in Table 12.

Table 15: Simulated regimes: correlation geometry and cumulative spillovers

Regime	$\rho_{12}, \rho_{13}, \rho_{23}$	$\Gamma_{r1 \leftarrow r3}$ (reverse)	$\Gamma_{r2 \leftarrow r1}$ (forward)
Symmetric (independent)	−0.50, +0.50, +0.50	+0.499	−0.496
Info. diffusion (independent)	−0.56, +0.65, +0.24	+0.802	−0.224
Liquidity delay (independent)	−0.54, +0.22, +0.54	+0.104	−0.936
Synthetic quotation	−0.70, +0.00, +0.71	+0.008	−0.987

Source: Misik and Szabó (2026)

Table 15 summarise the results. In these simulations, the synthetic quotation – or an economically close coefficient setting for liquidity delay – reproduces the regional "one weak, two strong" geometry and collapses the reverse spillover to zero. In the three independently quoted regimes the reverse spillover from the regional cross back to the global leg remains economically meaningful, ranging from about +0.10 under pronounced liquidity delay to +0.80 under information diffusion, even though the latter generates a strongly asymmetric lead – lag pattern. The synthetic regime collapses the reverse spillover to +0.008, and it is the only regime that reproduces the regional geometry (ρ12, ρ13, ρ23) ≈ (−0.70, 0.00, +0.71) observed in the data and derived in Appendix D. Information diffusion and liquidity delay alter the magnitude and timing of feedback but, consistent with their treatment as empirical implications, do not produce a terminal node.

Information diffusion models (B) readily allow for asymmetric leadership and uneven information diffusion across markets (Evans and Lyons (2002); Rime et al. (2010); Menkhoff et al. (2016); Babus and Kondor (2018)), but under independent quotation they do not mechanically imply the elimination of a feedback channel. In major triangles, feedback is large and immediate. In regional triangles, the independently quoted EUR-cross channel is economically negligible. This pattern is therefore more consistent with an effective reduction in the number of independent price-discovery channels than with temporary informational dominance.

If thinner liquidity (C) merely slowed price adjustment, reverse spillovers should emerge cumulatively over the two-second window on the empirical sample. The cumulative topology shows that they do not. Illiquidity may affect the timing and magnitude of price adjustment, but does not by itself imply the disappearance of an independently quoted market's capacity to transmit information (Payne (2003); Vayanos and Wang (2012)). The absence of economically meaningful cumulative reverse feedback is therefore difficult to reconcile with delayed adjustment alone. Across the parameter grid, the focal reverse coefficient ranges from 0.010 to 0.212. It enters the negligible region only when the delayed quote becomes tightly slaved to its arbitrage-consistent value. At that boundary, the model is no longer an economically distinct slow independent market; it is approaching a near-mechanical quote, which is observationally close to synthetic construction.

A further independent-quotation alternative deserves attention. If the USD cross dominates trading, a symmetric, Ciacci-style model (see: Chapter 5) can in principle generate a one-weak–two-strong correlation pattern without any synthetic construction, for example through a strong USD-buying or USD-selling trend. I therefore do not interpret the correlation geometry alone as evidence of synthetic quotation. The distinction lies in the accompanying market-structure implications: independent USD dominance should be associated with (i) materially larger USD-cross than EUR-cross trading volume, (ii) independent innovations in the USD cross, and (iii) non-negligible reverse spillovers from the USD cross into EURUSD or the EUR cross.

Figure 32: Generalised IRF of the global leg to a shock in the regional cross

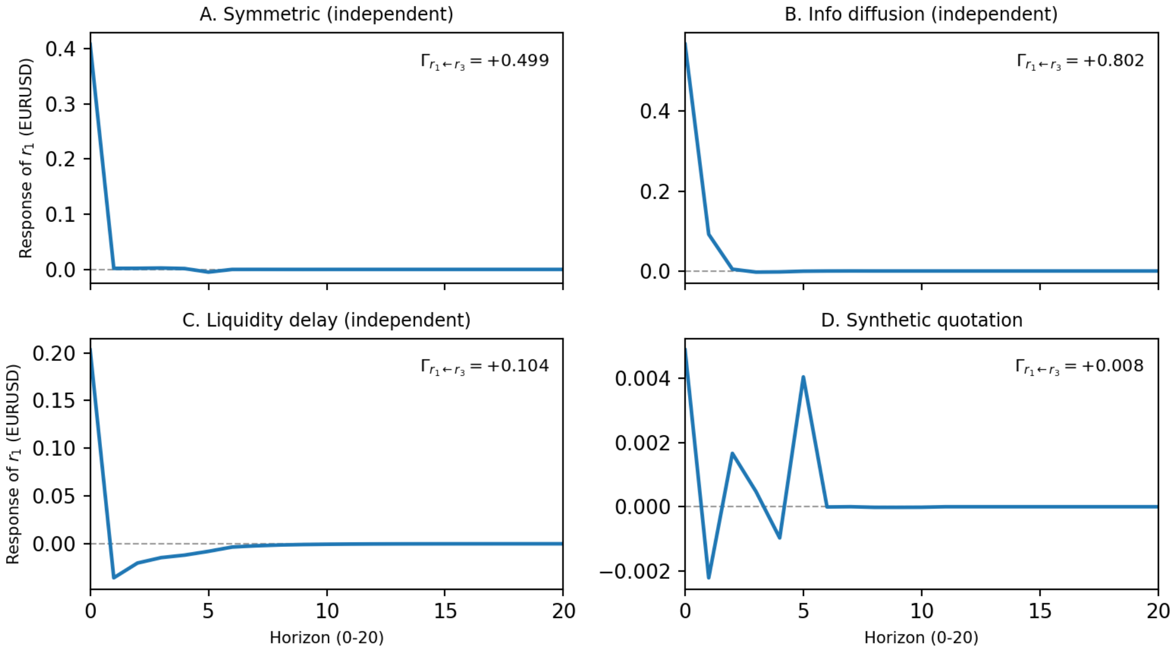

Source: Misik and Szabó (2026)

These features are plausible for major systems such as the USDJPY triangle, but much less so for USDHUF, USDCZK and USDPLN or for the regional Scandinavian systems. In the regional sample, the institutional evidence (Appendix C), the trading-share evidence (Chapter 2.5), and the weak reverse-spillover topology (Chapter 4.6) do not support USD-cross dominance as the operative mechanism.

4.7 Discussion and conclusions

This chapter examines the microstructural origins of exchange-rate correlations and seeks to explain why some currency triangles behave as highly integrated systems while others display persistent asymmetries in the transmission of shocks. The study is motivated by a fundamental question in foreign exchange market research: are the correlations observed between exchange rates the result of genuine bilateral information transmission, or are they largely a mechanical consequence of the quotation architecture imposed by triangular arbitrage? If some exchange rates are quoted synthetically rather than independently, the observed co-movements may reflect market design rather than economic information.

The chapter positions itself within the foreign exchange microstructure literature, which traditionally explains exchange-rate volatility and co-movements through liquidity provision, order flow, trading intensity, and information asymmetries. Although previous studies have documented asymmetric information transmission and heterogeneous quoting practices, relatively little attention has been paid to the possibility that the quotation architecture of a currency triangle itself may create systematic and lasting correlation patterns. If one leg is mechanically derived from the other two, its role within the network of spillovers should differ markedly from that of independently quoted exchange rates.

To investigate this issue, Chapter 4 employs ultra-high-frequency foreign exchange data obtained from Dukascopy and TrueFX. Because transaction-level quotes occur at irregular intervals, aggregation of observations to a uniform frequency of 100-200 milliseconds was needed. This approach preserves nearly all relevant information while making conventional time-series methods applicable. The empirical analysis combines vector autoregressive models, impulse-response functions, and local projection techniques. Together, these methods make it possible to trace the direction, magnitude, and persistence of spillovers among the exchange rates composing a currency triangle. Rather than examining a single lead-lag relationship, the chapter maps the complete three-by-three network of interactions within each system.

The empirical results reveal a clear distinction between major and regional currency triangles. In the benchmark EUR-USD-JPY system, shocks affecting any of the three exchange rates generate responses in the other two. Although the magnitude of the effects varies over time and across samples, the overall structure is characterised by dense, bidirectional spillovers and multiple independent channels of price discovery. Similar patterns are observed in the GBP and CHF currency triangles. These findings are also broadly consistent with institutional evidence from market participants, who report that the relevant exchange rates are usually quoted independently and aligned through arbitrage.

A fundamentally different pattern emerges in regional markets. In the EUR-USD-HUF triangle, as well as in the SEK- and NOK-based systems, the author documents a persistent one-weak-two-strong topology. The EURUSD exchange rate strongly influences the USD cross, while the independently quoted EUR cross primarily affects only the USD cross. Most importantly, reverse feedback from the regional EUR cross into the global EURUSD pair is either statistically insignificant or economically negligible. This result suggests that one leg of the triangle does not contribute an independent source of information and instead behaves as a mechanically derived exchange rate. Among the mechanisms examined, synthetic quotation is the explanation most consistent with the combined correlation, spillover, trading-volume, and practitioner evidence.

The chapter strengthens this interpretation through a series of robustness exercises. The results remain stable across different data providers, alternative samples, and estimation methodologies. Special attention is given to periods of Bank of Japan intervention, during which normal market dynamics are temporarily disrupted. While intervention episodes alter the intensity and direction of some spillovers, they do not overturn the broader conclusion that major currency systems remain fundamentally more symmetrical than regional ones. Local projection estimates confirm the same qualitative patterns obtained from vector autoregressions, further strengthening confidence in the findings.

H3 is supported by the available evidence, as the results demonstrate clear differences in the triangular dynamics of major and regional currency trios.

A major contribution of the chapter is the comparison between the synthetic quotation hypothesis and competing explanations from the microstructure literature. Considering whether information diffusion, liquidity frictions, or asymmetric market leadership could independently generate the observed correlation structures, simulation exercises demonstrate that these mechanisms may create lead-lag relationships and asymmetric responses, but they do not eliminate feedback channels between independently quoted exchange rates. Additionally, the "USD dominance" as another potential candidate is not supported based on the trading volume/market liquidity evidence. Out of the examined approaches only the synthetic quotation and the economically close liquidity delay with high partial adjustment coefficient frameworks are capable of reproducing both the near-zero reverse spillovers and the distinctive correlation geometry observed in regional currency triangles – **supporting the consistency of synthetic quotation with the observed spillover topology claimed by H3**.

The contribution of the chapter to the literature is therefore multifaceted. First, it provides a detailed empirical mapping of directional spillovers across a variety of currency triangles. Second, it introduces quotation architecture as an important determinant of exchange-rate correlation patterns. Third, it combines high-frequency econometric evidence with institutional information from market practitioners, thereby linking observable market dynamics to actual quoting practices. Finally, the chapter bridges the gap between foreign exchange microstructure and the broader literature on correlation regimes by demonstrating how differences in price-discovery mechanisms can shape long-run correlation structures.

The findings also have important practical implications. If correlations are partly generated by quotation practices rather than information transmission, conventional measures of financial integration and contagion may be misinterpreted. Risk managers, central banks, and portfolio managers should therefore be cautious when using correlation-based indicators in regional foreign exchange markets. Models that fail to account for synthetic quotation may overestimate the degree of interconnectedness between currencies and misjudge the transmission of shocks across markets.

Overall, the chapter concludes that quotation architecture represents a fundamental microstructural source of exchange-rate correlations. Major currency triangles are characterised by multiple independent price-discovery channels and balanced feedback mechanisms, whereas regional currency triangles display hierarchical structures consistent with synthetic quotation. By identifying this mechanism and documenting its empirical consequences, the chapter offers a novel explanation for the distinct correlation regimes observed across foreign exchange markets.

5. THE MICROSTRUCTURAL MODEL OF FOREIGN EXCHANGE MARKET CORRELATIONS

5.1 Introduction and research hypotheses

Considering the literature of correlation in finance, vast majority of the work treats correlation as an exogenous variable applied to models for various objectives – portfolio optimization, risk modelling etc.. Time-series modelling of correlation uses historical data and models certain properties of the observed correlation. In contrast any modelling that results certain outcomes for correlation between various assets by modelling certain underlying actions of market participants might be feasible to understand the underlying relationship and has the potential to explain the resulting correlation. By such model one could potentially improve by modelling the interaction between the relevant agents and therefore apply more realistic assumptions compared to a "naïve" time-series based approach.

The work of Ciacci et al. (2020) highlights the fact that the development of correlations in the foreign exchange market can be significantly influenced by the arbitrage mechanisms in the forex market and by the trend-following behaviour of market makers and lays the foundations for a possible model to explain these correlations. The Arbitrageur Model – ABM model, applied by the authors (see its introduction later) – surprisingly accurately replicates the levels of correlation observed in reality for the EUR-USD-JPY currency trio. This part of the thesis aims to demonstrate, by modifying the parameters and certain assumptions of the ABM model, that the model can provide an explanation for the correlation differences observed in specific markets (JPY, regional currencies, USD-driven currencies).
The ABM model also has limitations that inspires the fourth hypothesis of the thesis:

H4: Under empirically plausible quotation and market-maker parameters, the modified ABM reproduces the ordering and magnitudes, but not all signs, of regional FX-triangle correlations.

5.2 Related literature

Triangular arbitrage in the foreign exchange market means exploiting price differences between three currencies in order to achieve risk-free profits. BIS (2011) distinguishes between four types of arbitrage according to high frequency trading (HFT) strategies, of which "classic" arbitrage corresponds to the subject of this paper, i.e., triangular arbitrage.

Aiba et al. (2002) and Aiba and Hatano (2006) establishes a stochastic model based on triangular arbitrage in the foreign exchange market, where the activities of arbitrageurs in the three foreign exchange markets form a centre of gravity, which represents the arbitrage-free exchange rate triangle and acts as a kind of dynamic constraint on the three exchange rates. According to their theoretical model, triangular arbitrage can influence the behaviour of exchange rates, which can lead to negative autocorrelation in the short term.

One of the first studies on this topic Marshall et al. (2007), which goes beyond the Japanese yen, examines three currency trios: the EUR and USD, as well as the GBP, CHF, and JPY. According to their findings, although bid-ask spreads are wider during periods of illiquidity, and arbitrage opportunities are therefore less frequent, arbitrage profits are higher (in line with the Grossman-Stiglitz paradox).

Lyons and Moore, (2009) make several important observations. They introduce the concept of "vehicle currency," i.e., the possibility of trading a currency pair through the intermediation of a third currency. This so-called indirect trading also carries different types of information about the currency market, i.e., direct and indirect trading have different effects on the equilibrium exchange rate of the same currency pair. The authors also offer a possible explanation as to why certain exchange rates are not traded directly: if the information asymmetry is too great in a market, the players will not be sufficiently motivated to trade in it, and the market will effectively cease to exist.

Ito et al. (2012) Examining the period between 1999 and 2010, examining the EUR-USD-JPY currency trio, found that in 1999, the probability of triangular arbitrage opportunities persisting beyond one second was still 50%, whereas by 2009, this probability had fallen to just 10%. The reasons for this can be found in changes in market microstructure, with more and more market participants connecting directly to the foreign exchange market or through so-called prime brokers.

Kozhan and Tham (2012) argue that a large number of arbitrageurs does not necessarily help the market to clear and price efficiently, as intense competition for arbitrage opportunities greatly increases execution risk, which hinders arbitrage activity.

Based on the results mentioned above Chaboud, et al. (2014) show that the increase in algorithmic trading in the EUR-USD-JPY currency trio not only reduces the frequency of arbitrage opportunities, but also the autocorrelation of high-frequency returns, the latter being a kind of measure of market inefficiency. While in the case of triangular arbitrage, algorithmic traders take their share by trading at quoted prices, they contribute to the efficiency of individual markets through frequent price updates, in line with the findings of Hoffmann (2014).

In the foreign exchange market, arbitrage opportunities are typically very short-lived and mostly small in scale. These opportunities tend to occur during periods of high market liquidity, such as peak trading times, typically lasting less than 1 second and amounting to less than 1 basis point (Fenn, et al., 2009). However, according to Foucault et al. (2017) the presence of arbitrageurs reduces market liquidity by widening bid-ask spreads. There is agreement that arbitrage opportunities disappear quickly.

Gebarowski et al. (2019) uses multifractal trend-free correlation analysis to identify arbitrage opportunities. Based on their findings, the probability of arbitrage opportunities increases in the context of significant events (such as central bank interventions or Brexit), and correlations decrease over the shortest time periods.

Cui et al. (2020) formalize currency market arbitrage opportunities using a matrix of exchange rates, which is used in a computational framework for real-time detection and identification, taking transaction costs into account. Building on this approach, Maldonado et al. (2021) identify currency matrices as a useful analytical framework for understanding the dynamics of currency bubbles relative to fundamental exchange rates interpreted on the basis of purchasing power parity.

Ciacci et al. (2020) model the EUR-USD-JPY currency trio based on triangular arbitrage, according to which the correlation between currency pairs converges to the correlation levels observed in reality. This study forms an important basis for the doctoral dissertation, so the model is explained in detail in Chapter 5.3.

Wu et al. (2023) model the EURJPY exchange rate directly based on minute-by-minute closing rates (direct quotation) and indirectly via EURUSD-USDJPY (indirect quotation). They find that indirect quotation adapts to direct quotation to a lesser extent than vice versa, i.e., indirect quotation has more significant information content regarding the equilibrium exchange rate.

The literature on currency market arbitrage therefore focuses primarily on the EUR USD JPY currency trio, typically using data with a frequency of at least one second, draws conclusions based on market microstructure (market participants with different motivations), and examines "normal" and "exceptional" market events, such as central bank interventions, to a significant extent.

5.3 The benchmark agent-based model

The agent-based model presented here explains short-term correlations in the foreign exchange market through the behaviour of market makers and arbitrageurs. In the model, on the three foreign exchange markets formed from the three currencies, market makers exhibit the specific market-making behaviour characteristic of each market, one defining feature of which is short-term trend-following.

5.3.1 Market makers in the Arbitrageur Model

The i-th market maker on the ℓ-th market quotes at the $z_{i,\ell}$ mid-rate with a constant bid-ask spread (L_ℓ) – that is, all market makers apply the same spread for the given currency exchange rate. This assumption is consistent with the premise that the spread is the most important transaction cost in the foreign exchange market, meaning that significantly different spreads would result in considerable competitive advantages or disadvantages for the various participants. As a side note, the assumption of a constant spread is not generally realistic in the foreign exchange market, as in periods of high volatility market makers react by widening the spread to avoid unwanted losses. The dynamics of the quoted mid-rates are described by the following stochastic differential equation (Ciacci et al. (2020) p. 7):

$$\frac{dz_{i,\ell}(t)}{dt} = c_\ell \emptyset_{n,\ell}(t) + \sigma_\ell \epsilon_{i,\ell}(t) \qquad i = 1, \dots, N_l \tag{59}$$

where N_ℓ is the number of market makers on the ℓ-th market, σ_ℓ is the volatility factor, which represents the effect of random shocks on quoting (for example, order flows arriving in the given exchange rate), and $\epsilon_{i,\ell}$ is standard normally distributed white noise. According to the c_ℓ coefficient, market makers are sensitive to the short-term trend component $\emptyset_{n,\ell}$, which gives the exponentially weighted average of the last n exchange rate changes:

$$\emptyset_{n,\ell}(t) = \frac{\sum_{k=0}^{n-1}(p_\ell(g_{t,\ell}-k)-p_\ell(g_{t,\ell}-k-1))e^{-k/\xi}}{\sum_{k=0}^{n-1} e^{-k/\xi}} \tag{60}$$

where $g_{t,\ell}$ is the number of transactions executed in the interval [0,t], while ξ is the exponential weighting factor's scale parameter. According to the model of Ciacci et al. (2020), every market maker on a given market responds to both the trend component and the random market maker component to the same extent – that only comes into play when there is no price adjustment triggered by a transaction (see later). The c_ℓ trend coefficient can also take a negative value, which means that market makers on that market adjust their quotes against the short-term trend – thus following a so-called contrarian strategy. The market-making trend-following/contrarian strategy also appears in previous market-maker models, such as the Dealer model of Yamada et al. (2009), or the HFT model of Kanazawa et al. (2018), in which an explicitly trend-following market-making strategy reproduces the empirically observable market-maker behaviour.

It is also important to note that equation (59) describes the case of market maker activity where the market maker has not transacted but only responds to price movements observed in the market. If the quotes of two market makers enable a transaction – namely, if one market maker's bid price is higher than or equal to another's ask price – the model assumes that the transaction is executed at the average price of the two, and both market makers adjust their own mid-prices to this transaction price level. This mechanism is illustrated on Figure 33.

Figure 33: Limit order book – quotation realignment

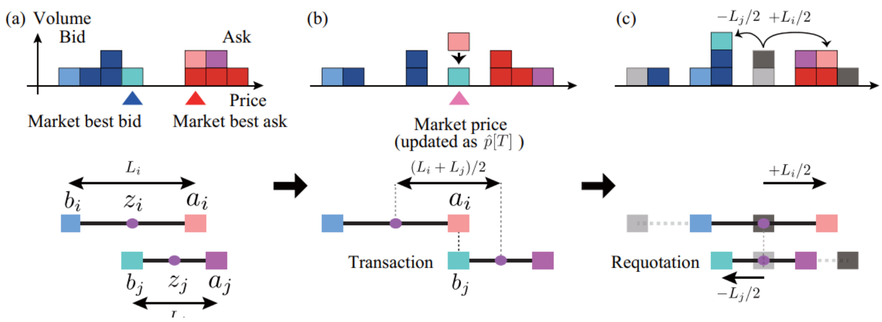

source: Kanazawa et al. (2018), p. 8

The figure illustrates the scenario where, moving from state "a" to state "b", the sell intention (a_i) of the i-th market maker meets the buy intention (b_j) of the j-th market maker. As a result, the i-th market maker adjusts their mid-quote upwards, while the j-th market maker adjusts theirs downwards, both precisely to the transaction price level ("c" state).

5.3.2 The role of the arbitrageurs

The model therefore assumes three independent market-making markets, which are interconnected by the activities of arbitrageurs. The arbitrageur is constantly seeking opportunities and by executing a buy/sell transaction in one currency pair and simultaneously entering into synthetic sell/buy transactions in the other two currency pairs, they can secure a positive profit. Using the example of the EUR-JPY-USD currency trio: if the highest EURJPY ask (sell intention) is lower than the product of the lowest USDJPY and EURUSD bid (buy intention), it is advantageous to buy EURJPY and sell USDJPY and EURUSD (Type 1 arbitrage). The relationship is reversed in the case of the EURJPY bid and the USDJPY, EURUSD ask (Type 2 arbitrage):

$$a_{EURJPY}(t) < b_{USDJPY}(t) \times b_{EURUSD}(t)$$
$$b_{EURJPY}(t) > a_{USDJPY}(t) \times a_{EURUSD}(t) \tag{61}$$

If the arbitrageur detects an arbitrage opportunity, the model assumes that they immediately execute the three transactions associated with the arbitrage cycle, and the affected market makers adjust their mid-prices to the transaction exchange rate.

5.3.3 The interaction between market makers and arbitrageurs

The ABM model attempts to simulate the correlation that develops within the currency trio, in which the so-called co-movement of trend states plays a key role. The trend state itself is determined by the sign of $\emptyset_{n,\ell}(t)$, meaning that for three currencies, $2^3 = 8$ different states can be described.

We can imagine a scenario where, for a short period, the possibility of arbitrage is "switched off". In this case, the three currency markets evolve independently of one another, and each market has an equal chance of moving into the same or an opposite trend – resulting in a total of 2^3, that is, 8 different cases, as shown in Figure 34.

Figure 34: Probability density of trend-alignments in ABM

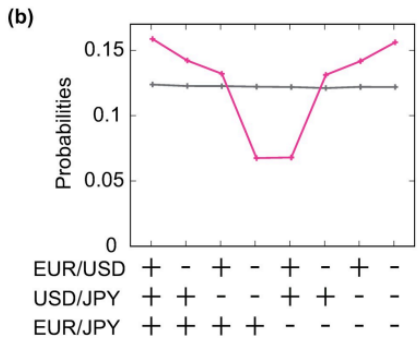

Source: Ciacci et al. (2020), p. 13

The grey line on Figure 34 indicates the typical 1/8 probability in the absence of arbitrage, while the purple line shows the scenario simulated by the authors where arbitrage is present. Without arbitrage, the expected value of correlations between the currencies is zero. When arbitrage is enabled, for example in the case of EURJPY-EURUSD, the probability of the trends aligning increases to 57%[3], so the modelled correlations shift in a positive direction consistent with these probabilities (Ciacci et al. (2020), p. 11).

