Optimising foreign exchange hedging strategies through cross-currency basis analysis

T.R. Janssen

452549

Master Thesis

Econometrics and Management Science

Quantitative Finance

Erasmus University Rotterdam [*]

Supervisor: dr. S.O.R. Lönn

Second assessor: prof. dr. M. van der Wel

Aegon Asset Management

Supervisor: drs. J.N. Schaap

January 2022

[*]The views stated in this thesis are those of the author and not necessarily those of the supervisor, second assessor, Erasmus School of Economics, Erasmus University Rotterdam or Aegon Asset Management.

Abstract

During recent years, the covered interest rate parity has been violated in the form of the cross-currency basis. This has led to rising costs of hedging against currency risk. The presence of the cross-currency basis has been widely established in the literature, but little quantitative analysis with respect to its dynamics has been performed. This paper aims to fill this gap by evaluating whether the cross-currency basis can be predicted at various forecast horizons and by investigating which hedging strategy could reduce hedging costs. Through ARIMA models, HAR models, neural networks and forecast combinations we make a prediction of the cross-currency basis. We find that a random walk is difficult to beat for one-day ahead forecasts, but most considered models provide superior forecasts regarding longer time horizons. Although we conclude that a model-driven hedging strategy would not be more profitable, we do find that funds which suffer from a large basis could reduce the absolute basis incurred by increasing the maturity of the traded forwards.

Contents

1 Introduction **1**

2 Literature review **4**

2.1 Macroeconomic background 4

2.2 Econometric time series modelling 5

2.3 Hedging strategy 7

3 Data **8**

4 Methodology **13**

4.1 Cross-currency basis analysis 13

4.1.1 Autoregressive models 13

4.1.2 Neural networks 14

4.1.3 Forecast combinations: linear model and neural network 16

4.1.4 Performance evaluation 17

4.2 Hedging strategy 20

4.2.1 Currency hedging in general 20

4.2.2 Current hedging strategy for UCITS 21

4.2.3 Alternative strategies 23

4.2.4 Performance evaluation 24

5 Results **27**

5.1 Model selection and calibration 27

5.2 Out-of-sample forecasts 31

5.3 Hedging strategies 38

6 Conclusion and discussion **45**

References **48**

1 Introduction

Recently, a phenomenon that influences the costs of hedging certain exchange rates has become more prominent: the cross-currency basis, also referred to as the basis. In short, this concerns the difference between the price of exchange rate forwards in the market and the theoretical exchange rate derived by interest rate differentials according to the covered interest rate parity (CIP) as outlined by Keynes (1923).

The basis arises due to a difference between supply and demand in the forward market of foreign exchange derivatives. Borio, McCauley, McGuire, and Sushko (2016) argue that this phenomenon is partially caused by recent financial regulation which requires financial intuitions to have certain amounts of particular currencies on their balance sheets at specific moments. As a consequence, apparent theoretical arbitrage opportunities cannot easily be capitalised on in practice. In this paper, we explore the features of the cross-currency basis time series and its predictability.

The cross-currency basis is relevant for asset managers who manage share classes for which investments are denoted in a different currency than the currency in which their clients participate. When one invests in a share class that has underlying assets denoted in a foreign currency, the performance of the investment should be caused by the performance of the underlying assets, rather than the appreciation or depreciation of the foreign currency in question. For this purpose, these share classes are hedged by means of exchange rate forwards.

Due to the presence of the basis, the pricing of these forwards is often not fully in line with what one expects according to the CIP. Hence, each time a trade is executed for the purpose of hedging against currency risk, an asset manager is confronted with additional costs or gains which influence the performance of the share class. The basis leads to the cost of hedging currency risk not being constant over time. Obtaining more information about the dynamics of the basis could enable us to develop hedging strategies that come with potentially lower trading costs.

The first part of the thesis aims to answer the following research question: 'Can we model and predict the cross-currency basis?' We analyse the cross-currency basis time series and evaluate which key features it exhibits. Based on that, we fit various linear models and learning methods, as well as a combination of them for both in-sample and out-of-sample analysis. A thorough analysis of the results is presented with a conclusion as to which models have the best predictive quality for various forecast horizons.

The second part of the thesis deals with the following research question: 'To what extent do alternative foreign exchange hedging strategies for share classes lead to a reduction in costs incurred from the basis and do these strategies generate additional performance?' A common hedging strategy in the industry is to use forwards maturing on a monthly basis, which are rolled at the end of the month. Historically, this has led to large transactions at times when the absolute basis tends to be relatively large.

We evaluate the consequences of applying various timings of full hedges, applying different hedge ratio bandwidths and hedging by using forwards with a different maturity. Furthermore, we discuss whether a hedging strategy can be developed based on the models that are considered in the first part of the thesis. Eventually, we propose an optimal currency hedging strategy for each share class under consideration.

From a practical perspective, asset managers aim to apply a hedging strategy which covers against exposure to currency risk, without creating additional performance. Trading at moments where the absolute basis is relatively large leads to some of such performance. However, maintaining a forward position that deviates substantially from the market value of the share class might lead to performance caused by exposure to currency risk.

Preferably, when a hedged share class is managed by an asset manager, the changes in the value of the share class are solely caused by market movements of the underlying assets rather than results from forward contracts which are unrelated to the appreciation or depreciation of the relevant currencies. It is interesting for asset managers to explore which strategy incurs as few absolute basis costs as possible, while not generating large performance due to exchange rate risk exposure. Modelling and predicting the basis is a logical starting point in order to develop such a strategy.

Relatively little quantitative analysis of the cross-currency basis has been performed thus far in the literature. This thesis aims to provide a starting point to discover time-series properties of the basis. Given that previous research has established the prominence of the basis, we evaluate which models are able to make accurate short-term predictions of the basis. For this purpose, we apply econometric tools which have been proved successful in forecasting time series of a different nature. We make an evaluation of the one-month cross-currency basis for three exchange rates: EUR/USD, EUR/GBP & EUR/CHF.

We find that it is difficult to provide significantly better 1-day ahead forecasts than a random walk. Nevertheless, for longer time horizons such as 5-day ahead (one week) and 22-day ahead (one month) forecasts, we are able to predict the basis more accurately than the random walk is able to. Although the models under consideration are chiefly able to capture the mean-reverting properties of the basis, they do not appear to have the ability to foresee sudden large changes in the basis.

As a consequence, applying a model-driven hedging strategy does not come with evident benefits. However, we do find that when the maturity of the forward contacts is extended from one month to three months, the absolute basis incurred is reduced substantially. This results in opportunities to reduce the costs of hedging against foreign exchange risk.

This paper continues with a Literature review which outlines the relevant literature for the macroeconomic background of the basis, econometric time series techniques and hedging strategies. Subsequently, the Data section presents the data that is used in this research. The Methodology section introduces the approach for both the forecasting exercise as well as the various evaluated hedging strategies, along with their respective performance evaluation. After applying the outlined methodologies, we present the results of the model selection and calibration, the out-of-sample forecast results and the results of the simulation of various hedging strategies in the Results section. In the Conclusion and discussion section, we reiterate our main findings and provide a discussion on their implications.

2 Literature review

2.1 Macroeconomic background

The covered interest parity (CIP) hypothesis was originally developed by Keynes (1923). This hypothesis states that the interest rate differential between two currencies equals the difference between the forward exchange rate and the spot exchange rate. In other words, the difference between a forward and a spot price ought to be fully explained by the difference between interest rates of the currencies. The CIP as introduced by Keynes can be denoted as follows:

$$\frac{F_t - S_t}{S_t} = \frac{r_t - r_t^*}{1 + r_t}, \tag{1}$$

where F_t equals the forward price in the units of the domestic currency versus the foreign currency, S_t equals the spot price in units of the domestic currency versus the foreign currency, r_t denotes the domestic risk-free interest rate and r_t^* equals the foreign risk-free interest rate. All variables are defined at a certain time t. The term of the interest rates corresponds with the time to maturity of the forward. If there is a violation of the parity outlined above, arbitrage is theoretically possible. The extent to which there is a deviation of the CIP is the definition of the basis. In practice, the following alternative version of Equation 1 is often used, which makes use of forward points rather than prices and applies continuous rather than discrete discounting:

$$FP_t = 10000 \times (S_t \times e^{\frac{cd_t}{36000}(r_t - (r_t^* + \frac{b_t}{100}))} - S_t), \tag{2}$$

where FP_t is the number of forwards points, S_t is the spot exchange rate, cd_t is the number of calendar days between the forward and the spot maturity date, r_t is the annualised interest rate of the domestic currency, r_t^* is the annualised interest rate of the foreign currency and b_t is the cross-currency basis. All variables are given on a certain day t. Trivially, the terms of the forwards and respective interest rates ought to be equal for an exact calculation of the basis. Forward points are the number of basis points added to or subtracted from the spot rate of a currency pair, resulting in the necessity to multiply by 10000. The number

36000 in Equation 2 originates from a day count convention of 360 days multiplied by 100, in order to obtain a percentage. We can rewrite Equation 2 by means of simple mathematical operations in order to retrieve a direct formula for the basis:

$$b_t = -100 \times (\ln(\frac{FP_t}{10000 \times S_t} + 1) \times \frac{36000}{cd_t} + r_t - r_t^*). \tag{3}$$

One could argue that the basis could theoretically be closed through arbitrage as expected according to the CIP and no-arbitrage theory. Borio et al. (2016) aim to explain the causes of the presence of the basis and elaborate on why it remains in place. They mention three sources of hedging demand, which are present regardless of the basis: demand from banks, institutional investors and non-financial firms. In particular, banks and institutional investors are faced with stringent regulation with respect to capital requirements. At some dates, notably at the end of months, quarters and years, these financial institutions are obliged to have a set amount of certain currencies on their balance sheets.

Borio et al. (2016) conclude that these constraints may result in tighter limits on arbitrage, raising a threshold to eliminate the basis immediately once it deviates from zero. Additionally, they give particular attention to the influence of central bank announcements. Borio et al. (2016) also note that around recent policy announcements of the Bank of Japan and the European Central Bank, the basis shows a pattern of diversion from its mean.

2.2 Econometric time series modelling

To start modelling the basis, we consider simple linear ARIMA models. Prominent empirical features we want to accommodate by means of these models are persistence and mean reversion. For ARIMA models to be valid, the time series must not contain a stochastic or deterministic trend, which can be tested through a Dickey and Fuller (1979) test. In addition, volatility clustering may arise in a time series, which is not something simple linear ARIMA models are able to account for, resulting in larger estimation errors.

A solution to this issue was presented by Corsi (2009), who introduces the heterogeneous autoregressive model. This model regresses a time series on the average of the values of that series up until particular lags. It has been proven to be quite successful at modelling realised volatility. Realised volatility and the cross-currency basis have in common that they tend to be persistent, mean-reverting and exhibit some sort of clustering. A difference, however, is that the basis can be both positive and negative, whereas volatility is non-negative by definition. It is interesting to evaluate whether the methodology of Corsi (2009) can help us model the basis more accurately than ARIMA models.

The aforementioned models are different examples of linear models. It is however most likely that the data generating process of the basis is not exclusively driven by linear components. Machine learning techniques could help us capture the potential non-linear features of the basis, in order to make more accurate forecasts. The concept of machine learning was first introduced by Turing (1948) and entails models and algorithms which have the property that computers learn while applying them.

Tang, De Almeida, and Fishwick (1991) have shown that regarding time series forecasting, neural networks have an advantage over the Box and Jenkins (1976) methodology applied for ARIMA models due to their flexibility for series with a short memory. Zhang and Hu (1998) present interesting findings by showing the superior predictive quality of neural networks for exchange rates as well as an evaluation as to how many input nodes and hidden nodes one should consider.

A useful tool which could help us improve the forecast quality of time series consists of forecast combinations. In general, forecast combinations yield a reduction in the forecast variance in comparison with individual forecasts, potentially at the cost of some additional bias. Rapach, Strauss, and Zhou (2010) and Stock and Watson (2004) successfully obtain improved predictions through forecast combinations by combining univariate forecasts. Interestingly, they all find that the optimal way to ascribe weights is by attributing each individual forecast a weight of $1/N$, rather than by introducing more complex weighting schemes.

Another approach regarding the weight allocation of individual forecasts within a forecast combination is proposed by Diebold and Shin (2019). They investigate whether it is beneficial to apply regularisation in this process. For this purpose, they develop a LASSO-based procedure which first regularises some of the individual forecast weights to zero and subsequently averages the weights of the remaining individual forecasts. This technique appears to be particularly useful in the context of a large set of potential regressors.

Additionally, Zhang (2003) presents a forecast combination approach which does not require a large set of regressors. He aims to forecast exchange rates using ARIMA models to solely capture the linear features of the time series. On top of that, he recognises the non-linear elements in the time series and aims to predict the residuals of the ARIMA models by means of a neural network. The eventual exchange rate forecast is presented by adding the forecasts of the linear and the non-linear components.

2.3 Hedging strategy

Eventually, our findings concerning the potential predictability of the basis are used to develop an optimal currency hedging strategy. Campbell, Serfaty-De Medeiros, and Viceira (2010) present a paper that investigates this subject. They show that certain currencies, such as the Australian and Canadian dollars, as well as the Japanese Yen and the British pound (GBP), correlate positively with world equity markets whereas American dollars (USD), euros (EUR) and the Swiss franc (CHF) have a negative correlation with equity markets.

Campbell et al. (2010) advocate that to minimise risk as a US-based equity investor, one should fully hedge all currency exposures apart from EUR and CHF. The necessity to hedge as a bond investor is not prevalent, as Campbell et al. (2010) show that there is little to no correlation between bond excess returns and currency excess returns. Lastly, they find no evidence that the forward premium puzzle can be ascribed to changing covariances of currencies with global equity returns.

