Erasmus University Rotterdam
Erasmus School of Economics

Master Thesis Quantitative Finance [FEM61008]

Dynamic Currency Characteristic Hedging: An Alternative to the Unconditional Mean Variance Approach

Koen Roskamp (622663)

Supervisor:	prof.dr. D.J.C. van Dijk
Second assessor:	dr. O. Kleen
Date final version:	30th April 2024

Dynamic Currency Characteristic Hedging

Koen Roskamp

30th April 2024

Abstract

In an increasingly globalized world, foreign exchange risk, or currency risk, is a significant concern for investors and businesses. We research an original model in the field of foreign exchange hedging and propose a novel hedging approach called Dynamic Currency Characteristic Hedging (DCCH). It is mainly based on theory by Brandt et al. (2009) and uses differentials in characteristics of exchange rates to determine optimal dynamic hedging weights in an overlaying currency portfolio, aiming to maximize returns without substantial additional risk. We find that our model yields positive results when tested in sample compared to traditional naive hedging methods. However, out-of-sample tests reveal difficulties within our model. Our investor can take on extreme hedging weights leading to high returns, but evenly large losses. This brings about undesirable levels of volatility. Suggestions for future research include incorporating more and higher-frequency data, implementing forecasting methods, and adjusting the investor's objectives. The study contributes to the literature by exploring dynamic hedging strategies based on currency characteristics, highlighting both potential benefits and limitations.

Contents

1 Introduction

Foreign exchange risk, also known as FX risk or currency risk, is a significant concern for investors and businesses operating in an increasingly globalized world. It arises from the potential volatility and uncertainty in exchange rates between different currencies. When a portfolio includes investments denominated in foreign currencies, the value of these assets can fluctuate due to changes in exchange rates, leading to potential gains or losses for the investor.

In this paper we present a new method for hedging foreign exchange risk that differs from more traditional hedging strategies. Common naive strategies are leaving currency exposure unhedged or the complete opposite, fully hedging away currency exposure. Then there is also the more active strategy of using mean variance optimization to assign weights to stocks and foreign currencies to minimize the risk of the total portfolio. An important example of this strategy is found in the highly cited paper 'Global Currency Hedging' by Campbell et al. (2010). However, the main focus there is minimizing the unconditional variance with fixed hedging weights, but we argue that currency rates vary enough that it would be beneficial to minimize the conditional variance and to also exploit the exchange rate fluctuations to maximize returns. Although there is research available on dynamic hedging models, none make use of the findings by Brandt et al. (2009) in optimizing portfolios. Here it was found that using a function of a firm's characteristics to assign portfolio weights can give a reliable performance in- and out-of-sample. Thus deviating from the traditional approach of first modeling the joint distribution of returns and consequently computing the portfolio weights. An opportunity to contribute to the literature presents itself by diverging from Campbell et al. (2010) and determine foreign currency positions that vary over time. To determine these dynamic positions we apply the findings of Brandt et al. (2009) and use exchange rate characteristics to form the currency portfolio for our investor. In doing so, we aim to yield higher returns without substantial additional risk. Furthermore, our research also differs from current related literature in that it utilizes data from a long time period and looks at its different sub-periods. This can provide insights into the existence and evolution of relationships between FX rates and our chosen characteristics.

Our proposed new hedging method, named Dynamic Currency Characteristic Hedging (DCCH), dynamically assigns weights to our investor's hedging portfolio based on the characteristics of exchange rates. These are factors that are believed to influence the movements of exchange rates. To form the currency portfolio that hedges an investor's global equity portfolio, we base the weights on the characteristics of each country pair. This includes the following differentials: Inflation rate, Interest rate, Unemployment rate, Trade balance, and Industrial production. We begin by taking into account an equity investor with a portfolio consisting of stock indices from 7 different economies who aims to maximize her portfolio's return while maintaining an acceptable unconditional variance. The main focus is testing our model against some naive hedging benchmark models and assess its performance by various metrics.

We assume a one-month investment horizon. Furthermore, it is assumed that the US investor has a predetermined long position in foreign equities and wants to manage her exposure to foreign exchange by establishing the best possible hedging positions in an overlaying currency portfolio using forward currency contracts. The portfolio consists of the following seven economies: Euro (EUR), Japanese yen (JPY), British pound sterling (GBP), Swiss franc (CHF), Canadian dollar (CAD), Australian dollar (AUD) and United States dollar (USD). Any of these are available for our investor to take a position in. These positions are first determined by running our model over the entire sample ranging from January 1975 until November 2023 with monthly sampling frequency. We evaluate the performance of our model against different unconditional hedging models and analyze the relation between the characteristics and exchange rate developments. Finally, the out-of-sample performance is tested to establish if our model is robust and could be a viable option to use in real life. For this we use estimation windows, sliding and expanding, where the training window consists of our sample's first 5 years and the out-of-sample period ranges from January 1980 until November 2023. The performance is measured by the out-of-sample Sharpe ratio, Certainty Equivalent (CEQ) return and the Information Ratio (IR), besides looking at utility, return and volatility.

Our main findings are that this method yields positive results when tested in-sample, but disappoints in the out-of-sample tests. The limited information our investor has each time-step she determines new hedging weights causes her to take very large over- or under-hedging positions. This in turn leads to some large negative returns and a significant volatility increase over the benchmark models.

Future extensions could improve the research by utilizing more and higher frequency data, implement a forecasting method for the characteristic differentials, and the investor's objective can be altered to shift more to risk minimizing instead of maximizing return like a traditional hedge strategy would.

This paper is organized as follows: In section 2, relevant literature is presented to provide context for the study. Section 3 gives a description of the data used in the research. Section 4 discusses our model extensively besides the benchmark models and the ways the models are compared. Section 5 outlines the main experimental results. Section 6 summarizes the main findings and provides a discussion of the limitations of the study and suggestions for future research.

2 Literature

This research paper is intricately linked to literature exploring global currency hedging. The ideal level of currency hedging is debatable and contingent upon the driving forces behind investors' currency demands. Examining the global average, 39% of investors implement a no-hedging policy, 34% opt for a 50% hedging policy, 14% adhere to a 100% hedging policy, and 13% employ a different hedge ratio. (Michenaud & Solnik, 2008). In addition to impacting portfolio risk, currency exposure also has an impact on returns to the extent that foreign currency returns do not equal zero.

There are multiple reasons for why an investor would wish to hedge her internationally diversified portfolio, but they can simply be ordered into two demands: risk management and speculative. Country-specific risks like changes in inflation or interest rates can be reduced. If an investor's goal is only to minimize risk and we assume foreign currencies are uncorrelated with equities, then hedging currency risk fully would be optimal (Solnik, 1974). As such in the literature, currency hedging is often declared as a 'free lunch', meaning it reduces risk without affecting portfolio returns (Perold & Schulman, 1988).

The possibility of increasing returns by hedging foreign exchange risk should be just as interesting to investors. Glen and Jorion (1993) examine both the speculative and risk minimization motives for adding currency hedges to international portfolios, where one-month forward contracts are the hedging instrument. Different approaches are considered and they find that conditional currency hedging strategies significantly improve the performance of stock and bond portfolios. They yield substantially higher returns without additional risk.

Now, our research aims to meet both risk management and speculative demands. For this we use a model that is initially derived from an important paper by Campbell et al. (2010). A key part of this paper is estimating unconditional currency demands for equity investors. Their U.S. investor holds a portfolio of domestic and foreign stocks and can change her currency exposure by using a zero investment currency portfolio. To determine the positions in this currency portfolio a mean-variance analysis is used that minimizes the risk of the total portfolio. Dependent on the correlation between stock returns and exchange rates, the investor can choose to over- or under-hedge her currency exposure besides a regular full hedge. However, these correlations can of course change over time making a static hedge not ideal. Therefore, we deviate from Campbell et al. (2010) by letting our model determine dynamic hedging positions by maximizing an investor's utility at different time-steps. The internet appendix of Campbell et al. (2010) is also a helpful source as it covers an equation for the arbitrarily hedged dollar portfolio return given a certain hedge ratio. We partly use this equation and formulation in setting up our model.

More motivation for using dynamic weights is obtained from a recent paper by Opie and Riddiough (2020). They study a method that uses time-varying hedging weights to generate a superior investment performance in terms of reducing risk and increasing return. They find currency returns to be predictable. Their method dynamically hedges foreign exchange exposure in international bond and equity portfolios and makes use of forecastable global risk factors to predict currency returns. The two global currency risk factors used are: *carry* and *dollar*. The *carry* factor corresponds to returns made on a currency carry trade, while the *dollar* factor reflects the average return of a collection of currencies against the U.S. dollar. Also, Campbell et al. (2010) do touch on conditional hedging, but their conditional model only depends on interest differentials, which would predict excess returns on currencies. This practice contributes to the returns of the currency carry trade. In the end, they find relatively weak evidence that only interest differentials provide the necessary conditioning information for currency risk management.

To build a model that determines dynamic hedging weights, inspiration is in our case taken from Brandt et al. (2009). In their paper they propose a novel approach for optimizing large-scale equity portfolios. A function of the firm's characteristics determines the portfolio weight in each stock. The function's coefficients are computed by maximizing the investor's average utility of the portfolio's return. The characteristics they include for firms are: market capitalization, book-to-market ratio, and lagged return. Their method gives reliable performance both in- and out-of-sample, is computationally easy, is simple to modify, and is easily extended. Even they themselves argue that a similar approach can be used to form currency portfolios based on the characteristics of each country pair. Here they already suggest the following possible characteristics: interest rate, inflation rate and trade balance.