The authors make three further important observations regarding the differences between markets without arbitrageurs and those with arbitrage (i.e., where arbitrageurs are active):

1. The expected duration of individual trend configurations decreases when arbitrage is enabled. Since arbitrage transactions typically slow down or reverse the trend of the affected exchange rate, arbitrage activity shortens the lifetime of certain trend configurations compared to the case without arbitrageurs.
2. The duration of some configurations increases compared to others. Certain configurations, such as a positive/negative trend in EURJPY and a negative/positive trend in EURUSD and USDJPY, can present arbitrage opportunities sooner than cases where all exchange rates show the same trend. Since the expected duration and probability of trend alignments are essentially two sides of the same coin, the configurations that persist longer determine the correlation between currency pairs.
3. The transition probabilities between different trend configurations vary, offering further explanation for the differing probabilities of each configuration occurring. Table 16 below clearly shows that from the {+,+,+} state, the market transitions to the {-,+,+} state with a 35.8% probability, and to the {+,+,-} state with only a 22.7% probability.

Table 16: Transition rates between two configurations

Table S5. Transition rates between two configurations

Configuration	{+, +, +}	{-, +, +}	{+, -, +}	{-, -, +}	{+, +, -}	{-, +, -}	{+, -, -}	{-, -, -}
{+, +, +}		0.358	0.336	0.010	0.227	0.028	0.026	0.014
{-, +, +}	0.375		0.027	0.218	0.011	0.329	0.014	0.026
{+, -, +}	0.364	0.026		0.217	0.012	0.014	0.341	0.026
{-, -, +}	0.035	0.302	0.295		0.006	0.028	0.030	0.305
{+, +, -}	0.304	0.030	0.028	0.007		0.294	0.302	0.035
{-, +, -}	0.026	0.340	0.014	0.012	0.220		0.026	0.363
{+, -, -}	0.028	0.013	0.330	0.011	0.217	0.027		0.374
{-, -, -}	0.015	0.027	0.027	0.226	0.010	0.335	0.359	

(The table shows the probability of transitioning from the row state to the column state, in the order {EURUSD, USDJPY, EURJPY}.)

Source: Ciacci et al. (2020), supplemental material, p. 6

[3] The overall probability of two exchange rate being in the same trend is the sum of the probabilities of trend-aligned states

Among the three exchange rates, the EURJPY exhibits the strongest trend component when normalised by exchange rate levels, $|\emptyset_{n,\ell}(t)|/p_\ell(t_0)$, which can also be considered as the model's resistance to trend changes. As a result, short-term trend reversals are less likely to occur than, for example, in the case of the EURUSD. This can be explained by the fact that, due to quoting conventions and arbitrage, the EURJPY tends to move in the same direction as both the EURUSD and USDJPY, while the EURUSD / USDJPY will move in the same direction as the EURJPY and in the opposite direction as USDJPY / EURUSD. For instance, if all three quotations are in an upward trend, the EURJPY can only reverse its trend with high probability if both of the other exchange rates also reverse their trends. If the EURUSD and USDJPY move slightly in opposite directions, it is less likely that the market will become arbitrageable and thus reverse the EURJPY trend. From an even more intuitive perspective, EURJPY is the product of EURUSD and USDJPY, meaning the aligned trends of EURUSD and USDJPY amplify and manifest more strongly in EURJPY. It is also noticeable how high the probability is of entering the {+,+,+} state from the {-,+,+} and {+,-,+} states (37.5% and 36.4%, respectively), and vice versa – in other words, the market easily gets stuck in these states.

It can also be established that the probabilities of type 1 and type 2 arbitrage strongly depend on the current trend configuration, shown in Figure 35 (black region: type 1 arbitrage, white region: type 2 arbitrage).

Figure 35: Probabilities of type 1 and type 2 arbitrage

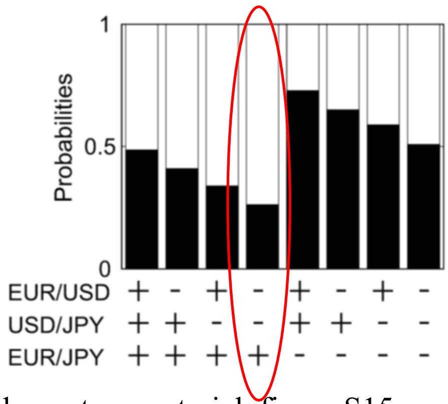

Source: Ciacci et al. (2020), supplementary material, figure S15

It can also be intuitively understood that, for example, when the EURJPY is in an upward trend while the other two exchange rates are in a downward trend, the likelihood increases significantly that selling EURJPY and buying EURUSD as well as USDJPY (type 2 arbitrage) offers arbitrage profit opportunities (a highlighted case in the figure).

As Figure 36 illustrates, in the ABM model, the combined dynamics of trend-following quoting and arbitrage among currency trios appears to support the empirically observed correlation levels in the EUR-USD-JPY currency trio.

Figure 36: Realized and simulated correlations

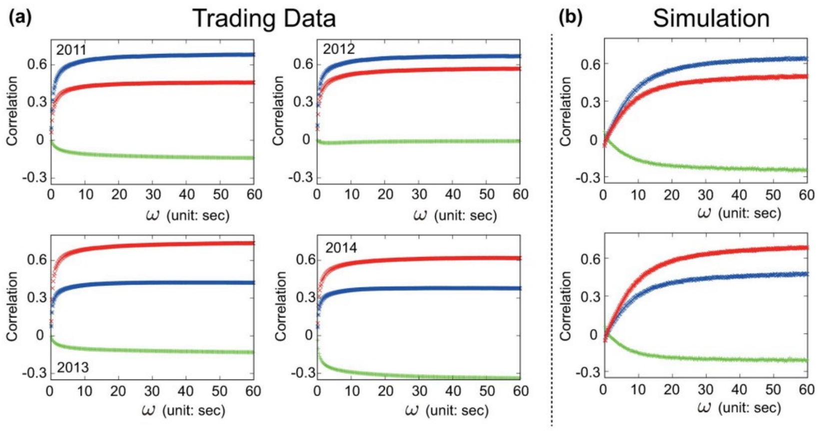

(Red: correlation between EURJPY and USDJPY, blue: correlation between EURJPY and EURUSD, green: correlation between EURUSD and USDJPY)

Source: Ciacci et al. (2020), p. 10

The simulated values, which are close to the empirical cases in the top and bottom rows, were achieved by the authors through the choice of the number of market makers: N_{EURUSD}, N_{USDJPY}, N_{EURJPY} = 35, 45, 25 in the upper case, and NEURUSD, NUSDJPY, NEURJPY = 50, 35, 25 in the lower case. All other parameters are identical in both cases: the strength of trend-following (c_l) is 0.8[4], the look-back period is n=15 periods, and ξ = 5 for all three currency pairs.

5.4 The ABM model in a symmetric currency market

In the following I reproduce the results of the ABM model based on Ciacci et al. (2020), to obtain a model that converges to the average correlation levels characteristic of CEE/Scandinavian markets, which differ from the yen market.

In the model I used, all three exchange rates are quoted by N=3 market makers. During the modelling process, in each step the market makers modify their mid-quote prices – as defined by the ABM model. If a transaction occurs within any of the currency markets (i.e., the bid

[4] The value of trend following coefficient is derived based on empirical findings – for further information see Supporting Information Ciacci et al. (2020)

price of one market maker overlaps with the ask price of another market maker), the market makers adjust their mid-quotes to the transaction price. This is then followed by an examination of arbitrage between markets. If arbitrage occurs, the affected market makers modify their mid-quote to the average of arbitraged exchange rate level. In the absence of arbitrage within the arbitrage loop, the market maker again modifies their quote according to relation (59). Thus, the model handles transactions both within and between currency exchange rates in two steps.

To investigate whether the ABM model can provide a microstructural explanation for the varying correlation levels observed within different currency trios, it is insightful to examine a scenario in which the currency trio is considered maximally homogeneous, meaning all three exchange rates are started from the same value ($p_{0,\ell}=1$), the bid-ask spread is identical for all exchange rates ($L_\ell=0{,}0001$), the volatility associated with quoting is identical for all ($\sigma_\ell = 6{,}32 * 10^{-5}$), and the trend-following coefficient is zero ($c_\ell = 0$). The magnitude of the bid-ask spread was chosen to be close to the currently observed empirical EURUSD spread, while the volatility is the proportion of an annual 10% volatility for one day, divided by one hundred: $0.00000632 = \frac{10\%}{\sqrt{250}}/100$. N=3 market makers are assumed in each market.

Figure 37 shows one realisation of the simulated correlation of log returns calculated over k=474 steps[5].

Figure 37: Simulated correlations – symmetric currency market

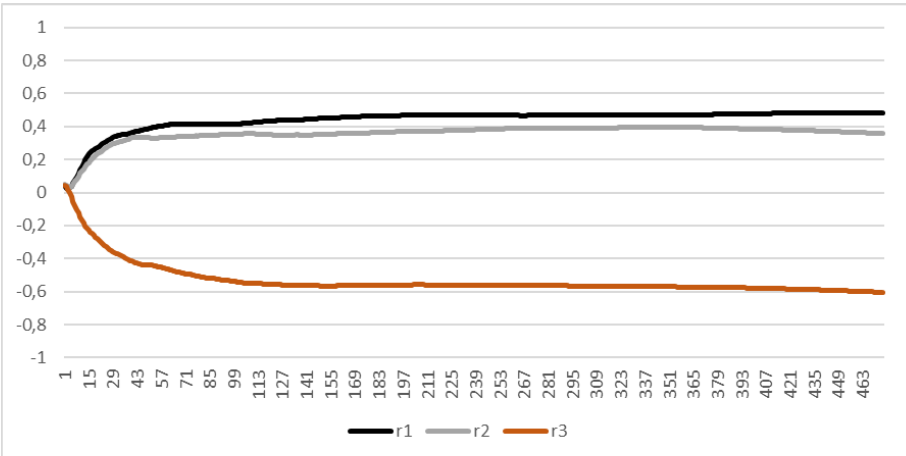

r1=EURUSD-EURJPY, r2=EURJPY-USDJPY, r3=EURUSD-USDJPY

Source: Author's own work

[5] The number of steps has been determined by the computing capacity behind the excel model.

In the symmetric model, therefore, the correlations clearly converge towards the values observed in the EUR-USD-JPY currency trio, and the probabilities of simultaneous or opposing movements are consistent with both the strength and the sign of the correlations – see earlier chapter. Based on 250 runs of the model describing the EUR-USD-JPY trio, the 90% confidence interval for the individual correlations is shown in Figure 38.

Figure 38: Confidence intervals of simulated correlations – symmetric currency market

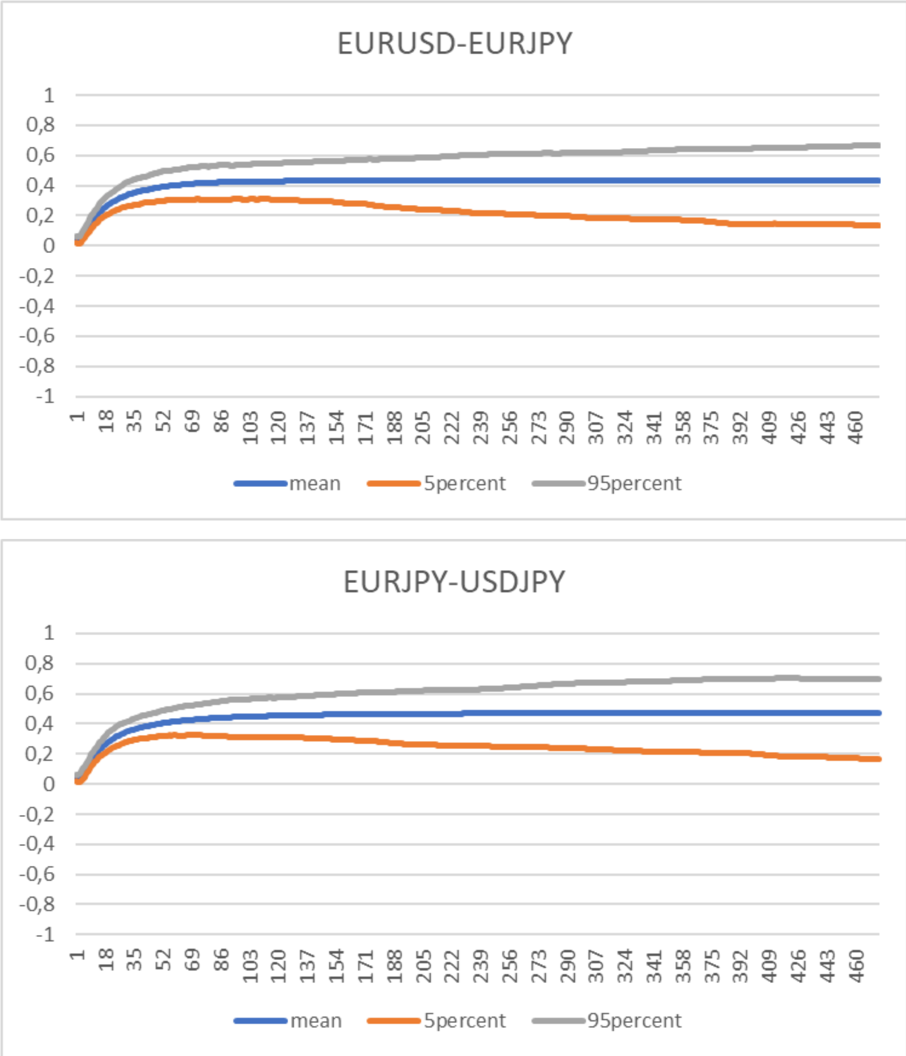

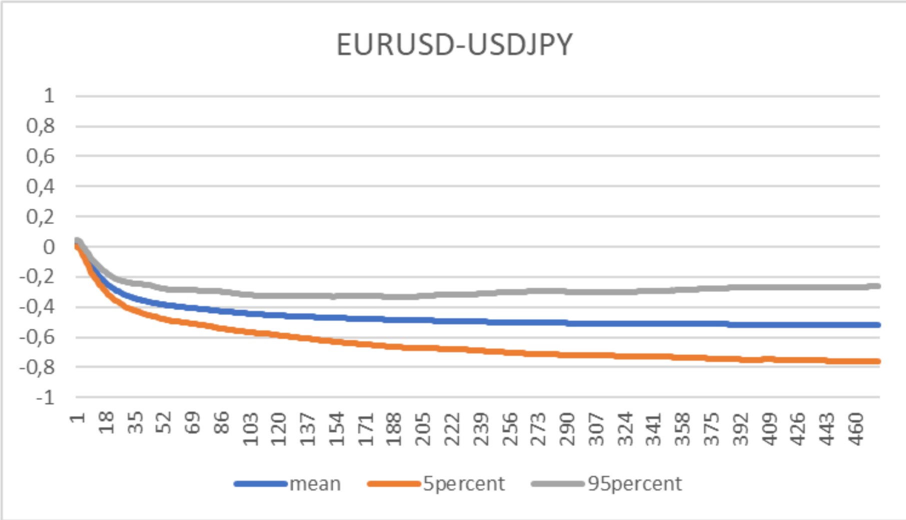

Source: Author’s own work

The confidence interval marks a rather wide correlation range, which is consistent with what has been observed empirically – see Table 1 for the range of each currency correlation.
Based on the model runs, I also calculated the probability of each trend alignment in a manner similar to Figure 34. In the case where the trend-following coefficient, $c_\ell = 0$, a picture emerges that is similar to, but not entirely symmetric with, Ciacci et al. (2020), as shown in Figure 39.

Figure 39: Probability density of trend-alignments in symmetric currency market (C=0)

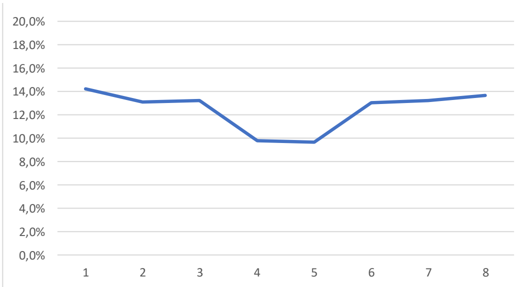

EURUSD	+	-	+	-	+	-	+	-
USDJPY	+	+	-	-	+	+	-	-
EURJPY	+	+	+	+	-	-	-	-

Source: Author’s own work

In this case, the probability of the EURUSD-EURJPY and EURJPY-USDJPY trend alignment is 54%-54% (cf. 57% in Ciacci et al. (2020) when C=0.8). If we assume $c_\ell = 0.8$, then the probability of trend alignment for EURUSD-EURJPY and EURJPY-USDJPY changes to 56% and 57%, respectively – which is in agreement with Ciacci et al. (2020), as shown in Figure 40.

Figure 40: Probability density of trend-alignments in symmetric currency market (C=0.8)

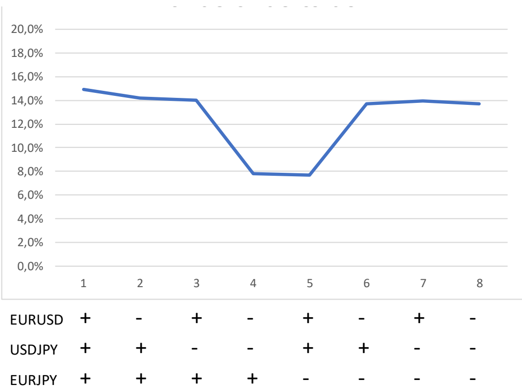

Source: Author's own work

The trend-following coefficient, c_ℓ, of the ABM model markedly alters the levels of correlation within the model, especially when, for one or more exchange rates, the coefficient is significantly different from the others. If we assume identical trend-following coefficients across all three currency markets, we obtain correlation levels close to those found in Ciacci et al. (2020): that is, we see two moderately strong positive correlations (EURJPY-EURUSD and EURJPY-USDJPY) and one moderately strong negative correlation (EURUSD-USDJPY). Among the different combinations of trend-following coefficients, a particularly interesting result emerges when a value of 1.5 is applied to USDJPY while the other two currencies are set to zero, see Figure 41.

Figure 41: Simulated correlations (C1=C2=0, C3=1.5)

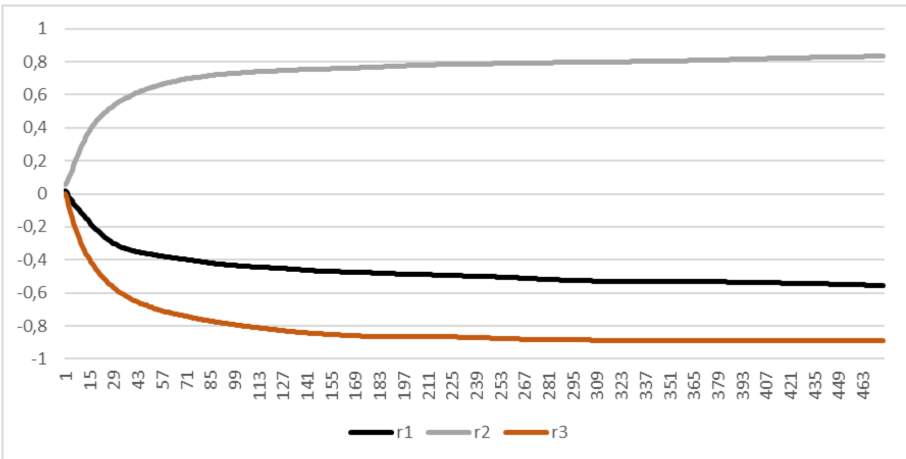

r1=EURUSD-EURJPY, r2=EURJPY-USDJPY, r3=EURUSD-USDJPY

Source: Author's own work

Under this assumption, the probabilities of the trend configurations also change significantly, as shown in Figure 42.

Figure 42: Probability density of trend-alignments (C1=C2=0, C3=1.5)

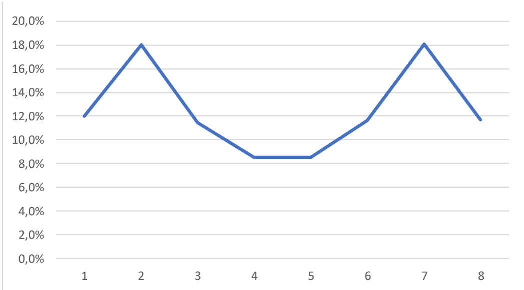

EURUSD	+	-	+	-	+	-	+	-
USDJPY	+	+	-	-	+	+	-	-
EURJPY	+	+	+	+	-	-	-	-

Source: Author's own work

The probability of trend alignment between EURUSD and EURJPY falls to 47%, which is consistent with a weak negative correlation. For EURUSD-USDJPY, this figure is 41%, indicating a moderately strong negative correlation, while the probability of trend alignment between EURJPY and USDJPY rises to 60%, suggesting a moderately strong positive correlation – in line with the correlation chart (Figure 41). In this version, the introduction of a high trend-following coefficient for USDJPY dramatically increases the likelihood of USDJPY and EURJPY moving in the same trend, meaning that USDJPY is less likely to break out of an upward or downward price trend. Through the arbitrage process, this forces EURJPY to move together with USDJPY, while EURUSD moves in the opposite direction (due to inverse quoting). As a reminder: in a strong positive USDJPY upward trend, the arbitrage opportunity manifests as USDJPY sale – EURJPY purchase – EURUSD sale in the arbitrage circuit, meaning that an increase in USDJPY forces an increase in EURJPY and a decrease in EURUSD in the other two markets. It is easy to see that the opposite occurs in a strong USDJPY downward trend.

In this special case, the model produces precisely the correlation levels also observed in the forint market, despite the fact that, none of the forint market (CEE regional) quoting practices detailed by banks have been incorporated into the model. Upon replacing JPY with HUF, it becomes apparent that the EURUSD-EURHUF correlation falls into a weak negative range, the EURHUF-USDHUF correlation is strongly positive, while the EURUSD-USDHUF correlation shows a strongly negative level.

In summary, trend-following and market quoting conventions play a key role in the ABM model of Ciacci et al. (2020). In the symmetric version – where the trend-following coefficients are identical for each exchange rate – it can be concluded that increasing the trend-following coefficient increases the absolute value of the correlations. Furthermore, due to the quoting convention and arbitrage, the exchange rate acting as a product (EURJPY) is more prone to becoming "stuck" in a short-term trend, which probabilistically leads to co-movement, meaning a positive correlation with both EURUSD and USDJPY. In the same ABM model, there is an asymmetric case (C_{EURUSD}=0, C_{EURJPY}=0, C_{USDJPY}=1.5) where USDJPY is forced into strong trend following, thereby closely approximating the correlation levels observed in the EUR-USD-HUF currency trio. (Although not elaborated in the thesis, with the condition set as (C_{EURUSD}=0.8, C_{EURJPY}=0.8, C_{USDJPY}=1.5), the correlation patterns of the forint market can also be observed.)

This state is therefore created by overriding the trend component arising from the market's characteristics with a highly significant trend coefficient. Although the asymmetric ABM model almost perfectly describes the correlation patterns observed in the forint market, there is no known empirical evidence for such pronounced differences in the trend-following coefficient. For this reason, in the next chapter, I examine a version of the ABM model that aligns with the quoting practices described by bank market makers – see the results of the qualitative questionnaire.