Although the aforementioned papers on hedging strategies are insightful, they cannot help that much in the context of this research. Unlike Campbell et al. (2010), we do not consider the question whether certain currency risk should be hedged and if so to what extent. We investigate what the cheapest trading strategy is, given the objective to minimise exposure to currency risk.

3 Data

We aim to analyse the effect of various hedging strategies for UCITS share classes. UCITS, which stands for Undertakings for the Collective Investment in Transferable Securities, is a European regulatory framework that harmonises regulations for the sale and management of certain mutual funds. UCITS share classes are open-end funds with one valuation point per day. Fund managers may enable clients to participate in a UCITS share class in their domestic currency, while the actual investments of the fund are made in assets denoted in a particular foreign currency. In the context of this research, we refer to the domestic currency in which a client makes an investment as the share class currency. The foreign currency in which the actual investments in assets are denoted is called the target currency.

Making investments in a fund that invests in assets that are denoted in a non-domestic currency leaves investors exposed to currency risk. An appreciation or depreciation of the target currency influences the performance of the fund, regardless of the performance of the underlying assets. Such currency risk can be eliminated by means of hedging. This can be achieved by trading foreign exchange forwards, which are over-the-counter financial derivatives in which two parties agree to buy or sell an amount of a certain currency against a particular exchange rate at the end of the contract.

We consider four hedged UCITS share classes which are managed by Aegon Asset Management. The first one is the Emerging Market Debt (EMD) UCITS, which has a target currency USD and is hedged toward its share class currency EUR. This fund is the largest share class that is hedged towards EUR. Additionally, we consider Asset-Backed Securities (ABS) UCITS which have target currency EUR. ABS UCITS are the largest share classes which are hedged towards USD, GBP and CHF.

For the out-of-sample period July 2018 to June 2021, we backtest various hedging strategies regarding the four share classes described above. For this purpose, we use data on the market value of the fund, the daily inflows and outflows due to subscriptions and redemptions and the inflows and outflows due to forward exchange rate results. Based on that we calculate the daily return of the underlying assets through the following formula:

$$R_t = \frac{MV_t}{MV_{t-1} + FX_t + FL_t} \times 100\%. \tag{4}$$

Here, R_t is the daily return of the underlying assets, MV_t equals the market value of the fund, FX_t is the result of the foreign exchange forward and FL_t shows flows due to subscriptions and redemptions. Each variable is determined on each day t and expressed in the target currency. Equation 4 is a simple return where we correct for inflows and outflows due to either the results of the forwards or deposit and withdrawal decisions by clients. The formula deduces what the return of the underling assets had been if none of such flows had occurred. All variables are expressed in the target currency. Using these calculated returns, we estimate the daily returns of the share classes for the simulated strategies.

In order to investigate what the best hedging strategies would be for the four share classes mentioned above, we calculate the cross-currency basis of the relevant exchange rates. In practice, we consider three exchange rates: EUR/GBP, EUR/CHF and EUR/USD, since the USD/EUR exchange rate implicitly lies within the EUR/USD exchange rate. The period under consideration in this research is July 2006 to June 2021, since from this moment onwards the basis tends to deviate from zero. This period is subdivided into an in-sample period of the first twelve years (July 2006 to June 2018) and an out-of-sample period of the last three years (July 2018 to June 2021). The cross-currency basis is not a time series that can be extracted directly from available databases. It can be calculated at a particular point in time with the five variables displayed in Equation 3. These variables are retrieved from Bloomberg.

To start with, for each time t, we retrieve the spot rate S_t of EUR/USD, EUR/GBP and EUR/CHF. Furthermore, we obtain the forward points FP_t of each of the exchange rates for each day t. For the forecasting exercise and the estimation of parameters in the in-sample period, we consider one-month forwards at each time t. For the hedging strategy simulation, we consider a strategy that uses forwards with a maturity at the end of each month, at the end of each quarter, at the fifteenth day of each month and at the fifteenth day of the last month of each quarter. The number of calendar days cd_t between the forward and spot maturity dates at time t is also retrieved from Bloomberg.

As displayed in Equation 3, we also need the relevant annualised risk-free interest rates of the share class currency r_t and the relevant annualised risk-free interest rates of the target currency r_t^*. For each currency, we consider the rate which comes the closest to being a risk-free rate. For EUR, this is EONIA, for USD, this is the OIS, for GBP this is the SONIA rate and for CHF this is the SARON rate. Unfortunately, the SARON rate is only available from 2009 onwards. Therefore, for calculation of the basis for the EUR/CHF rate, we use the EUR and CHF LIBOR rates. Using LIBOR rates is less ideal since it contains a default risk premium, but it gives us the possibility to consider the same time period for a CHF related share class as well. Apart from this, since the LIBOR rates are subtracted in the calculation of the basis, the effect of the default risk premium is reduced since there is such a premium for both rates, albeit not fully equal.

For the forecasting exercise with one-month forwards, we use the relevant annualised one-month rate for each day. With regards to the hedging strategy simulation, we use the annualised one-month or three-month rates for forwards that are rolled monthly and quarterly respectively. Evidently, the time to maturity is not constant when we consider forwards that are rolled either monthly or quarterly. We use a one-month and three-month rate with a constant time to maturity as not all terms are available in the database. It must be noted that, due to the difference of the terms between the forwards and risk-free rates, we do not make an exact calculation of the basis for the hedging strategy simulation. Nevertheless, this appears to be the best option given the data availability constraints. On the day of the roll, which is when the largest transaction occurs, the terms of the forward and the maturities are equal such that the calculation of the basis is exact on this date.

Calculating the basis for a particular exchange rate with a particular forward requires five data points per day, while the number of data points that can be retrieved from Bloomberg is limited. Therefore we calculate the basis for the forwards with the four maturities mentioned above exclusively for the test period July 2018 to June 2021. For the period July 2006 to June 2021, which we consider for the forecast exercise, we only calculate the basis based on one-month forwards for the EUR/USD, EUR/GBP and EUR/CHF exchange rates.

Figure 1 shows a plot of the three basis series. In Table 1, we display the descriptive statistics of these one-month basis series: mean, median, volatility, skewness, autocorrelation and Dickey and Fuller (1979) statistic. We observe that on average the EUR/USD basis is

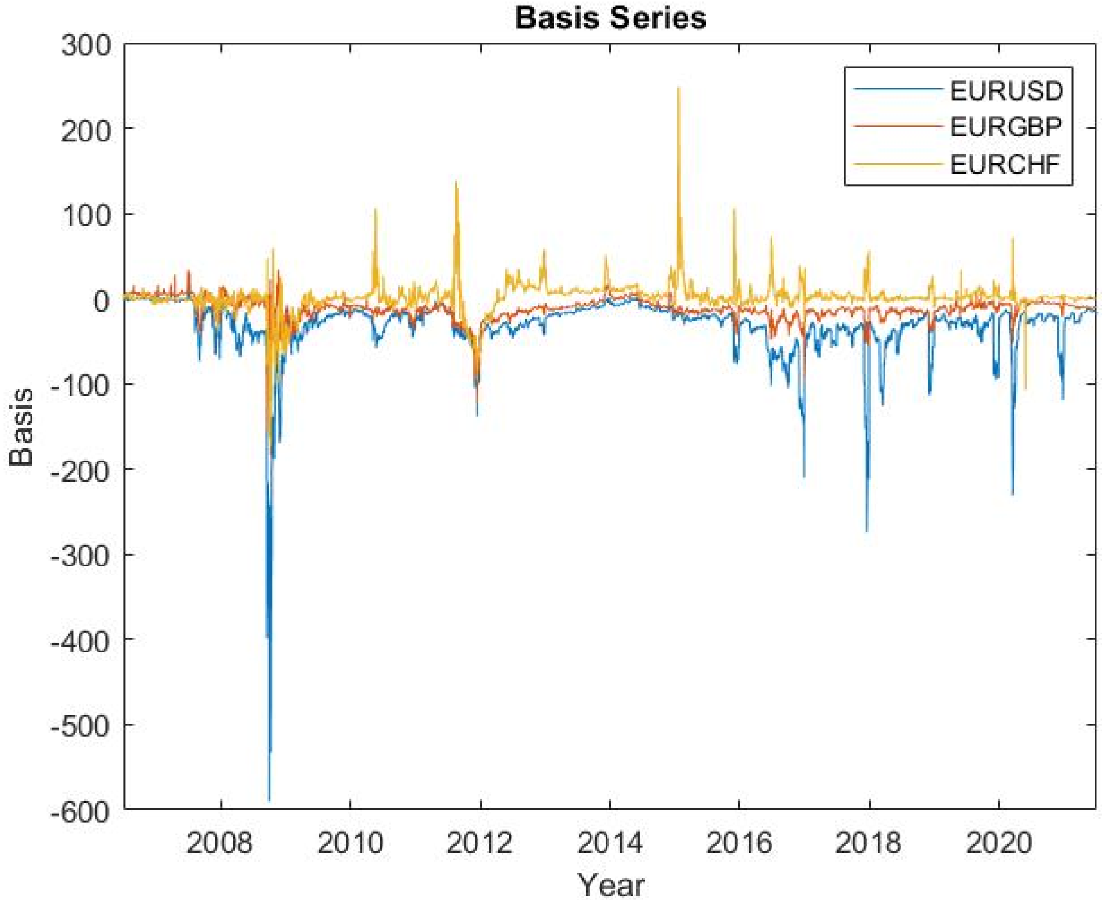

Figure 1: Basis series. *Note*: This figure shows the one-month EUR/USD, EUR/GBP and EUR/CHF basis expressed in basis points for the period July 2006 to June 2021.

Table 1: Descriptive statistics

	EUR/USD	**EUR/GBP**	**EUR/CHF**
Mean	-32.335	-12.215	1.579
Median	-24.820	-11.343	1.167
Standard deviation	35.740	13.277	18.952
Skewness	-5.672	-2.588	1.064
Autocorrelation	0.958	0.894	0.877
Dickey-Fuller statistic	-6.777*	-10.736*	-15.920*

Note: This table displays the descriptive statistics of the one-month EUR/USD, EUR/GBP and EUR/CHF cross-currency basis during the period July 2006 to June 2021. We display the mean, median, standard deviation, skewness, first-order autocorrelation and the Dickey and Fuller (1979) statistic. A * indicates significance at a 1% level.

negative, the most volatile and skewed downwards. The EUR/GBP basis is also negative, but to a lesser extent and also less volatile than the EUR/USD basis. The EUR/CHF basis shows a different pattern and is slightly positive on average. All basis series show persistence with a first-order autocorrelation that is higher than 0.85. Lastly, the three series do not contain a stochastic or deterministic trend, which can be concluded from the Dickey and Fuller (1979) statistics which are significant at a 1% level for all series under consideration.

4 Methodology

4.1 Cross-currency basis analysis

4.1.1 Autoregressive models

When we observe the three basis series in Figure 1, we note that they appear to have some mean-reverting properties. For this reason, it seems logical to start with the set of ARIMA models. An ARIMA(p,d,q) model is defined as follows,

$$\Delta^d b_t = \beta_0 + \beta_1 \Delta^d b_{t-1} + ... + \beta_p \Delta^d b_{t-p} + \epsilon_t + \theta_1 \epsilon_{t-1} + ... + \theta_q \epsilon_{t-q}. \tag{5}$$

Here p is the number of autoregressive (AR) terms, q is the number of moving average (MA) terms, d is the level of differencing. Furthermore, b_t is the basis at time t and ϵ_t is the error of the model at time t. The coefficients are estimated through maximum likelihood estimation. A common strategy to select the optimal hyperparameters is investigating their autocorrelation function (ACF) and partial autocorrelation function (PACF) as originally proposed by Box and Jenkins (1976). We do this for the in-sample period July 2006 to June 2018 for both the undifferenced series ($d = 0$) and the differenced series ($d = 1$).

Corsi (2009) presents an alternative to simple ARIMA models in the form of heterogeneous autoregressive (HAR) models. As can be derived from the name, this is an autoregressive model with heterogeneous lags. It was originally developed as a model for realised volatility and proved to be successful in that context. Realised volatility and the cross-currency basis have in common that they tend to be mean-reverting processes which exhibit some form of clustering as can be noted in Figure 1. It is interesting to evaluate whether this model has some predictive power with regard to the basis as well. The specific model proposed by Corsi (2009) is the HAR(3) model, which is defined as follows:

$$b_t = c + \beta b_{t-1} + \beta^{(w)} b_{t-1}^{(w)} + \beta^{(m)} b_{t-1}^{(m)} + \omega_t, \tag{6}$$

with

$$b_{t-1}^{(w)} = \frac{1}{5}\sum_{i=0}^{5} b_{t-i} \tag{7}$$

and

$$b_{t-1}^{(m)} = \frac{1}{22}\sum_{i=0}^{22} b_{t-i}. \tag{8}$$

Here c is a constant, β, $\beta^{(w)}$, $\beta^{(m)}$ are coefficients, b_t is the basis on day t and ω_t is the error term on day t. Equation 6 is estimated through ordinary least squares. Based on the in-sample analysis, we determine an appropriate set of lags to consider for the out-of-sample prediction.