3 Data

All data series used are available at a monthly frequency. To test the robustness of our results, we want to look at multiple different sub-periods. Therefore, the aim is to get the largest possible sample period ranging from 1975 until present day 2023. Of course, prior to 1999 the Euro was not in use yet. Therefore, we use Germany and the Deutsche Mark as proxy for earlier years. We base this on results in Campbell et al. (2010). Here, before 1999 they define 'Euroland' as a value-weighted stock basket that includes Germany, France, Italy and the Netherlands. But they also conduct their analysis including only Germany in Euroland and find similar results. From now on we will simply refer to Europe or EU.

To test our hedging model we firstly create a portfolio consisting of equity from the following seven economies: Euro, Japanese yen, British pound sterling, Swiss franc, Canadian dollar, Australian dollar and U.S. dollar. The equity portfolio consists of stock indices from Morgan Stanley Capital International (MSCI). The historical data for these equity indices are obtained via the International Monetary Fund (IMF). Another obvious necessity are the foreign exchange rate data between our domestic currency the USD and the other 6 economies. These are also from MSCI and obtained via the IMF. The returns across FX, equity, and money markets are summarized in table 1 over the sample period. There is considerable variability in risk-free rates across the different countries, with most notably the British pound and the Australian dollar giving significantly higher interest rates in comparison to the Swiss franc and Japanese yen. Excess stock returns are highest for Switzerland and the United States, followed by Europe. We see that Australia displays the lowest returns which is to be expected with its high risk-free rate. Uncovered Interest

Parity (UIP) suggests that changes in the exchange rate should exactly offset variations in risk-free rate differentials (Engel, 2014). We see that this is somewhat the case. We also see that the countries with the lowest excess currency returns, Japan and Switzerland, also have the lowest interest rates. They have appreciated the most against the US dollar as seen in the exchange rate changes. Meanwhile the countries with the highest interest rate, namely Canada, Australia and the United Kingdom, have depreciated against the US dollar.

Table 1: Summary Statistics

	United States	Europe	Japan	United Kingdom	Canada	Australia	Switzerland
				Risk-free rate(%, annualized)			
mean	4.62	3.78	2.60	6.06	5.67	7.13	2.41
std	3.94	3.13	3.34	4.64	4.51	4.82	2.76
				Excess stock returns in local currency (%, annualized)			
mean	3.81	2.78	1.93	1.95	1.23	0.13	3.92
std	54.7	63.0	63.8	61.4	56.8	61.3	54.4
				Exchange rate returns (%, annualized)			
mean	-	3.07	2.06	-0.83	-0.43	-0.86	2.76
std	-	60.6	39.0	35.3	23.6	38.2	40.3
				Excess currency returns (%, annualized)			
mean	-	-0.56	-1.80	1.24	0.83	2.21	-1.93
std	-	5.59	4.33	3.70	2.44	4.11	4.49

All values are given in percent per annum. The data sample ranges from January 1975 until September 2023

Next, for each of the seven economies we need the characteristics that potentially explain the value of a currency compared to others. The motivation for the choice of these is discussed later in section 4.2.1, as well as their expected relation to exchange rates. The characteristics include:

- Inflation rates
- Interest rates
- Industrial production
- Unemployment rate
- Trade Balance

We use industrial production instead of Gross Domestic Product (GDP), because GDP data is often only available quarterly or annually. The historical data for these characteristics are obtained by the OECD, IMF, Federal Reserve Economic Data (FRED) and via Eikon/Datastream provided by the Erasmus University. Due to the unavailability of monthly data, Switzerland uses quarterly data for Industrial production, as well as Australia using quarterly data for Industrial production and CPI.

As seen later on, before using these in our model, the characteristics are standardized cross-sectionally. This is done to make sure our currency overlay portfolio is a zero investment portfolio. This does come with one interesting consequence. When taking the differential of each time-series with respect to the USD and standardizing afterwards, you get a dataset with the same values as you would get when standardizing the dataset immediately without taking the differential. This fortunately does not cause a problem, because it just means the relations between our home currency and the foreign currencies are already embedded in the standardized dataset.

Figure 1 describes the five exchange rate characteristics for each economy individually. It is important to note that this concerns the non-standardized characteristics, so not the form in which they are used in the DCCH model. We see clear correlation between all characteristics across economies, but figure 1e stands out because of its extreme outliers. Just as we do with CPI and industrial production, we take the annual growth rate for trade balance to get a percentage that is comparable across economies. This allows us to analyze how changes in trade balance relate to changes in exchange rates. Although the average value is around 60%, we unfortunately get these ugly extreme outliers. It remains to see how this will affect our results.

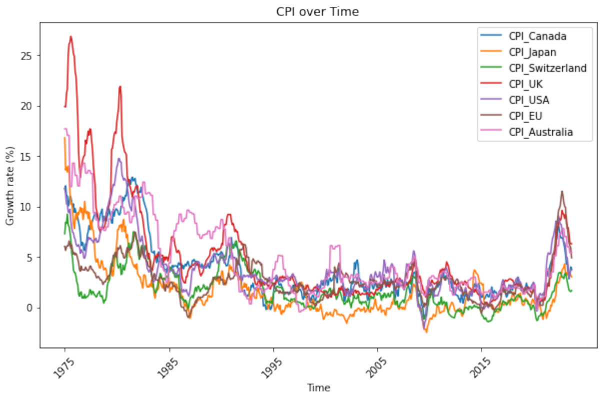

(a) The annual growth rate of the CPI of our seven economies, from January 1975 until September 2023

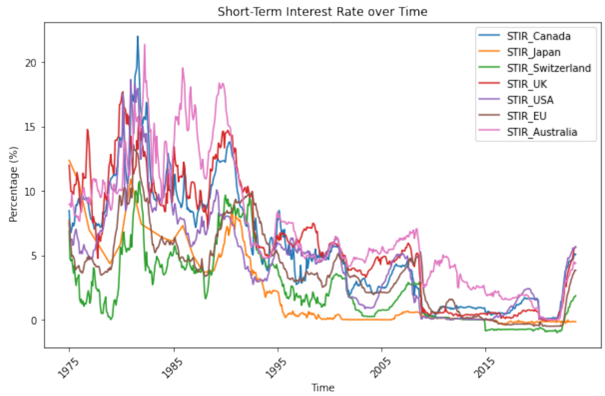

(b) The 3-month interest rate of our seven economies, from January 1975 until September 2023

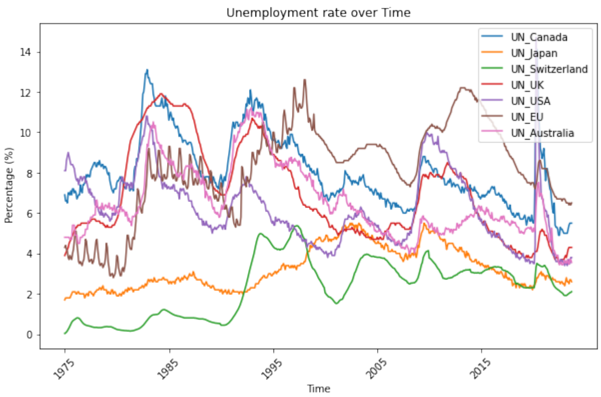

(c) The monthly unemployment rate of our seven economies, from January 1975 until September 2023

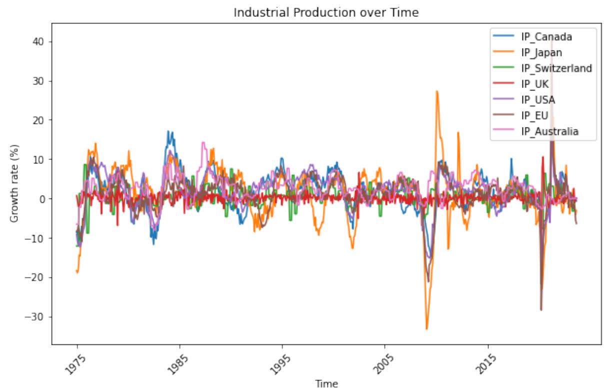

(d) The annual growth rate of the industrial production of our seven economies, from January 1975 until September 2023

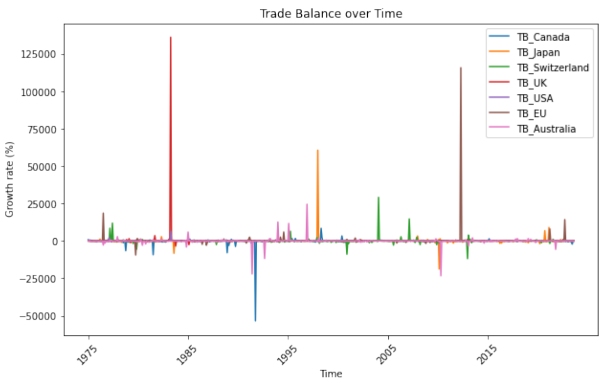

(e) The annual growth rate of the trade balance of our seven economies, from January 1975 until September 2023

Figure 1: Time-series of the five exchange rate characteristics for the period from January 1975 until September 2023

4 Methodology

4.1 General framework for arbitrarily hedged portfolio returns

We take an existing portfolio where the stocks have constant equal weights. Our US investor wants to hedge her FX exposure in her portfolio by overlaying a zero investment currency portfolio of foreign and domestic bills or, likewise, enter in a number of forward currency contracts.