5.5 The ABM in the Forint market

The next step is to assume that the volume or significance of one currency (HUF) is negligible compared to the other two currencies (EUR, USD). Although all three exchange rates are involved in transactions during arbitrage, from a modelling perspective, this is equivalent to the outcome of a forint market arbitrage transaction having a negligible effect on the EURUSD rate. That is, in the model, arbitrage does not influence the EURUSD exchange rate, which thus evolves entirely independently. In terms of the model, this means that EURUSD is excluded from the arbitrage circuit, and it is as if deals are made only in EURHUF and USDHUF, which then move the exchange rates. The modification causes a slight shift, but the modelled correlation regime still more closely resembles that of a major currency trio. In the next step, I took into account the qualitative survey result which indicated that, according to most banks, relatively illiquid markets are quoted in an arbitrage-free manner from the outset, meaning that USDHUF = EURHUF/EURUSD. This modification is applied to individual market makers, i.e., each market maker adjusts their USDHUF quoting according to the EURUSD and EURHUF rates quoted by their own bank. However, this also implies that arbitrage opportunities may remain within the arbitrage circuit, meaning that triangular arbitrage can still operate between banks.

In this case, as a result of the simulation, the model will converge to correlation levels much closer to those observed in the forint market: a moderately strong positive correlation (r2) was modelled between EURHUF and USDHUF, an almost zero correlation (r1) between EURHUF and EURUSD, and a moderately strong negative correlation (r3) between EURUSD and USDHUF, see Figure 43.

Figure 43: Confidence intervals of simulated correlations – forint market

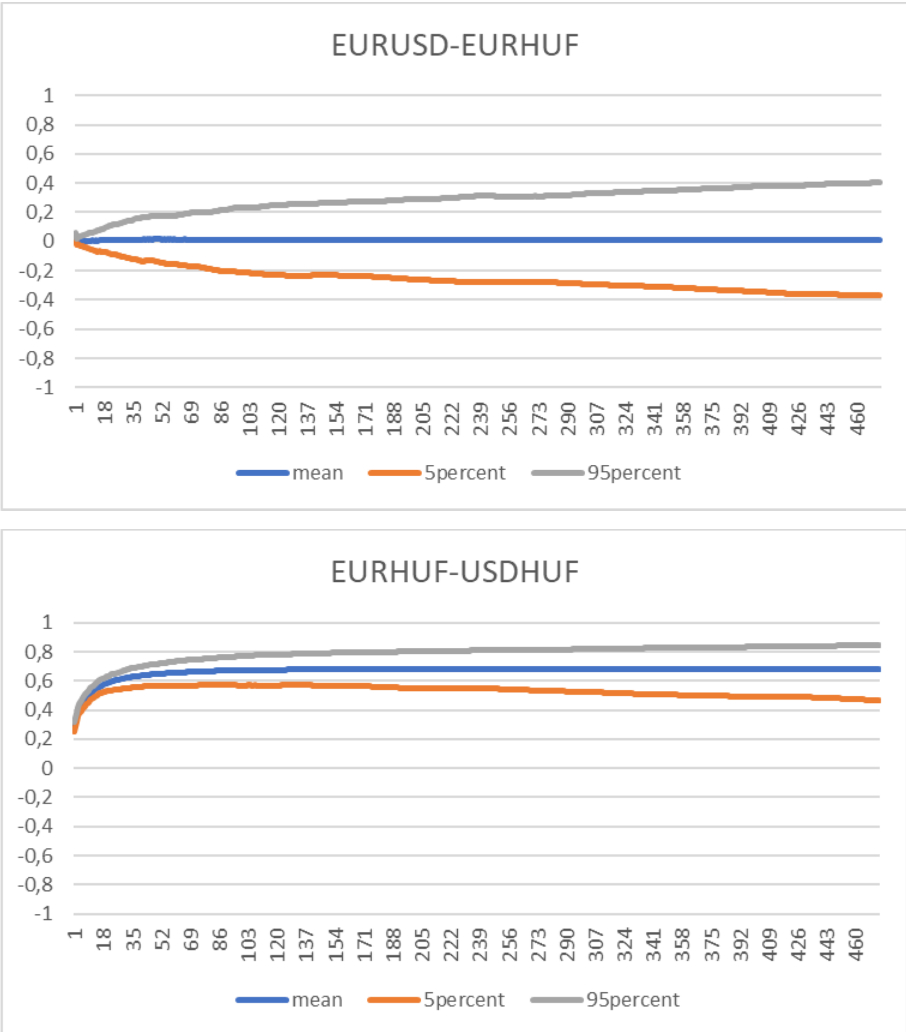

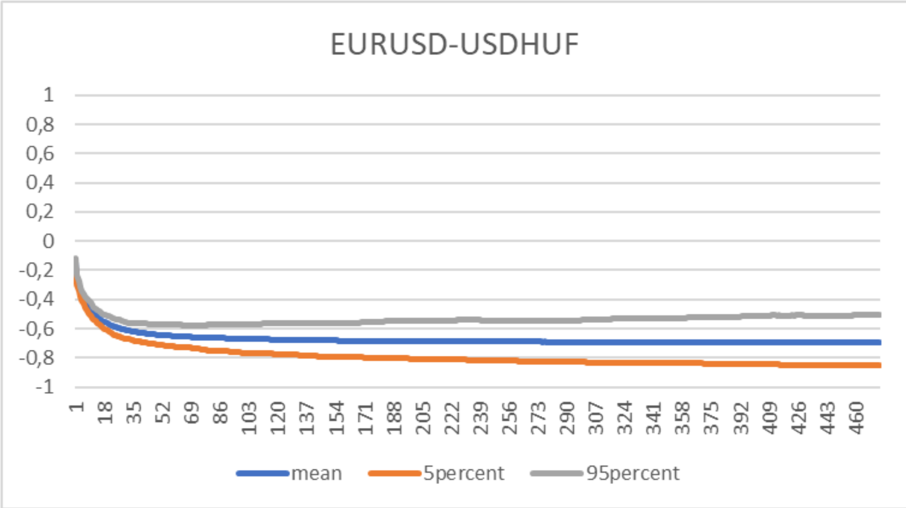

Source: Author's own work

Examining the probability of trend alignments, Figure 44 presents results of the modelling the HUF market.

Figure 44: Probability density of trend-alignments in forint currency market (C=0)

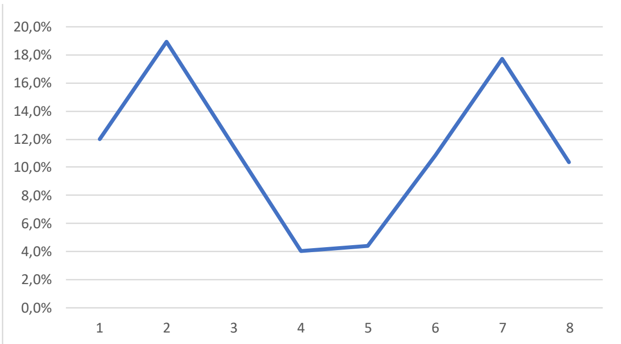

EURUSD	+	-	+	-	+	-	+	-
USDHUF	+	+	-	-	+	+	-	-
EURHUF	+	+	+	+	-	-	-	-

Source: Author's own work

Figure 44 is reminiscent of the symmetric model with c1=0, c2=0, c3=1.5 version.
The probability of the EURUSD and EURHUF moving in the same trend is 45 percent, while EURUSD and USDHUF align in the same trend in 31 percent of cases, and USDHUF and EURHUF align in the same trend in 59 percent of cases. All this is consistent with a near-zero EURUSD-EURHUF correlation, a strong negative EURUSD-USDHUF correlation, and a strong positive EURHUF-USDHUF correlation.

It is intuitively clear that, for example, it is difficult for USDHUF to move in the same trend as EURUSD and in the opposite trend to EURHUF (cases 4 and 5 in Figure 44) when USDHUF quoting in the market is continuously aligned to the EURHUF/EURUSD ratio. On the other hand, the two opposite cases (2 and 7) are more probable.

With this modification, I have essentially reached a currency correlation regime that can be achieved endogenously by the model parameters, and which is already very close to the empirically observed regional currency correlation regime.

As discussed previously, the key to achieving the model that reproduces the correlation levels typical of CEE/Scandinavian currencies was that the illiquid (USDHUF), or by market convention, indirectly quoted (USDNOK) exchange rate was set to be arbitrage-free from the outset.

This modified version of the ABM closely reproduces the correlation levels observed in Scandinavian and CEE currency trios. Apart from certain unrealistic specifications of the ABM model, so far I have not been able to simulate a market setup that reproduces the weak negative correlation between EURUSD and EURCCY exchange rates, and which can also be empirically supported based on quoting practices. One explanation for this may be that the phenomenon is not microstructural in nature, but rather related to market risk/pricing. In practice, it is a well-known phenomenon that “risk-on/risk-off” market conditions regularly coincide with a weakening dollar – strengthening CEE currencies, or a strengthening dollar – weakening CEE currencies market.

It is easy to imagine a scenario – consistent with the impulse response analysis presented in Chapter 4.5.2, which found that USDHUF also has a positive effect on EURHUF, as well as the banking practice described in the Appendix C, where in illiquid markets, market makers often hedge their positions in the two other (liquid) markets – where, in a significant EURUSD trend, the flow into USDHUF is hedged by market makers through EURHUF and EURUSD. Because synthetic hedging of USDHUF always involves trades in the opposite direction for EURUSD and EURHUF, this may intensify the negative relationship between EURUSD and EURHUF. In any case, attempts to modify the ABM model in this direction have not been successful; that is, I have not managed to reproduce the weak negative correlation between EURUSD and EURHUF.

5.6 The ABM in USD-driven markets

The open question regarding the ABM model and the empirically observed correlation levels in currency markets concerns those currencies where the dominant exchange rate is USDCCY, i.e., where the EURCCY rate is quoted via EURUSD and the USDCCY rates. Good example of this is the Mexican peso, the correlations of which are shown in Figure 45.

Figure 45: Correlations in Mexican peso currency trio

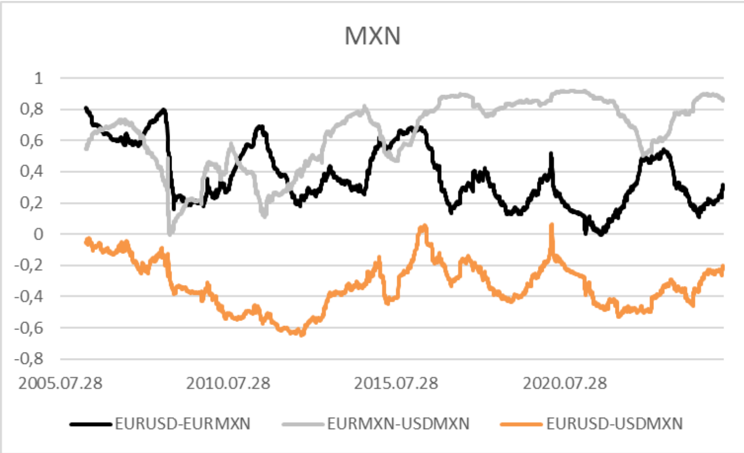

Source: Refinitiv, Author's own work

With a further modification of the ABM model – where the EURCCY exchange rate is quoted indirectly – I modelled the correlations shown in Figure 46.

Figure 46: Confidence intervals of simulated correlations – USD-driven market

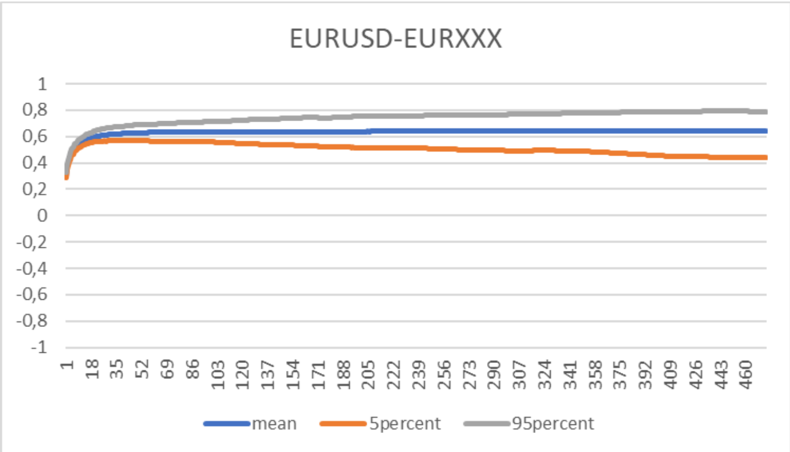

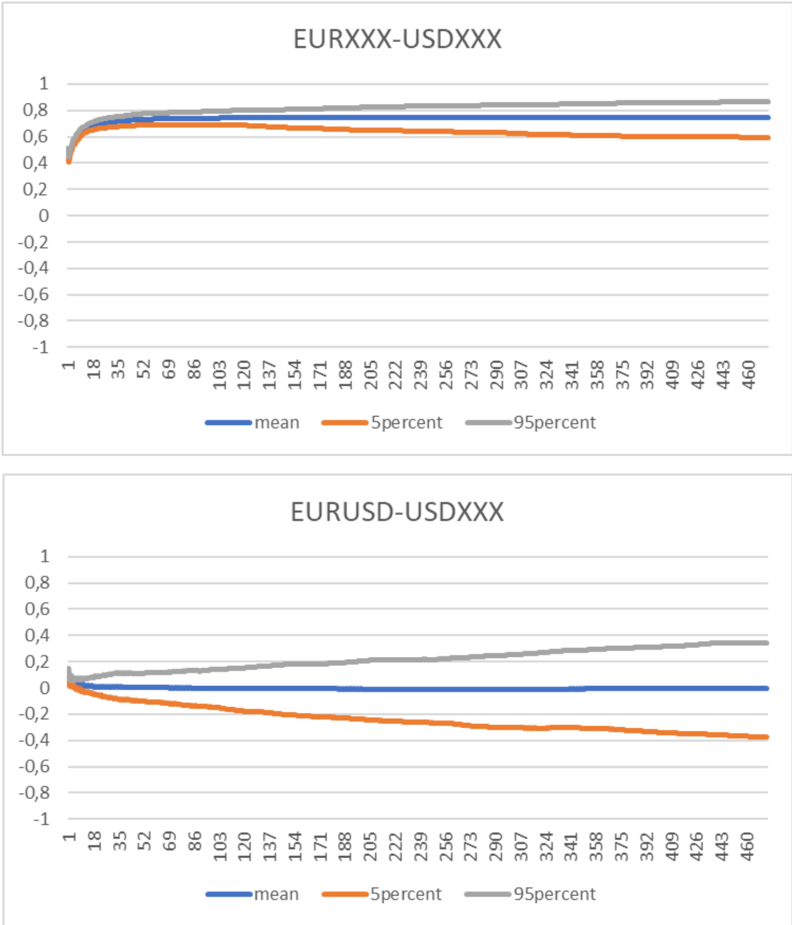

Source: Author's own work

The probability of trend alignments are illustrated in Figure 47.

Figure 47: Probability density of trend-alignments in USD-driven currency market (C=0)

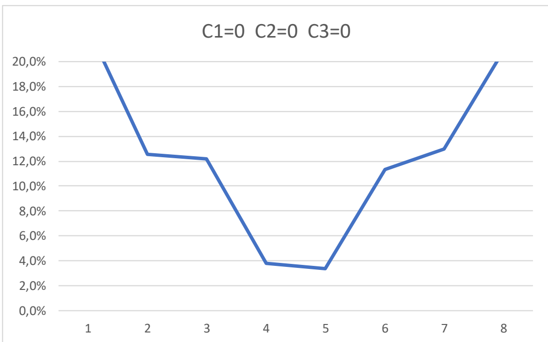

EURUSD	+	-	+	-	+	-	+	-
USDXXX	+	+	-	-	+	+	-	-
EURXXX	+	+	+	+	-	-	-	-

Source: Author's own work

The probability of EURUSD and EURCCY moving in the same trend is 67 percent, EURUSD and USDCCY align in the same trend in 51 percent of cases, while USDCCY and EURCCY move together in 69 percent of cases. All this is consistent with a near-zero EURUSD-USDCCY correlation, and a strong positive correlation between EURUSD-EURCCY and EURCCY-USDCCY. Intuitively, since the EURCCY quote is the product of EURUSD and USDCCY, a positive correlation and frequent trend alignment can be expected with both rates. The main difference between the modelled and empirically observed correlations, is the correlation between the two exchange rates assumed to be independent does not converge to zero but shifts into the negative range.

5.7 Discussion and conclusions

The ABM model and its various modifications are able to reproduce the correlation levels observed in different currency trios along fairly realistic market/quoting assumptions:

- The relative correlation levels within currency trios are positioned similarly to those observed empirically, both in terms of their strength and in most cases their sign.
- The differences in correlation levels between different currency trios are reproduced in line with the empirically observed differences.

Based on the ABM model, we can identify three groups of currency trios, while noting that in reality currency trios rarely operate purely according to one regime or another, but rather as some combination of three (or typically two) regimes:

1. In a symmetric currency trio, all exchange rates are similarly liquid and essentially independently quoted and are held together by triangular arbitrage. In this trio, typically two moderately strong positive and one moderately strong negative correlation emerge.
2. In some currency markets, the EURCCY exchange rate is significantly more liquid (CEE region), or by tradition, the USDCCY rate is quoted via EURCCY and EURUSD (Scandinavian markets). In these cases, there are effectively two independently quoted rates (EURUSD and EURCCY) and one synthetically quoted rate (USDCCY). In these markets, the correlation between EURUSD and EURCCY is close to zero (empirically weakly negative), strong positive between EURCCY and USDCCY, and strong negative between EURUSD and USDCCY.
3. In other currency markets, USDCCY is the dominant exchange rate (e.g. Mexican peso), meaning these currencies are primarily measured against the USD on the FX market, and the EURCCY rate is merely a derivative of the other two rates in the trio. In these markets, a near-zero (slightly negative) correlation develops between EURUSD and

USDCCY, while in the other two pairs (EURUSD-EURCCY and EURCCY-USDCCY), a moderately strong positive correlation can be measured.

According to the ABM model, in all three cases, the resulting correlation levels are essentially probabilistic in nature, determined by the probabilities of formation and transitions of same and opposite directional trends over the short term. We have seen that, in different market structures, the various trend alignments occur with different probabilities, which is consistent with the differing correlation levels. In the symmetric currency trio (most closely resembling JPY), the probability of same/opposite trends is essentially determined only by quoting convention (based on numerator/denominator currency; for example, EURJPY tends to move together with EURUSD and USDJPY, while USDJPY tends to move contrary to EURUSD). In the CEE region, for instance, the EUR rate trends together with USD rate. The reason for this is that, in CEE (and Scandinavian) currency trios, triangular arbitrage is reflected in the USDCCY rate, and EURUSD-EURCCY rates remain independent of each other. Thus, while in the symmetric currency trio the quoting method (numerator/denominator currency) determines the probability of trend alignment, in asymmetric currency trios the dominant exchange rate (directly quoted currency) shapes the relative position, sign, and strength of the correlations.

It was also established that, the model is able to reproduce the weak negative correlation between EURUSD and EURCCY/USDCCY in each currency trio, but under realistic assumptions it cannot fully do so. The thesis does not attempt to answer this question, but merely raises the possibility that this negative relationship is primarily not due to quoting reasons, but may have risk-perception (“risk-on”/“risk-off”) or risk-pricing causes.

The results support H4: the ABM model cannot fully reproduce the correlation structure observed in the regional currency trios under realistic assumptions.

Impulse response analysis based on ultra-high-frequency observations of currency trios, as well as responses from surveyed international banks, also support the existence of different dynamic relationships and quoting practices between currency trios. These empirical findings also support the observation that currency trios do not “purely” exhibit different structures, but rather that these structures vary by market maker and over time. This is also reflected in the temporal development of the correlation coefficients measured in individual currency trios, which typically move within a (sometimes quite wide) band.

The primary aim of this chapter is to provide an explanation to the differences in correlations between currency trios, using the ABM model as a tool. The three-market-maker version I employed and its various modifications proved suitable for achieving this main objective, but there remain open questions which could be addressed by further modifying the ABM model. Without claiming to be exhaustive, further refinements of the arbitrageur model might include:

1. Non-homogeneous market makers: The literature has extensively examined the structure of foreign exchange market makers (e.g., Hagströmer and Menkveld (2019)), which supports the notion that market maker markets are not homogeneous; a kind of centre and periphery formation can be observed, as well as specialisation among individual market makers. Based on this, it may be assumed that certain market makers behave differently. According to the findings of Babus and Kondor (2018), central market makers necessarily become better informed than those on the periphery. Therefore, it is conceivable to develop an ABM version in which individual market makers behave differently both within and across currency pairs.
2. The degree and direction of market maker trend-following: The literature widely investigates feedback in market maker behaviour, during which price movements themselves generate order flow (either as positive or negative feedback). In the ABM model (Ciacci, et al., 2020) a trend-following coefficient of 0.8 is applied, but the choice of this measure is not strongly substantiated. In an interesting version of the ABM model, a coefficient greater than 1 introduced for USDCCY, interestingly, reproduces the correlation levels observed on the forint market – including the weakly negative correlation between EURUSD and EURHUF.
3. Extension of the three-currency market maker model to four or more currencies: The main question is whether, in every case, the quotations against EUR and USD are dominant, or whether minor cross rates also affect currency trios – and, through this, the correlations within the trio. Breedon, et al (2018) examined the impact of the Swiss National Bank's 2015 decision, when the bank unexpectedly abandoned the 1.2 exchange rate limit. By including JPY in the EUR-USD-CHF analysis, they did not find a significantly higher number of arbitrage opportunities during this period. Another interesting aspect of the multi-currency model is that, for example, CEE regional currencies are often treated as a group on the market, meaning there is a strong positive correlation observed between the HUF, CZK, and PLN exchange rates. While, when evaluating the individual currency markets separately, their insignificant trading volumes compared to EURUSD meant that their feedback effect on EURUSD was switched off in the modelling, in certain market situations – if the market moves simultaneously in all three regions – it is conceivable that arbitrage activity could feed back into EURUSD as well.

6. CONCLUSIONS

6.1 Summary and main findings

The thesis captures a specific feature of foreign exchange market correlations, namely that currency trios (currency triangles) formed from three currencies are interconnected due to the possibility of spot FX market arbitrage. Because of this special relationship, the implied correlations, which essentially constitute the only forward-looking correlation forecasts reflecting market expectations, can be clearly calculated using implied volatilities derived from currency option pricing. After presenting the mathematical properties of currency correlations, the predictive power of implied and traditional time series-based correlation estimates was tested on the regional and major FX market. Based on the results, the implied correlation estimate – consistent with the international literature – cannot be considered a clearly more accurate forecast.

A detailed examination of the forecast errors led to several further interesting findings:

- In the regional markets the largest forecast errors occur in the correlations between EURUSD and EURCCY, where essentially the relationship is between two independent exchange rates.
- Although there is a functional relationship between FX volatilities and correlations, the various forecasting methods do not provide the same ranking of accuracy in forecasting the two measures. The method that forecasts volatility with the smallest error (ATM implied) is not the most accurate for estimating correlation (EWMA). This is because for volatility forecasting, the absolute value of volatility must be predicted most accurately, while for correlation estimation, it is the volatility ratios that matter.
- Further breaking down the correlation estimation error by volatility ratios, it is found that in ATM volatility-based correlation estimation, the largest contribution to the estimation error comes from the volatility estimate of USDCCY, consistent with the ranking observed in volatility estimation, where, the volatility estimate for USDCCY is the least certain.

The results support the H1 the implied correlation does not provide more accurate predictions than time-series based estimates. However, based on the results, H2 is rejected: despite the intuition volatility and correlation forecasting methods are not resulting the same ranking in the regional currency trios.