4.1.2 Neural networks

We also forecast the cross-currency basis by means of artificial neural networks (ANNs). In this part of the research, we follow a methodology comparable to the approach of Zhang and Hu (1998). We consider a single hidden-layer feedforward neural network, which takes the following form:

$$b_t = \beta_0 + \sum_{j=1}^{q} \beta_j h(\gamma_{ji} + \sum_{i=1}^{p} \gamma_{j1} b_{t-i}) + \epsilon_t, \tag{9}$$

with a p-dimensional input of previous values of the basis $(b_{t-1}, b_{t-2}, .., b_{t-p})$ and q nodes in the hidden layer. We use a hyperbolic tangent activation function $h(x)$ which is defined as:

$$h(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}. \tag{10}$$

The actual estimation of the parameters is done through the damped least squared method which was introduced by Levenberg (1944) and Marquardt (1963). This method is also called the Levenberg-Marquardt algorithm. The idea behind this algorithm is to combine two other numerical minimisation algorithms. On the one hand, it uses the gradient descend method, which reduces the sum of the squared errors by updating the parameters in the direction of the steepest descent. On the other hand, it uses the Gauss-Newton method, which reduces the sum of the squared errors though assuming that the least squares function is locally quadratic in the parameters and subsequently minimising this quadratic.

In the case that the parameters are far from their optimal values, the Levenberg-Marquandt algorithm behaves like a gradient descend method, whereas it behaves more like the Gauss-Newton method when parameters are near their optimal value. A more thorough explanation of the algorithm is provided by Gavin (2019). We use a package from the Deep Learning Toolbox from MATLAB to implement the Levenberg-Marquardt alogrithm.

The neural networks are initialised according to the Nguyen and Widrow (1990) algorithm. This algorithm chooses values such that the active region of each neuron in the layer is distributed virtually evenly across the input space of the layer. When initialising neural networks, some randomness is inevitable. We use the same seed for each estimation such that randomness does not play a role in the differences among the evaluated model specifications.

It would be ideal to have multiple initialisations, but this requires substantial computation power, in particular for the cross-validation. We train the neural network during a training period and stop the training if the Root Mean Squared Error (RMSE) of the forecasts in the validation period does not improve for six steps. With the trained model, we eventually construct out-of-sample predictions.

To select the appropriate hyperparameters for the number of lags p and hidden nodes h, we ought to perform cross-validation. Since we are considering a time series, the data are ordered. Subdividing individual data points within the in-sample period randomly into a training set, a validation set and a test set would lead to the loss of important dynamics in the data. For this reason, we apply a nested cross-validation approach as presented by Varma and Simon (2006). According to them, this approach yields an almost unbiased error for time series which exhibit autocorrelation.

For the nested cross-validation, we divide the in-sample period into a training, validation and test set for each of the five folds. The training set always starts at the beginning of July 2006 and finishes at the end of June 2011 for the first fold. The validation set lasts one year and spans from July 2011 to June 2012 and the test set spans from July 2012 to June 2013 for the first fold. For the consecutive folds, we extend the training period with one year and shift both the validation and test set one year to the future. A visualisation of the concept

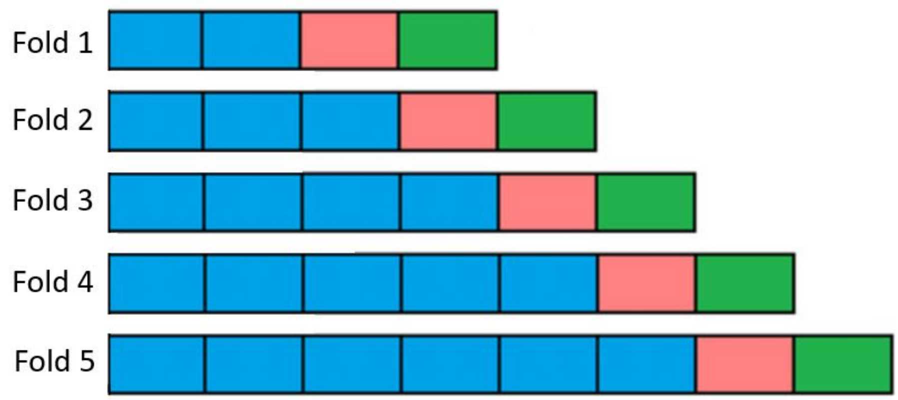

Figure 2: Concept of nested 5-fold cross-validation. *Note*: For each of the five folds, we show the training set in blue, the validation set in pink and the test set in green.

of the applied nested cross-validation approach is provided in Figure 2.

After performing the nested 5-fold cross-validation, where we evaluate $p = 1, ..., 30$ and $q = 1, ..., 50$, we calculate the mean of the RMSE for the five test sets for each set (p, q) of the hyperparameters. We select the set (p, q) which yields the lowest average RMSE over the five folds. For the eventual out-of-sample forecasts during the period July 2018 to June 2021, we use the training period July 2006 to June 2017 and the validation period July 2017 to June 2018.

4.1.3 Forecast combinations: linear model and neural network

Forecast combinations are known for improved forecasts as well as a reduction in the forecast variance as previously discussed in the Literature review section. In our application, we follow the approach of Zhang (2003). This paper considers a time series that is subdivided into a linear and a non-linear component:

$$b_t = l_t + n_t + \kappa_t. \tag{11}$$

Here, b_t is the basis, l_t is the linear component an n_t is the non-linear component and κ_t the error term. Every variable is determined on day t. The linear component is estimated through a linear ARIMA model. The fitted value according to the most appropriate ARIMA model as determined during the in-sample analysis is denoted by $\hat{l}_t$. The non-linear component of the basis time series is based on the residuals of the ARIMA model, which are constructed as follows:

$$e_t = b_t - \hat{l}_t. \tag{12}$$

These residuals e_t are modelled by using ANNs, which are a universal approximator to non-linear problems. For p input nodes, this model can be displayed as follows:

$$n_t = f(e_{t-1}, e_{t-2}, ..., e_{t-p}) + \upsilon_t. \tag{13}$$

Here f is the nonlinear function which is determined by the ANN, $(e_{t-1}, e_{t-2}, ..., e_{t-p})$ is the p-dimensional input of the residuals as calculated in Equation 12 and υ_t is the error term of the ANN on time t. Fitting the ANN eventually yields the fitted value of the non-linear component $\hat{n}_t$. The estimation and cross-validation of the ANNs are performed in the same manner as when the basis itself is modelled with an ANN. The eventual forecast $\hat{b}_t$ of the basis for this forecast combination approach is given by:

$$\hat{b}_t = \hat{l}_t + \hat{n}_t. \tag{14}$$

4.1.4 Performance evaluation

To evaluate the out-of-sample quality of all 1-day ahead, 5-day ahead and 22-day ahead predictions, we calculate the Root Mean Squared Error ($RMSE$) and Mean Absolute Error (MAE):

$$RMSE = \sqrt{\frac{\sum(b_t - \hat{b}_t)^2}{P}} \tag{15}$$

and

$$MAE = \frac{\sum |b_t - \hat{b}_t|}{P}. \tag{16}$$

Here b_t is the observed basis on day t, $\hat{b}_t$ is the predicted basis according to a model and P equals the size of the out-of-sample period. Eventually, we want to evaluate whether the models under consideration have significantly better predictive qualities than a benchmark model, which is the random walk. We use a Diebold and Mariano (2002) test to evaluate whether a model yields a significant improvement of the forecast accuracy. This test evaluates whether the forecast errors resulting from the models under consideration differ significantly from one another. First, we calculate the following statistic, which is called the loss differential, for each day t:

$$d_{t+h} = e^2_{H0,t+h|t} - e^2_{H1,t+h|t}, \tag{17}$$

where d_{t+h} is the loss differential for an h-day ahead prediction, $e_{H0,t+h|t}$ is the forecast error of the less complex model for an h-day ahead forecast and $e_{H1,t+h|t}$ is the corresponding forecast error for the tested model. The Diebold and Mariano (2002) statistic is subsequently calculated as follows:

$$DM = \frac{\bar{d}_h}{\sqrt{V(\hat{d}_{t+h})/P}}, \tag{18}$$

with

$$V(\hat{d}_{t+h}) = \frac{1}{P-1} \sum_{t=T+1-h}^{T+P-h} (d_{t+h} - \bar{d}_h)^2. \tag{19}$$

Here $\bar{d}_h$ is the average of d_{t+h}, $V(\hat{d}_{t+h})$ is the sample variance of the loss differential, P is the number of observations during the out-of-sample period and T denotes the last observation in the in-sample period. Under the null hypothesis that the more complex model does not yield a different forecast quality than the benchmark model, DM follows a standard normal distribution. A negative DM statistic implies that the model under consideration has better predictive accuracy than the benchmark model.

It is important to note that the DM statistic only gives a judgement about the forecast quality of a certain model in comparison with the quality of a benchmark model. The statistic does not imply anything about the absolute forecast quality. For the latter, we can compare the $RMSE$ and MAE. Another interesting feature about the forecast quality of a model is whether it is able to predict the direction of the movement of a time series.

Pesaran and Timmermann (1992) present a performance measure which gives an indication about this feature. By introducing a non-parametric test, they enable us to evaluate whether a model is able to accurately predict the direction of the movement of a particular time series. In contrast to the Diebold and Mariano (2002) statistic, the test statistic does not ascribe any value to the quantitative proximity of the forecast to the realised value. It solely considers the forecast qualitatively by evaluating whether or not the direction of the prediction is correct.

The share of correctly predicted directions during the out-of-sample period is denoted $\hat{Q}$. The share of predictions of increases of the basis during this period is given by $\hat{Q}_{\hat{b}}$ and the actual share of increases during this period is given by $\hat{Q}_b$. We now introduce $\hat{Q}_*$:

$$\hat{Q}_* = \hat{Q}_b + \hat{Q}_{\hat{b}} + (1 - \hat{Q}_b)(1 - \hat{Q}_{\hat{b}}). \tag{20}$$

Additionally we have:

$$\hat{\text{var}}(\hat{Q}) = \frac{1}{P}\hat{Q}_*(1 - \hat{Q}_*) \tag{21}$$

and

$$\hat{\text{var}}(\hat{Q}_*) = \frac{1}{P}(2\hat{Q}_{\hat{b}} - 1)^2\hat{Q}_b(1 - \hat{Q}_b) + \frac{1}{P}(2\hat{Q}_b - 1)^2\hat{Q}_{\hat{b}}(1 - \hat{Q}_{\hat{b}}) + \frac{4}{P^2}\hat{Q}_b\hat{Q}_{\hat{b}}(1 - \hat{Q}_b)(1 - \hat{Q}_{\hat{b}}), \tag{22}$$

with P the number of days in the out-of-sample period. We calculate the PT statistic as follows:

$$PT = \frac{\hat{Q} - \hat{Q}_*}{\sqrt{\hat{\text{var}}(\hat{Q}) - \hat{\text{var}}(\hat{Q}_*)}}. \tag{23}$$

Under the null hypothesis of no ability to forecast the direction of the basis, the PT statistic follows a standard normal distribution.

4.2 Hedging strategy

4.2.1 Currency hedging in general

To eliminate risk due to currency exposure, we aim to hedge the investments in a share class such that any result of the fund can be ascribed to the performance of the underlying assets rather than changes in the exchange rate. To hedge currency risk, we must take a position in a forward which sells a certain amount of the target currency versus the share class currency for the value of the investment in the share class.

As a consequence, changes in the exchange rate of the target currency versus the share class currency do not affect the overall value of the portfolio: appreciation of the target currency increases the value of the investment in the share class, which is offset by an equal decrease of the value of the short forward position. A depreciation of the target currency decreases the value of the underlying assets in the share class which is mitigated by an increase in the value of the short forward position.

However, there are practical issues that lead to the possibility of losing a fully hedged position. This may lead to exchange rate risk exposure at some point in time. In that case, the theoretical situation described above may not hold anymore. To illustrate this, we introduce the following formula:

$$MV_t = I_t + UF_t, \tag{24}$$

where I_t denotes the value of investments in the share class, UF_t denotes the unrealised result of the forward contract and MV_t is the market value of the fund. The variables are expressed in the share class currency and are determined for a day t.

Suppose there is a certain amount of money invested in the share class at the beginning of day t and we assume that the investment in the share class is fully hedged. For this reason, changes in the exchange rate during the day do not affect the market value of the portfolio MV_t. Naturally, the value of the investments in the fund changes during the day due to market movements. Yet, the forward contract is fixed on a set amount of the target currency. A change in the value of the underlying assets of the share class thus results in the investment not being fully hedged anymore.

Another cause of a discrepancy between the value I_t and the set amount of the target currency in the forward position are inflows and outflows for the fund. If a client wants to subscribe to or redeem investments from a share class, the amount invested in the target currency changes, whereas the set amount of the target currency in the forward contract is not automatically altered. This changes the extent to which the share class is hedged, such that a change in the relevant exchange rate does not have an equal opposite effect on I_t and UF_t.

The two challenges discussed above centre around changes in the value of the investments. There is, however, also a practical issue for the forward contract. A forward contract has a finite maturity date at which the underlying asset of the forward contract is actually delivered. In this case, this concerns a particular amount of the share class currency. We must maintain a hedged position after the maturity date of the forward contract. For this reason, at some point before maturity, we must close current forward contracts and open ones with a later maturity date.

4.2.2 Current hedging strategy for UCITS

For hedged UCITS, the strategy which is the current industry common practice consists of both a daily process as well as a monthly process to resolve some of the challenges outlined above. In this paragraph, we explain which hedging process is currently performed and which transactions are executed at which moment. We start by introducing the daily hedging process. On a daily basis, a UCITS encounters flows. This can for example be due to subscriptions and redemptions by clients. The policy is that corrections in the hedging positions due to inflows and outflows are executed on the same day, simultaneously with the investment or divestment in the share class.

If there is a net inflow of the share class currency of amount X, we buy the corresponding amount of the target currency according to the spot rate such that the money can be invested in assets that are denoted in the target currency. To hedge this inflow immediately, we sell a forward of the target currency which corresponds to an amount X of the share class currency. If there is a net outflow, the buy becomes a sell and vice versa. The procedure of buying/selling a spot and selling/buying a forward simultaneously is called a swap.