Before our new model is discussed, it is necessary to define the way arbitrarily hedged portfolio returns are calculated. The basis we use for this lies in the 'Internet Appendix to Global Currency Hedging' of Campbell et al. (2010). First we define some variables. Let $R_{c,t+1}$ represent the return in domestic currency units resulting from holding stocks from country c from the beginning of the period t until the end $t+1$. Similarly, $S_{c,t+1}$ is the foreign currency spot exchange rate at $t+1$, denoted as the dollar price of foreign currency c at $t+1$. Because the USD is our domestic currency and it is given by $c=1$, $S_{1,t+1}=1$ for all t. In an unhedged portfolio our investor exchanges a dollar for $1/S_{c,t}$ foreign currency units at time t, which she then invests in the same foreign stock market. The return after one period $t+1$ is then exchanged back to $S_{c,t+1}$ dollars. This leads to the following unhedged portfolio return:

$$R^{uh}_{p,t+1} = \boldsymbol{R}'_{t+1}\boldsymbol{\omega}_t(\boldsymbol{S}_{t+1} \div \boldsymbol{S}_t) \tag{1}$$

Here, $\boldsymbol{\omega}_t$ is the (7×7) diagonal matrix of weights at time t. Because we use 7 different stock indices and $\sum_{c=1}^{7} \omega_{c,t} = 1$ for all t, each weight equals $1/7$. $\boldsymbol{R}_{t+1}$ and $\boldsymbol{S}_{t+1}$ are both (7×1) vectors.

Next, we take a look at hedged portfolio returns. Here we need the one-period forward exchange rate which is denoted by $F_{c,t}$ in dollars per foreign currency. Another new variable is $\Psi_{c,t}$. This is the dollar value of the amount of forward exchange rate contracts, that the investor enters into at time t for every dollar invested in her stock portfolio. Our investor is able to exchange $\Psi_{c,t}/S_{c,t}$ units of the return in foreign currency back into USD at a rate of $F_{c,t}$. The rest is exchanged at the standard spot exchange rate $S_{c,t+1}$ Therefore, a hedged portfolio return is given by:

$$R^{h}_{p,t+1} = \boldsymbol{R}'_{t+1}\boldsymbol{\omega}_t(\boldsymbol{S}_{t+1} \div \boldsymbol{S}_t) - \boldsymbol{\Psi}'_t(\boldsymbol{S}_{t+1} \div \boldsymbol{S}_t) + \boldsymbol{\Psi}'_t(\boldsymbol{F}_t \div \boldsymbol{S}_t), \tag{2}$$

where $\boldsymbol{F}_t$ is the vector of forward exchange rates and $\boldsymbol{\Psi}_t$ the vector of hedge ratios. One more step is needed before we can use this formula. We need to assume covered interest rate parity (CIRP). CIRP implies that interest rate differences between two nations should be equal to the difference between the forward and spot exchange rates in an efficient market with no arbitrage opportunities and minimal transaction costs (Sarno & Taylor, 2003). According to covered interest rate parity:

$$F_{c,t} = \frac{S_{c,t}(1 + I_{1,t})}{(1 + I_{c,t})}, \tag{3}$$

where $I_{1,t}$ is the domestic short-term interest rate, or the risk free rate, and $I_{c,t}$ is the short-term interest rate for the corresponding country c, both at the end of period t. This means our equation (2) can be rewritten as:

$$R^h_{p,t+1} = \boldsymbol{R}'_{t+1}\boldsymbol{\omega}_t(\boldsymbol{S}_{t+1} \div \boldsymbol{S}_t) - \boldsymbol{\Psi}'_t(\boldsymbol{S}_{t+1} \div \boldsymbol{S}_t) + \boldsymbol{\Psi}'_t\left[(1 + \boldsymbol{I}^d_t) \div (1 + \boldsymbol{I}_t)\right], \tag{4}$$

where $\boldsymbol{I}^d_t = I_{1,t}\mathbf{1}$ and $\boldsymbol{I}_t$ is a vector of the short-term interest rates. From this equation it is easy to see that a portfolio with $\Psi_{c,t} = 0$ leads to the unhedged portfolio in equation (1). Meanwhile, a fully hedged portfolio means $\psi_{c,t} = \omega_{c,t}$. Finally, because the currency overlay portfolio is a zero investment portfolio, our currency exposure has to add up to zero: $\boldsymbol{\omega}'_t\mathbf{1} - \boldsymbol{\Psi}'_t\mathbf{1} = 0$ or $\boldsymbol{\Psi}'_t\mathbf{1} = 1$. So, our domestic currency exposure is equal to the negative sum of foreign currency exposures and the hedge ratios add to one. The only way our investor can establish a long position in a currency is by borrowing in her domestic currency and using those funds to purchase bonds denominated in that currency.

4.2 DCCH model

The main approach to our model follows Brandt et al. (2009). Their basic idea is to configure the optimal weights for their equity portfolio as a function of every stocks' characteristics, like market capitalization and book-to-market ratio, such that the expected utility of the portfolio's return $r_{p,t+1}$ is maximized. In our case we also want to maximize the expected utility of our portfolio's return, but this portfolio consists of a constant equally weighted equity portfolio and a zero investment portfolio where the weights can be adjusted. The currency portfolio is configured by finding the optimal weights in forward currency contracts as a function of their characteristics. All calculations and tests for our model are done in Python.

For our hedged portfolio return, we use our equation (4) from the previous section. The main problem for our investor is choosing the hedge ratios Ψ_t that maximize the expected utility of its return $R^h_{p,t+1}$:

$$\begin{aligned}\max_{\Psi_t} \mathbb{E}_t \left[u(R^h_{p,t+1})\right] = \\ \mathbb{E}_t \left[u\left(\boldsymbol{R}'_{t+1}\boldsymbol{\omega}_t(\boldsymbol{S}_{t+1} \div \boldsymbol{S}_t) - \boldsymbol{\Psi}'_t(\boldsymbol{S}_{t+1} \div \boldsymbol{S}_t) + \boldsymbol{\Psi}'_t\left[(1+\boldsymbol{I}^d_t) \div (1+\boldsymbol{I}_t)\right]\right)\right]\end{aligned} \tag{5}$$

The optimal hedge ratios are given as a function of FX exposure in the main portfolio and a deviation determined by the currencies' characteristics:

$$\psi_{c,t} = \omega_{c,t} + \frac{1}{N_t}\theta^\top x_{c,t} \tag{6}$$

Here, $\omega_{c,t}$ is the weight of currency c at date t in our main portfolio. Because we use an equally weighted portfolio with indices from seven different currencies, every value is 1/7. θ is a vector of coefficients to be estimated when maximizing the utility function. The characteristics of currency c are represented in $x_{c,t}$, and they are standardized cross-sectionally to have zero mean and unit standard deviation. The standardization ensures we retain a zero-investment portfolio. Finally, the normalization term $\frac{1}{N_t}$ allows for the function to be used for a random and even possibly time-varying number of foreign currencies.

As mentioned in section 4.1, a fully hedged portfolio means $\psi_{c,t} = \omega_{c,t}$. Thus, formula (6) is essentially calculating optimal hedging weights as deviation from a full hedge, which is traditionally the optimal strategy. This will likely lead to so-called over-hedging or under-hedging in certain currencies.
The coefficients θ are found by maximizing the coinciding sample analogue:

$$\max_{\theta} \frac{1}{T}\sum_{t=0}^{T-1} u(R^h_{p,t+1}) \tag{7}$$

for a prestated utility function like constant relative risk aversion (CRRA). The utility functions that were chosen to generate results for are discussed in subsection 4.2.2.

4.2.1 Choosing the relevant characteristics

We want our model to make the best dynamic hedging decisions, therefore the characteristics chosen to predict the exchange rate are crucial. However, researching which characteristics have the most significant correlation with FX rates is a subject extensive enough for its own paper. Therefore, we have chosen characteristics that have been used in prior papers researching this subject (Engel & West, 2005) (Meese & Rogoff, 1983). The chosen characteristics are: Inflation rates, short-term interest rates (3-month), trade balance, industrial production and unemployment rate. The short-term interest rate is sometimes also referred to as the risk-free rate. During testing it will become clear which characteristics can have the most positive influence on our results.

Before conducting our research we have expectations on the relations between the characteristics and movements in exchange rates. These expectations are based on previous publications, mainly the textbook 'The Economics of Exchange Rates' by Sarno and Taylor (2003). There are two important theories that aim to explain movements in exchange rates. The first one comes forth from the efficient market hypothesis. This suggests that the expected FX gain from holding a foreign currency instead of the domestic currency should be equal to the difference in interest rates over the same period, the interest differential. This condition is known as the uncovered interest rate parity (UIP). We take the USD/EUR FX rate for illustration. If the interest rates in the EU rise while the rates in the US stay the same, it leads to an appreciation of the euro against the dollar. Our investor wants to hedge their EUR exposure against unexpected FX movements, but can also under-hedge if she anticipates that the EUR will appreciate leading to extra returns. Nevertheless, foreign exchange market data has been known to reject the efficient market hypothesis and thus the (un)covered interest parity. This is regardless of the quality of data or econometric techniques employed.

The second theory is Purchasing Power parity (PPP). The exchange rate under PPP is the rate at which the national price levels of two currencies become equivalent when denominated in a shared currency. So, in both economies the purchasing power of one currency unit would be equal. The foundation of the PPP condition lies in the law of one price (LOOP). This essentially means a good should have the exact same price across economies if denoted in the same currency. According to PPP, if inflation is higher in one country than in another, the currency of the country with lower inflation should appreciate to maintain parity or the currency of the country with higher inflation should depreciate. Our investor could (over)hedge the currency with higher inflation or/and (under)hedge the currency with lower inflation. Testing of the validity of the PPP hypothesis yields that it might be seen as a legit long-run parity condition among major economies.