Partly based on the results of the correlation estimation and partly on the observation of differing correlation levels across currency trios (more specifically, currency regions) over longer periods, Chapter 4 of the thesis seeks to answer why distinctly different correlation levels arise despite the fact that most currency trios use the same quoting conventions: EURUSD = value of 1 euro in dollars, EURCCY = value of 1 euro in currency CCY, and USDCCY = value of 1 dollar in currency CCY. The literature mainly focuses on correlation risk and its pricing, while empirical evidence shows that, in Scandinavian and CEE regional currency trios, the pairwise correlations are at very similar levels. In Chapter 2.5, based partly on previous literature, the relative trading volumes of FX markets are explored, laying the groundwork for the idea that, in certain currency trios, the trading volumes and thus the relative importance of specific exchange rates differ significantly. **This implies that, in many currency trios, only two of the three exchange rates genuinely play a price-determining role, while the third (EURCCY or USDCCY) acts as a passive participant, essentially as a function of the movements of the other two rates, supporting the 3rd hypothesis of the thesis (H3), pointing to a major difference in sub-second dynamics between currency trios.**

In Chapter 4.5, the dynamics of several currency trios are mapped using ultra-high-frequency data, which highlights the observation that individual exchange rates influence the other two in different ways. In the yen and, to varying degrees, other major currency trios, impulse response functions show a kind of symmetry during normal periods, i.e., each exchange rate affects the other two. In contrast, in the regional markets, the EUR rate does not influence EURUSD, and the effect of USD rate is likely primarily attributable to the short-term autocorrelation of EURUSD. The results of the vector autoregressive (VAR) and local projection (LP) analysis also revealed that the dynamic relationships are not homogeneous in time and space (i.e., across market makers). To further improve the robustness of the findings directional spillover matrices are calculated pointing to the same conclusion as clear distinction can be done between regional and major currency trios where the former appear to consist of two independent FX rates with a third synthetic leg limiting the price discovery channels within the trio and the later shows more symmetric topology. Provided the argument of the thesis that "synthetic quotation" is probably the root cause of the different correlation regimes some microstructural alternatives are also tested if they can reproduce the observed spillover topology and they fail to do so.

The results of the VAR modelling are supported by the responses of the surveyed international banks as follows:

- Bank quoting practices can differ across currency trios, primarily due to the relative importance of each exchange rate, which stems partly from trading volumes and partly from historical reasons.
- Market makers have a uniform view regarding quoting in the CEE region, but their opinions are less homogeneous regarding Scandinavian and the major currencies regarding the synthetic/independent quotational practices.

An important finding is that quoting and hedging conventions are not uniform across institutions, currencies, or time periods. Consequently, observed volatility and correlation patterns should be interpreted as outcomes of evolving market practices rather than immutable characteristics of specific currencies.

Following both quantitative and qualitative empirical analysis, an arbitrageur model and its various modifications are presented in Chapter 5, with which it was possible, using realistic and empirically consistent model assumptions, to model three distinct currency trio setups and the correlations that develop within them:

- In the symmetric currency trio, all three exchange rates are of equal/similar importance and independently quoted, linked by triangular arbitrage. In this setup, correlations converge in the long run to the levels observed in the yen market.
- In the euro-centric currency trio, the two independently traded exchange rates are EURUSD and EURCCY, while USDCCY is determined by the other two, i.e., it is quoted arbitrage-free. In this model, correlations converge towards the levels observable in the regional currencies. It should be noted, however, that the model only finds a mathematically meaningful solution for the weak negative correlation seen between EURUSD and EURCCY, and, under realistic assumptions, only manages to reproduce a correlation close to zero. The results suggest that the weak negative correlation may have a risk evaluation/pricing explanation, rather than a market microstructural one.
- The third model variant aims to represent those currency trios where the two independently traded rates are EURUSD and USDCCY, while EURCCY is determined by the other two and thus quoted arbitrage-free. In this case, correlations similar to those observed for, among others, the Mexican peso is reflected. Similarly to the euro-centric currencies, the empirical evidence here also differs in that USDCCY exhibits not a zero, but rather a slightly negative long-term average correlation with EURUSD.

The results of Chapter 4 and Chapter 5 support the 3rd hypothesis of thesis that, in the FX market, the level of correlation between currency pairs is largely explained by market quoting conventions and the triangular arbitrage relationships between exchange rates, where the major drivers are the relative importance of the respective currency rates. Also, Hypothesis 4 (H4) is accepted, as I did not find a realistic ABM model that converges exactly to the observable levels – more specifically the weak negative correlation between the independently quoted FX rate and the EURUSD was not reproduced.

6.2 Directions for future research

One key finding of the thesis is the heterogeneity of quotation practices among different market makers. The VAR modelling in Chapter 4.5 was conducted based on two free data-sources. Both Dukascopy and TrueFX are known names in the currency market as platforms however the related literature mostly applies data from CLS, EBS, Refinitiv Matching that are considered as major venues of currency trading.

Although the thesis is concentrating on separate currency trios, in reality all currency trios are part of the currency market and arbitrage could arise when a fourth or even more currencies are added to the ecosystem. Therefore, I see some potential in terms of deeper understanding the correlation and volatility structure of the currency market within both the VAR and ABM framework.

Due to the limitations of the thesis the question of slightly negative but close to zero correlations with the EURUSD remained open, as I did not find a microstructural answer under realistic assumptions. Further investigation of this phenomenon is needed.

The thesis and the main source of the ABM Ciacci et al. (2020) does not put emphasis on explaining the actual value of the trend-following parameters in the arbitrageur model. As under certain trend-following assumptions I found major shift in correlation structures it is important to investigate further what might be a realistic range and how it affects the ABM model and the correlation structure.

Finally, the main finding of the thesis – namely there is a certain level of pairwise correlations explained by the quotational architecture – may provide an interesting extension to the literature investigating the dispersion of currency correlations with relation to various risk factors – like Mueller et al. (2017).

References

Adammer, P. (2019). *lpirfs: An R Package to Estimate Impulse Response Functions by Local Projections.* [Online]
Available at: https://journal.r-project.org/articles/RJ-2019-052/index.html
[Accessed 20 10 2025].

Aiba, Y. and Hatano, N. (2006). A Microscopic Model of Triangular Arbitrage. *Physica A,* 371(2), pp. 572-584. https://doi.org/10.1016/j.physa.2006.05.046

Aiba, Y., Hatano, N., Takayasu, H., Marumo, K. ans Shimizu, T. (2002). Triangular Arbitrage as an Interaction Among Foreign Exchange Rates. *Physica A,* 310(3-4), pp. 467-479. https://doi.org/10.1016/S0378-4371(02)00799-9

Babus, A. and Kondor, P. (2018). Trading and Information Diffusion in Over-the-Counter Markets. *Econometrica,* 86(5), pp. 1727-1769. https://doi.org/10.3982/ECTA12043

Baillie, R. T., Booth, G. G., Tse, Y., and Zabotina, T. (2002). Price discovery and common factor models. *Journal of financial markets*, 5(3), pp. 309–321.

Barunik, J., and Vacha, L. (2018). Do Co-Jumps Impact Correlations in Currency Markets?. *Journal of Financial Markets,* Volume 37, pp. 97-119. DOI: 10.1016/j.finmar.2017.11.004

Bauwens, L., & Xu, Y. (2023). DCC- and DECO-HEAVY: Multivariate GARCH models based on realized variances and correlations. *International Journal of Forecasting*, 39(2), 938–955. https://doi.org/10.1016/j.ijforecast.2022.03.005

Benavides, G., & Capistrán, C. (2012). Forecasting exchange rate volatility: The superior performance of conditional combinations of time series and option implied forecasts. *Journal of Empirical Finance*, 19(5), 627–639. https://doi.org/10.1016/j.jempfin.2012.07.001

Berger, D. W., Chaboud, A. P., Chernenko, S. V., Howorka, E., and Wright, J. H. (2008). Order flow and exchange rate dynamics in electronic brokerage system data. *Journal of international Economics*, 75(1), pp. 93–109.

Berlinger, E., Dömötör, B., Illés, F. and Váradi, K. (2016). Stress Indicator for Clearing Houses. *Central European Business Review,* 5(4), pp. 47-60. DOI: 10.18267/j.cebr.166

BIS (2011). *High-frequency Trading in the Foreign Exchange Market,* Basel, Switzerland: Bank for International Settlements.

BIS (2022). *Triennial Central Bank Survey - OTC Foreign Exchange Turnover in April 2022,* https://www.bis.org/statistics/rpfx22_fx_annex.pdf: Bank for International Settlements.

BIS (2025). *Triennial Central Bank Survey of Foreign Exchange and Over-the-counter (OTC) Derivatives Markets in 2025,* https://www.bis.org/statistics/rpfx25.htm: Bank for International Settlements.

Blair, B. J., Poon, S.-H., & Taylor, S. J. (2001). Forecasting S&P 100 volatility: The incremental information content of implied volatilities and high-frequency index returns. *Journal of Econometrics*, 105(1), 5–26. https://doi.org/10.1016/S0304-4076(01)00068-9

Bodurtha, J. N. and Shen, Q. (1995). Historical and Implied Measures of ́'Value at Risk': The DM and Yen Case ́. *Manuscript, Georgetown Univerity.*

Bollerslev, T. (1986). Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics,* 31(3), pp. 307-327. https://doi.org/10.1016/0304-4076(86)90063-1

Bollerslev, T. and Engle, R. F. (1993). Common Persistence in Conditional Variances. *Econometrica,* 61(1), pp. 167-186. https://doi.org/10.2307/2951781

Breedon, F., Chen, L., Ranaldo, A. and Vause, N. (2018). Judgement Day: Algorithmic Trading Around the Swiss Franc Cap Removal. *Bank of England,* Volume Staff Working Paper 711.

Buss, A. and Vilkov, G. (2012). Measuring Equity Risk with Option-Implied Correlations. *Review of Financial Studies,* 25(10), pp. 3113-3140.

Campa, J. M. and Chang, P. H. K. (1998). The Forecasting Ability of Correlations Implied in Foreign Exchange Options. *Journal of International Money and Finance,* 17(6), pp. 855-880. https://doi.org/10.1016/S0261-5606(98)00031-X

Carr, P. and Madan, D. (2002). Toward a Theory of Volatility Trading. *Risk Publications,* Volume Volatility, pp. 417-427.

Carr, P. and Wu, L. (2009). Variance Risk Premiums. *The Review of Financial Studies,* 22(3), pp. 1311-1341.

Castrén, O. and Mazzotta, S. (2005). Foreign Exchange Option and Returns Based Correlation Forecasts. Evaluation and Two Applications. *ECB Working Paper,* Volume 447., pp. 1-48.

Chaboud, A., Chiquoine, B., Hjalmarsson, E. and Vega, C. (2014). Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market. *The Journal of Finance,* 69(5), pp. 2045-2084. https://doi.org/10.1111/jofi.12186

Chaboud, A., Hjalmarsson, E. and Zikes, F. (2021). The Evolution of Price Discovery in an Electronic Market. *Journal of Banking and Finance,* Volume 130.
DOI: 10.1016/j.jbankfin.2021.106171

Chernov, M. (2007). On the Role of Risk Premia in Volatility Forecasting. *Journal of Business and Economic Statistics,* 25(4), pp. 411-426. DOI: 10.1198/073500106000000350

Chernov, M., Haddad, V. and Itskhoki, O. (2024). What do Financial Markets Say About the Exchange Rate?. *National Bureau of Economic Research,* Volume Working Paper 32436.

Chong, J. (2004). Value at Risk from Econometric Models and Implied from Currency Options. *Journal of Forecasting,* 23(8), pp. 603-620. DOI: 10.1002/for.934

Christensen, B. J., & Prabhala, N. R. (1998). The relation between implied and realized volatility. *Journal of Financial Economics*, 50(2), 125–150. https://doi.org/10.1016/S0304-405X(98)00034-8

Ciacci, A., Sueshige, T., Takayasu H., Christensen, K. and Takayasu, M. (2020). The Microscopic Relationships Between Triangular Arbitrage and Cross-Currency Correlations in a Simple Agent Based Model of Foreign Exchange Markets. *Plos one,* 15(6):e0234709.

CLS (2023). *CLS.* [Online]
Available at: https://www.cls-group.com/insights/data-analysis/bank-of-japan-intervention-clsmarketdata-insights/ [Accessed 10 03 2025].

CNB (2025). *CNB.* [Online]
Available at: https://www.cnb.cz/en/financial-markets/foreign-exchange-market/foreign-exchange-market-turnover/ [Accessed 12 08 2025].

Cohen, B. H. and Shin, H. S. (2003). Positive Feedback Trading Under Stress: Evidence from the US Treasury Securities Market. *Bank for International Settlements,* Volume BIS Working Papers 122.
Corsi, F. (2009). A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics*, 7, pp. 174-196.

Cui, Z., Qian, W., Taylor, S. M. and Zhu, L. (2020). Detecting and Identifying Arbitrage in the Spot Foreign Exchange Market. *Quantitative Finance,* 20(1), pp. 119-132.
https://doi.org/10.1093/jjfinec/nbp001

Daníelsson, J. and Love, R. (2006). Feedback Trading. *International Journal of Finance & Economics,* 11(1), pp. 35-53. https://doi.org/10.1002/ijfe.286

Demeterfi, K., Derman, E., Kamal, M. and Zou , J. (1999). *More than You Ever Wanted to Know* About Volatility Swaps.* [Online]
Available at: https://emanuelderman.com/wp-content/uploads/1999/02/gs-volatility_swaps.pdf [Accessed 02 09 2025].

DeMiguel, V., Plyakha, Y., Uppal, R. and Vilkov, G. (2013). Improving Portfolio Selection Using Option-Implied Volatility and Skewness. *Journal of Financial and Quantitative Analysis,* 48(6), pp. 1813-1845.

Dömötör, B. and Havran, D. (2011). *Risk Modeling of EUR/HUF Exchange rate Hedging Strategies.* s.l., ECMS, pp. 269-274.

Drehmann, M. and Suhsko, V. (2022). The Global Foreign Exchange Market in a Higher Volatility Environment. *BIS Quarterly Review,* Volume December, pp. 33-48.

Driessen, J., Maenhout, P. J. and Vilkov, G. (2009). The Price of Correlation Risk: Evidence from Equity Options. *Journal of Finance,* 64(3), pp. 1377-1406. https://doi.org/10.1111/j.1540-6261.2009.01467.x

Dufour, A. and Engle, R. F. (2000). Time and the Price Impact of a Trade. *The Journal of Finace,* 55(6), pp. 2467-2498. https://doi.org/10.1111/0022-1082.00297

Engle, R. F. (2000). The Econometrics of Ultra-High-Frequency Data. *Econometrica,* 68(1), pp. 1-22. https://doi.org/10.1111/1468-0262.00091

Engle, R. F. (2002). Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models. *Journal of Business & Economic Statistics,* 20(3), pp. 339-350. https://doi.org/10.1198/073500102288618487

Engle, R. F. and Granger, C. W. J. (1987). Co-Integration and Error Correction: Representation, Estimation, and Testing. *Econometrica,* 55(2), pp. 251-276. https://doi.org/10.2307/1913236

Engle, R. F. and Patton, A. J. (2004). Impacts of trades in an error-correction model of quote prices. *Journal of Financial Markets,* 7(1), pp. 1-25.

Esposito , M. and Laruccia, E. (1999). *Exchange Rates Statistical Properties Implied in FX Options.* [Online]
Available at: https://group.intesasanpaolo.com/content/dam/portalgroup/repository-documenti/investor-

relations/Contenuti/RISORSE/Documenti%20PDF%20Storici/PDF_studi_eng/CNT-04-000000001D646.pdf [Accessed 22 09 2025].

Evans, M. D. and Lyons, R. K. (2002). Order Flow and Exchange Rate Dynamics. *Journal of Political Economy,* 110(1), pp. 170-180. https://doi.org/10.1086/324391

Fenn, D. J. Howison, S. D., McDonald, M., Williams, S., and Johnson, N. F. (2009). The Mirage of Triangular Arbitrage in the Spot Foreign Exchange Market. *International Journal of Theoretical and Applied Finance,* 12(8), pp. 1105-1123.
DOI: 10.1142/S0219024909005609

Foucault, T., Kozhan, R. and Tham, W. W. (2017). Toxic Arbitrage. *The Review of Financial Studies,* 30(4), pp. 1053-1094.

Gebarowski, R., Oswiecimka, P., Watorek, M. and Drozdz, S. (2019). Detecting Correlations and Triangular Arbitrage Opportunities in the Forex by Means of Multifractal Detrended Cross-Correlations Analysis. *Nonlinear Dynamics,* Volume 98, pp. 2349-2364.
https://doi.org/10.1007/s11071-019-05335-5

Gereben, Á. and Pintér, K. (2005). Devizaopciókból számolt implikált volatilitás: érdemes-e vizsgálni?. *MNB-tanulmányok,* Volume 39.

Green, T. C. (2004). Economic News and the Impact of Trading on Bond Prices. *The Journal of Finance,* 59(3), pp. 1201-1233. https://doi.org/10.1111/j.1540-6261.2004.00660.x

Hagströmer, B. and Menkveld, A. J. (2019). Information Revelation in Decentralized Markets. *The Journal of Finance,* 74(6), pp. 2751-2787. https://doi.org/10.1111/jofi.12838

Hansen, P. R., Huang, Z., & Shek, H. H. (2012). Realized GARCH: A joint model for returns and realized measures of volatility. *Journal of Applied Econometrics*, 27(6), 877–906. https://doi.org/10.1002/jae.1234

Hasbrouck, J. (1991). Measuring the Information Content of Stock Trades. *The Journal of Finance,* 46(1), pp. 179-207. https://doi.org/10.1111/j.1540-6261.1991.tb03749.x

Hasbrouck, J. and Levich, R. M. (2019). FX Liquidity and Market Metrics: New Results Using CLS Bank Settlement Data. . *Manuscript,*
available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2912976

Hoffmann, P. (2014). A Dynamic Limit Order Market with Fast and Slow Ttraders. *Journal of Financial Economics,* 113(1), pp. 156-169. https://doi.org/10.1016/j.jfineco.2014.04.002

Huang, W., Ranaldo, A., Schrimpf, A. and Somogyi, F. (2025). Constrained liquidity provision in currency markets. Journal of Financial Economics, 167 https://doi.org/10.1016/j.jfineco.2025.104028

Ito, T., Yamada, K., Takayasu, M. and Takayasu, H. (2012). Free Lunch! Arbitrage Opportunities in the Foreign Exchange Markets. *National Bureau of Economic Research,* Volume Working Paper 18541.

Johansen, S. (1988). Statistical Analysis of Cointegration Vectors. *Journal of Economic Dynamics and Control,* 12(2-3), pp. 231-254. https://doi.org/10.1016/0165-1889(88)90041-3

Jordá, O. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review,* 95(1), pp. 161-182. https://doi.org/10.1257/0002828053828518

Jorion, P. (1995). Predicting volatility in the foreign exchange market. *The Journal of Finance*, 50(2), 507–528. https://doi.org/10.1111/j.1540-6261.1995.tb04793.x

Kanazawa, K., Sueshige, T., Takayashu, H. and Takayashu, M. (2018). Derivation of the Boltzmann Equation for Financial Brownian Motion: Direct Observation of the Collective Motion of High-Frequency Traders. *Physical Review Letters,* Volume 120. DOI: https://doi.org/10.1103/PhysRevLett.120.138301

Karolyi, G. A. and Stulz, R. M. (1996). Why Do Markets Move Together? An Investigation of U.S.- Japan Stock Return Comovements. *Journal of Finance,* 51(3), pp. 951-986.

Keskin, M., Deviren, B. and Kocakaplan, Y. (2011). Topology of the Correlation Networks among Major Currencies using Hierarchical Structure Methods. *Physica A: Statistical Mechanics and its Applications,* 390(4), pp. 719-730. DOI: 10.1016/j.physa.2010.10.041

King, M. R., Osler, C., and Rime, D. (2012). Foreign exchange market structure, players, and evolution. *Handbook of exchange rates*, pp. 1–44.

Kinkyo, T. (2022). The Intermediating Role of the Chinese Renminbi in Asian Currency Markets: Evidence from Partial Wavelet Coherence. *The North American Journal of Economics and Finance,* Volume 59. DOI: 10.1016/j.najef.2021.101598

Kourtis, A., Markellos, R. N. and Symeonidis, L. (2016). An International Comparison of Implied, Realized and GARCH Volatility Forecasts. *Journal of Futures Markets,* 36(12), pp. 1164-1193.

Kozhan, R. and Tham, W. W. (2012). Execution Risk in High-Frequency Arbitrage. *Management Science,* 58(11), pp. 2131-2149. https://doi.org/10.1287/mnsc.1120.1541

Lee, B. (2010). *MIT.* [Online]
Available at: https://ocw.mit.edu/courses/15-450-analytics-of-finance-fall-2010/511a32446b77d2566dae5d97253e83c9_MIT15_450F10_rec03.pdf

Li, B. and Liao, Z. (2020). Finding Changes in the Foreign Exchange Market from the Perspective of Currency Network. *Physica A: Statistical Mechanics and its Applications,* Volume 545. DOI: 10.1016/j.physa.2019.123727

Longin, F. and Solnik, B. (1995). Is the Correlation in International Equity Returns Constant: 1960–1990?. *Journal of International Money and Finance,* 14(1), pp. 3-26.

Love, R. and Payne, R. (2008). Macroeconomic News, Order Flows, and Exchange Rates. *The Journal of Financial and Quantitative Analysis,* 43(2), pp. 467-488.

Lütkepohl, H. (1991). *Introduction to Multiple Time Series Analysis.* 1 ed. Heidelberg: Springer Berlin.

Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis.* Berlin: Springer Science & Business Media.

Lyons, R. K. (1996). Foreign Exchange Volume: Sound and Fury Signifying Nothing?. *National Bureau of Economic Research,* Volume NBER 11365.

Lyons, R. K. and Moore, M. J. (2009). An Information Approach to International Currencies. *Journal of International Economics,* 79(2), pp. 211-221.

Maldonado, W. L., Egozcue, J. J. and Pawlowsky-Glahn, V. (2021). Compositional Analysis of Exchange Rates. In: A. Daouia and A. Gazen, eds. *Advances in Contemporary Statistics and Econometrics.* Cham, Switzerland: Springer, pp. 489-507. DOI: 10.1007/978-3-030-73249-3_25

Mancini, L., Ranaldo, A., and Wrampelmeyer, J. (2013). Liquidity in the foreign exchange market: Measurement, commonality, and risk premiums. *The Journal of Finance,* 68(5), pp. 1805–1841. https://doi.org/10.1111/jofi.12067

Marshall, B. R., Treepongkaruna, S. and Young, M. . (2007). *Exploitable Arbitrage Opportunities Exist in the Foreign Exchange Market.* Discussion Paper: Massey University, Palmerston North.

Medvegyev, P. and Száz, J. (2010). *A meglepetések jellege a pénzügyi piacokon.* Budapest: Bankárképző.

Menkhoff, L., Sarno, L., Schmeling, M., and Schrimpf, A. (2016). Information flows in foreign exchange markets: Dissecting customer currency trades. *The Journal of Finance*, 71(2), pp. 601–634. https://doi.org/10.1111/jofi.12378

Misik, S. (2023). Korrelációbecslés a forintpiacon. *Közgazdasági szemle,* Volume 7-8, pp. 772-794. https://doi.org/10.18414/KSZ.2023.7-8.772

Misik, S. and Szabó, D. Z. (2026). Triangular Arbitrage, Market Microstructure, and Correlation Dynamics in Global and Regional Currency Trios. *Manuscript,* Available at: *https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7134042*

Mizuno, T., Kurihara, S., Takayasu, M., Takayasu, H. (2004). Time-scale dependence of correlations among foreign currencies. In: Takayasu, H. (eds) *The Application of Econophysics*. Springer, Tokyo. https://doi.org/10.1007/978-4-431-53947-6_3.

MOF (2024). *Ministry of Finance Japan.* [Online]
Available at:
https://www.mof.go.jp/english/policy/international_policy/reference/feio/monthly/20240731e.html [Accessed 30 06 2025].

Mueller, P., Stathopoulos, A. and Vedolin, A. (2017). International Correlation Risk. *Journal of Financial Economics,* 126(2), pp. 270-299. https://doi.org/10.1016/j.jfineco.2016.09.012

NBH (2025). *NBH.* [Online]
Available at: https://statisztika.mnb.hu/publikacios-temak/kamatlabak_-penz-es-tokepiaci-adatok/devizapiaci-forgalmi-statisztikak-/devizapiaci-forgalmi-statisztikak
[Accessed 01 09 2025].

NBP (2022). *NBP.* [Online]
Available at: https://nbp.pl/wp-content/uploads/2022/10/Poland2022.pdf
[Accessed 12 08 2025].