In addition to this process, during every trading day of the month, we monitor whether the forward position remains close enough to the market value of the fund denoted in the target currency. This might have changed due to the performance of the underlying assets of the fund. For this purpose, we monitor the hedge ratio:

$$HR_t = \frac{HP_t}{MV_t} * 100\%. \tag{25}$$

Here HR_t denotes the hedge ratio, HP_t is the hedge position which is equal to the net forward position denoted in the target currency and MV_t entails the market value of the fund expressed in the target currency. All variables are determined on each day t. When on a certain day this hedge ratio deviates from the bandwidth of 96% to 104%, a hedging process is triggered which brings back the hedge ratio to 100%. This process consists of increasing or decreasing the forward position depending on the difference between MV_t and HP_t. If MV_t-HP_t, is positive, we sell MV_t-HP_t of the forward of the target currency versus the share class currency. If MV_t-HP_t is negative, we buy this amount of the forward of the target currency versus the share class currency.

Now that we have described the daily process, we move on to the monthly process of the foreign exchange hedging strategy. In short, this procedure consists of rolling the forward position to the next month. After this has been processed, we steer the hedge ratio back to 100% at the end of the current month and on the first trading day of the next month.

Two days before the last trading day of the month, the forward position for the current month is closed. This is done by buying a forward in the target currency versus the share class currency for the net amount of sold units of the target currency. At the same time, we take the same net forward short position as before, but now with a maturity at the end of the following month. This process is called rolling the forwards.

Closing the forward positions of the current month yields a certain result which is set from the day after the roll onwards. The cash results of the forward are set at a fixed value of the share class currency. At this date, there is an equally long and short forward position that matures during the current month. New fluctuations in the exchange rate do therefore not influence the value of the forward positions for this month overall.

At maturity, which is the last trading day of the month, the result of the forward currency is settled in the share class currency. If this is a positive result, we use this result to make an additional investment in the fund and use the share class currency to buy the target currency at the spot rate while selling an equal forward denoted in the target currency versus the share class currency to maintain a hedged position. In the case of a negative result, we withdraw a certain amount out of the fund to cover the loss on the forward position and buy a forward denoted in the target currency versus the share class currency which matures the next month to maintain the hedged position. The way we treat the forward result is thus equal to a regular inflow and outflow on a random day.

On the last day of the month, we bring back the hedge ratio to 100%. This procedure is equal to the previously explained hedging procedure which is triggered when the hedge ratio diverges from the bandwidth 96% to 104%. We sell (buy) a forward worth MV_t-HP_t of the target currency versus the share class currency if this number is negative (positive). Unfortunately, not all information on the market movements of the underlying assets is known at this point, since this hedge is performed before the closure of the market during the last day of the month. For this purpose, we trigger another full hedge one day after the end of the month, to ensure a 100% hedge position at the beginning of the month.

A disadvantage of the current strategy is that for some funds the trading costs in the form of basis tend to be substantial during the end of the month, which is the time at which the largest transactions take place. For other funds, this phenomenon implies a positive result as for them the basis is generally positive. As a consequence, the hedging strategy tends to cause some performance, either negative or positive. All in all, four parameters can be altered to develop a new strategy: the hedge ratio bandwidth, the moments when a full hedge is performed, the moment of the roll of the forward and the choice of maturity dates of the forward contracts.

4.2.3 Alternative strategies

We evaluate the effect of alternative hedging strategies on several factors. For this purpose we discard the influence of inflows and outflows due to subscriptions and redemptions respectively. When such a flow takes place, there is an inevitable spot trade and a forward can be immediately traded as well via a swap. It is also uncertain if particular inflows or outflows would be the same if another hedging strategy was implemented, while it does not influence the eventual basis cost and exposure to currency risk to a large extent.

In practice, the most costs incurred by the basis are caused by rolling the forward during the end of the month. This is the largest transaction since it entails approximately the full value of the share class investments. On top of that, it is in particular during the end of the month when the forwards are rolled, that the basis tends to deviate from the mean as can be observed in Figure 1. Potential explanations for this phenomenon are given in the Literature review section.

We evaluate a strategy that rolls the forwards in the middle of the month, on its fifteenth day, such that we avoid substantial trades at the end of the month. Furthermore, we evaluate whether it may be beneficial to trade in forwards with a longer maturity of three months, instead of one month. This would lead to fewer forward rolls in general, which could potentially lead to a lower exposure to the basis.

Furthermore, we investigate the implications of eliminating the monthly hedge and only hedge when a certain threshold is exceeded. Regarding the threshold, we also evaluate alternative bounds of the hedge ratios to be maintained: 98%-102% and 99%-101%. We do not consider a wider bandwidth than the current one, since regulation forbids a deviation from the 95%-105% region when a hedged share class product is offered to the public.

All in all, for each UCITS under consideration, we evaluate four maturity dates: end of each month, halfway each month, at the end of each quarter and halfway the last month of each quarter. We consider the situation with a monthly full hedge and without such a hedge and the current and two alternative hedge ratio bounds mentioned above. This results in 24 potential strategies for each fund.

4.2.4 Performance evaluation

We simulate the various hedging strategies outlined above for each of the four UCITS described in the Data section during the period July 2018 to June 2021. For each strategy i for each fund, we calculate four performance measures: alpha (α_i), exposure (EXP_i), share of days on which a trade occurs and average daily basis points incurred (AB_i).

In order to determine the α_i of a hedging strategy, we need to subtract the underlying asset return UAR from the hedged share class return HSR_i according to strategy i. The hedged share class return for strategy i in the out-of-sample period is calculated as follows:

$$HSR_i = ((\frac{MV_e^i * S_e}{MV_b^i * S_b})^{1/Y} - 1) \times 100\%. \tag{26}$$

Here HSR_i is the annualised share class return in percentages, MV_t^i is the market value of the share class investments denoted in the target currency and S_t is the spot rate at time t. The letters b and e refer to the beginning and end of the out-of-sample period respectively. Y is the number of years, which is three since we evaluate the period July 2018 to June 2021. Consequently, we calculate the annualised return of the underlying assets for a particular share class as follows:

$$UAR = ((\prod_{t=1}^{P}(\frac{R_t}{100} + 1))^{1/Y} - 1) \times 100\%. \tag{27}$$

Here, R_t for each day t are the daily returns of the investments as calculated in Equation 4. P is the number of observations in the out-of-sample period and Y is the number of years. After applying the aforementioned formulas, we can calculate the hedging α_i of a particular strategy i as follows:

$$\alpha_i = HSR_i - UAR. \tag{28}$$

This performance measure displays to what extent the hedging strategy adds performance to the share class. It displays the excess return of the hedged share class over the returns of the underlying assets. There are two factors which might influence α_i: exposure to the target currency and the basis. Exposure to the target currency can yield either a profit or a loss based on changes in the spot rate, whereas the basis leads to a riskless profit or loss.

The average exposure EXP_i according to strategy i is defined as the average deviation of the hedge ratio from 100% and indicates to what extent a fund is exposed to currency risk on average according to a particular strategy. It is calculated as follows:

$$EXP_i = \frac{1}{P}\sum_{t=1}^{P}|100\% - HR_{i,t}|. \tag{29}$$

Here P is the number of observations in the out-of-sample period and $HR_{i,t}$ is the hedge ratio according to strategy i at day t for which the calculation is displayed in Equation 25.

Another statistic we calculate is the share of days on which a trade occurs. Generally, executing a trade comes with some costs, regardless of whether the basis is present or not. The fewer trades, the fewer of such costs are made. To get more insight into this, we calculate the share of days in which a trade occurs in the out-of-sample period according to a certain strategy i.

Some of the performance of α_i may arise due to a currency appreciation or depreciation which occurs while the hedge ratio is not equal to 100%. Another part of α_i is the incurred basis. If we want to evaluate the costs or gains received through the basis purely, we can calculate the average daily basis incurred AB_i for a particular fund according to strategy i. This statistic displays the average daily amount of basis points of the share class currency which are gained or lost due to the presence of the basis. It is calculated as follows:

$$AB_i = \frac{1}{P} \sum_{t=1}^{P} \frac{ft_{i,t}}{MV_{i,t}} \times b_t. \tag{30}$$

Here, $ft_{i,t}$ is the size of the forward trade in the share class currency according to strategy i at time t and $MV_{i,t}$ is the value of the share class investments according to strategy i at time t denoted in the share class currency. b_t is the relevant basis on day t and P is the size of the out-of-sample period. A positive AB_i implies a gain and a negative AB_i implies a loss from the perspective of the investor.

5 Results

5.1 Model selection and calibration

In this section, we make an in-sample analysis of the one-month EUR/USD, EUR/GBP and EUR/CHF basis during the period July 2006 to June 2018. We start with an autocorrelation analysis, which serves primarily for the lag selection of the various autoregressive models. Additionally, we present the results of the cross-validation which is used to select the hyperparameters of the neural networks under consideration.

In the autocorrelation graphs in Figure 3, we display the sample autocorrelation and the sample partial autocorrelation of the one-month EUR/USD, EUR/GBP and EUR/CHF basis. All three series show a decaying autocorrelation function, albeit not fully quadratic. On the other hand, the partial autocorrelation is significant at a 5% level at various lags. Considering the sample autocorrelation and partial autocorrelation graphs for the three series, we conclude that the model should probably contain an AR component according to the approach proposed by Box and Jenkins (1976).

To select the correct lag order, we consider the sample partial autocorrelations of the series. We observe that each of the three basis series has a significant partial autocorrelation at a 5% level up until lag 2. This indicates that an AR(2) model might be an appropriate choice. We also consider an AR(1) model to evaluate whether an AR(2) model indeed has a better forecast quality. On top of that, there are significant lags after lag 2 for each series under consideration, albeit at different lags per series. For this reason, it could be interesting to evaluate models with a larger amount of lags. We therefore also consider the intuitive amount of lags for a week (5 lags) and a month (22 lags). All in all, we consider the following four AR models: AR(1), AR(2), AR(5) and AR(22).

Additionally, the feasibility of the HAR model as introduced by Corsi (2009) is evaluated. The potential advantage of this model is that it considers a mean of the weekly and monthly lags, thereby smoothing the effect of outliers while still giving larger lags some influence. Since the number of significant partial autocorrelation lags is relatively scarce within the frame of one month, we also consider a version of the HAR model which does not consider a monthly component but solely a daily and a weekly one. We thus evaluate both the HAR(2)