The three remaining characteristics do not have a prominent theory that directly links them to exchange rates, but are of course quite interconnected with inflation rates, interest rates and economies themselves. An increase in trade balance means an economy's export increases relative to their import, leading to an appreciation of their currency. High unemployment rates typically indicate a weaker economy and a lower value of currency. In this case, central banks may also choose to lower interest rates, which affects FX rates. Strong industrial production indicates economic growth and a healthy economy. Higher economic growth can attract foreign investment, leading to an appreciation of the country's currency. Again, central banks may adjust interest rates to control inflation. Section 5 will show if the results are in line with our expectations.

4.2.2 Choosing the objective function

An important aspect of our model is the investor's objective function that ultimately decides the hedging weights. Because we base our model on Brandt et al. (2009), an obvious choice would be to copy the objective function they use. In this case, this concerns the standard constant relative risk aversion (CRRA) utility function, which collapses to log-utility in the case of $\gamma = 1$:

$$u(R^h_{p,t+1}) = \begin{cases} \frac{(1+R^h_{p,t+1})^{1-\gamma}}{1-\gamma} \text{ for } \gamma > 0, \gamma \neq 1 \\ log(1 + R^h_{p,t+1}) \text{ for } \gamma = 1 \end{cases} \tag{8}$$

However, this seems counter-intuitive as we are building a hedging model and the CRRA function does not include an explicit variance term. For our model we want to capture the fact that as opposed to fully hedging our investor can benefit from speculative returns besides only focusing on risk management. Therefore, we also run tests where currency hedges are selected to maximize a mean-variance investor's utility arising from the quadratic utility, according to the objective function:

$$u(R^h_{p,t+1}) = \mu_{p,t+1} - \frac{\gamma}{2}\sigma^2_{p,t+1} \tag{9}$$

For γ we consider three different values: 1, 5, 10. These correspond to low, medium and high levels of risk aversion an investor could have. Now, we have two different objective functions to test against certain benchmark models. These are discussed in the next section.

4.3 Benchmark models

To be able to make proper model evaluations, which will be discussed in section 4.4, it is necessary to define naive benchmark models to make a fair comparison.

Once again looking at formula (6), we find the main benchmark to consider to be the fully hedged portfolio. As our model takes this as its basis and determines desired deviations. Fully hedging all foreign exchange exposure is researched in classic papers like Solnik (1974) and Eun and Resnick (1988). Here it is argued that when the only goal is to minimize risk a full hedge is optimal, given that currencies and equities are uncorrelated. This strategy provides certainty in cash flows and protects against any unfavorable currency movements. However, it comes with the expense of hedging instruments.

Besides the main benchmark we want to include other hedging strategies for comparison. If we would not include these, we could not be sure of the practicality of our model. For example, if an unhedged strategy turns out to perform the best, a suggestion could be to change formula (6) into one that determines deviations from an unhedged currency portfolio. Thus, a second naive benchmark is to leave the portfolio entirely unhedged. Some investors, particularly those with a long-term investment horizon, may choose to leave their portfolios unhedged. They believe that over time, exchange rates display mean reversion, and the impact on the portfolio's performance becomes less significant (Froot, 1993). This approach is common when management of a company thinks they can withstand the effects of currency fluctuations or when they feel the cost of hedging is too high.

The last approach for completeness is partial hedging. Applying a 50/50 hedged/unhedged policy, in which 50% of total foreign exchange risk is hedged at all times, is a common industry practice. It was shown by Michenaud and Solnik (2008) that this currency hedging policy is the simplest that deals with possible investor's 'regret'. It allows them to participate in favorable currency movements to some extent while still being protected against significant adverse movements.

It is possible that relations between exchange rates and their characteristics change over time. Thus, our results will be evaluated over multiple sub-periods. For the chosen periods we look at key events that have had a significant influence on FX markets and financial markets in general. The sub-periods are:

- 01/1975 to 01-1985
- 01-1985 to 01-1999
- 01-1999 to 01-2007
- 01-2007 to 01-2020

The first period starts at the beginning of our dataset and cuts off before 1985. This is the year The Plaza Accord was signed. This accord meant the USD was to be depreciated in relation to the French franc, the German Deutsche Mark, the Japanese yen and the British pound sterling. This was done because of the USD appreciating 44 percent against other major currencies in the five years before and the trade balance had fallen to record lows (Frankel, 2015). The second period continues on from 1985 and ends when the euro is introduced. This period includes the signing of the Louvre Accord in 1987, where it was agreed to stabilize exchange rates after the Plaza Accord's intended depreciation of the USD. The third period encompasses the introduction of the euro as the official currency in 1999. This period reflects the initial stages of the euro's existence and its impact on FX rates. Furthermore, in section 3 we mention that before 1999 the Deutsche Mark is used as proxy for the euro. The last period begins in 2007, the year in which the Global Financial Crisis started. This includes the 'surprising' appreciation of the US dollar caused by a flight to safety into US treasury bills and reversal of carry trades (McCauley & McGuire, 2009). We end our final period before 2020. This means we do not include the official start of the Brexit and the COVID-19 pandemic. Naturally, in history there have been many more events that significantly impacted the dollar value with respect to other certain economies. We have tried to pick out the most noteworthy worldwide events.

4.4 Performance measures

This section looks at different tests used to evaluate the models and compare our model to the benchmark models. Given our utility framework and the objective functions found in section 4.2.2, the first performance measure is comparing portfolios based on their utility values. A higher utility value indicates a more preferred portfolio according to the investor's preferences. Other obvious points of interest are the annualized return and volatility. Some additional performance measures can be useful to add to our results. The following tests have been chosen based on related literature (Brandt et al., 2009)(Opie & Riddiough, 2020).

4.4.1 Sharpe ratio

One of the first performance measures we look at to compare our model to the benchmark models is the Sharpe ratio (SR). The out-of-sample Sharpe ratio is a frequently used metric for assessing a portfolio's statistical performance. It offers insights into how much return a portfolio generates compared to its risk (Sharpe, 1998). The difference in Sharpe Ratios is given by:

$$\Delta = SR_i - SR_j = \frac{\mu_i}{\sigma_i} - \frac{\mu_j}{\sigma_j}, \tag{10}$$

where μ_i is the out-of-sample average return of portfolio i over the risk free rate and σ_i is the out-of-sample standard deviation.

Besides comparing the Sharpe ratios of different portfolios, we need to test if their difference is significant. We follow the method presented in (Ledoit & Wolf, 2008). First, we set $\hat{\Delta} = \hat{SR}_i - \hat{SR}_j = f(\hat{v})$, with $\hat{v} = (\hat{\mu_i}, \hat{\mu_j}, \hat{\gamma_i}, \hat{\gamma_j})$. Here, $\hat{\mu}$ are the estimated first moments and $\hat{\gamma}$ are the estimated second moments. This means $\hat{\Delta} = f(\hat{v})$ with:

$$f(a, b, c, d) = \frac{a}{\sqrt{c - a^2}} - \frac{b}{\sqrt{d - b^2}} \tag{11}$$

Now with the assumption that $\sqrt{T}(\hat{v} - v) \xrightarrow{d} N(0, \Psi)$, we can find $\hat{\Psi}$, the estimated symmetric positive semi-definite matrix of $\hat{v}$. Then, using the Delta method we find:

$$\sqrt{T}(\hat{\Delta} - \Delta) \xrightarrow{d} N(0; \nabla' f(\hat{v})\Psi \nabla f(v)), \tag{12}$$

and the related standard error for $\hat{\Delta}$ is given by:

$$s\left(\hat{\Delta}\right) = \sqrt{\frac{\nabla' f(\hat{v})\hat{\Psi}\nabla f(\hat{v})}{T}} \tag{13}$$

We use the estimated covariance matrix of the return series as a plug in estimate of $\hat{\Psi}$. Conclusively, a two sided p-value for the null hypothesis $H_0 : \Delta = 0$ is given by:

$$\hat{p} = 2\Phi\left(-\frac{|\hat{\Delta}|}{s(\hat{\Delta})}\right) \tag{14}$$

4.4.2 Certainty Equivalent

The Certainty Equivalent (CE) return is the lowest risk-less return an investor accepts instead of a risky return. Thus, this metric can be seen as an expression of the price one pays for taking on risk. It accounts for both the variability in a portfolio's returns and the individual investor's risk aversion level, denoted as γ. Under low order Taylor approximations we can treat CE as:

$$CE \approx \mu_p - \frac{1}{2}\gamma\sigma_p^2 \tag{15}$$

When comparing two portfolios of investors with the same risk aversion, a higher CE indicates a greater preference for the associated strategy, as it represents the guaranteed amount that makes an investor indifferent between the certain outcome and the risky one.

4.4.3 Information Ratio

Another way to compare our model to the benchmark models is through the Information Ratio (IR). The information ratio serves as a performance metric for our method when compared to a benchmark method. A higher Information Ratio is generally desirable, as it indicates that the portfolio is generating higher excess returns than a benchmark for each unit of risk taken. The definition of the out-of-sample information ratio is as follows:

$$IR_b = \frac{\mu_p - \mu_b}{\sigma_{p-b}} = \frac{\textit{Active Return}}{\textit{Tracking Error}}, \tag{16}$$

where $\hat{\mu}_b$ is the out-of-sample average return of one of the benchmark portfolios and $\hat{\sigma}_{p-b}$ is the standard deviation of the difference in the two portfolios' returns.