Newey, W. K., and West, K. (1987). A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica,* Volume 55, pp. 703-708. https://doi.org/10.2307/1913610

Norges Bank (2024). *Norway's Financial System,* https://www.norges-bank.no/en/news-events/publications/norways-financial-system/2024-nfs/web-report-2024-nfs/: s.n.

Noureldin, D., Shephard, N., & Sheppard, K. (2012). Multivariate high-frequency-based volatility (HEAVY) models. *Journal of Applied Econometrics*, 27(6), 907–933. https://doi.org/10.1002/jae.1260

Payne, R. (2003). Informed trade in spot foreign exchange markets: an empirical investigation. *Journal of International Economics*, 61(2), pp. 307–329. DOI: 10.1016/S0022-1996(03)00003-5

Pesaran, H. H. and Shin, Y. (1998). Generalized Impulse Response Analysis in Linear Multivariate Models. *Economics Letters,* 58(1), pp. 14-29. https://doi.org/10.1016/S0165-1765(97)00214-0

Ranaldo, A. and Somogyi, F. (2021). Asymmetric information risk in FX markets. *Journal of Financial Economics,* 140(2), pp. 391-411. https://doi.org/10.1016/j.jfineco.2020.10.016

Rebonato, R. (1999). *Volatility and Correlation - In The Pricing of Equity, FX and interest - Rate Options.* s.l.:John Wiley & Sons .

Riksbank (2025). *Riksbank,* https://www.riksbank.se/en-gb/statistics/turnover-statistics/fx-market: s.n.

Rime, D., Sarno, L., and Sojli, E. (2010). Exchange rate forecasting, order flow and macroeconomic information. *Journal of International Economics*, 80(1), pp. 72–88. https://doi.org/10.1016/j.jinteco.2009.11.003

RiskMetrics (1996). *RiskMetrics Technical Document, fourth edition,* New York: J.P.Morgan/Reuters.

Rousseeuw, P. J. and Molenberghs, G. (1994). The Shape of Correlation Matrices. *The American Statistician,* 48(4), pp. 276-279. https://doi.org/10.1080/00031305.1994.10476081

Siegel, A. F. (1997). International Currency Relationship Information Revealed by Cross-Option Prices. *Journal of Futures Markets,* 17(4), pp. 369-384.

Sims, C. A. (1980). Comparison of Interwar and Postwar Cycles: Monetarism Reconsidered. *National Bureau of Economic Research,* Volume Working Papaer 430.

Szüle, B. (2012). Indefinit korrelációs mátrixok korrekciós módszerei. *Statisztikai Szemle,* 90(2-3), pp. 144-164.

Walter , C. and Lopez, J. A. (2000). Is Implied Correlation Worth Calculating? Evidence from Foreign Exchange Option Prices. *Journal of Derivatives,* 7.(3.), pp. 65-82..

Walter, C. A. and Lopez, J. A. (1999). The Sahpe of Things in a Currency Trio. *Federal Reserve Bank of San Francisco ,* Volume Working Paper.

Wu, Z.-X., Gau, Y.-F. and Chen, Y.-L. (2023). Price Discovery and Triangular Arbitrage in Currency Markets. *Journal of International Money and Finance,* 137(C). https://doi.org/10.1016/j.jimonfin.2023.102912

Xu, L. and Kinkyo, T. (2022). Asian Currency Market Integration. *The Kokumin-Keizai Zasshi,　(Journal of Economics & Business Administration),* 225(2), pp. 29-44.

Yamada, K., Tkayasu, H., Ito, T. and Takayasu, M. (2009). Solvable stochastic dealer models for financial markets. *Phisical Review E,* 79(5). https://doi.org/10.1103/PhysRevE.79.051120

Zheng, Y., Osmer, E. and Liu, J. (2014). The Persistency of Correlation between Currency Futures: A Macro Perspective. *International Journal of Economics and Finance,* 6(5), pp. 1-17. DOI:10.5539/ijef.v6n5p17

Zivot, E. and Wang, J. (2006). *Modelling Financial Time Series with S-PLUS.* [Online] Available at: https://faculty.washington.edu/ezivot/econ589/manual.pdf
[Accessed 10 07 2025].

Appendices

Appendix A: descriptive statistics of log returns and correlations

Figure A1: Daily log returns of CEE currency trios

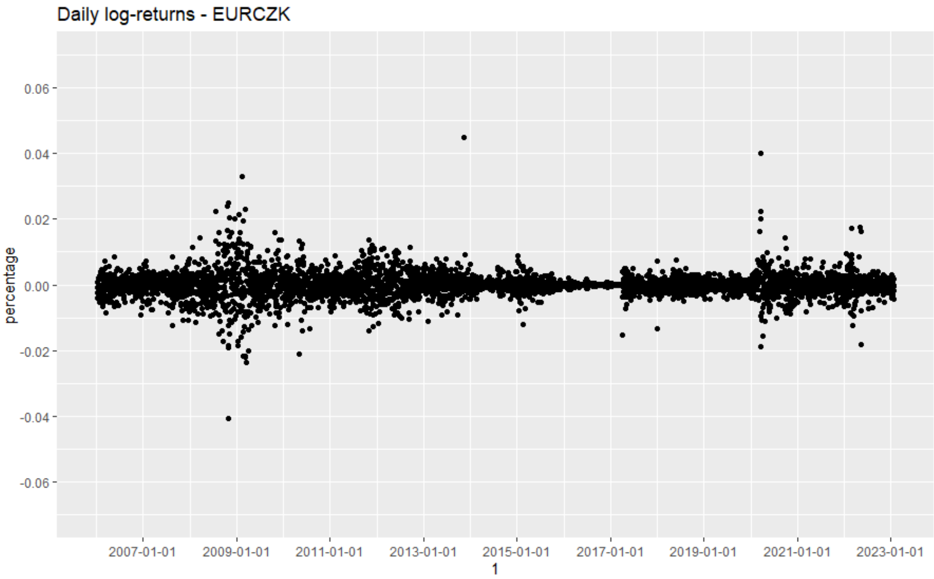

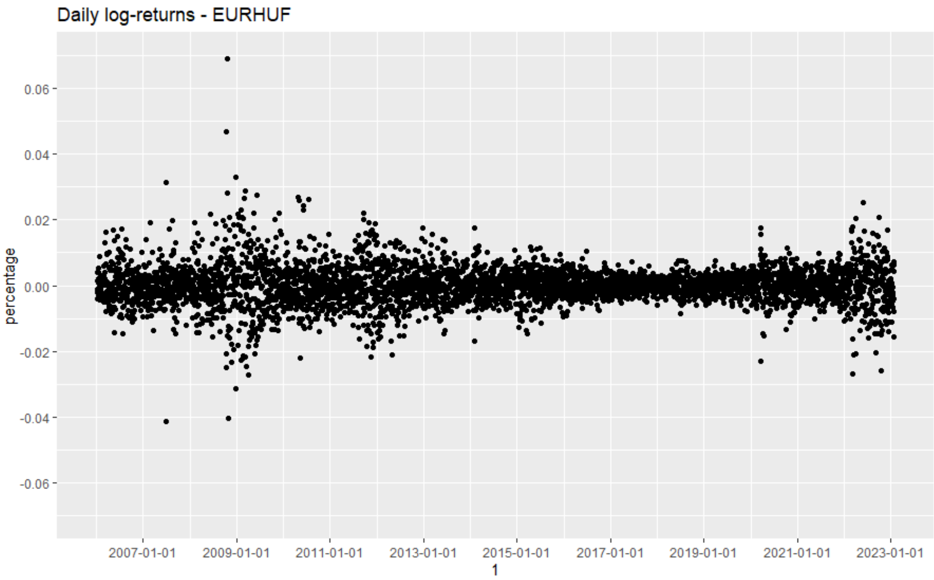

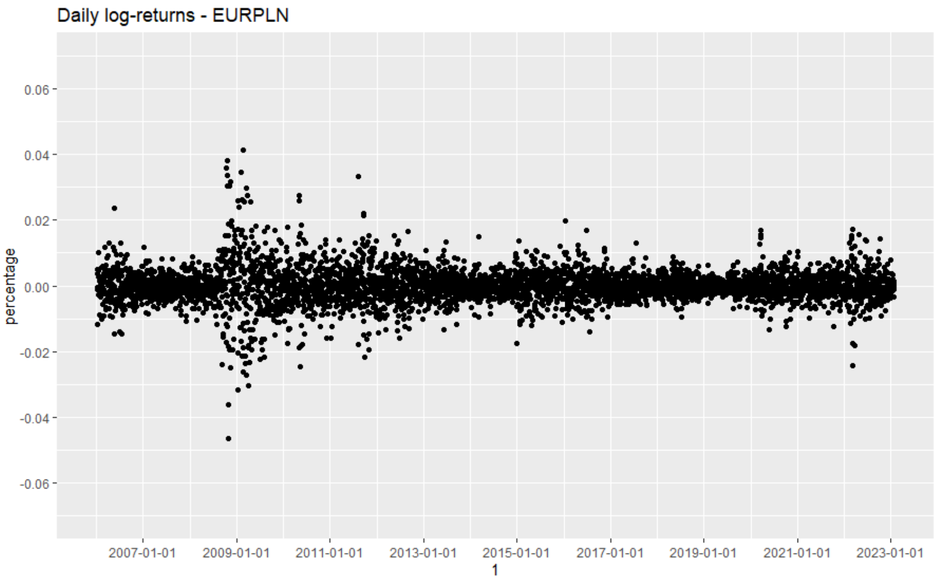
Daily log-returns - EURPLN
percentage
0.06
0.04
0.02
0.00
-0.02
-0.04
-0.06
2007-01-01
2009-01-01
2011-01-01
2013-01-01
2015-01-01
2017-01-01
2019-01-01
2021-01-01
2023-01-01
1

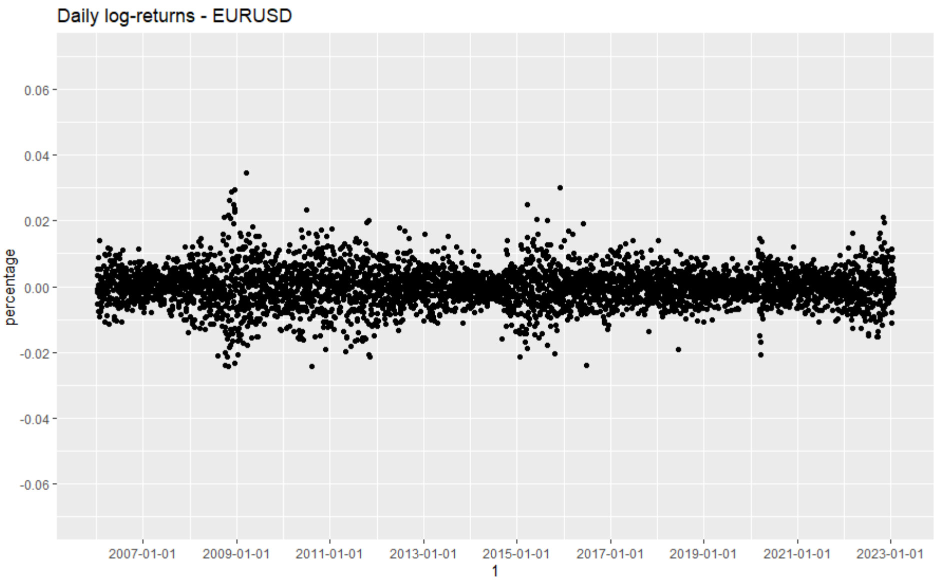
Daily log-returns - EURUSD
percentage
0.06
0.04
0.02
0.00
-0.02
-0.04
-0.06
2007-01-01
2009-01-01
2011-01-01
2013-01-01
2015-01-01
2017-01-01
2019-01-01
2021-01-01
2023-01-01
1

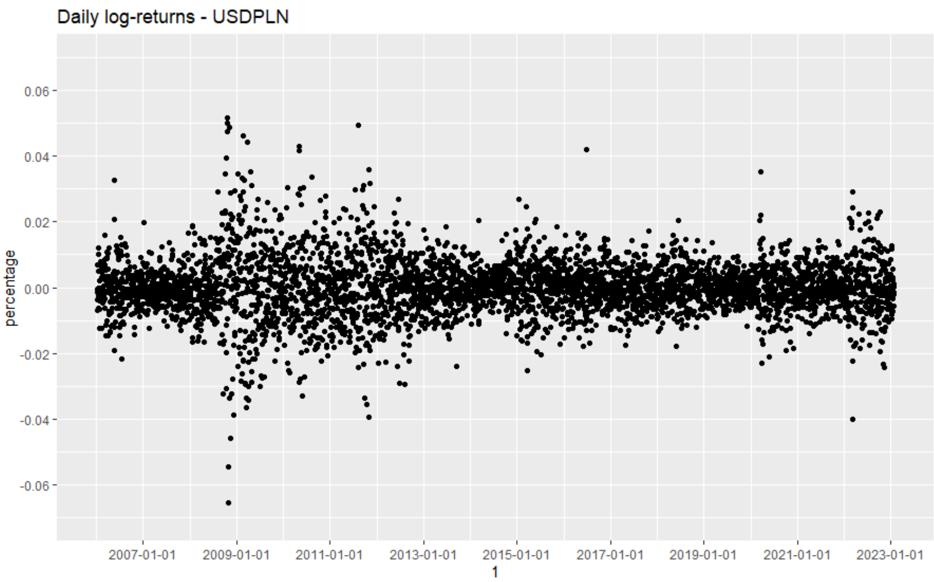
Daily log-returns - USDPLN
percentage
0.06
0.04
0.02
0.00
-0.02
-0.04
-0.06
2007-01-01
2009-01-01
2011-01-01
2013-01-01
2015-01-01
2017-01-01
2019-01-01
2021-01-01
2023-01-01
1

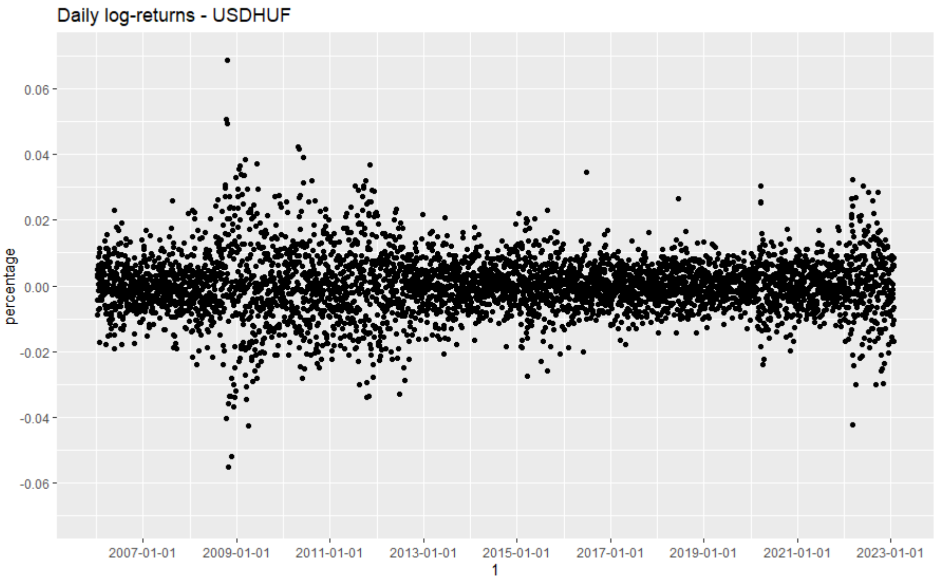

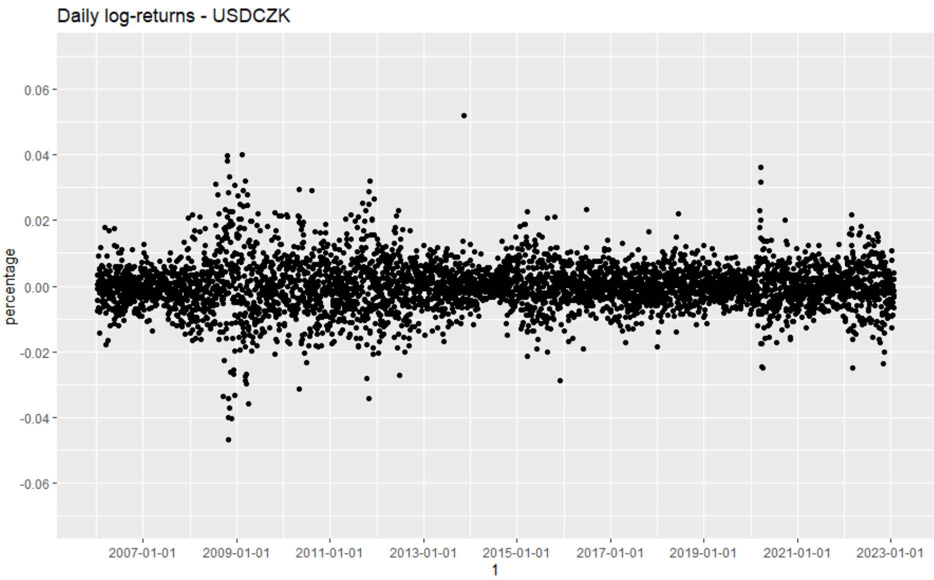

Table A1: Statistics of 12-month correlations

currency trio	pair	min	max	mean	n	Newey-West standard error	t_stat	p_value
HUF	EURHUF_EURUSD	-0,56	0,20	-0,23	5170	0,0106	-21,89	0,0000
	EURHUF_USDHUF	0,37	0,91	0,74	5170	0,0080	92,93	0,0000
	USDHUF_EURUSD	-0,94	-0,58	-0,80	5170	0,0049	-162,51	0,0000
CZK	EURCZK_EURUSD	-0,53	0,19	-0,16	5170	0,0085	-18,40	0,0000
	EURCZK_USDCZK	0,02	0,84	0,58	5170	0,0117	49,19	0,0000
	USDCZK_EURUSD	-1,00	-0,65	-0,88	5170	0,0045	-195,38	0,0000
PLN	EURPLN_EURUSD	-0,64	0,27	-0,22	5170	0,0114	-19,28	0,0000
	EURPLN_USDPLN	0,35	0,91	0,72	5170	0,0074	97,55	0,0000
	USDPLN_EURUSD	-0,91	-0,62	-0,81	5170	0,0041	-197,05	0,0000
NOK	EURNOK_EURUSD	-0,38	0,32	-0,02	5170	0,0108	-2,18	0,0292
	EURNOK_USDNOK	0,39	0,91	0,69	5170	0,0075	91,91	0,0000

	USDNOK_EURUSD	-0,89	-0,31	-0,72	5170	0,0069	-104,71	0,0000
	EURSEK_EURUSD	-0,38	0,40	-0,08	5170	0,0117	-6,76	0,0000
SEK	EURSEK_USDSEK	0,21	0,84	0,65	5170	0,0081	79,86	0,0000
	USDSEK_EURUSD	-0,92	-0,61	-0,80	5170	0,0043	-187,32	0,0000
	EURJPY_EURUSD	0,05	0,78	0,48	5170	0,0116	41,48	0,0000
JPY	EURJPY_USDJPY	0,12	0,84	0,63	5170	0,0088	71,19	0,0000
	USDJPY_EURUSD	-0,68	0,19	-0,35	5170	0,0119	-29,41	0,0000
	EURGBP_EURUSD	-0,12	0,74	0,36	5170	0,0132	26,98	0,0000
GBP	EURGBP_GBPUSD	-0,83	0,02	-0,45	5170	0,0135	-33,11	0,0000
	GBPUSD_EURUSD	0,39	0,86	0,65	5170	0,0067	96,53	0,0000
	EURCHF_EURUSD	-0,30	0,58	0,22	5170	0,0128	17,07	0,0000
CHF	EURCHF_USDCHF	-0,04	0,95	0,42	5170	0,0120	34,60	0,0000
	USDCHF_EURUSD	-0,98	-0,11	-0,76	5170	0,0106	-71,63	0,0000

Source: Bloomberg, Author's own work

Table A2: Statistics of 3-month correlations

currency trio	pair	min	max	mean	n	Newey-West standard error	t_stat	p_value
	EURHUF_EURUSD	-0,72	0,43	-0,23	5368	0,0129	-17,80	0,0000
HUF	EURHUF_USDHUF	0,22	0,96	0,74	5368	0,0093	79,51	0,0000
	USDHUF_EURUSD	-0,96	-0,28	-0,80	5368	0,0064	-125,39	0,0000
	EURCZK_EURUSD	-0,71	0,52	-0,15	5368	0,0117	-12,77	0,0000
CZK	EURCZK_USDCZK	-0,47	0,91	0,55	5368	0,0135	40,64	0,0000
	USDCZK_EURUSD	-1,00	-0,51	-0,88	5368	0,0046	-190,86	0,0000
	EURPLN_EURUSD	-0,77	0,49	-0,21	5368	0,0142	-14,97	0,0000
PLN	EURPLN_USDPLN	-0,04	0,94	0,70	5368	0,0093	75,95	0,0000
	USDPLN_EURUSD	-0,96	-0,50	-0,81	5368	0,0054	-150,86	0,0000
	EURNOK_EURUSD	-0,62	0,59	-0,03	5368	0,0149	-1,96	0,0498
NOK	EURNOK_USDNOK	0,18	0,94	0,68	5368	0,0095	71,19	0,0000
	USDNOK_EURUSD	-0,93	0,07	-0,72	5368	0,0085	-84,65	0,0000
	EURSEK_EURUSD	-0,61	0,53	-0,09	5368	0,0145	-6,07	0,0000
SEK	EURSEK_USDSEK	-0,07	0,91	0,64	5368	0,0103	62,57	0,0000
	USDSEK_EURUSD	-0,95	-0,27	-0,80	5368	0,0060	-132,94	0,0000
	EURJPY_EURUSD	-0,26	0,92	0,47	5368	0,0149	31,68	0,0000
JPY	EURJPY_USDJPY	-0,36	0,91	0,60	5368	0,0122	49,21	0,0000
	USDJPY_EURUSD	-0,85	0,48	-0,37	5368	0,0154	-23,84	0,0000
	EURGBP_EURUSD	-0,58	0,82	0,36	5368	0,0151	23,92	0,0000
GBP	EURGBP_GBPUSD	-0,95	0,20	-0,43	5368	0,0148	-29,19	0,0000
	GBPUSD_EURUSD	0,09	0,91	0,65	5368	0,0090	71,98	0,0000
	EURCHF_EURUSD	-0,53	0,75	0,19	5368	0,0170	11,09	0,0000
CHF	EURCHF_USDCHF	-0,33	0,98	0,38	5368	0,0142	26,53	0,0000
	USDCHF_EURUSD	-1,00	0,12	-0,80	5368	0,0104	-76,74	0,0000

Source: Bloomberg, Author's own work

Appendix B: Impulse-response functions of various FX triangles

Figure B1: IRFs in JPY Trio – bid/ask - 11/07/2024 – intervention – Dukascopy – second aggregation