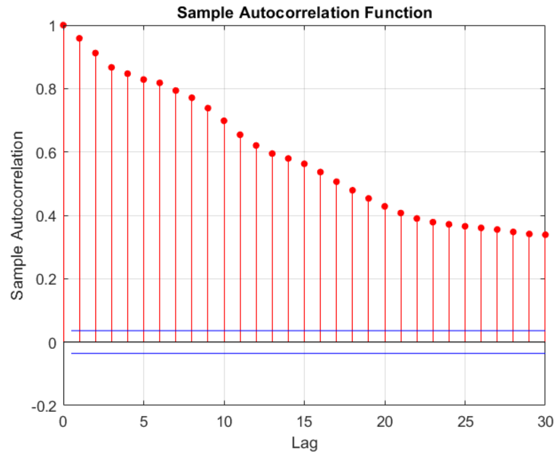

(a) Autocorrelation EUR/USD basis

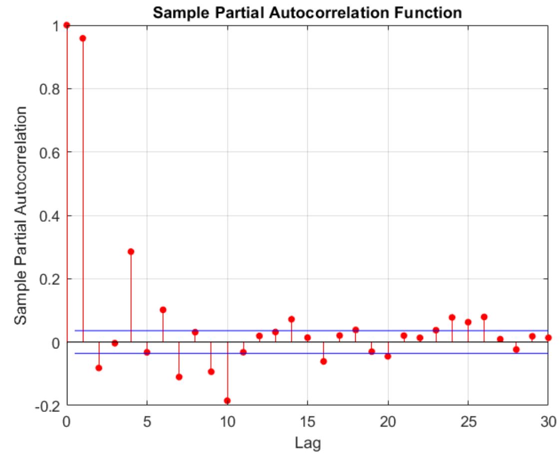

(b) Partial autocorrelation EUR/USD basis

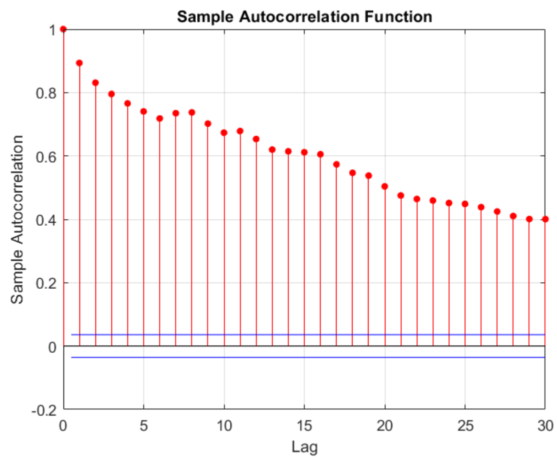

(c) Autocorrelation EUR/GBP basis

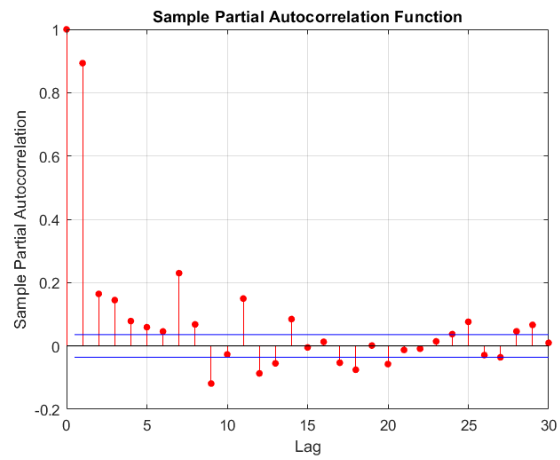

(d) Partial autocorrelation EUR/GBP basis

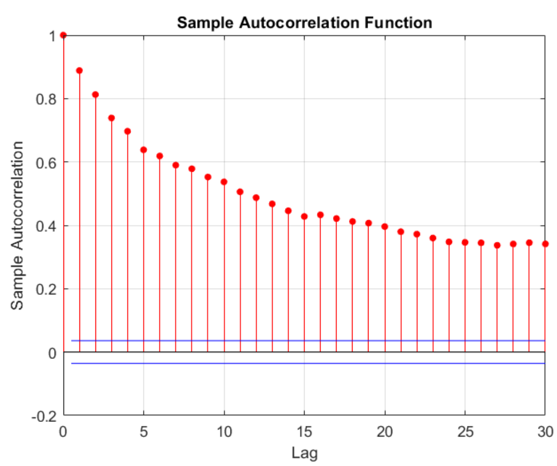

(e) Autocorrelation EUR/CHF basis

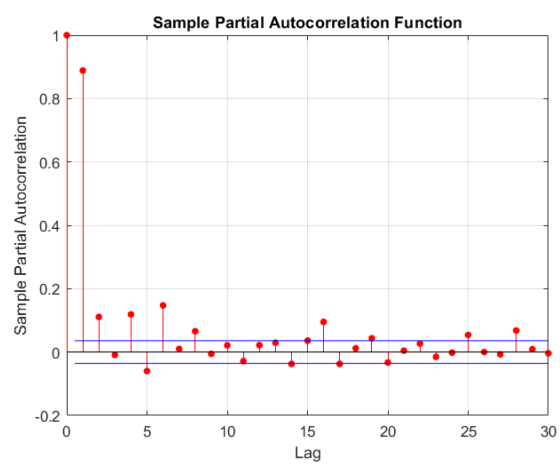

(f) Partial autocorrelation EUR/CHF basis

Figure 3: Correlation analysis of the one-month basis. *Note*: The blue line indicates significance at a 5% level.

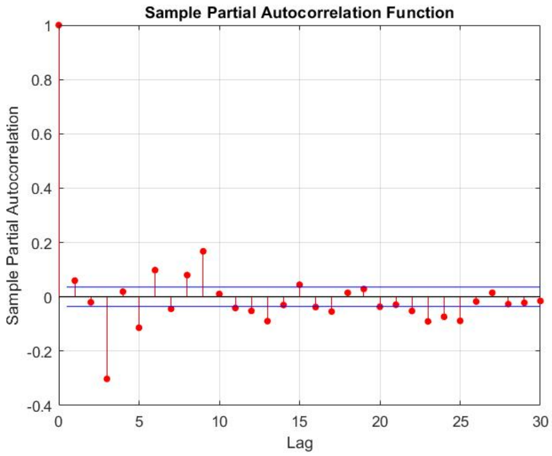

(a) Autocorrelation EUR/USD basis

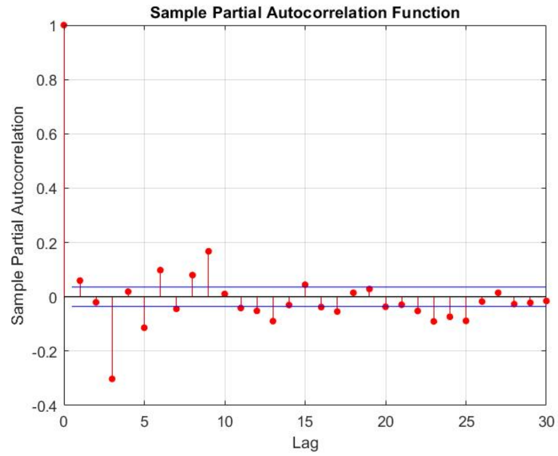

(b) Partial autocorrelation EUR/USD basis

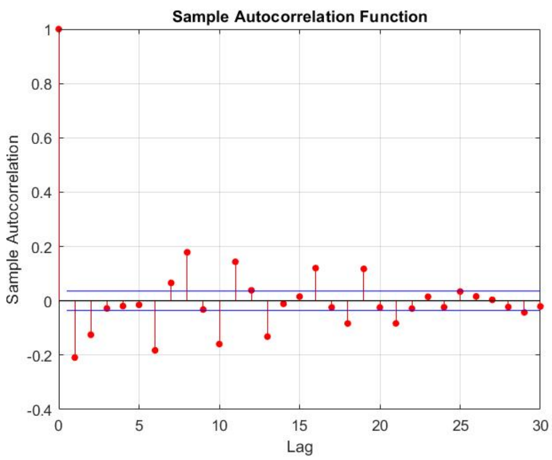

(c) Autocorrelation EUR/GBP basis

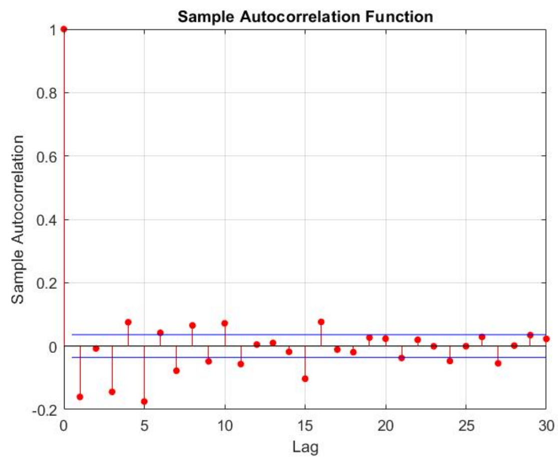

(d) Partial autocorrelation EUR/GBP basis

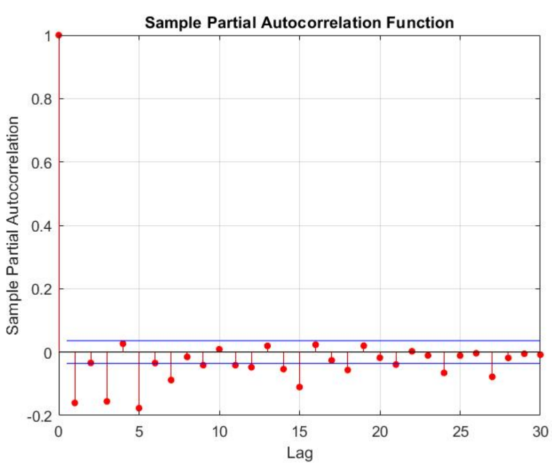

(e) Autocorrelation EUR/CHF basis

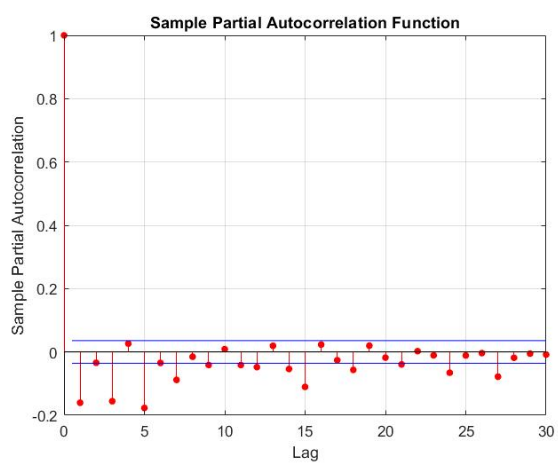

(f) Partial autocorrelation EUR/CHF basis

Figure 4: Correlation analysis of the first differences of the one-month basis. *Note*: The blue line indicates significance at a 5% level.

and a HAR(3) model.

We also consider the autocorrelation graphs for first differences, which are included in Figure 4. According to the Box and Jenkins (1976) methodology, (partial) autocorrelation series which show decaying features point into the direction of a linear autoregressive or moving average model. In Figure 4, we do not observe any of such decaying (partial) autocorrelations for the three series under consideration.

There is however one particular important model which is based on the first differences: the random walk. A random walk model, which assumes that the next data point is equal to the previous data point is equal to an ARIMA(0,1,0) model with no constant. This model is used as the benchmark with which we compare the forecasts that are based on the various models.

In Table 2, we display the results of the application of the nested 5-fold cross-validation in a time-series context as outlined in the Methodology section. For the linear model of the forecast combination, we choose the random walk in line with Zhang and Hu (1998). The residuals of this linear model (RLM) are equal to the increments of the observed basis,

Table 2: Optimal hyperparameters according to the nested 5-fold cross validation

		1-day ahead		**5-day ahead**		**22-day ahead**	
		Lags	**Hidden**	**Lags**	**Hidden**	**Lags**	**Hidden**
Basis	**EUR/USD**	1	17	18	5	22	16
	EUR/GBP	5	2	1	11	22	7
	EUR/CHF	6	7	2	1	24	2
RLM	**EUR/USD**	10	2	17	6	15	5
	EUR/GBP	5	3	3	4	15	5
	EUR/CHF	5	8	15	2	11	41

Note: This table shows the results of the nested 5-fold cross-validation. We use the five last years of the in-sample period July 2006 to June 2018 as the five test sets, the year before the particular year as validation set and the preceding data in the sample as training set. The optimal hyperparameters are selected by choosing the hyperparameter set which yields the lowest average RMSE among the five years in the test sets. EUR/USD, EUR/GBP and EUR/CHF are the one-month basis series under consideration. Basis implies that the neural networks are applied on the series itself, whereas RLM implies that they are applied on the residuals of the linear model, in this case the random walk. For 1-day, 5-day and 22-day ahead forecasts, we present the optimal hyperparameter selection for the amount of lags and the number of hidden units.

as these are by definition equal to the residuals of the random walk. We note that the appropriate number of lags and hidden nodes differ substantially among forecast horizon, basis series and whether the level of the basis itself or the residuals of the linear model are considered. To construct out-of-sample forecasts, we use the hyperparameters displayed in Table 2 for each respective series, model and forecast horizon.

5.2 Out-of-sample forecasts

We now evaluate the forecast quality during the out-of-sample period. First, we present and discuss the forecast quality for each forecast horizon. Afterwards, we make some more general remarks and discuss the extent to which these forecasts are useful for the development of hedging strategies.

In Table 3, we observe the out-of-sample performance of the 1-day ahead forecasts. In general, we note that the EUR/USD basis appears to be more difficult to forecast than the EUR/CHF basis, which in turn is more difficult to predict than the EUR/GBP basis. No model can outperform the random walk for the EUR/USD basis and EUR/CHF basis at a 10% significance level. For the EUR/GBP basis the AR(2), AR(5), HAR(2), NN and NN-LM models outperform the random walk at a 10% significance level. For each series, there is a PT statistic that is significant at a 5% level for each series. The models which have the best predictive power for the direction of the basis are HAR(3), AR(5) and HAR(2) for the EUR/USD, EUR/GBP and EUR/CHF basis respectively.

In Table 4, we observe the out-of-sample performance of the 5-day ahead forecasts. We note that for the autoregressive models, indirect iterated forecasts result in more accurate predictions than direct forecasts. This is not something to be expected upfront, since usually, an iterated forecast implies an accumulation of individual forecast errors. Nevertheless, an explanation for this phenomenon could be that the mean-reverting feature of the autoregressive models suits well with the behaviour of the considered time series. The mean reversion taken into account by the forecast is stronger when individual forecasts are iterated than when a direct forecast is made.

Some models outperform the random walk at the 5-day ahead forecast horizon for the EUR/USD and EUR/GBP at a 5% significance level, whereas no model can do this for the

Table 3: 1-day ahead forecasts

		RSME	MAE	DM	PT
EUR/USD	**AR(1)**	7.616	2.953	-0.650	0.924
	AR(2)	7.532	2.953	-1.362	1.655*
	AR(5)	7.934	3.143	1.132	1.439
	AR(22)	8.037	3.276	1.403	1.235
	HAR(2)	7.668	2.957	-0.235	1.678*
	HAR(3)	7.667	2.947	-0.288	2.150**
	NN	8.521	3.258	1.599	-1.019
	NN-LM	8.179	3.181	2.281	0.636
EUR/GBP	**AR(1)**	2.869	1.560	-1.282	2.018**
	AR(2)	2.786	1.523	-1.802*	3.180**
	AR(5)	2.786	1.494	-1.717*	4.001**
	AR(22)	2.959	1.613	0.202	2.887**
	HAR(2)	2.783	1.492	-1.749*	4.086**
	HAR(3)	2.803	1.499	-1.443	2.404**
	NN	2.806	1.545	-1.758*	3.078**
	NN-LM	2.842	1.470	-1.818*	3.855**
EUR/CHF	**AR(1)**	6.681	2.014	-1.521	4.653**
	AR(2)	6.479	1.966	-1.368	6.417**
	AR(5)	6.554	2.047	-1.202	5.461**
	AR(22)	6.677	2.234	-0.930	4.666**
	HAR(2)	6.488	1.972	-1.454	5.468**
	HAR(3)	6.477	1.984	-1.477	4.949**
	NN	6.699	2.201	-0.661	-0.204
	NN-LM	7.755	2.664	1.769	-0.968

Note: This table contains performance measures of the 1-day ahead forecasts of the one-month EUR/USD basis, EUR/GBP basis and EUR/CHF basis during the out-of-sample period July 2018 to June 2021, which contains 783 observations. All estimations are made with an expanding window which starts at July 2006. The performance measures displayed are the Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Diebold and Mariano (2002) statistic (DM) and the Pesaran and Timmermann (1992) statistic (PT). These statistics follow a standard normal distribution under the null hypothesis. * implies significance at a 10% level, whereas ** implies significance at a 5% significance level. The autoregressive AR(1), AR(2), AR(5) and AR(22) models are estimated by maximum likelihood, while the heterogeneous autoregressive HAR(2) and HAR(3) models as introduced by Corsi (2009) are estimated via ordinary least squares. The NN is an ANN model which according to a previous 5-fold cross validation is based on the hyperparameters displayed in Table 2 for EUR/USD, EUR/GBP and EUR/CHF respectively. The model is trained with the Levenberg–Marquardt algorithm with data from July 2006 to June 2017 and validated based on July 2017 to June 2018, whereby we stop the training if six iterations do not lead to an improvement in the validation set. The NN-LM model combines the forecast of a NN with the random walk as linear model.