4.5 Out-of-Sample performance

Our model outperforming classic strategies such as fully hedging in-sample is one thing, but to actually establish if our model would perform well in practice we need to perform out-of-sample tests to mimic a real-life scenario. This means that rather than estimating coefficients θ and consequently the hedging weights based on our entire sample from 01/1975 until 10/2023, we estimate θ based on sub-samples that determine the hedging weights for the following period. We do this in two ways, using a sliding window and an expanding window. The issue of structural changes in the characteristics can be effectively addressed through the utilization of a sliding window. This method relies on the latest observations to estimate our weights. Conversely, the expanding window approach uses information from all observations, making use of the entirety of available data up until $t+1$. Classic macroeconomics literature tends to prefer the expanding window for estimating parameters over the sliding window preferred by finance literature, see e.g. (Stock & Watson, 2003). Nonetheless, it remains interesting to compare the average θ coming out of both methods. Multiple runs with different sized training windows will be performed and compared. When assessing out-of-sample performance, the training window is not included in the results.

Because our results vary depending on what kind of window is used and the training window size, we have to beware of data snooping. Thus, when we compare our out-of-sample test with the most positive results to our benchmark, we need to check if this over- or under-performance is statistically significant. For this purpose we use White's Reality Check (White, 2000). H. White provides a procedure that tests the null hypothesis that the best model identified in a specification search does not offer any predictive advantage over a benchmark model, thereby mitigating the impact of data snooping. Or in our case, we test the null hypothesis that a portfolio strategy does not outperform a benchmark. We take:

$$\hat{f}_{t+1} = log[1 + R^{DCCH}_{p,t+1}] - log[1 + R^{fh}_{p,t+1}], \tag{17}$$

where we compare the DCCH portfolio returns with the fully hedged portfolio returns and test the null hypothesis: $H_0 : E(f^*) \leq 0$. For $R^{DCCH}_{p,t+1}$ we use our best out-of-sample results, which vary due to choice between an expanding or sliding window and training window size.

First, we use the stationary bootstrap method by Politis and Romano (1994) to create 1000 simulated series for the alternative data. It is required to choose the average block size for the Politis Romano stationary bootstrap. We use the rule of thumb $b = n^{1/3}$ for a balance between variance and bias, where n is the number of observations or the number of months in our case. This empirically derived rule offers a compromise between capturing temporal dependence and computational feasibility (Hall et al., 1995). Next, $\bar{V}_l$ is calculated as the square root of the number of observations multiplied by the mean of the loss function:

$$\bar{V}_l = n^{1/2} \bar{f} \tag{18}$$

Similarly, for each bootstrap iteration $\bar{V}_l^*$ is calculated as:

$$\bar{V}_l^* = n^{1/2} (\bar{f}^* - \bar{f}), \tag{19}$$

where $\bar{f}^*$ is the loss function of a bootstrap. Finally, the p-value is calculated by comparing $\bar{V}_l^*$ values against $\bar{V}_l$. The proportion of $\bar{V}_l^*$ values greater than $\bar{V}_l$ gives the p-value. Unlike previous calculations using self-written Python code, for the Reality Check we use a Matlab code sourced from Tilgenkamp (2013).

5 Results

For our first results, the DCCH model found in section 4.2 is compared against all naive hedging methods mentioned in 4.3. Our investor has a risk aversion of $\gamma = 5$ and has a mean-variance utility function. Return graphs of the benchmark models and the DCCH model can be found in appendix A.1. The first table 2 compares utility values across portfolios. Tables 3, 4 and 5 compare the annualized volatility, Sharpe ratio and return, respectively. We do this over our full period 01/1975 - 10/2023 and over the four sub-periods:

- I: 01/1975 to 01-1985
- II: 01-1985 to 01-1999
- III: 01-1999 to 01-2007
- IV: 01-2007 to 01-2020

The main performance measure in table 2 shows DCCH obtains a higher value for utility than all other strategies over the entire sample. This means a mean-variance investor with an average risk aversion of $\gamma = 5$ would prefer the DCCH model over the benchmarks. However, over the sub-periods there is variation in which strategy gives the highest utility. The remaining tables could provide more explanation for this phenomenon. Table 3 shows the volatility of portfolio returns. First thing we notice is that the naive hedging methods of half hedging and fully hedging successfully reduce the volatility over our whole sample and over all sub-periods. Also, it is clear that the DCCH does not achieve the volatility levels obtained by the static hedging strategies. Even having no hedge at all gives less volatility. Our mean-variance investor places importance on maximizing mean returns alongside the goal of minimizing variance, giving this as consequence. Table 4 indeed shows that even though our volatility is higher, we still outperform every benchmark when it comes to Sharpe ratio. This is the case for every sub-period. The explanation for this is found in table 5. Every benchmark is outperformed in annualized return by a significant percentage. The DCCH model successfully exploits opportunities that arise from varying exchange rates. We also see that the half hedge and full hedge surrender some return in favor of reducing volatility.

However, in sub-period II an unhedged strategy even outperforms the DCCH method return-wise. This is the period starting with the Plaza accord and two years later the Louvre accord which aimed to stabilize international currency markets. The USD was relatively stable during most of this period which could explain the higher returns in an unhedged portfolio. There are less swings in exchange rates for the DCCH method to exploit. Conversely, sub-period I sees more extreme fluctuations which were aimed to be solved by the mentioned Plaza accord. Sub-period III saw the USD depreciate, because of the introduction of the euro. Finally, sub-period IV dealt with the aftermath of the global financial crisis in 2008 and saw the USD appreciate with respect to most economies.

Table 2: Utility of Hedged Global Equity Portfolios

	No hedge	Half hedge	Full hedge	DCCH
Full Period	2.8×10^{-3}	3.5×10^{-3}	3.6×10^{-3}	5.4×10^{-3}
Subperiod I	5.2×10^{-3}	7.3×10^{-3}	9.0×10^{-3}	1.4×10^{-2}
Subperiod II	7.9×10^{-3}	6.5×10^{-3}	4.5×10^{-3}	4.7×10^{-3}
Subperiod III	2.6×10^{-3}	2.5×10^{-3}	2.0×10^{-3}	1.0×10^{-2}
Subperiod IV	-2.6×10^{-3}	-1.3×10^{-3}	-3.9×10^{-4}	-2.7×10^{-3}

Investor has a MV utility function with $\gamma = 5$

Table 3: Annualized Volatility (%) of Hedged Global Equity Portfolios

	No hedge	Half hedge	Full hedge	DCCH
Full Period	15.9	14.4	13.7	16.9
Subperiod I	15.5	13.9	13.0	14.9
Subperiod II	14.9	14.2	14.6	17.1
Subperiod III	13.6	12.5	12.1	13.8
Subperiod IV	17.2	15.1	13.4	19.2

Investor has a MV utility function with $\gamma = 5$

Table 4: Annualized Sharpe Ratio of Hedged Global Equity Portfolios

	No hedge	Half hedge	Full hedge	DCCH
Full Period	0.61	0.65	0.66	0.80
Subperiod I	0.79	0.98	1.15	1.47
Subperiod II	1.01	0.90	0.73	0.76
Subperiod III	0.58	0.56	0.50	1.25
Subperiod IV	0.25	0.27	0.30	0.31

Investor has a MV utility function with $\gamma = 5$

Table 5: Annualized Return (%) of Hedged Global Equity Portfolios

	No hedge	Half hedge	Full hedge	DCCH
Full Period	9.75	9.36	8.96	13.6
Subperiod I	12.2	13.6	15.0	22.0
Subperiod II	15.1	12.9	10.7	12.9
Subperiod III	7.87	6.97	6.08	17.3
Subperiod IV	4.24	4.11	3.99	5.98

Investor has a MV utility function with $\gamma = 5$

Now, we take a closer look at the specific outcomes of our DCCH model. Table 6 shows the results for running our model over the whole sample using a mean-variance utility function. Table 7 shows the same results but uses the CRRA utility function. We analyze and compare the two results. First of all, we direct our attention to the thetas. They decide the over- or under-hedging of each currency, relative to the full hedged portfolio. We see that these coefficients lie very close to each other for both utility functions. The cross-sectional standardization of the characteristics allows for a direct comparison of the magnitudes of the coefficients. When the risk aversion coefficient γ is small, the estimated thetas for short-term interest rate, consumer price index, trade balance, unemployment and industrial production all have a large absolute value. Especially the short-term interest rate is a clear number one. This coefficient together with the coefficient for trade balance start as a negative value and increase as our investor's risk aversion increases.

This means that for maximizing the return our investor under-hedges currencies that have a higher short-term interest rate and higher trade balance than our home currency USD. But, as variance becomes more important with a higher level of risk aversion the coefficients approach zero. This implies a connection between these characteristics and both average returns and risk. As risk aversion rises, the investor places more emphasis on these characteristics in relation to risk and reduces their significance. The remaining three thetas are positive values. This implies that our US investor wants to over-hedge currencies from countries which have a higher inflation rate, unemployment rate, and industrial production. Same as with the negative thetas, a higher risk aversion means these too approach zero in an effort to reduce variance. Calling back on theory from section 4.2.1, the signs of the estimated thetas are to a large extent in line with our expectations. We expected that we would under-hedge economies with a higher interest rate and trade balance and we would over-hedge economies with a higher inflation rate and unemployment rate. However, unlike our findings we predicted our model to under-hedge economies with a higher industrial production as their currency would appreciate. The positive sign shows that this is not the case. It is difficult to pinpoint what causes this disparity, but one possibility is the relation between industrial production and inflation rates. A higher industrial production leads to higher inflation which according to our theory results in a benefit from over-hedging. On average, the thetas lead to our US investor over-hedging their exposure to the dollar, the euro, the Swiss franc and the yen. The remaining currencies are under-hedged. However, these positions vary greatly over time and $\psi_{c,t}$ takes on large positive and negative positions for each currency. Further details on exposure per economy are found in appendix A.2.