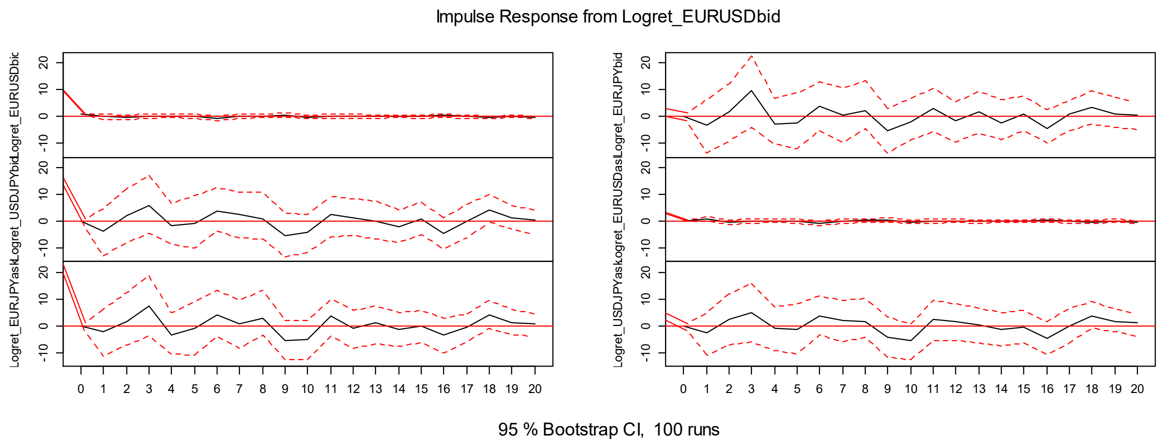

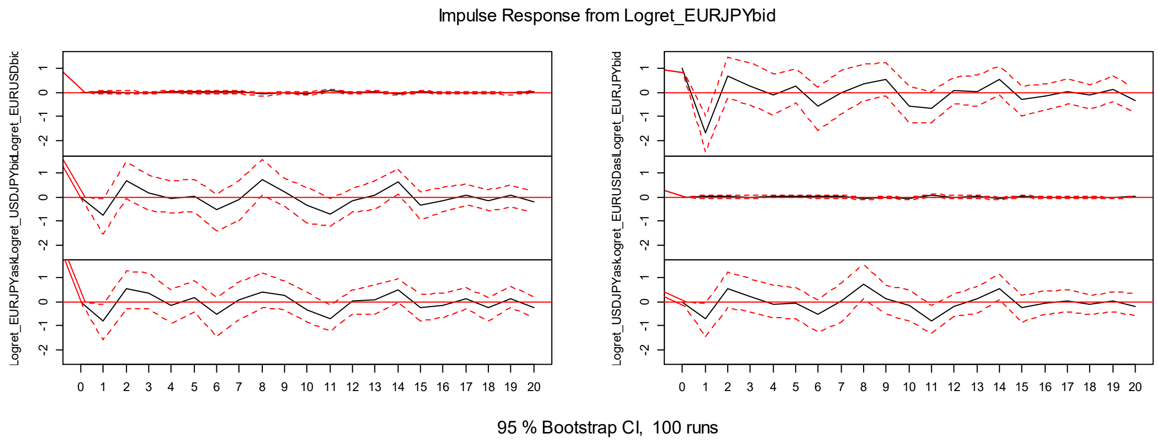

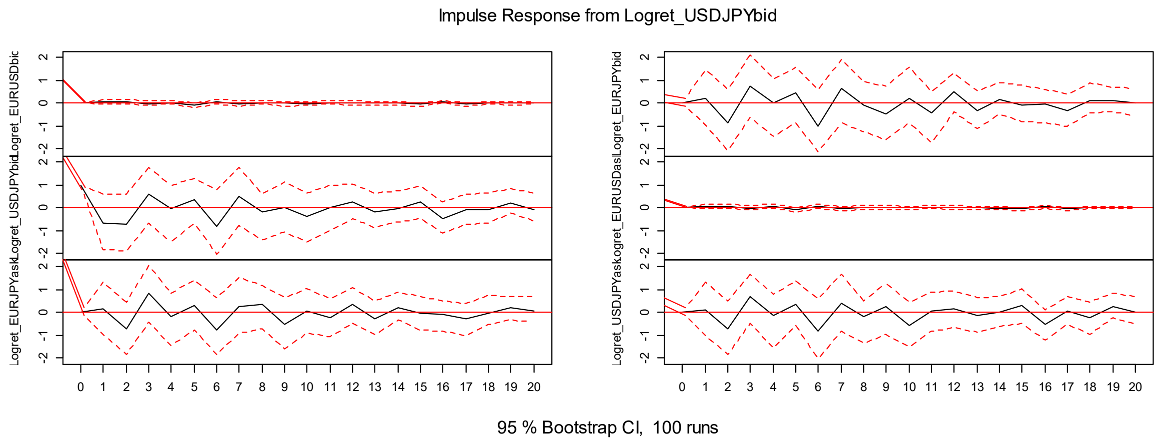

Impulse Response from Logret_EURUSDask

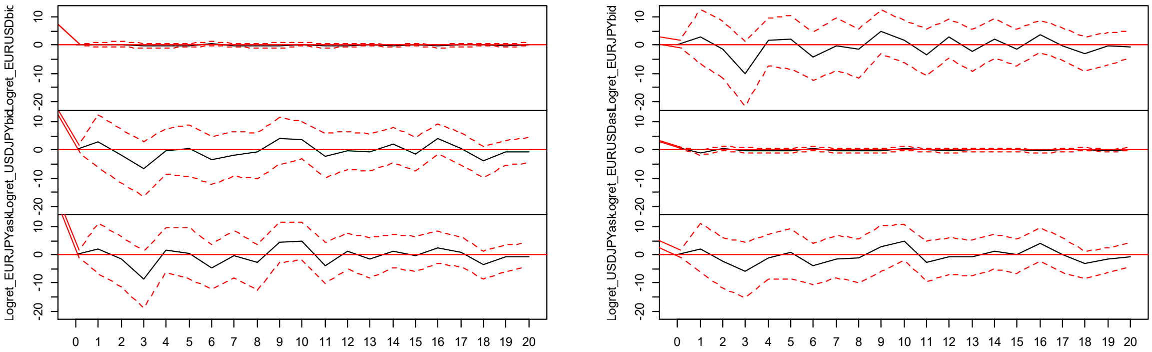

95 % Bootstrap CI, 100 runs

Impulse Response from Logret_EURJPYask

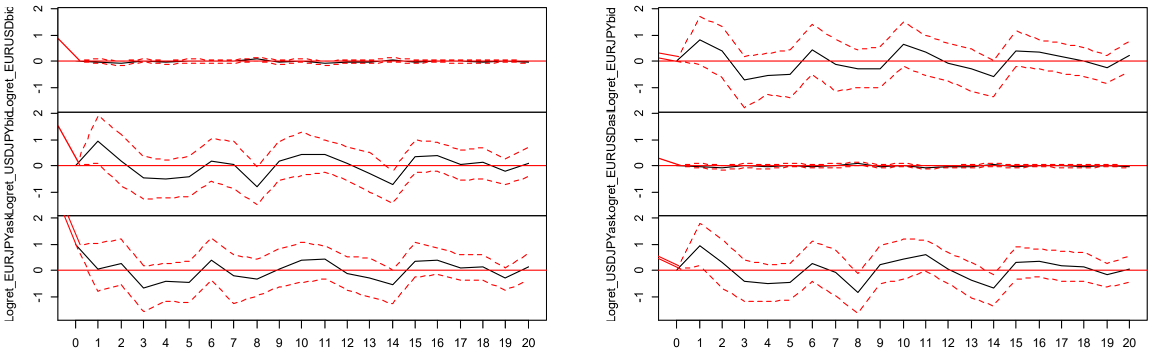

95 % Bootstrap CI, 100 runs

Impulse Response from Logret_USDJPYask

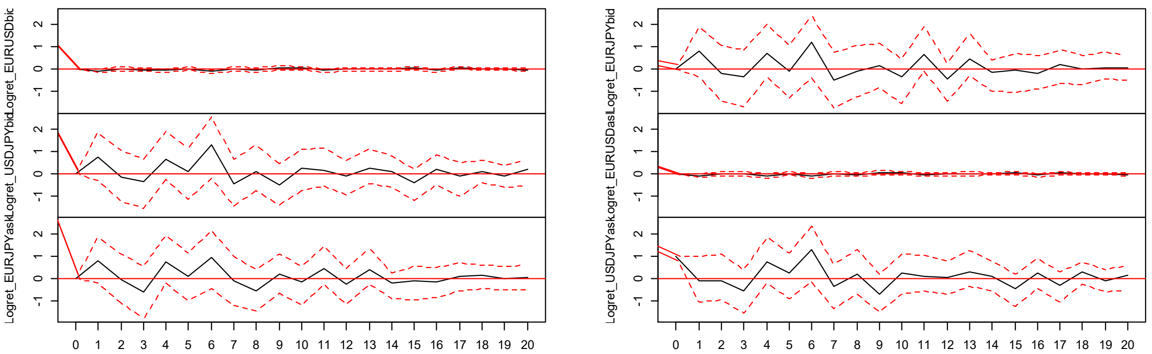

95 % Bootstrap CI, 100 runs

Figure B2: IRF in JPY Trio – bid/ask - 24/10/2022 – pre-intervention

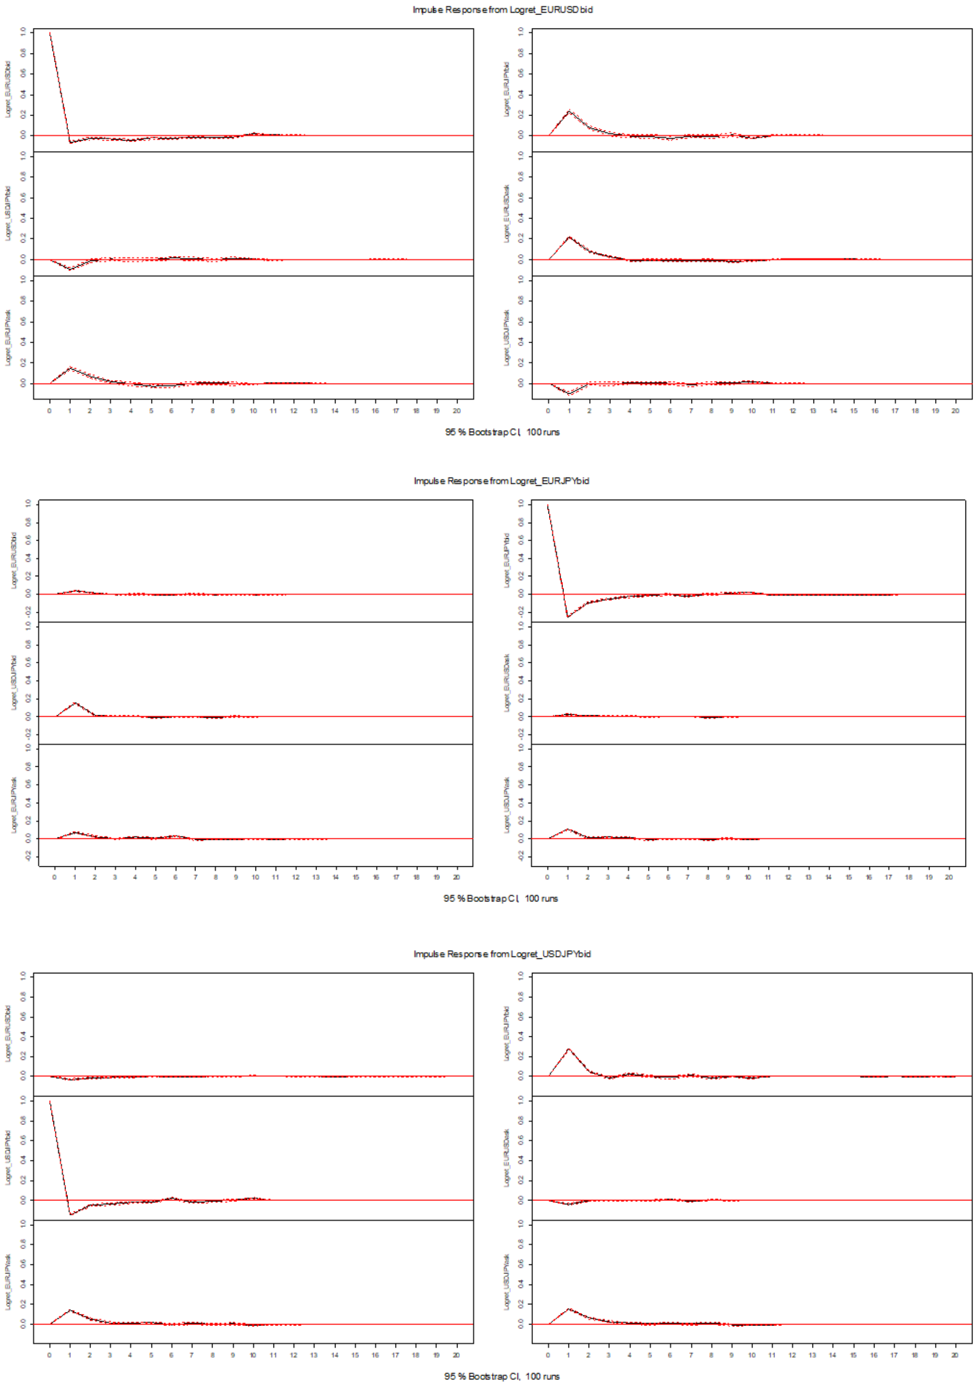

Figure B3: IRF in JPY Trio – bid/ask - 24/10/2022 – intervention

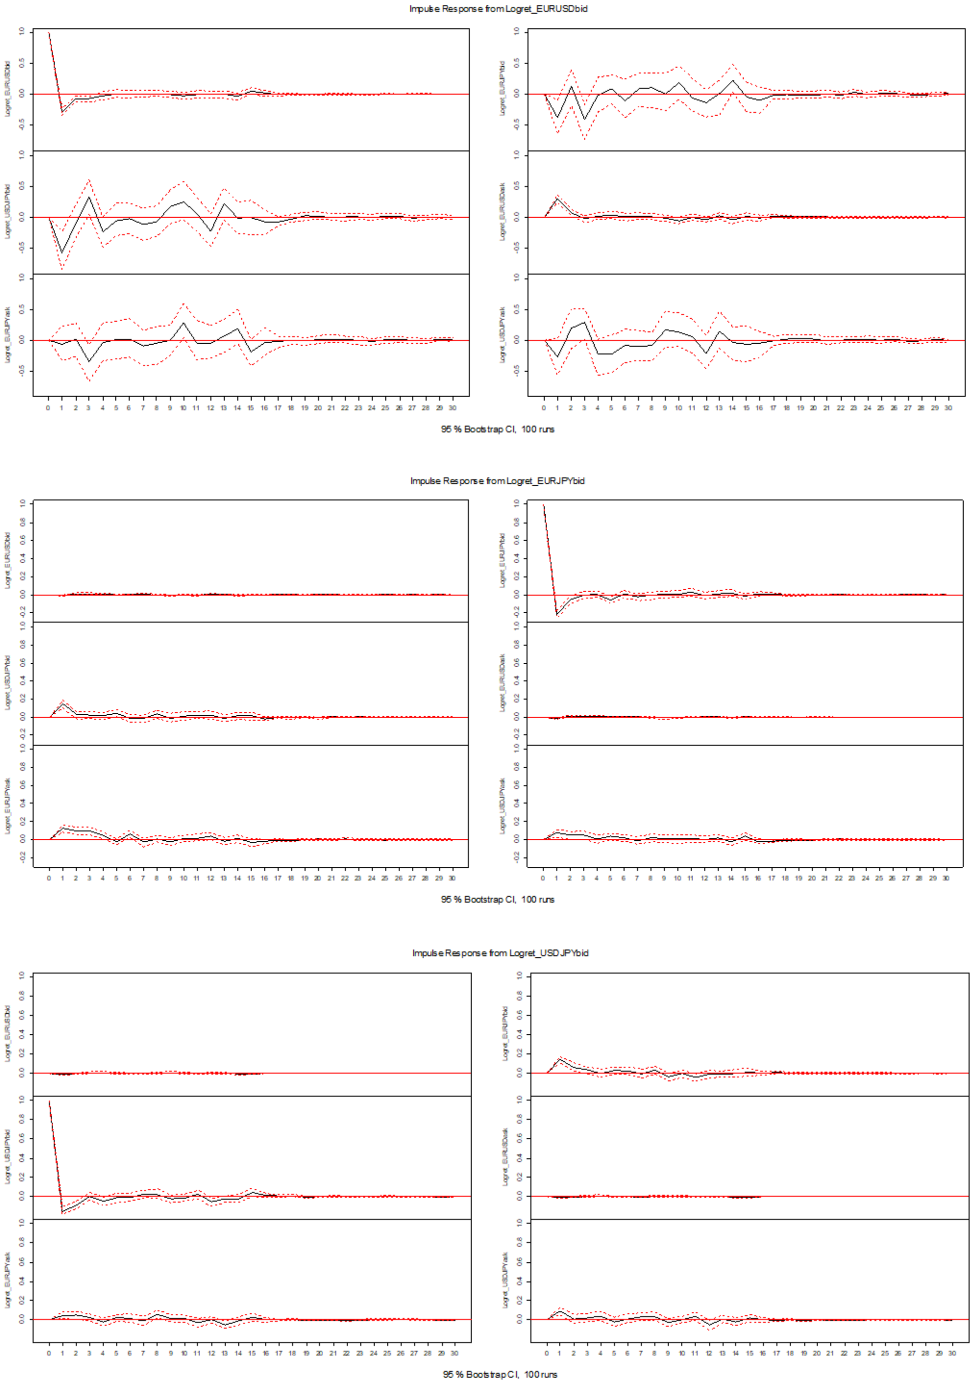

Figure B4: IRF in JPY Trio – bid/ask - 21/10/2022 – intervention

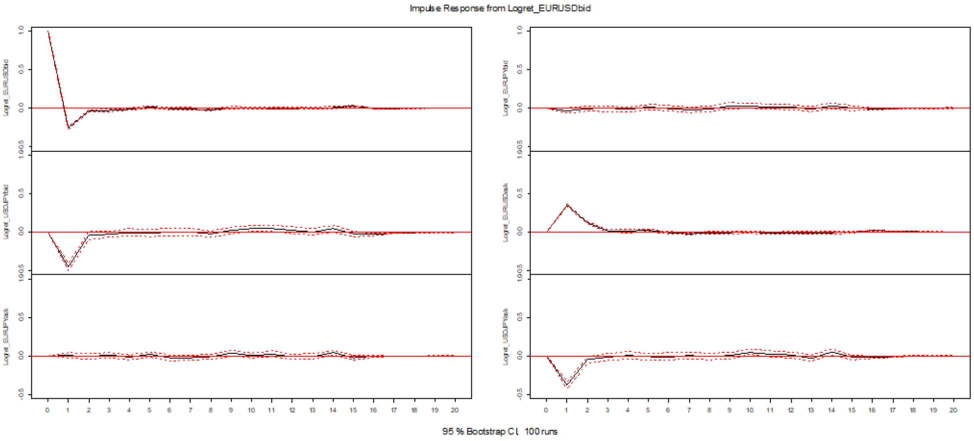

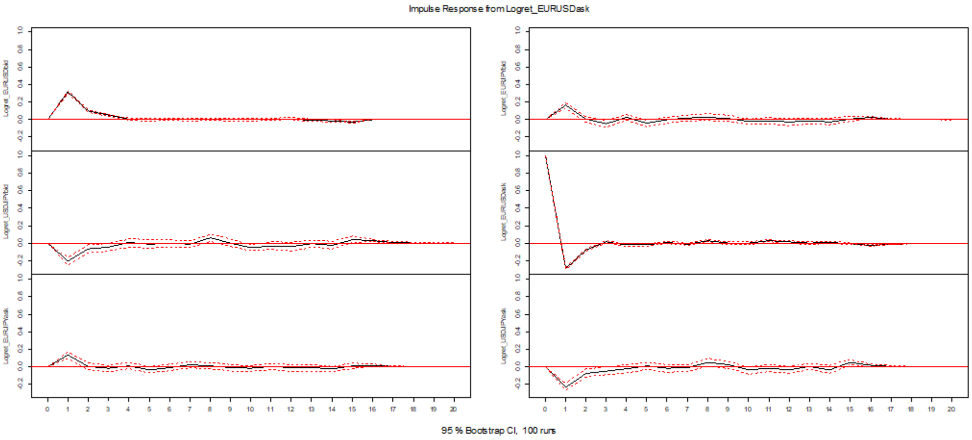

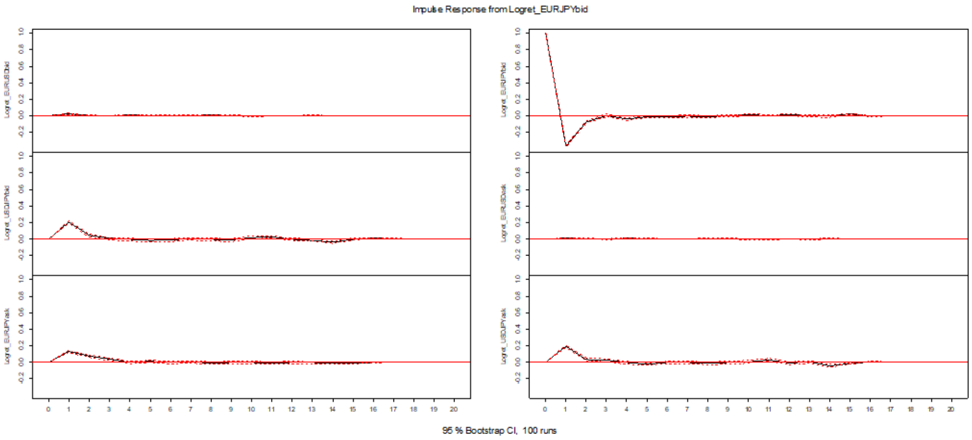

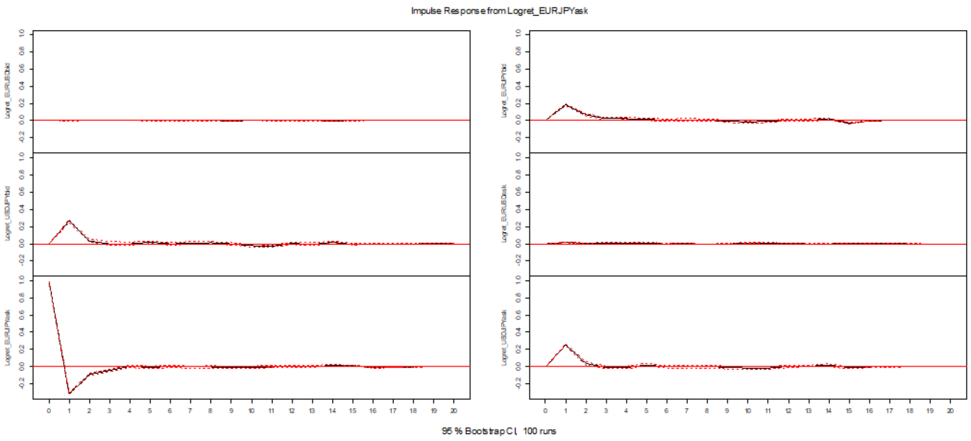
Impulse Response from Logret_EURJPYask
95 % Bootstrap CI, 100 runs

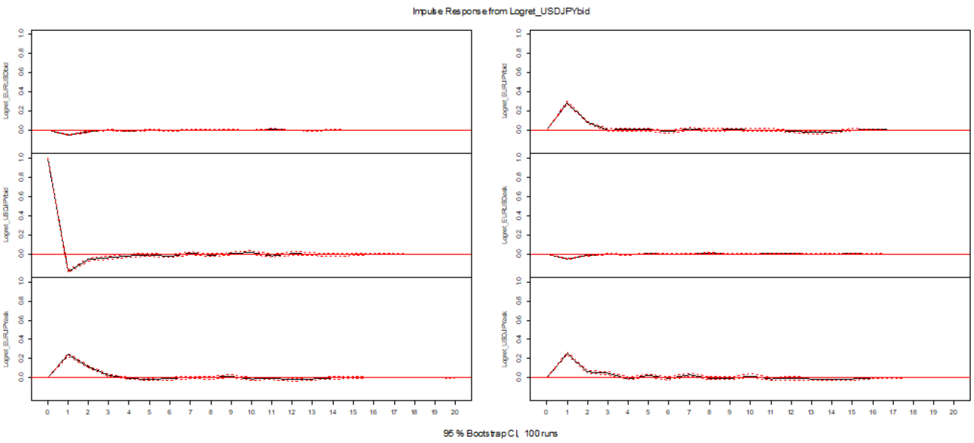
Impulse Response from Logret_USDJPYbid
95 % Bootstrap CI, 100 runs