Table 4: 5-day ahead forecasts

		Indirect				Direct			
		RSME	**MAE**	**DM**	**PT**	**RSME**	**MAE**	**DM**	**PT**
EUR/USD	**AR(1)**	17.951	7.923	-2.899**	4.007**	23.416	11.268	4.014	2.379**
	AR(2)	17.810	7.980	-2.888**	3.558**	23.707	11.453	4.241	2.060**
	AR(5)	18.365	8.137	-1.751*	4.031**	23.920	11.543	4.321	1.493
	AR(22)	18.533	8.145	-1.656*	4.328**	18.227	8.249	-2.230**	3.258**
	HAR(2)	18.055	7.941	-2.712**	4.448**	18.408	8.193	-1.566	3.945**
	HAR(3)	17.966	7.816	-3.436**	4.243**	18.384	8.156	-1.896*	4.609**
	NN	21.519	10.242	4.437	-3.356	18.386	7.838	-2.617**	0.474
	NN-LM	19.779	8.567	1.202	1.607	19.511	8.361	1.683	-0.383
EUR/GBP	**AR(1)**	4.571	3.163	-0.600	5.834**	5.373	3.399	2.752	6.189**
	AR(2)	4.413	2.924	-1.910*	6.015**	5.384	3.375	2.775	5.513**
	AR(5)	4.377	2.727	-2.302**	6.088**	5.555	3.418	3.272	5.361**
	AR(22)	4.715	2.876	0.354	5.034**	4.678	2.876	0.126	5.210**
	HAR(2)	4.379	2.730	-2.272**	6.265**	4.526	2.757	-0.854	6.842**
	HAR(3)	4.457	2.766	-1.403	5.438**	4.528	2.762	-0.844	6.878**
	NN	4.435	2.951	-1.614	5.928**	4.575	2.978	-1.224	5.110**
	NN-LM	4.587	2.782	-0.939	2.405**	4.640	2.819	-0.904	3.699**
EUR/CHF	**AR(1)**	6.966	2.823	-1.959*	7.816**	7.238	3.103	-1.101	7.226**
	AR(2)	6.994	2.805	-1.958*	7.048**	7.015	3.097	-1.342	6.496**
	AR(5)	6.988	2.807	-1.939*	7.265**	6.936	3.114	-1.420	6.265**
	AR(22)	6.798	2.858	-1.824*	6.646**	6.786	2.867	-1.817*	6.426**
	HAR(2)	6.911	2.778	-1.958*	7.518**	6.685	2.751	-1.917*	7.399**
	HAR(3)	6.805	2.787	-1.908	7.400**	6.690	2.766	-1.910*	6.763**
	NN	7.344	3.294	-1.437	-0.955	7.094	2.942	-1.713*	5.525**
	NN-LM	9.394	6.093	1.857	1.011	8.344	3.255	-0.753	3.845**

Note: This table contains performance measures of the 5-day ahead forecasts of the one-month EUR/USD basis, EUR/GBP basis and EUR/CHF basis during the out-of-sample period July 2018 to June 2021, which contains 783 observations. We display both direct forecasts and indirect forecasts, consisting of five consecutive iterated 1-day ahead forecasts. All estimations are made with an expanding window which starts at July 2006. The performance measures displayed are the Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Diebold and Mariano (2002) statistic (DM) and the Pesaran and Timmermann (1992) statistic (PT). These statistics follow a standard normal distribution under the null hypothesis. * implies significance at a 10% level, whereas ** implies significance at a 5% significance level. The autoregressive AR(1), AR(2), AR(5) and AR(22) models are estimated by maximum likelihood, while the heterogeneous autoregressive HAR(2) and HAR(3) models as introduced by Corsi (2009) are estimated via ordinary least squares. The NN is an ANN model which according to a previous 5-fold cross validation is based on the hyperparameters displayed in Table 2 for EUR/USD, EUR/GBP and EUR/CHF respectively. The model is trained with the Levenberg–Marquardt algorithm with data from July 2006 to June 2017 and validated based on July 2017 to June 2018, whereby we stop the training if six iterations do not lead to an improvement in the validation set. The NN-LM model combines the forecast of a NN with the random walk as linear model.

Table 5: 22-day ahead forecasts

		Indirect				Direct			
		RSME	**MAE**	**DM**	**PT**	**RSME**	**MAE**	**DM**	**PT**
EUR/USD	**AR(1)**	27.213	16.044	-7.182**	11.394**	27.477	16.341	6.729	4.078**
	AR(2)	26.480	15.746	-7.167**	11.394**	27.488	16.356	6.742	4.118**
	AR(5)	28.514	16.680	-6.864**	11.652**	27.487	16.377	6.777	3.810**
	AR(22)	27.258	16.001	-7.357**	12.130**	25.927	16.023	6.552	4.116**
	HAR(2)	27.700	16.301	-7.086**	11.342**	26.903	16.049	6.418	2.919**
	HAR(3)	27.080	15.991	-7.348**	11.394**	27.100	16.422	6.560	3.156**
	NN	44.088	30.676	13.123	-6.578	29.033	18.385	7.890	3.143**
	NN-LM	37.137	20.863	0.970	-7.814	40.040	37.865	26.472	2.099**
EUR/GBP	**AR(1)**	6.616	4.970	-3.834**	9.976**	6.707	4.581	6.437	7.182**
	AR(2)	6.462	4.799	-4.539**	10.070**	6.699	4.543	6.401	7.655**
	AR(5)	6.335	4.556	-5.710**	10.259**	6.690	4.491	6.294	7.556**
	AR(22)	6.539	4.581	-5.799**	10.236**	6.440	4.273	5.787	7.083**
	HAR(2)	6.342	4.566	-5.694**	10.259**	6.402	4.496	5.989	5.049**
	HAR(3)	6.413	4.461	-5.950**	10.456**	6.357	4.441	5.818	5.368**
	NN	6.366	4.581	-4.838**	10.466**	6.627	4.339	6.225	6.766**
	NN-LM	7.503	4.885	-2.312**	-1.092	10.281	9.276	24.456	4.590**
EUR/CHF	**AR(1)**	6.853	3.160	-2.457**	12.262**	7.194	3.473	-2.155**	11.380**
	AR(2)	6.874	3.166	-2.471**	12.167**	7.143	3.454	-2.193**	11.663**
	AR(5)	6.890	3.171	-2.479**	12.301**	7.119	3.417	-2.207**	12.647**
	AR(22)	7.066	3.312	-2.430**	11.458**	7.119	3.421	-2.179**	11.449**
	HAR(2)	6.904	3.179	-2.478**	12.220**	7.083	3.344	-2.404**	11.425**
	HAR(3)	7.073	3.318	-2.429**	11.289**	7.035	3.312	-2.434**	11.190**
	NN	11.585	8.490	2.477	-7.604	6.838	3.160	-2.443**	12.114**
	NN-LM	20.857	19.242	21.165	0.982	10.070	4.824	3.008	3.693**

Note: This table contains performance measures of the 22-day ahead forecasts of the one-month EUR/USD basis, EUR/GBP basis and EUR/CHF basis during the out-of-sample period July 2018 to June 2021, which contains 783 observations. We display both direct forecasts and indirect forecasts, consisting of twenty-two consecutive iterated 1-day ahead forecasts. All estimations are made with an expanding window which starts at July 2006. The performance measures displayed are the Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Diebold and Mariano (2002) statistic (DM) and the Pesaran and Timmermann (1992) statistic (PT). These statistics follow a standard normal distribution under the null hypothesis. * implies significance at a 10% level, whereas ** implies significance at a 5% significance level. The autoregressive AR(1), AR(2), AR(5) and AR(22) models are estimated by maximum likelihood, while the heterogeneous autoregressive HAR(2) and HAR(3) models as introduced by Corsi (2009) are estimated via ordinary least squares. The NN is an ANN model which according to a previous 5-fold cross validation is based on the hyperparameters displayed in Table 2 for EUR/USD, EUR/GBP and EUR/CHF respectively. The model is trained with the Levenberg–Marquardt algorithm with data from July 2006 to June 2017 and validated based on July 2017 to June 2018, whereby we stop the training if six iterations do not lead to an improvement in the validation set. The NN-LM model combines the forecast of a NN with the random walk as linear model.

EUR/CHF basis. It is noteworthy that for the EUR/CHF basis the forecast quality of the 5-day ahead forecasts is not substantially worse than the quality of 1-day ahead forecasts displayed in Table 3.

In Table 5, we observe the out-of-sample performance of the 22-day ahead forecasts. For this forecast horizon we can draw similar conclusions as for the 5-day ahead forecasts, since all autoregressive model under consideration with iterated forecasts significantly outperform the random walk at a 5% significance level. Taking into account all forecasts at different time horizons, we note that it is possible to outperform the random walk for the basis series, albeit at different significance levels for the different series and forecast horizons.

The question is to what extent the models under consideration can help create an improved hedging strategy. A model can succeed in developing a new hedging strategy if it can predict a sudden increase or decrease in the basis. One must note that, due to the strong mean-reverting properties of the basis, the random walk is a relatively easy model to beat at a larger forecast horizon. The superiority of our forecasts at this forecast horizon compared to the random walk arises from predicting that the basis goes back to the mean, rather than forecasting sudden deviations from the mean. This is displayed in Figure 5, 6 and 7.