Next, we see that the optimal hedging weight distribution is also subject to change depending on the level of risk aversion. Our investor takes extreme hedging weights when $\gamma = 1$. When the risk aversion increases to a more realistic level, the hedging weights also become more reasonable. Again, these weights do not differ much for both objective functions.

It comes as no surprise that the differences in optimal hedging weights result in equally notable differences in the distribution of the hedged portfolio returns and its standard deviation. For our lowest risk aversion the annualized average return is the highest but so is the annualized average volatility. Both of these decrease with the investor's level of risk aversion. The highest Sharpe ratio is found when $\gamma = 5$. This counts for both utility functions.

In this instance, the certainty equivalent return directly represents the utility function of mean-variance investors. Therefore, the next line which depicts the difference in utility from a fully hedged portfolio is also the difference in CEQ return. We get a positive increase for both investors across different levels of risk aversion. Lastly, the information ratio relative to the full hedge benchmark also displays a positive value emphasizing the outperformance of the DCCH method.

Table 6: Performance of MV utility function $\mu_{p,t} - \frac{\gamma}{2}\sigma^2_{p,t}$

	$\gamma = 1$	$\gamma = 5$	$\gamma = 10$
θ_{STIR}	-21.3	-3.99	-1.83
θ_{CPI}	10.9	2.70	1.67
θ_{TB}	-5.40	-1.14	-0.61
θ_{unemp}	3.22	0.64	0.31
θ_{IP}	3.51	0.94	0.62
max $\psi_{c,t}$	7.88	1.81	1.06
min $\psi_{c,t}$	-7.78	-1.44	-0.67
mean (%)	37.3	13.6	10.6
std (%)	56.0	16.9	14.0
SR	0.67	0.80	0.76
CE (%)	1.80	0.54	0.06
$\Delta u_{fullyhedged}$	1.1×10^{-2}	1.8×10^{-3}	1.0×10^{-3}
IR$_{fullyhedged}$	0.14	0.06	0.02

Table 7: Performance of CRRA utility function $\frac{(1+\mu_{p,t})^{1-\gamma}}{1-\gamma}$

	$\gamma = 1$	$\gamma = 5$	$\gamma = 10$
θ_{STIR}	-19.2	-3.92	-1.77
θ_{CPI}	9.98	2.82	1.89
θ_{TB}	-4.98	-1.14	-0.64
θ_{unemp}	2.81	0.69	0.46
θ_{IP}	3.62	1.14	0.95
max $\psi_{c,t}$	7.20	1.87	1.18
min $\psi_{c,t}$	-7.02	-1.41	-0.73
mean (%)	34.4	13.3	10.2
std (%)	50.8	16.6	13.9
SR	0.68	0.80	0.73
CE (%)	1.79	0.53	0.05
$\Delta u_{fullyhedged}$	1.0×10^{-3}	3.4×10^{-3}	1.0×10^{-3}
IR$_{fullyhedged}$	0.14	0.06	0.02

CRRA utility function with $\gamma = 1$ corresponds to log-utility

The increase in annualized Sharpe ratio is hopeful, but it is necessary to test if the difference from the benchmarks is significant. Following the method presented in section 4.4.1 yields the results presented in table 8, based on the full period Sharpe ratios. As evident from the row of p-values, the null hypothesis $H_0 : \Delta = 0$ is rejected when we compare the DCCH model to every benchmark.

Table 8: Significance test of Sharpe ratio's

	No hedge	Half hedge	Full hedge
$\hat{\boldsymbol{\Delta}}$	0.19	0.15	0.14
$\mathbf{s}\left(\hat{\boldsymbol{\Delta}}\right)$	0.03	0.03	0.03
$\hat{\mathbf{p}}$	0.0	0.0	0.0

To actually establish if our model would perform well in practice we also need to evaluate the out-of-sample results. Like mentioned before, these results are obtained by both a sliding and an expanding window. Multiple training window sizes have been tested on both methods. The results for the expanding window do not change significantly when the window size changes, besides the changes resulting from having a different out-of-sample period. Concerning the sliding window, a smaller window gives larger average returns but also more extreme volatility. Here we present the results for a mean-variance investor with a risk aversion of 5 using a training window of 5 years or 60 months. Results of out-of-sample tests using different sized training windows can be found in appendix A.3.

The results are shown below in table 9. A few things stand out. When we look at the average coefficients $\bar{\theta}$, there are some differences between them and the coefficients in table 6 with the in-sample results. For the sliding window, we get a negative coefficient for industrial production. A closer look at the development of the coefficients over time provides an answer.

This development is shown in figure 3b. The theta for industrial production fluctuates around zero and there does not seem to be a clear pattern. The same goes for θ_{TB} which also fluctuates around zero, but takes on larger values than the coefficient for industrial production. Both average coefficients lie close to zero in table 9, thus their significance remains in doubt. The second thing that catches our attention is the negative correlation between the coefficient for short-term interest rate and the coefficient for inflation rate. All in all, the optimal coefficients vary quite a lot when using a sliding window which leads to the non-optimal results in table 9. The large hedging weights lead to an extremely high average volatility. The negative certainty equivalent return and negative difference in utility particularly stand out, meaning our investor dislikes this strategy as opposed to a simple full hedge.

The out-of-sample performance improves when evaluated with an expanding window. This time something striking is the average coefficient for trade balance. It is positive, while the coefficient in table 6 over the whole sample is negative. We find our explanation in figure 3a. θ_{TB} does indeed start positive, but ends as the negative value we find in table 6. This indicates a change in the relationship between exchange rates and trade balance differentials or a weak relationship seeing the value is so close to zero. The most clear relations to exchange rates come from the short-term interest rate, inflation rate and unemployment differentials, although the coefficient for unemployment does drop down to near zero from 1999 onwards.

The Sharpe ratio and return are more favorable this time, but the standard deviation is still too high and again the negative certainty equivalent and utility difference are discouraging. In both cases it seems our model takes extreme hedging positions that cause the high returns and high volatility.

From figure 2 and table 9 we can see which characteristics are the main cause of these extreme weights. These characteristics also seem to be correlated with each other. Therefore, we run the out-of-sample tests with an expanding window again while each time omitting one of these a characteristics. The results are shown in table 10. As was expected, our model produces less extreme hedging weights leading to a lower average return and volatility. Unfortunately, the investor's utility is still lower than for a fully hedged portfolio. The signs of the coefficients are consistent across tests with the exception of θ_{CPI}, which changes from positive to negative when θ_{STIR} is omitted. The correlation between these two characteristics is probably what causes this. This shows that because the characteristics are economically connected, the relation between them and exchange rates can change based on which are included in the model.

Table 9: Out-of-sample performance using a sliding and expanding window

	Sliding window	Expanding window
$\bar{\theta}_{STIR}$	-6.71	-10.0
$\bar{\theta}_{CPI}$	1.24	5.96
$\bar{\theta}_{TB}$	-0.65	0.37
$\bar{\theta}_{unemp}$	5.53	3.97
$\bar{\theta}_{IP}$	-0.52	1.18
max $\psi_{c,t}$	9.72	7.47
min $\psi_{c,t}$	-9.69	-6.73
mean (%)	18.4	19.4
std (%)	40.1	28.9
SR	0.46	0.67
CE (%)	-1.82	-0.13
$\Delta u_{fullyhedged}$	-2.1 $\times 10^{-2}$	-4.0 $\times 10^{-3}$
$IR_{fullyhedged}$	0.06	0.09

Investor has a MV utility function with $\gamma = 5$

Table 10: Out-of-sample performance using an expanding window

	Without $\bar{\theta}_{STIR}$	Without $\bar{\theta}_{CPI}$	Without $\bar{\theta}_{unemp}$
$\bar{\theta}_{STIR}$	-	-4.49	-6.64
$\bar{\theta}_{CPI}$	-1.27	-	5.41
$\bar{\theta}_{TB}$	-0.58	-0.36	-0.08
$\bar{\theta}_{unemp}$	1.03	3.46	-
$\bar{\theta}_{IP}$	0.22	0.83	1.06
max $\psi_{c,t}$	2.42	4.46	4.90
min $\psi_{c,t}$	-1.38	-4.25	-5.72
mean (%)	8.30	16.3	12.4
std (%)	16.4	23.7	22.8
SR	0.51	0.69	0.54
CE (%)	0.13	0.19	-0.05
$\Delta u_{fullyhedged}$	-1.3 $\times 10^{-3}$	-9.0 $\times 10^{-4}$	-3.0 $\times 10^{-3}$
$IR_{fullyhedged}$	0.00	0.08	0.04