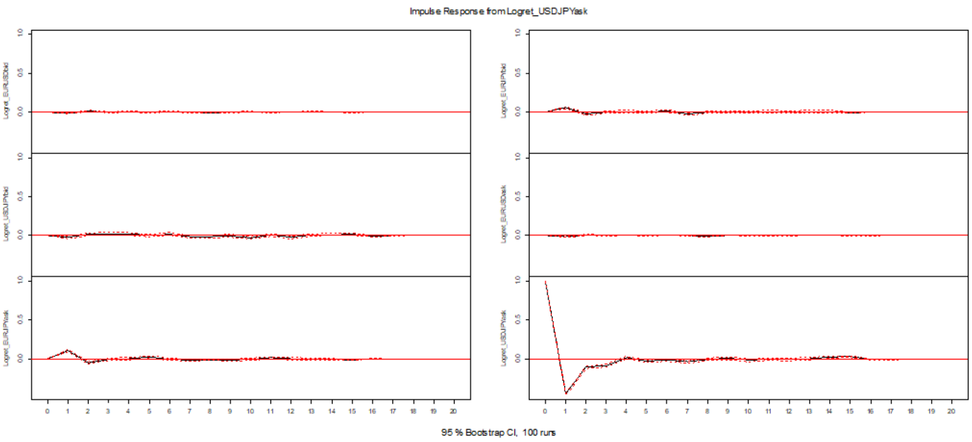
Impulse Response from Logret_USDJPYask
95 % Bootstrap CI, 100 runs