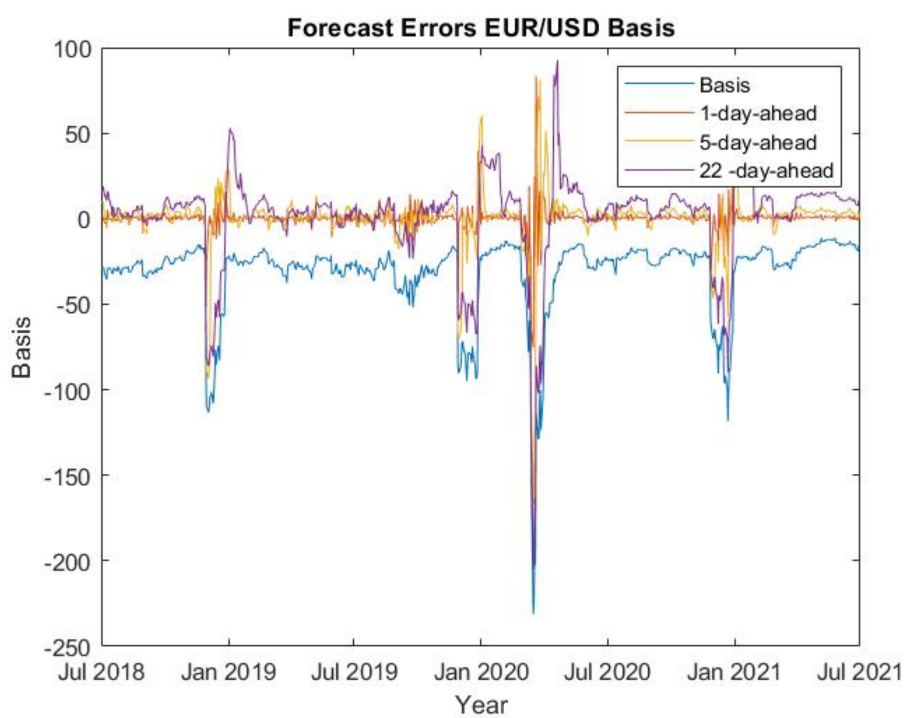

Figure 5: Forecast errors of the EUR/USD models with the lowest DM-statistic

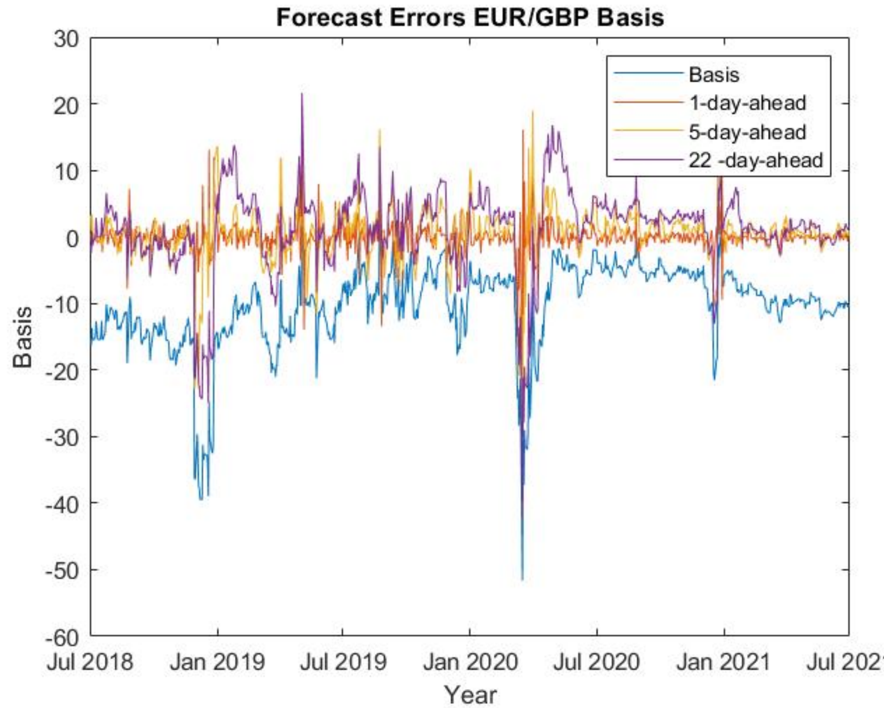

Figure 6: Forecast errors of the EUR/GBP models with the lowest DM-statistic

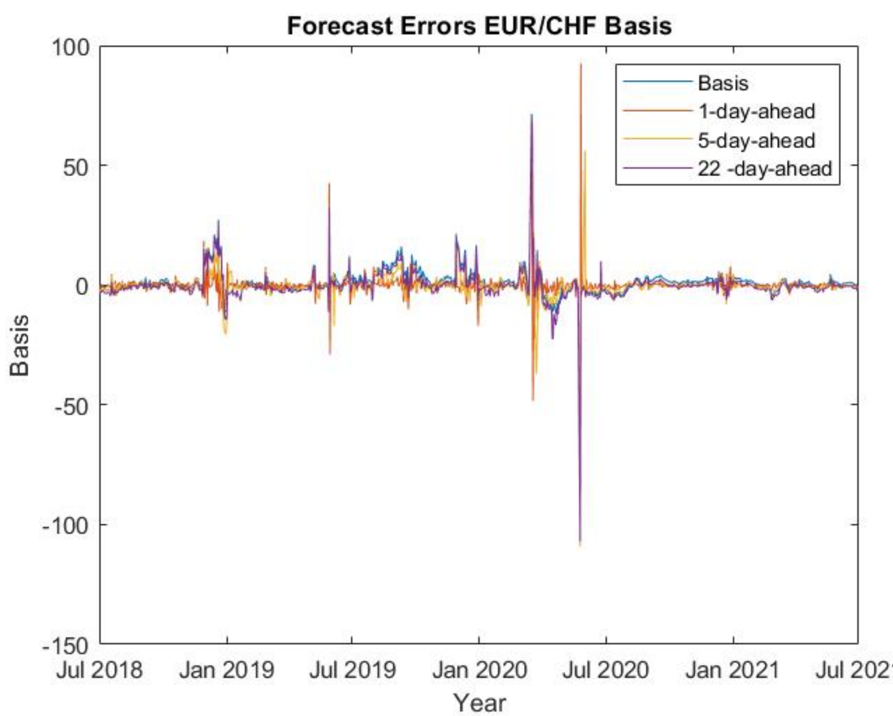

Figure 7: Forecast errors of the EUR/CHF models with the lowest DM-statistic

When we consider the best performing model for each series at each forecast horizon, as displayed in Figure 5, 6 and 7, we note that the models are not particularly successful in predicting certain spikes. Neither the random walk, nor the models we apply foresee sudden increases or decreases in the basis. However, when the basis is already at a point that is far away from the mean, the considered models outperform the random walk. A random walk makes the naive prediction that the next value is equal to the previous value. The mean-reverting AR and HAR models, and in some cases the neural networks, are more effective in predicting mean-reverting features.

If certain spikes in the basis can be predicted, one could close the forward position and roll the forward to a later maturity at an earlier moment, when the prediction shows an decrease of the basis. In that manner, a trade concerning the full amount of investments in the fund takes place at a less expensive moment. Unfortunately, such spikes cannot be predicted as outlined above. This leaves us with the question whether we can profit from the ability to model the mean-reverting nature of the basis, after a negative spike has occurred. In the case of a large predicted negative basis, one would like to postpone a roll, whereas in the case of a large predicted positive basis one would prefer to let the roll take place earlier.

However, a roll momentarily occurs at the latest possible moment before maturity. At some point, due to the imminent maturity of the forward, further postponement is not possible. The only way in which we could profit from knowing that the basis goes back to its mean is by delaying a full hedge. The size of these trades is however minor compared to the value of the portfolio such that the average basis incurred compared to the share class value as a whole would be relatively small. Potential profits are minimal, whereas the implementation would add quite some complexity.

Taking everything into account, we must conclude that we have no indication that a model-based hedging approach results in a strategy that substantially reduces the costs incurred by the cross-currency basis. Nevertheless, the insight that the basis is indeed mean-reverting could be very helpful in another context, which we elaborate upon in the Conclusion and discussion section.

5.3 Hedging strategies

In Tables 6, 7, 8 and 9, we present the results of various hedging strategies for the four UCITS share classes under consideration. For each fund, we consider four maturities, three hedge ratio bounds and whether or not a monthly hedge is performed as outlined in the Methodology section. Within a particular share class, the factor having the largest effect on the average basis incurred appears to be whether the forwards have a maturity of one month or three months. Having to roll a forward maturing monthly, rather than quarterly, substantially increases the absolute average basis costs incurred.

These costs are largely incurred during the roll of the forward two days before the maturity date since the trade at this time has the size of the full value of the invested amount. A quarterly roll implies a quarter less of these trades. Although the absolute amount of basis points incurred per trade is larger for a roll of a three-month forward, this is more than made up for by the reduction of the number of such trades. Broadly speaking, it appears to be beneficial for funds that suffer from the cross-currency basis to roll quarterly, whereas for funds that profit from it, one should maintain a one-month forward.

The differences of the average incurred basis costs between a roll at the end of the month or halfway through the month appear to be minor. Based on Figure 1, during the in-sample period, one notices that the absolute basis tends to deviate from its mean around most year-ends and on some quarter ends. However, during the considered out-of-sample period July 2018 to June 2021, this was apparently not the case on a structural basis.

It appears that when the absolute basis deviates from the mean, this usually occurs around the end of a quarter, but being in the end of a quarter does not necessarily imply a deviation of the basis from its mean. During the out-of-sample period, it seems that rolling halfway through the month even leads to a marginally higher absolute incurred basis. This would imply that the funds suffering from the basis should maintain a roll at the end of the month and the funds profiting from it could consider rolling halfway through the month. The gains of this strategy during the last three years were, however, not substantial.

Apart from the average basis incurred, we present three performance statistics of the strategy in each of the Tables 6, 7, 8 and 9. The alpha of the hedging strategy is calculated by determining the annualised return of the fund, minus the annualised return of the underlying

Table 6: Emerging Market Debt UCITS (USD/EUR)

Maturity	Bounds	Hedge	Alpha (%)	Exposure (%)	Trade (%)	Average Basis
1m	**96-104**	**Yes**	-1.243	1.059	15.052	1.151
	98-102	**Yes**	0.276	0.787	16.885	1.154
	99-101	**Yes**	0.387	0.476	22.251	1.160
	96-104	**No**	-1.117	1.653	11.126	1.142
	98-102	**No**	-0.401	0.949	13.874	1.157
	99-101	**No**	-0.347	0.508	18.848	1.159
1m-half	**96-104**	**Yes**	-1.531	0.936	14.791	1.262
	98-102	**Yes**	-0.873	0.681	17.147	1.269
	99-101	**Yes**	-2.052	0.485	21.990	1.271
	96-104	**No**	-1.454	1.652	10.864	1.258
	98-102	**No**	-0.882	0.986	14.136	1.266
	99-101	**No**	-1.214	0.512	19.241	1.272
3m	**96-104**	**Yes**	-2.619	1.366	6.283	0.430
	98-102	**Yes**	-1.755	0.898	8.770	0.434
	99-101	**Yes**	-1.734	0.492	14.791	0.435
	96-104	**No**	-2.519	1.686	4.843	0.426
	98-102	**No**	-1.764	0.935	7.592	0.441
	99-101	**No**	-1.729	0.506	13.351	0.435
3m-half	**96-104**	**Yes**	-1.491	1.447	6.152	0.415
	98-102	**Yes**	-1.563	0.809	9.031	0.423
	99-101	**Yes**	-1.539	0.514	14.136	0.424
	96-104	**No**	-1.447	1.654	4.712	0.419
	98-102	**No**	-1.572	0.959	7.853	0.423
	99-101	**No**	-1.547	0.512	13.482	0.424

Note: This table displays the performance of various hypothetical hedging strategies for the Emerging Market Debt UCITS in USD hedged towards EUR from July 2018 to June 2021. In the 'Maturity' column we display what the maturity of the traded forwards is. We consider a monthly roll on the end of the month (1m) and halfway the month (1m-half). Additionally, we consider a quarterly roll on the end of the quarter (3m) and halfway the last month of the quarter (3m-half). 'Bounds' indicates the hedge ratio which must be maintained. If this is violated, the hedge ratio is brought back to 100%. If 'Hedge' is 'Yes' we bring the hedge ratio back to 100% after the roll of the forward and if this is 'No', we only perform a full hedge when the 'Bounds' are violated. 'Alpha' equals the annualised return of the hedged share class minus the annualised return of the underlying assets. 'Exposure' is the daily average absolute deviation of the hedge ratio from 100%. 'Trades' is the share of days in which a trade occurs following the strategy expressed in percentages. 'Average Basis' is the average number of basis points gained per day according to the particular strategy. A positive number indicates a gain and a negative number implies a loss.

Table 7: Asset-Backed Securities UCITS (EUR/GBP)

Maturity	**Bounds**	**Hedge**	**Alpha (%)**	**Exposure (%)**	**Trades (%)**	**Average Basis**
1m	**96-104**	**Yes**	-0.973	0.182	14.136	-0.457
	98-102	**Yes**	0.555	0.167	14.267	-0.457
	99-101	**Yes**	0.555	0.153	14.660	-0.457
	96-104	**No**	0.585	0.922	9.686	-0.458
	98-102	**No**	0.552	0.679	9.948	-0.457
	99-101	**No**	0.539	0.411	10.602	-0.457
1m-half	**96-104**	**Yes**	0.912	0.231	14.267	-0.525
	98-102	**Yes**	0.868	0.189	14.398	-0.524
	99-101	**Yes**	1.250	0.150	14.921	-0.524
	96-104	**No**	0.861	0.919	9.817	-0.524
	98-102	**No**	0.837	0.697	10.209	-0.524
	99-101	**No**	0.820	0.408	10.864	-0.524
3m	**96-104**	**Yes**	0.859	0.431	4.712	-0.236
	98-102	**Yes**	0.809	0.347	5.105	-0.235
	99-101	**Yes**	0.794	0.287	5.497	-0.235
	96-104	**No**	0.819	0.916	3.534	-0.236
	98-102	**No**	0.791	0.699	3.927	-0.236
	99-101	**No**	0.776	0.414	4.581	-0.235
3m-half	**96-104**	**Yes**	0.765	0.446	4.843	-0.232
	98-102	**Yes**	0.729	0.308	5.105	-0.231
	99-101	**Yes**	0.713	0.275	5.628	-0.231
	96-104	**No**	0.742	0.917	3.534	-0.231
	98-102	**No**	0.725	0.708	3.927	-0.231
	99-101	**No**	0.702	0.414	4.581	-0.231

Note: This table displays the performance of various hypothetical hedging strategies for the Asset-Backed Securities UCITS in EUR hedged towards GBP from July 2018 to June 2021. In the 'Maturity' column we display what the maturity of the traded forwards is. We consider a monthly roll on the end of the month (1m) and halfway the month (1m-half). Additionally, we consider a quarterly roll on the end of the quarter (3m) and halfway the last month of the quarter (3m-half). 'Bounds' indicates the hedge ratio which must be maintained. If this is violated, the hedge ratio is brought back to 100%. If 'Hedge' is 'Yes' we bring the hedge ratio back to 100% after the roll of the forward and if this is 'No', we only perform a full hedge when the 'Bounds' are violated. 'Alpha' equals the annualised return of the hedged share class minus the annualised return of the underlying assets. 'Exposure' is the daily average absolute deviation of the hedge ratio from 100%. 'Trades' is the share of days in which a trade occurs following the strategy expressed in percentages. 'Average Basis' is the average number of basis points gained per day according to the particular strategy. A positive number indicates a gain and a negative number implies a loss.

Table 8: Asset-Backed Securities UCITS (EUR/CHF)

Maturity	Bounds	Hedge	Alpha (%)	Exposure (%)	Trades (%)	Average Basis
1m	**96-104**	**Yes**	-0.282	0.164	14.267	0.081
	98-102	**Yes**	-1.078	0.154	14.267	0.082
	99-101	**Yes**	-0.279	0.143	14.660	0.079
	96-104	**No**	-0.296	1.157	9.555	0.080
	98-102	**No**	-1.015	0.662	10.079	0.080
	99-101	**No**	-0.280	0.338	10.995	0.079
1m-half	**96-104**	**Yes**	-0.034	0.213	14.398	0.066
	98-102	**Yes**	-0.035	0.179	14.398	0.065
	99-101	**Yes**	0.360	0.144	14.791	0.066
	96-104	**No**	-0.058	1.153	9.555	0.065
	98-102	**No**	0.648	0.657	10.079	0.065
	99-101	**No**	-0.032	0.337	10.995	0.065
3m	**96-104**	**Yes**	-1.937	0.339	4.712	0.035
	98-102	**Yes**	-1.222	0.286	4.974	0.034
	99-101	**Yes**	-0.521	0.229	5.497	0.032
	96-104	**No**	-0.528	1.195	3.272	0.033
	98-102	**No**	-1.234	0.651	3.796	0.034
	99-101	**No**	-0.520	0.339	4.712	0.032
3m-half	**96-104**	**Yes**	-0.024	0.420	4.974	0.035
	98-102	**Yes**	-0.035	0.343	4.974	0.034
	99-101	**Yes**	-0.025	0.232	5.759	0.034
	96-104	**No**	-0.051	1.158	3.272	0.034
	98-102	**No**	0.655	0.663	3.796	0.034
	99-101	**No**	-0.025	0.337	4.712	0.034

Note: This table displays the performance of various hypothetical hedging strategies for the Asset-Backed Securities UCITS in EUR hedged towards CHF from July 2018 to June 2021. In the 'Maturity' column we display what the maturity of the traded forwards is. We consider a monthly roll on the end of the month (1m) and halfway the month (1m-half). Additionally, we consider a quarterly roll on the end of the quarter (3m) and halfway the last month of the quarter (3m-half). 'Bounds' indicates the hedge ratio which must be maintained. If this is violated, the hedge ratio is brought back to 100%. If 'Hedge' is 'Yes' we bring the hedge ratio back to 100% after the roll of the forward and if this is 'No', we only perform a full hedge when the 'Bounds' are violated. 'Alpha' equals the annualised return of the hedged share class minus the annualised return of the underlying assets. 'Exposure' is the daily average absolute deviation of the hedge ratio from 100%. 'Trades' is the share of days in which a trade occurs following the strategy expressed in percentages. 'Average Basis' is the average number of basis points gained per day according to the particular strategy. A positive number indicates a gain and a negative number implies a loss.