Investor has a MV utility function with $\gamma = 5$

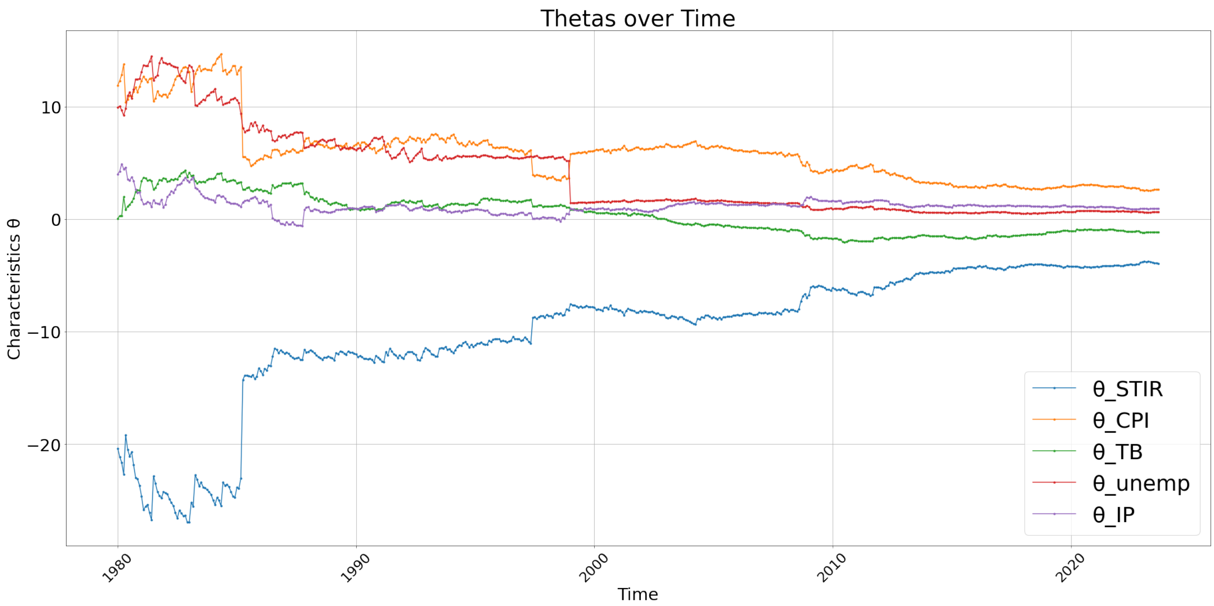

(a) The characteristic coefficients over time using an expanding window, from January 1980 until September 2023

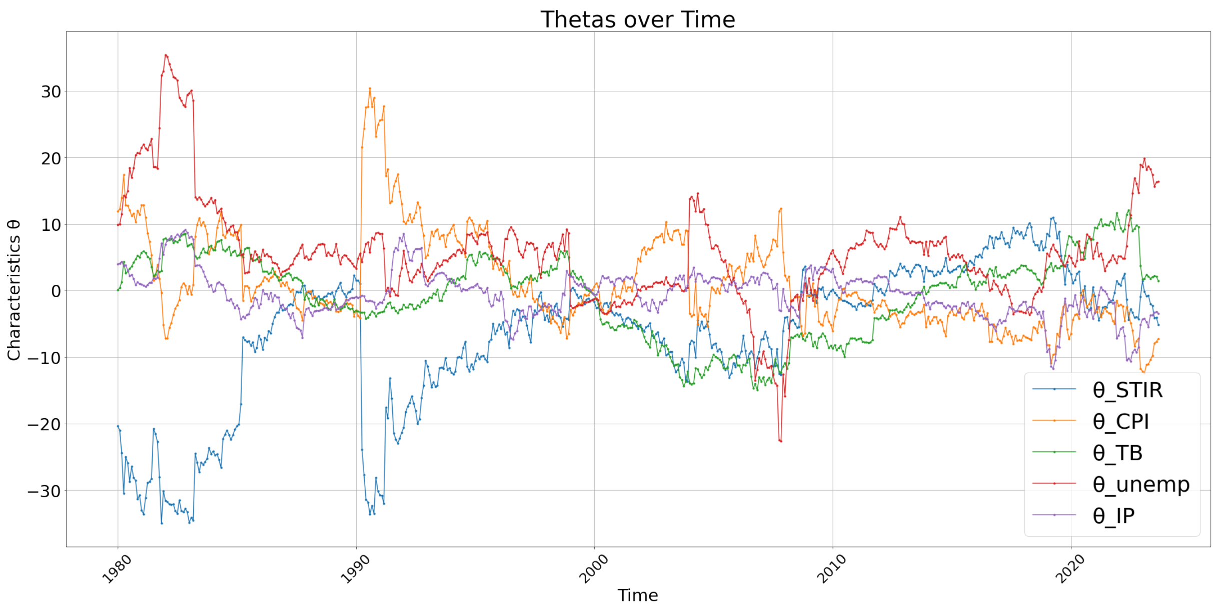

(b) The characteristic coefficients over time using a sliding window , from January 1980 until September 2023

Figure 2: Time-series of the five coefficients for the period from January 1980 until September 2023

We have yet to closely compare the out-of-sample performance of the DCCH method to the full hedge benchmark over the out-of-sample period 01-1980 until 09-2023. Tables 11, 12, 13 and 14 compare the utility, annualized volatility, Sharpe ratio and return, respectively. The in-sample performance of the DCCH method is added as additional benchmark. The out-of-sample DCCH uses an expanding window with a training window of 60 months. Lastly, sub-period I is omitted as half of it consists of the in-sample period. Here, the root of the poor out-of-sample performance is again made clear. If we would only compare the Sharpe ratios, one could conclude DCCH even outperforms a full hedge out-of-sample. However, the high returns are paired with even higher volatilities. For an investor with an average risk aversion this gives a low utility and a dislike towards this method with respect to a full hedge.

Finally, we want to check for data snooping by performing White's Reality Check. Following the steps presented in section 4.5, we test the null hypothesis that we have not found a model that outperforms the full-hedge benchmark in terms of out-of-sample return. Our 'best' model in this case is the DCCH using an expanding window with a training window size of 60 months. The out-of-sample period has $n = 525$ observations, so the average block size used is $b \approx 8$. We obtain a p-value of $p \approx 0.20$, so we fail to reject the null that our model does not outperform a fully hedged portfolio in out-of-sample returns. This is in line with the results presented in table 14. This also means that the observed performance difference between our model and the benchmark is unlikely to have occurred by random chance alone. Thus, for the difference in volatility this is sadly also the case.

Table 11: Utility of Hedged Global Equity Portfolios

	Full hedge	DCCH in-sample	DCCH out-of-sample
Full Period	2.9×10^{-3}	4.2×10^{-3}	-1.3×10^{-3}
Subperiod II	4.5×10^{-3}	4.7×10^{-3}	-9.5×10^{-3}
Subperiod III	2.0×10^{-3}	1.0×10^{-2}	1.0×10^{-2}
Subperiod IV	-3.9×10^{-4}	-2.7×10^{-3}	-6.8×10^{-3}

Investor has a MV utility function with $\gamma = 5$

Table 12: Annualized Volatility (%) of Hedged Global Equity Portfolios

	Full hedge	DCCH in-sample	DCCH out-of-sample
Full Period	13.6	17.0	28.9
Subperiod II	14.6	17.1	32.8
Subperiod III	12.1	13.8	17.8
Subperiod IV	13.4	19.2	24.3

Investor has a MV utility function with $\gamma = 5$

Table 13: Annualized Sharpe Ratio of Hedged Global Equity Portfolios

	Full hedge	DCCH in-sample	DCCH out-of-sample
Full Period	0.59	0.72	0.67
Subperiod II	0.73	0.76	0.47
Subperiod III	0.50	1.25	1.12
Subperiod IV	0.30	0.31	0.27

Investor has a MV utility function with $\gamma = 5$

Table 14: Annualized Return (%) of Hedged Global Equity Portfolios

	Full hedge	DCCH in-sample	DCCH out-of-sample
Full Period	8.03	12.3	19.4
Subperiod II	10.7	12.9	15.6
Subperiod III	6.08	17.3	20.0
Subperiod IV	3.99	5.98	6.56

Investor has a MV utility function with $\gamma = 5$

6 Conclusion

To conclude this research paper we review the performance of our new hedging method. First, we used our DCCH model over the entire sample to generate coefficients related to each characteristic that chose our time-varying hedging weights. For an average mean-variance investor this unfortunately does not decrease the annualized volatility. The increase from an unhedged strategy is not much, but a fully hedged portfolio does perform significantly better in this department. When we look at the speculative gains, the DCCH model does give a notable increase. In summary, this results in a relatively higher utility for our investor and a substantial improvement in the Sharpe ratio, with statistically significant findings.

However, the outcome of the out-of-sample tests tells a different story. The limited information our investor has each time-step she determines new hedging weights causes her to take large over- or under-hedging positions. Of course, this sometimes does result in high returns, but it evenly brings about large losses. So, the rate of return does not increase by the desired percentage in proportion to the rising volatility. This is why we get a negative certainty equivalent and negative difference in utility, which means this method is undesirable for our mean-variance investor.

Even though the performance of our DCCH model is disappointing, we do find that our chosen characteristics show the possibility of aiding in choosing hedging weights through their relation with exchange rates. Positive trade balance differentials and short-term interest rate differentials signal to an investor to generally under-hedge these foreign currencies, and vice versa. On the flip side, positive unemployment, industrial production and inflation rate differentials should be recognized by investors as cues for over-hedging in these non-domestic currencies. It should be noted that these judgements could be different for investors whose only goal in FX hedging is to decrease risk. Also, due to the interconnection of the characteristics the relationships between them and FX rates can change based on which of them are included.