Figure B5: IRF in JPY Trio – bid/ask - 11/07/2024 – intervention - Dukascopy

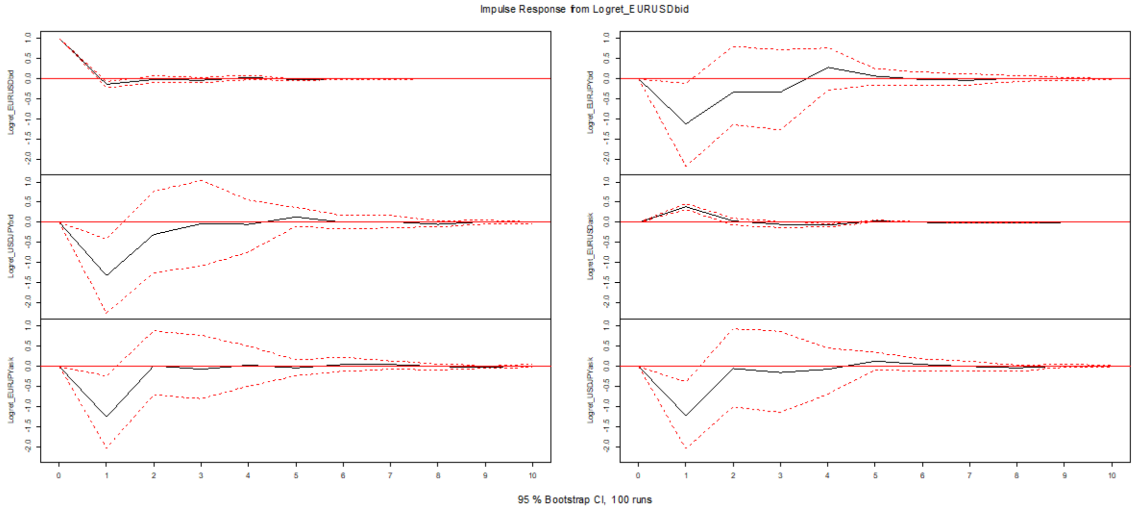

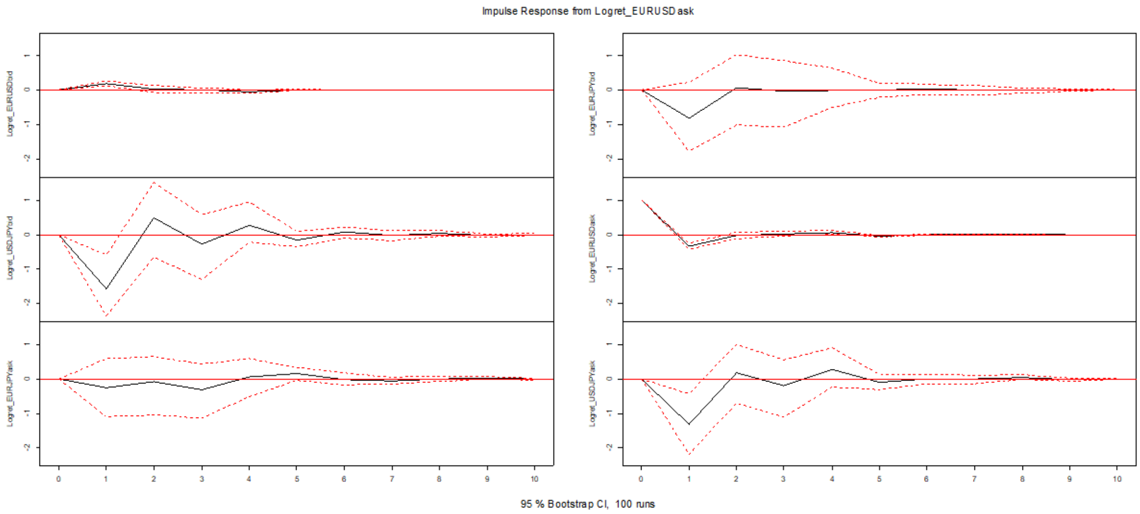

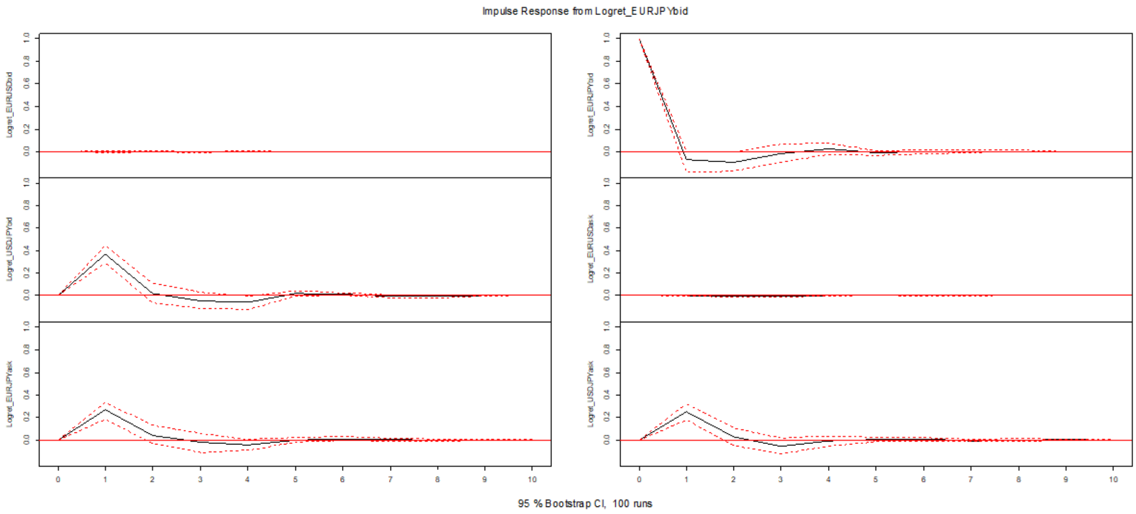

Impulse Response from Logret_EURJPYask

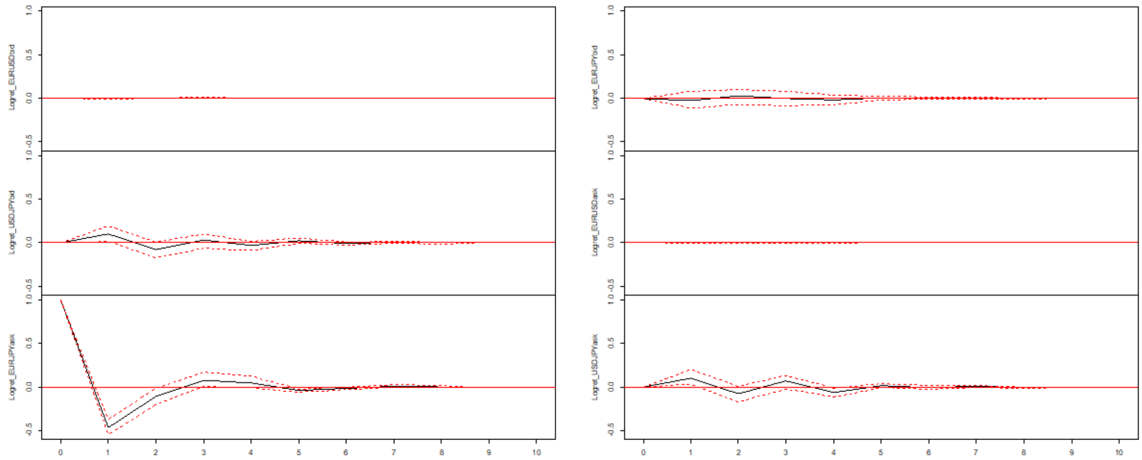

95 % Bootstrap CI, 100 runs

Impulse Response from Logret_USDJPYbid

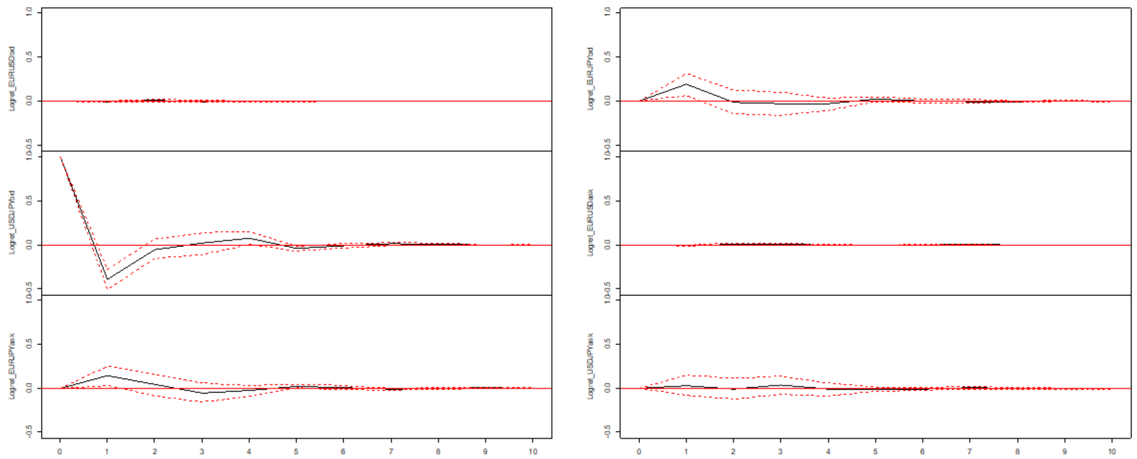

95 % Bootstrap CI, 100 runs

Impulse Response from Logret_USDJPYask

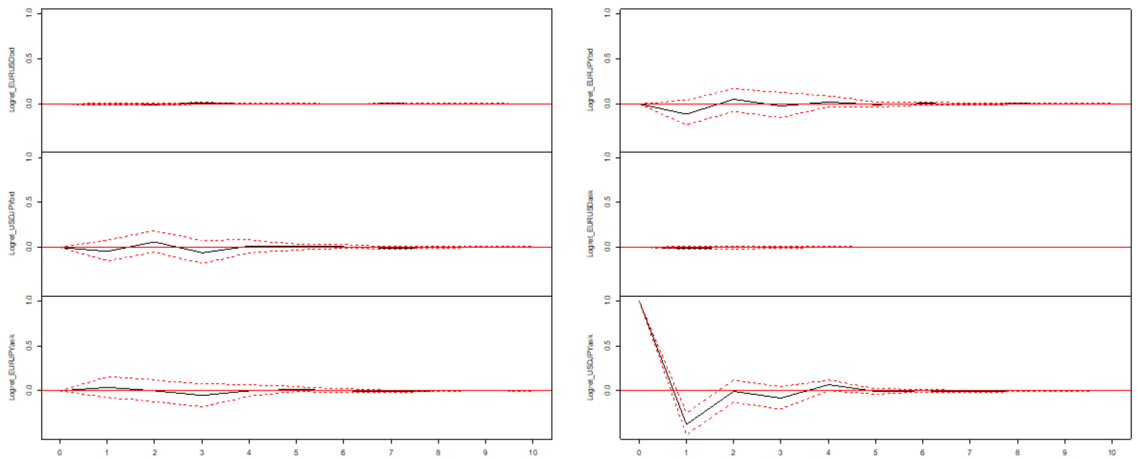

95 % Bootstrap CI, 100 runs

Figure B6: IRF in JPY Trio – bid/ask - 11/07/2024 – intervention - TrueFX

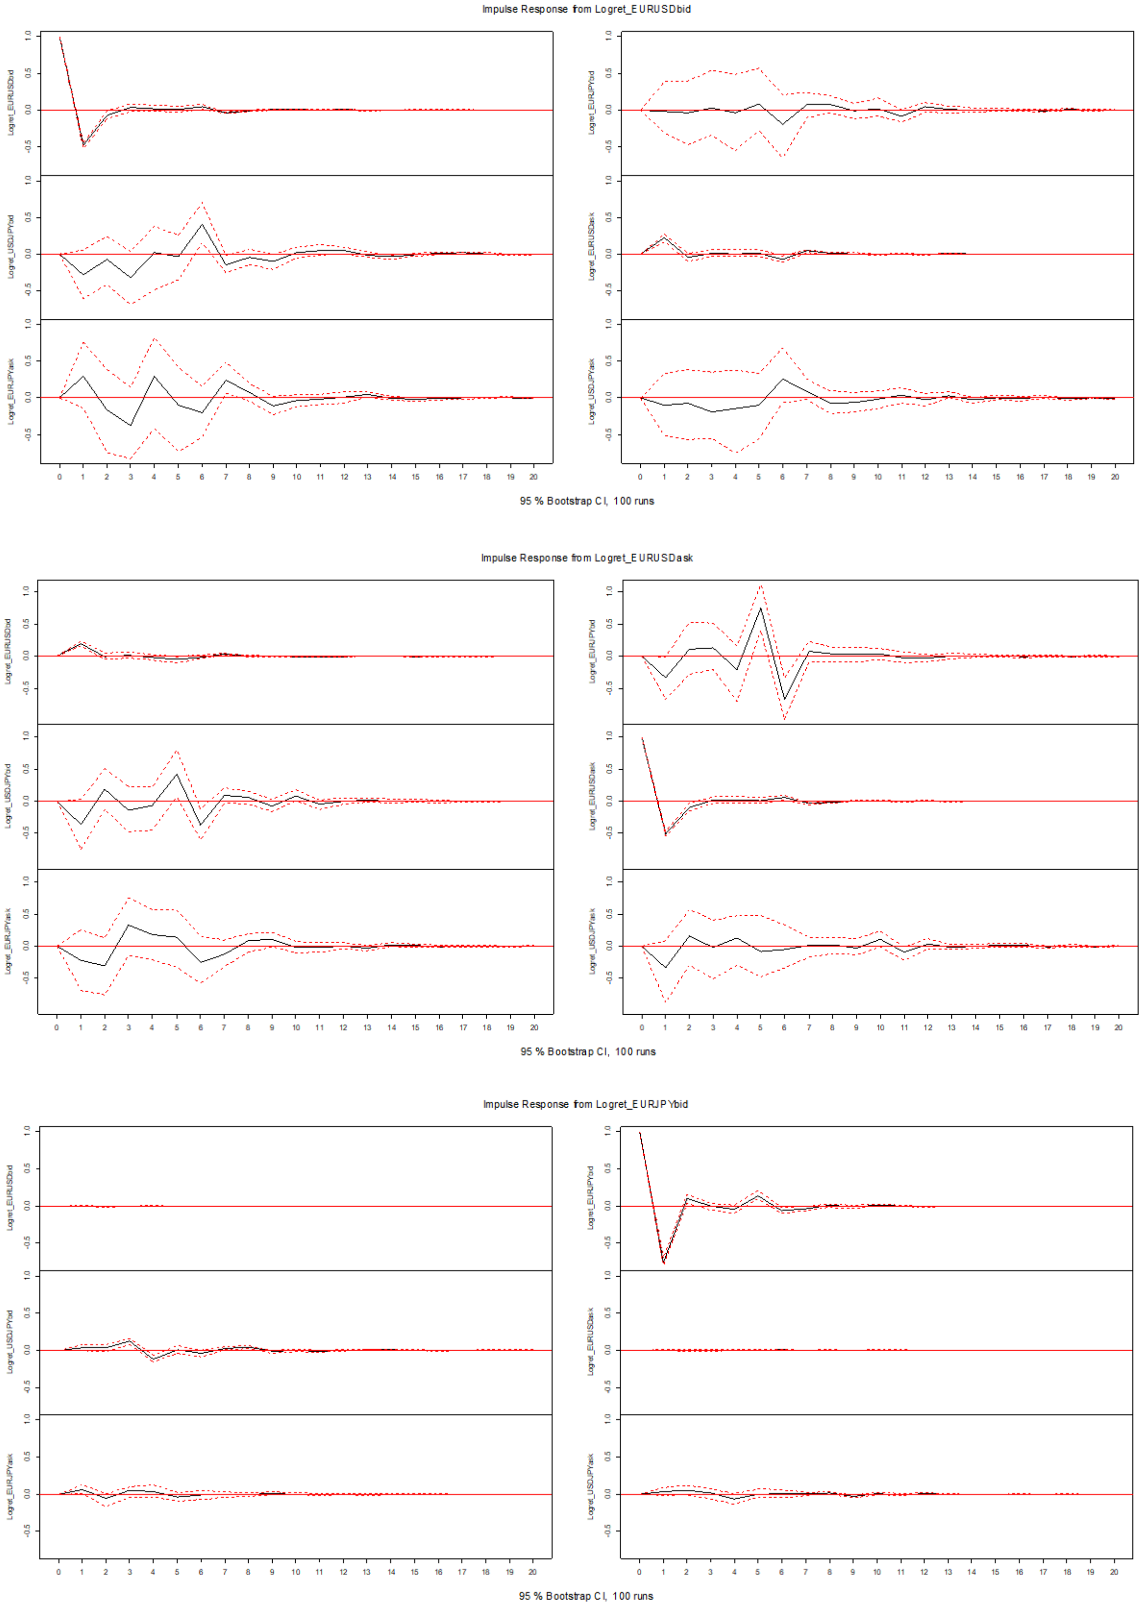

Impulse Response from Logret_USDJPYbid

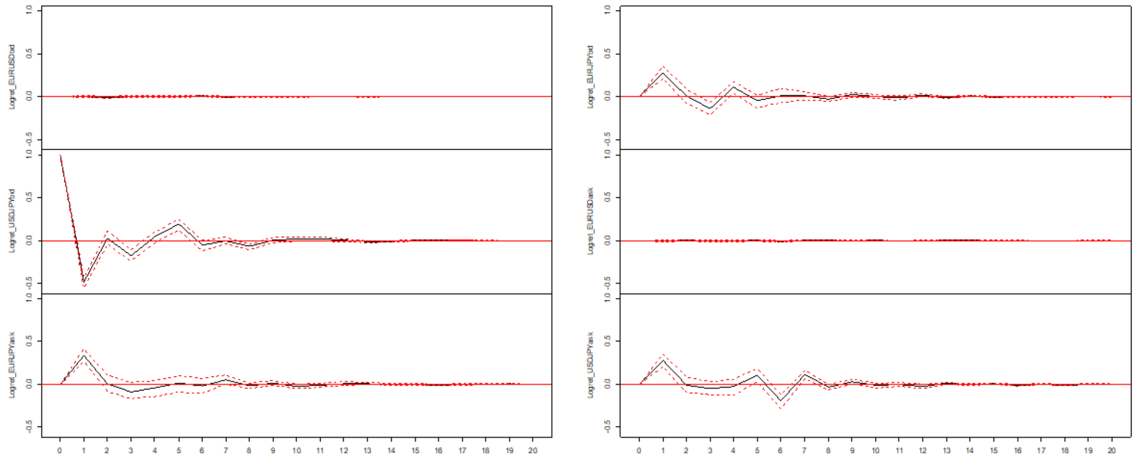

95 % Bootstrap CI, 100 runs

Impulse Response from Logret_EURJPYask

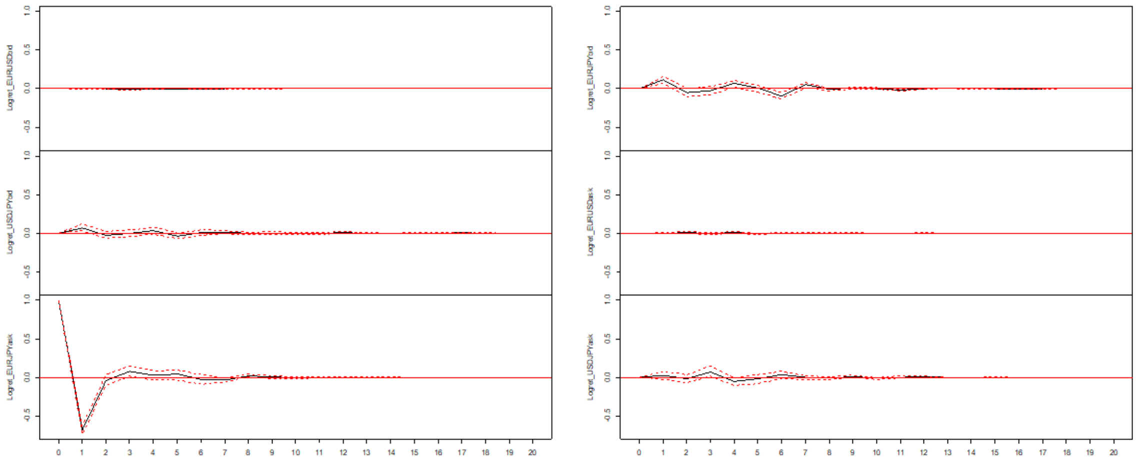

95 % Bootstrap CI, 100 runs

Impulse Response from Logret_USDJPYask

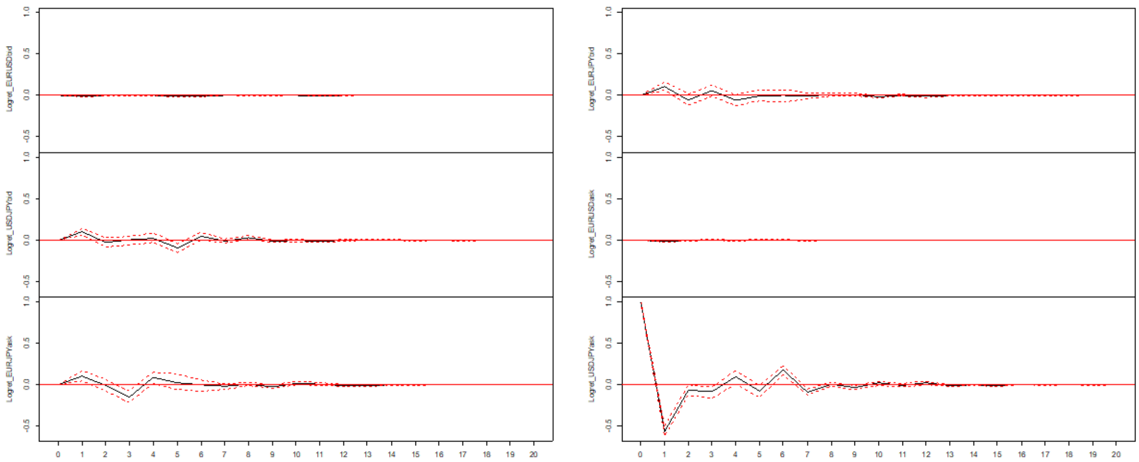

95 % Bootstrap CI, 100 runs

Figure B7: LP IRFs in JPY trio 20/01/2025 Dukascopy – 1st 200 000

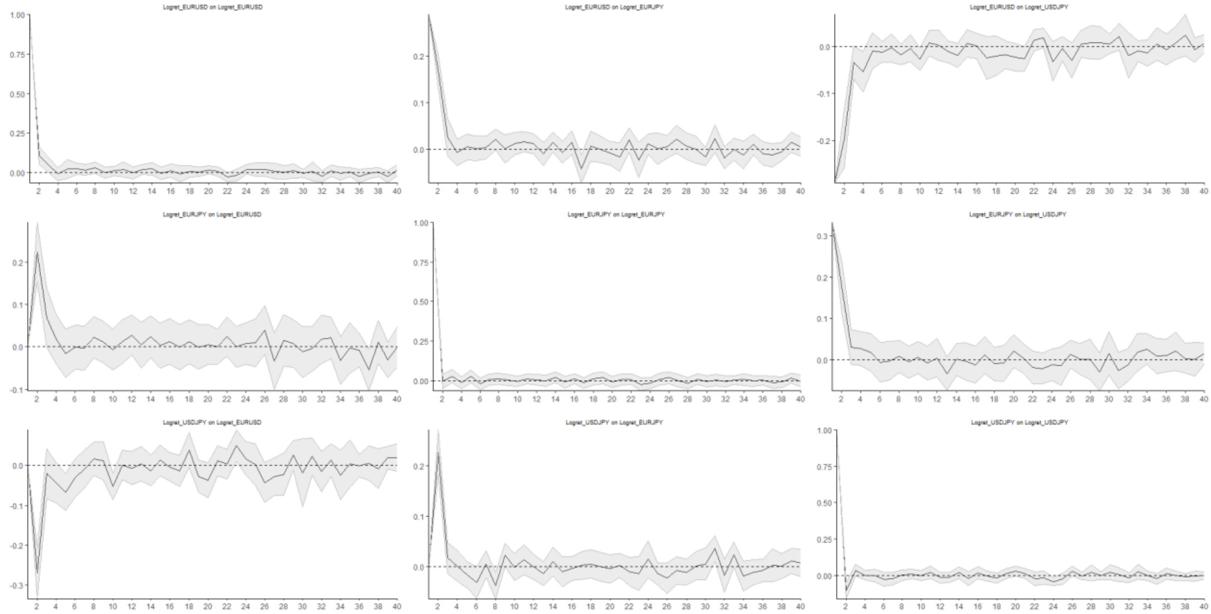

Figure B8: LP IRFs in JPY trio 20/01/2025 Dukascopy – 3rd 100 000 sample

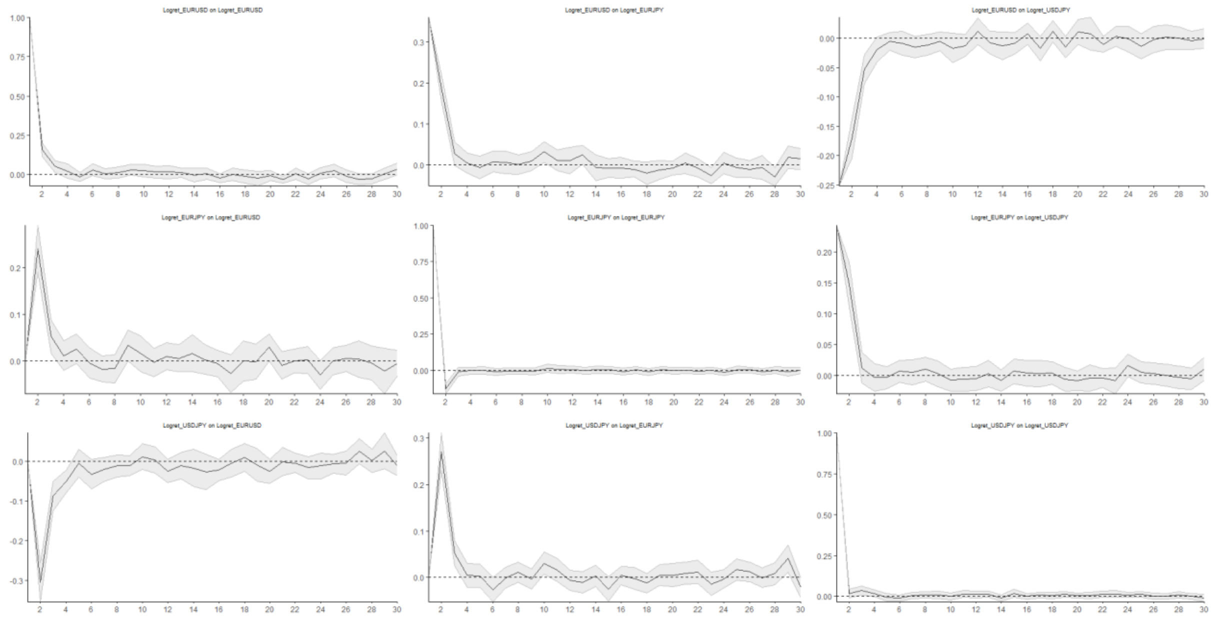

Figure B9: LP IRFs in HUF trio 20/01/2025 Dukascopy – 1st 200 000 sample

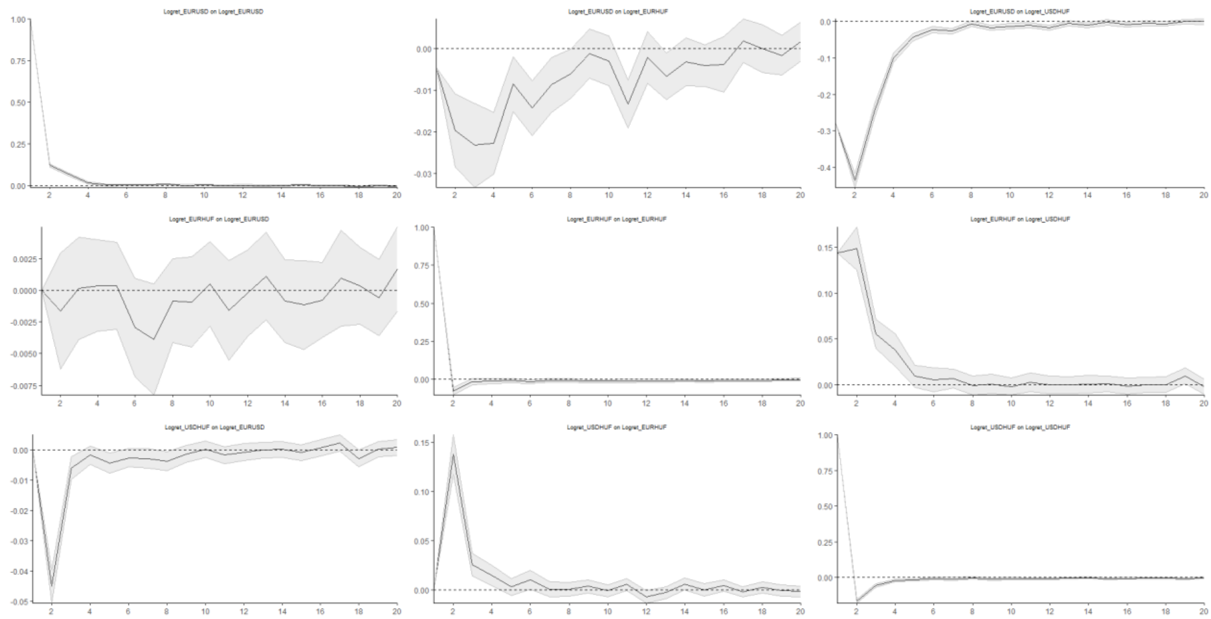

Appendix C: Bank questionnaire on quoting differences among currency trios

Although the econometric analysis conducted on the available ultra-high frequency data provides evidence of differing dynamic relationships within individual currency triangles, the conclusions can only be generalised to a limited extent due to the limited number of data sources and samples. In order to support or refuse the results obtained through quantitative methods, a questionnaire survey was carried out, enabling us to obtain responses from the foreign exchange trading departments of four major international financial institutions. Given the potential sensitivity of the subject, the four institutions requested anonymity, so the responses are presented anonymously. The questionnaire aimed to explore whether, according to quoting banks, there are differences in the nature of quoting among the various currency trios. This, in turn, serves to underpin the potential model variations discussed in Chapter 5, which appear to support the different correlation levels observed in currency trios.

The questions and the answers received are as follows (one bank provided a summary, which is included after the questions without modification):

1. According to your knowledge/experience/practice considering EUR-USD-JPY FX trio:
    a. all three FX rates are quoted independently

b. EURUSD and USDJPY are quoted independently and EURJPY is always quoted based on EURUSD and USDJPY
c. The answer depends on market circumstances – handling of large order flow e.g. (please explain):
d. market works another way (please explain):

Bank 1: *EURUSD / USDJPY and EURJPY are all directly quoted crosses.*

Bank 2: *EURUSD and USDJPY are quoted independently and EURJPY is always quoted based on EURUSD and USDJPY*

Bank 3: *all three FX rates are quoted independently*

Bank 5: *All three pairs are well established with good liquidity – can trade either pair independent of the others*

Bank 6*: EURJPY is very liquid, and one could easily hedge using EURJPY rather than via the crosses, however this would depend on market liquidity and size of flow. For larger transaction (e.g >50mio), most would look to leg out part if not all of the flow, while executing in clips*

2. According to your knowledge/experience/practice considering EUR-USD-GBP FX trio:
   a. all three FX rates are quoted independently
   b. EURUSD and EURGBP are quoted independently and GBPUSD is always quoted based on EURUSD and EURGBP <u>or</u> GBPUSD is independent and EURGBP is indirectly quoted
   c. The answer depends on market circumstances – handling of large order flow e.g. (please explain):
   d. market works another way (please explain):

Bank 1: *EURUSD / EURGBP and GBPUSD are all directly quoted crosses.*

Bank 2: *EURUSD and EURGBP are quoted independently and GBPUSD is always quoted based on EURUSD and EURGBP <u>or</u> GBPUSD is independent and EURGBP is indirectly quoted*

Bank 3: *all three FX rates are quoted independently*

Bank 5: *All three pairs are well established with good liquidity – can trade either pair independent of the others*

Bank 6: *EURGBP is incredibly liquid (arguable more liquid than GBPUSD. In addition, volatility is lower in EURGBP than GBPUSD, due to the correlation between eur and gbp. As such. Most traders, if given a EURGBP exposure, would look to hedge with EURGBP only. Indeed if given GBPUSD exposure, traders may consider hedging using EURGBP and EURUSD alongside GBPUSD*

3. According to your knowledge/experience/practice considering EUR-USD-CEE(HUF, PLN, CZK) FX trio:
    a. all three FX rates are quoted independently
    b. EURUSD and EURCEE are quoted independently and USDCEE is always quoted based on EURUSD and EURCEE
    c. The answer depends on market circumstances – handling of large order flow e.g. (please explain):
    d. market works another way (please explain):

Bank 1: *HUF, PLN & CZK are mainly quoted via the EUR. There is a small amount of direct USDHUF, USDPLN, USDCZK trading (guess, 10% of volume).*

Bank 2: *EURUSD and EURCEE are quoted independently and USDCEE is always quoted based on EURUSD and EURCEE*

Bank 3: *EURUSD and EURCEE are quoted independently and USDCEE is always quoted based on EURUSD and EURCEE*

Bank 5: *All pairs here are frequently quoted and well established however we would typically expect better liquidity on the EUR/Skandi crosses as opposed to vs the USD.*

Bank 6: *EURCEE is much more liquid than USDCEE, benefiting again from the lower vol seen in EURGBP vs GBPUSD. For small flows / limit orders, or v. liquid times, traders may consider USDCEE, however chatting to our trader, he nearly always hedges a USDCEE exposure through the legs (e.g EURCEE and EURUSD)*

4. According to your knowledge/experience/practice considering EUR-USD-Skandie (NOK, SEK,DKK) FX trio:
    a. all three FX rates are quoted independently
    b. EURUSD and EURSkandie are quoted independently and USDSkandie is always quoted based on EURUSD and EURSkandie
    c. The answer depends on market circumstances – handling of large order flow e.g. (please explain):
    d. market works another way (please explain):

Bank 1: *Almost all (if not all) DKK flow is via EURDKK*
Would estimate 70 – 80% of SEK and NOK flow is quoted via the EUR.

Bank 2: *Skandies are different, very low liquidity and very important the locals … sometimes liquidity is in EUR/XXX and some others the best liquidity is in USD/XXX*

Bank 3: *EURUSD and EURSkandie are quoted independently and USDSkandie is always quoted based on EURUSD and EURSkandie*

Bank 5: *All pairs here are frequently quoted and well established however we would typically expect better liquidity on the EUR/CEE crosses as opposed to vs the USD.*

Bank 6: *USDNOK and USDSEK are just as liquid as EURNOK / EURSEK. Therefore, when trading these currency pairs vs the USD, most would trade directly vs USD, rather than legging out via EUR. DKK would be the exception, given it's pegged to EUR, there is much greater* *liquidity vs EUR. All traders, when given USDDKK exposure, would hedge using EURUSD and EURDKK*

5. If there is any difference in quotation practices between different currency trios what circumstances/factors explain them - in you opinion?

Bank 1: *Perhaps the less active CEE currencies have stayed being mainly quoted via the EUR due to lower liquidity.*

Bank 2: *The EUR/XXX is normally the cross and the best liquidity is in the USD legs*

Bank 3: *The main driver as you mention is the traded volume – this is typically driven by trade flows between the countries if there are large trade flows then the pair will be quoted separately (i.e. Europe and UK have large flows, as do Europe and US and US and UK so all three pairs are quoted in their own right), whereas if the trade flows are not significant then the 3rd pair is generally only quoted off the legs.*
I would note however that any triangular arbitrage between pairs will only be exhibited for extremely short periods even for pairs where all 3 legs are quoted of their own right. As arbitragers will move the rates back into equilibrium.

Bank 5: *Not that I'm aware of*

Bank 6: *For all pairs that are hedged via the legs, the bid offer will be somewhere between the max. spread out of the two legs, and the sum spread of the two legs. Where the pricing stands between the two will be driven by perceived correlation. Alongside this, you may find that at times USDxxx may be wider (or tighter), resulting from specific market conditions at the time of trade. Ie if market volatility is being driven by US specific concerns, volatility of non-USD pairs will likely be lower, therefore the spread will be tighter. This is driven more by market sentiment than volumes & liquidity.*

6. Do you see any evidence that the quotation practices to have changed at any currency trio in the past 20 years (referring specifically the change in directly and indirectly quoted FX rates)?

Bank 1: *In past times, EBS and Reuters saw most of the flow. Liquidity in the CME has increased a lot in recent years, and perhaps helped drive an increase in direct USD crosses.*

Bank 2: *GBP/USD has become more liquid than EUR/GBP in the last few years*

Bank 3: *A good example of this is after the CHF de-peg on 15th January 2015. Following this event the market went from being quoted mainly via EURCHF (which had been very stable prior to the event) (~ 80:100 USDCHF:EURCHF volume ratio) to being dominated by USDCHF volumes (~ 180:100 today).*

Jan 2015 – EURCHF collapses after SNB intervention to support the pair ends

Bank 5: *I wouldn't expect much material change other than increased liquidity over the years as technology and access to liquidity improves.*

Bank 6: *Anecdotally it seems as though there is better liquidity in non USD crosses, specifically G7. This may be due to many more brokers (electronic vs voice & EBS/Reuters that would have been used 20yrs ago), much higher e-flows, and larger internalization from banks, meaning they are more likely to go to market with the remaining cross post internalization. In addition, the growth of dark mid-pools makes it relatively easy for banks to leave offers / bids in crosses with limited market impact, again likely improving liquidity. However, it remains that most banks would hedge currencies in the most liquid crosses.*

Bank 4:
Not all FX Triangles are created equal:

*While FX pairs within a triangle are mathematically linked (e.g. EURJPY = EURUSD * USDJPY), this relationship doesn't translate into identical correlation behaviour across triangles.*
In triangles where all three pairs are actively traded (e.g. EURUSD, USDJPY, and EURJPY), each leg responds to its own macro drivers, capital flows, rate differentials, market sentiment etc.
In triangles where one leg is structurally less liquid, that leg tends to follow pricing in the more liquid legs. In these pairs, triangular arbitrage becomes the dominant driver, and there is less room for independent price action that translates into stronger correlations.
*Irrespective of whether pairs in a triangle are quoted independently, or derived from the legs, triangular arbitrage enforces a mathematical constraint (e.g. EURJPY = EURUSD * USDJPY) if this does not hold arbitrageurs would step in and instantaneously close any price misalignments).*
Examples:
EUR-USD-JPY // EUR-USD-GBP // EUR-USD-Skandie
all three FX rates are quoted independently

- *These triangles are made up of deeply liquid, independently quoted pairs all pairs are listed on primary liquidity platforms.*
- *All legs in these triangles exhibit their own supply/demand dynamics, and thus observed correlations between the legs are not purely mechanical.*
- *Dealers typically cover in the cross itself - unless they have a directional bias in one of the legs.*

EUR - USD - CEE:
market works another way:

- *This triangle exhibits a clear asymmetry in liquidity roughly 80% of the CEE liquidity is concentrated in EUR based pairs.*
- *USDCEE are still independently quoted, but as liquidity is thinner, independent price action is constrained, and therefore correlations to their EUR based counterparts are high.*
- *The price in USDCEE crosses largely reflects its need to align with its more liquid EURCEE counterparts, given triangular arbitrage dynamics.*

Bottom Line:

- *When all three pairs in an FX triangle are quoted independently, each will reflect its own price drivers, and thus correlations of pairs within FX triangles can vary. When one leg's liquidity is meaningfully lower, its price behaviour is relatively less important, and thus correlations become more mechanically consistent.*

Other contributing factors:

- *Platform listing: when a cross is quoted on a major liquidity provider, like EBS or Reuters, it often has very good liquidity, given natural interbank flow. While not a hard rule, this tends to correlate with independent price formation rather than mechanical derivation from legs.*
- *Regional trade patterns: i.e. roughly 80% of the CEE liquidity is concentrated in EUR based pairs.*
- *Trade behaviour: irrespective of liquidity, a dealer may cover a position in the legs as opposed to the crosses if they have a directional bias or for execution advantages.*

Evolution of quotation practices:

- *Electronic market making and algorithmic arbitrage has virtually zeroed out price inefficiencies (situations whereby AAABBB ≧ AAACCC * CCC*BBB).*

Increased FX volumes mean many crosses that may have been derived from the legs are now independently quoted, given the depth of liquidity now in these pairs.

Appendix D.: Derivation of correlations in the three specific cases: CEE/Scandinavian, USD-driven and symmetric currency trio

The thesis started with some math that described the specifics of currency correlations considering the triangular arbitrage relationship. As a major finding of the thesis based on VAR modelling, qualitative research and the ABM model we can assume that in many cases there are two independently quoted FX rates in the currency trios and the third is quoted based on the first two.

As seen in chapter 2.2, considering $S_1 = EURHUF$, $S_2 = USDHUF$, $S_3 = EURUSD$ as spot FX rates we can express S_2 as $S_1/S3$ meaning $USDHUF = \frac{EURHUF}{EURUSD}$. Based on Ito's Lemma one can derive the change of S_2 and rate of change of S_2 of as follows (based on Lee (2010)):

$$dS_{2,t} = S_{2,t}(\mu_1 - \mu_3 - \rho_{1,3}\sigma_1\sigma_3 + \sigma_3^2)dt + S_{2,t}\sigma_1 dw_t^1 - S_{2,t}\sigma_3 dw_t^3 \quad (1)$$

$$\frac{dS_{2,t}}{S_{2,t}} = (\mu_1 - \mu_3 - \rho_{1,3}\sigma_1\sigma_3 + \sigma_3^2)dt + \sigma_1 dw_t^1 - \sigma_3 dw_t^3 \quad (2)$$

where μ_i is the drift of the i-th currency rate, σ_i is the volatility of the i-th currency rate and w_t^i is the realization of Brownian motion. Considering the quadratic variation as

$$\left(\frac{dS_{2,t}}{S_{2,t}}\right)^2 = (\sigma_1 dw_t^1 - \sigma_3 dw_t^3)^2 = \left(\sigma_1{}^2 + \sigma_3{}^2 - \rho_{1,3}\sigma_1\sigma_3\right)dt \tag{3}$$

we can derive the volatility of $\frac{dS_{2,t}}{S_{2,t}}$:

$$\sigma_2 = \sqrt{\sigma_1{}^2 + \sigma_3{}^2 - 2\rho_{1,3}\sigma_1\sigma_3} \tag{4}$$

Let's assume the correlation between EURHUF and EURUSD to be zero:

$$\sigma_2 = \sqrt{\sigma_1{}^2 + \sigma_3{}^2} \tag{5}$$

In a currency set where the volatility and correlation between two independently traded currencies are known we can express the pairwise correlations of the third currency as follows:

$$\rho_{1,2} = \frac{\sigma_1^2 + \sigma_2^2 - \sigma_3^2}{2\sigma_1\sigma_2} \tag{6}$$

Utilizing the relationship between volatilities:

$$\rho_{1,2} = \frac{\sigma_1^2 + \sigma_1{}^2 + \sigma_3{}^2 - \sigma_3^2}{2\sigma_1\sqrt{\sigma_1{}^2 + \sigma_3{}^2}} = \frac{\sigma_1}{\sqrt{\sigma_1{}^2 + \sigma_3{}^2}} \tag{7}$$

similarly

$$\rho_{2,3} = \frac{-\sigma_3}{\sqrt{\sigma_1{}^2 + \sigma_3{}^2}} \tag{8}$$

Based on the derivation we can say if two FX rates can be considered as independent, the correlations with the third currency rate will be driven solely by the volatility ratios. In case the volatilities are similar, we will find correlations as $\rho_{1,2} = 0.707$ and $\rho_{2,3}$ =-0.707 – values being fairly close to those what was modelled with the modified ABM in Chapter 5.5.

Let's assume now the EURHUF is quoted as the multiple of EURUSD and USDHUF: $S_1 = S_2 \times S_3$, and EURUSD and USDHUF are independent thus the correlation between them is zero. In this particular case Ito's Lemma would lead to:

$$\sigma_1 = \sqrt{\sigma_2{}^2 + \sigma_3{}^2 + 2\rho_{2,3}\sigma_2\sigma_3} \tag{9}$$

and since $\rho_{2,3} = 0$ we again derive:

$$\sigma_1 = \sqrt{\sigma_2{}^2 + \sigma_3{}^2} \tag{10}$$

Considering the correlation when the multiplication is applied:

$$\rho_{1,2} = \frac{\sigma_1^2+\sigma_2^2-\sigma_3^2}{2\sigma_1\sigma_2} \tag{11}$$

$$\rho_{1,2} = \frac{2{\sigma_2}^2}{2\sqrt{{\sigma_2}^2+{\sigma_3}^2}\sigma_2} = \frac{\sigma_2}{\sqrt{{\sigma_2}^2+{\sigma_3}^2}} \tag{12}$$

and

$$\rho_{1,3} = \frac{\sigma_3^2+\sigma_1^2-\sigma_2^2}{2\sigma_1\sigma_3} \tag{13}$$

resulting

$$\rho_{1,3} = \frac{2{\sigma_3}^2}{2\sqrt{{\sigma_2}^2+{\sigma_3}^2}\sigma_3} = \frac{\sigma_3}{\sqrt{{\sigma_2}^2+{\sigma_3}^2}} \tag{14}$$

Again, if the two volatilities of the two independently traded currencies are identical, we will find $\rho_{1,2} = 0.707$ and $\rho_{1,3}$ =0.707 close to the convergence values of the ABM model applied when USDCCY rate is considered to be independently traded.

Considering the symmetric case we consider all three currency pairs are quoted independently, and they are bound together by the arbitrage. In this case we cannot assume any of the FX rates are always the ratio/multiple of the two others (this will be satisfied at a long enough time horizon). We can assume however in a symmetric setup that the absolute value of the correlation coefficients will be equal:

$$\rho_{1,3} = \frac{\sigma_3^2+\sigma_1^2-\sigma_2^2}{2\sigma_1\sigma_3} = \rho_{1,2} = \frac{\sigma_1^2+\sigma_2^2-\sigma_3^2}{2\sigma_1\sigma_2} = \mathrm{c} \tag{15}$$

$$\rho_{2,3} = \frac{\sigma_1^2-\sigma_2^2-\sigma_3^2}{2\sigma_2\sigma_3} = -c \tag{16}$$

this will only be satisfied when $\sigma_1 = \sigma_2 = \sigma_3$, and will result $\rho_{1,3} = \rho_{1,2}$ =0.5 and $\rho_{2,3}$ =-0.5. These are close to the convergence values of the ABM model in the symmetric case.