Table 9: Asset-Backet Securities UCITS (EUR/USD)

Maturity	Bounds	Hedge	Alpha (%)	Exposure (%)	Trades (%)	Average Basis
1m	**96-104**	**Yes**	0.969	0.166	14.136	-1.352
	98-102	**Yes**	0.957	0.152	14.267	-1.350
	99-101	**Yes**	0.560	0.140	14.529	-1.348
	96-104	**No**	-0.452	0.958	9.555	-1.347
	98-102	**No**	1.005	0.688	9.948	-1.352
	99-101	**No**	0.617	0.410	10.733	-1.348
1m-half	**96-104**	**Yes**	1.353	0.214	14.267	-1.467
	98-102	**Yes**	1.357	0.190	14.398	-1.465
	99-101	**Yes**	1.341	0.153	14.791	-1.463
	96-104	**No**	1.352	0.907	9.686	-1.458
	98-102	**No**	1.382	0.686	9.948	-1.462
	99-101	**No**	1.343	0.409	10.733	-1.462
3m	**96-104**	**Yes**	2.323	0.376	4.712	-0.536
	98-102	**Yes**	1.591	0.326	4.843	-0.531
	99-101	**Yes**	2.339	0.238	5.628	-0.530
	96-104	**No**	0.959	0.951	3.272	-0.532
	98-102	**No**	2.367	0.694	3.665	-0.534
	99-101	**No**	2.344	0.409	4.450	-0.529
3m-half	**96-104**	**Yes**	1.401	0.389	4.843	-0.503
	98-102	**Yes**	1.406	0.353	5.105	-0.502
	99-101	**Yes**	1.377	0.294	5.497	-0.501
	96-104	**No**	1.389	0.951	3.403	-0.500
	98-102	**No**	1.408	0.684	3.665	-0.501
	99-101	**No**	1.380	0.408	4.450	-0.502

Note: This table displays the performance of various hypothetical hedging strategies for the Asset-Backed Securities UCITS in EUR hedged towards USD from July 2018 to June 2021. In the 'Maturity' column we display what the maturity of the traded forwards is. We consider a monthly roll on the end of the month (1m) and halfway the month (1m-half). Additionally, we consider a quarterly roll on the end of the quarter (3m) and halfway the last month of the quarter (3m-half). 'Bounds' indicates the hedge ratio which must be maintained. If this is violated, the hedge ratio is brought back to 100%. If 'Hedge' is 'Yes' we bring the hedge ratio back to 100% after the roll of the forward and if this is 'No', we only perform a full hedge when the 'Bounds' are violated. 'Alpha' equals the annualised return of the hedged share class minus the annualised return of the underlying assets. 'Exposure' is the daily average absolute deviation of the hedge ratio from 100%. 'Trades' is the share of days in which a trade occurs following the strategy expressed in percentages. 'Average Basis' is the average number of basis points gained per day according to the particular strategy. A positive number indicates a gain and a negative number implies a loss.

assets as shown in Equation 28. The alpha is caused by two factors: exposure to the target currency resulting in a certain positive or negative performance and the basis incurred. In general, the aim is to keep the currency exposure as close to zero as possible and to minimise the number of trades, without creating additional performance due to the hedging strategy.

In Table 6, 7, 8 and 9, maintaining a narrower hedge ratio and not hedging every month yields more currency exposure, which is something to be expected upfront. On the other hand, using three-month forwards, wider hedge ratio bounds and not implementing a monthly hedge, reduces the number of trades. The absolute alpha of a strategy with a three-month forward tends to be larger than the alpha of strategies implementing one-month forwards. Hence, the reduction of the average daily absolute basis incurred by switching from one-month to three-month forwards comes at the cost of slightly more absolute alpha created as a result of additional currency exposure.

It is interesting to note that a generally positive average basis does not necessarily imply a positive alpha, as can be observed in Table 6 for the USD share class hedged towards EUR. The contrasting also holds for the EUR share class hedged towards USD, as can be seen in Table 9. This result is somewhat counter-intuitive since a positive basis has a positive effect on alpha and a negative basis has a negative effect on alpha.

Nevertheless, the fact that the basis actually influences alpha can be observed as well. For instance, in Table 6 the basis is generally larger for the one-month forward strategies compared to the three-month forward strategies, as well as the alphas of the respective strategies. A similar result holds for Table 9. The apparent discrepancy is caused by the fact that according to the hedging strategies for Table 6, the exposure to depreciating currencies tends to be relatively high, which results mostly negative alphas. On the other hand, in Table 9 the exposure to a appreciating currencies is relatively high, resulting in mostly positive alphas.

Two notable results are the negative single alphas in Tables 7 and 9. When looking into the cause of these results, we note that at one point a daily hedge does not take place in these two instances, whereas for all other strategies such a hedge is triggered. This happens during March 2021, which was during unrest on the financial markets due to the Covid-19 outbreak. This led to a relatively large basis, large market movements in the underlying assets of the

UCITS and large movements in the exchange rates of various currencies.

As a consequence of the hedge not having taken place, there was quite a large exposure to EUR during a substantial depreciation of the currency. This exposure leads to a larger loss compared with all other strategies. This result is useful in the sense that it reminds us to be careful in drawing sharp conclusions on optimal strategies. Noise from an individual data point may have a substantial effect on the results of the strategy, which does not necessarily point into the direction of the strategy itself being the driving force behind that result.

6 Conclusion and discussion

All in all, we consider three forecast horizons for the one-month EUR/USD, EUR/GBP and EUR/CHF basis during the period July 2018 to June 2021. With regard to 1-day ahead predictions, no model can provide a forecast that performs better than a random walk at a 5% significance level. Only the EUR/GBP basis can be predicted more accurately than than the random walk benchmark at a 10% level, through the AR(2), AR(5), HAR(2) and NN models. Notably, 5-day ahead predictions, as well as 22-day ahead predictions, are more accurate through iterated 1-day ahead forecasts than direct forecasts. Broadly speaking, at these time horizons with iterated forecasts, the random walk can be beaten at a 5% significance level.

Although for the longer forecast horizons, some models make a significantly better prediction than the random walk. This does not necessarily imply that they can be profitably used within a model-based hedging strategy. In particular, the considered models tend to have some predictive quality regarding the mean reversion. Hence, the superior performance compared to the random walk for the longer time horizons is chiefly caused by the random walk being a naive model in this context, rather than the AR models, HAR models, neural networks and forecast combination being that sophisticated.

When we apply the current hedging strategy, which is common practice in the industry, major trades often occur at times when the absolute basis is relatively large. A model-driven hedging strategy only has added value when a model can foresee a sudden growth of the absolute basis, such that a forward roll can ensure that the trade takes place at a moment when the deviation of the CIP is smaller. With respect to more efficient hedging strategies, the evaluated models are unfortunately not so helpful, since neither the linear models nor the neural network models or their combination appear to have the ability to forecast such spikes in the basis.

Taking everything into consideration, trading forwards with a maturity of three months rather than those with a one-month maturity appears to yield a lower absolute incurred basis. It would thus be advisable to use longer maturities for funds that generally suffer from a negative basis, such as the Asset-Backed Securities UCITS with the GBP and USD as share class currency. The EUR/CHF primarily centres around zero and there is no good argumentation to deviate from the current strategy.

For the Emerging Market Debt UCITS, the basis is generally positive. We would advise maintaining a strategy based on one-month forwards, since its positive average basis incurred is larger than for a three-month forward strategy. Eliminating the monthly full hedge or changing the hedge ratio bounds does not have a substantial effect on the incurred basis. Based on the preferences of the asset manager, a trade-off must be made between costs of additional trades and risk caused by currency exposure, which can lead to either an increase or decrease of the value of the investments.

The fact that the models do not appear to be helpful in developing a cheaper hedging strategy, does not imply that the findings with regard to the cross-currency basis and the models do not give interesting insights. It appears that the basis has strong mean-reverting properties. Although it can increase or decrease substantially, after a spike, the basis usually reverts back to the mean and rarely deviates from its mean substantially for a longer period of time.

Being aware of this feature, one can choose to invest in forwards with a longer time to maturity and then in principle roll them a few weeks or month before their maturity. This would create the option to postpone a roll of the forward if the absolute basis is high at a particular point of time. A disadvantage of this alternative strategy is that forwards with a longer time to maturity tend to suffer from a larger absolute basis. Therefore, with additional flexibility comes a higher absolute basis incurred upfront. Apart from that, trading in forwards with a longer time to maturity in practice comes with some complexities concerning margin costs as well as collateral issues.

This research can be extended in several manners. One could also forecast the basis of different currency pairs and maturities, as well as investigate the correlations between the basis of different exchange rates. A major challenge arises from the enormous amount of data points that is required for this forecasting exercise. To determine the cross-currency basis at one day for a currency pair with a particular forward maturity, we require five data points: a spot rate, forward points, calendar days between the maturity of the spot and forward and two risk-free interest rates.

Another interesting extension of this research could be applying more dynamic strategies, where the hedging strategy itself is also subject to change depending on the short-term development of the basis. This also comes with the necessity of huge data collection, which may be costly. In addition to that, more complex hedging strategies must be sufficiently conceptually understandable in order to be able to offer the share class to the public. This may be quite a challenge when one makes use of a dynamic strategy.

Regarding the forecast exercise, it appears that the ANN models are not constantly better performing than traditional AR models. Perhaps other machine learning models which are not evaluated in this thesis are more successful in making predictions. However, one must note that such models could be even more complex than ANN models. Long-term short memory (LTSM) models might be more successful in foreseeing sudden spikes in the basis. Nevertheless, these models come with an increase in complexity and even more hyperparameters to select.

References

Borio, C. E., McCauley, R. N., McGuire, P., & Sushko, V. (2016). Covered interest parity lost: understanding the cross-currency basis. *BIS Quarterly Review September*.

Box, G. E., & Jenkins, G. M. (1976). Time series analysis. forecasting and control. *Holden-Day Series in Time Series Analysis*.

Campbell, J. Y., Serfaty-De Medeiros, K., & Viceira, L. M. (2010). Global currency hedging. *The Journal of Finance*, *65*(1), 87–121.

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, *7*(2), 174–196.

Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American statistical association*, *74*(366a), 427–431.

Diebold, F. X., & Mariano, R. S. (2002). Comparing predictive accuracy. *Journal of Business & economic statistics*, *20*(1), 134–144.

Diebold, F. X., & Shin, M. (2019). Machine learning for regularized survey forecast combination: Partially-egalitarian lasso and its derivatives. *International Journal of Forecasting*, *35*(4), 1679–1691.

Gavin, H. P. (2019). The levenberg-marquardt algorithm for nonlinear least squares curve-fitting problems. *Department of Civil and Environmental Engineering, Duke University*, 1–19.

Keynes, J. M. (1923). *A tract on monetary reform.* London, Macmillan.

Levenberg, K. (1944). A method for the solution of certain non-linear problems in least squares. *Quarterly of applied mathematics*, *2*(2), 164–168.

Marquardt, D. W. (1963). An algorithm for least-squares estimation of nonlinear parameters. *Journal of the society for Industrial and Applied Mathematics*, *11*(2), 431–441.

Nguyen, D., & Widrow, B. (1990). Improving the learning speed of 2-layer neural networks by choosing initial values of the adaptive weights. In *1990 ijcnn international joint conference on neural networks* (pp. 21–26).

Pesaran, M. H., & Timmermann, A. (1992). A simple nonparametric test of predictive performance. *Journal of Business & Economic Statistics*, *10*(4), 461–465.

Rapach, D. E., Strauss, J. K., & Zhou, G. (2010). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *The Review of Financial Studies*, *23*(2), 821–862.

Stock, J. H., & Watson, M. W. (2004). Combination forecasts of output growth in a seven-country data set. *Journal of forecasting*, *23*(6), 405–430.

Tang, Z., De Almeida, C., & Fishwick, P. A. (1991). Time series forecasting using neural networks vs. box-jenkins methodology. *Simulation*, *57*(5), 303–310.

Turing, A. M. (1948). *Intelligent machinery.* NPL. Mathematics Division.

Varma, S., & Simon, R. (2006). Bias in error estimation when using cross-validation for model selection. *BMC bioinformatics*, *7*(1), 1–8.

Zhang, G. P. (2003). Time series forecasting using a hybrid arima and neural network model. *Neurocomputing*, *50*, 159–175.

Zhang, G. P., & Hu, M. Y. (1998). Neural network forecasting of the british pound/us dollar exchange rate. *Omega*, *26*(4), 495–506.