As we conclude this study, it is essential to acknowledge and critically reflect upon its potential limitations and to look ahead and consider avenues for future research and extensions. To start, our main goal was to improve hedged portfolios in terms of decreasing risk and increasing return, but we mainly increased returns while still maintaining or even increasing the level of risk. Additional tests could be done where we take an investor that only focuses on minimizing risk and see if then our method does improve upon classic hedging approaches. A limitation in our study was the amount of data. For our portfolio consisting of equity from foreign economies we used stock indices. Though they provide a good indication of the average performance of every economy, when comparing portfolios' returns and their volatilities more data gives results that are more statistically significant. For a larger dataset one could, for example, use returns from 100 prominent stocks of each economy. Furthermore, data of higher frequency such as daily could also prove to be beneficial. In closing, a logical way out-of-sample decisions on hedging weights can be improved is by trying to forecast the characteristics that ultimately decide the hedging weights. For instance, a forecasting method using a machine learning technique like a K-Nearest Neighbours model (KNN).

References

Brandt, M. W., Santa-Clara, P., & Valkanov, R. (2009). Parametric Portfolio Policies: Exploiting Characteristics in the Cross-Section of Equity Returns. *The Review of Financial Studies*, *22*(9), 3411–3447.

Campbell, J. Y., SERFATY-DE MEDEIROS, K., & VICEIRA, L. M. (2010). Global currency hedging. *The Journal of Finance*, *65*(1), 87–121.

Engel, C. (2014). Exchange rates and interest parity. *Handbook of international economics*, *4*, 453–522.

Engel, C., & West, K. (2005). Exchange rates and fundamentals. *Journal of Political Economy*, *113*(3), 485–517.

Eun, C. S., & Resnick, B. G. (1988). Exchange rate uncertainty, forward contracts, and international portfolio selection. *The Journal of Finance*, *43*(1), 197–215.

Frankel, J. (2015). *The plaza accord, 30 years later* (tech. rep.). National Bureau of Economic Research.

Froot, K. (1993). Currency hedging over long horizons. *National Bureau of Economic Research, Inc, NBER Working Papers*.

Glen & Jorion. (1993). Currency hedging for international portfolios. *The Journal of Finance*, *48*(5), 1865–1886.

Hall, P., Horowitz, J. L., & Jing, B.-Y. (1995). On blocking rules for the bootstrap with dependent data. *Biometrika*, *82*(3), 561–574.

Ledoit, O., & Wolf, M. (2008). Robust performance hypothesis testing with the sharpe ratio. *Journal of empirical finance*, *69*(15), 860–859.

McCauley, R. N., & McGuire, P. (2009). Dollar appreciation in 2008: Safe haven, carry trades, dollar shortage and overhedging. *BIS Quarterly Review December*.

Meese, R. A., & Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics*, *14*(1), 3–24.

Michenaud, S., & Solnik, B. (2008). Applying regret theory to investment choices: Currency hedging decisions. *Journal of International Money and Finance*, *27*(5), 677–694.

Opie, W., & Riddiough, S. J. (2020). Global currency hedging with common risk factors. *Journal of Financial Economics*, *136*(3), 780–805.

Perold, A. F., & Schulman, E. C. (1988). The free lunch in currency hedging: Implications for investment policy and peformance standards. *Financial Analysts Journal*, *44*(3), 45–50.

Politis, D. N., & Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, *89*(428), 1303–1313.

Sarno, L., & Taylor, M. P. (2003). *The economics of exchange rates*. Cambridge University Press.

Sharpe, W. F. (1998). The sharpe ratio. *Streetwise–the Best of the Journal of Portfolio Management*, *3*, 169–85.

Solnik, B. H. (1974). An equilibrium model of the international capital market. *Journal of Economic Theory*, *8*(4), 500–524.

Stock, J. H., & Watson, M. W. (2003). Forecasting output and inflation: The role of asset prices. *Journal of economic literature*, *41*(3), 788–829.

Tilgenkamp, A. (2013, October). White reality check. https://nl.mathworks.com/matlabcentral/fileexchange/34306-white-reality-check

White, H. (2000). A reality check for data snooping. *Econometrica*, *68*(5), 1097–1126.

A Appendix

A.1 DCCH vs Benchmark returns over time

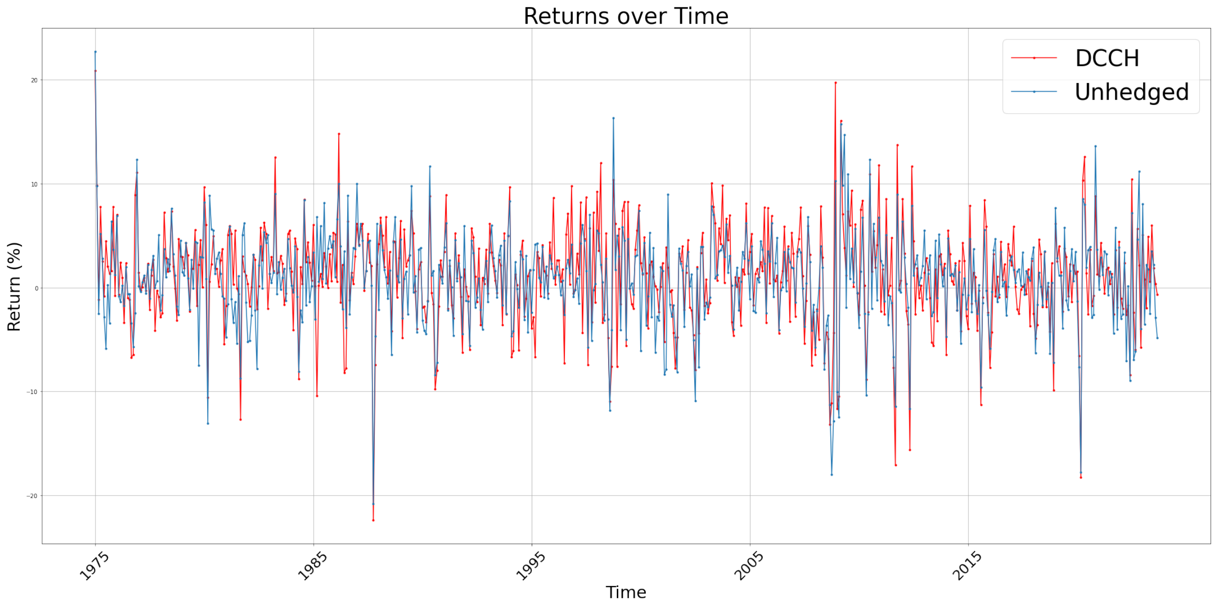

(a) Returns (%) over time, Unhedged & in-sample DCCH,
from January 1975 until September 2023

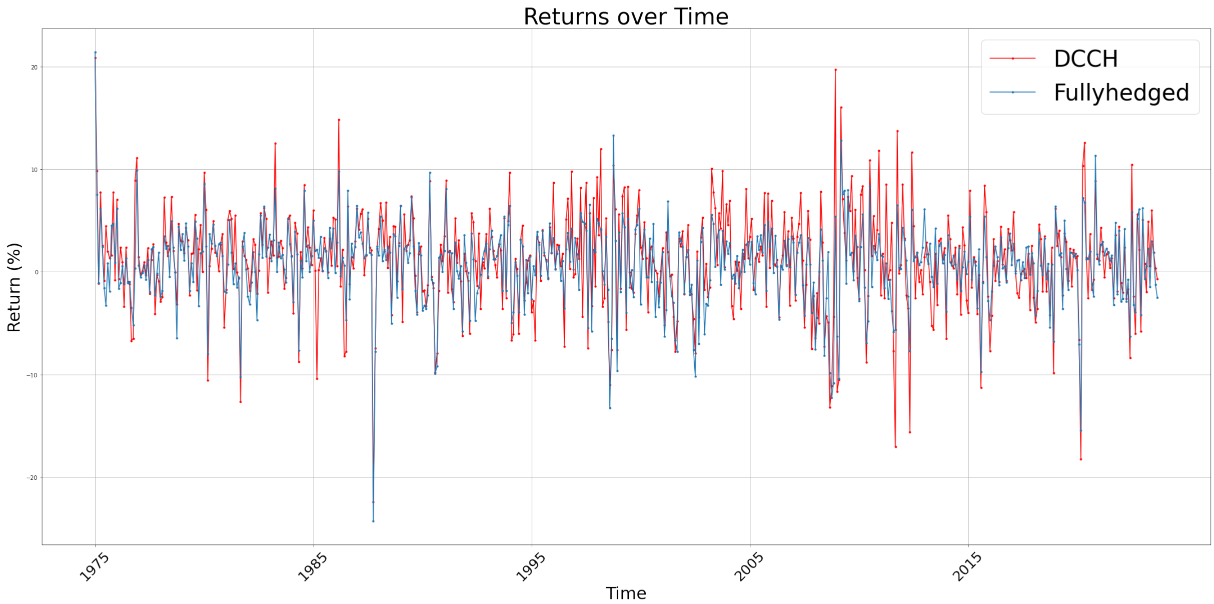

(b) Returns (%) over time, Fully hedged & in-sample DCCH,
from January 1975 until September 2023

Figure 3: Time-series of returns, the period from January 1975 until September 2023

A.2 Hedging ratios over time

Table 15: Average Hedging Ratios

	United States	Europe	Japan	United Kingdom	Canada	Australia	Switzerland
$\bar{\psi}_t$	0.31	0.39	0.19	0.01	0.10	-0.21	0.22
$max\ \psi_t$	1.32	1.81	1.33	1.23	1.07	0.94	1.15
$min\ \psi_t$	-0..85	–0.95	-1.15	-1.33	-1.44	-1.35	-0.71

The data sample ranges from January 1975 until September 2023

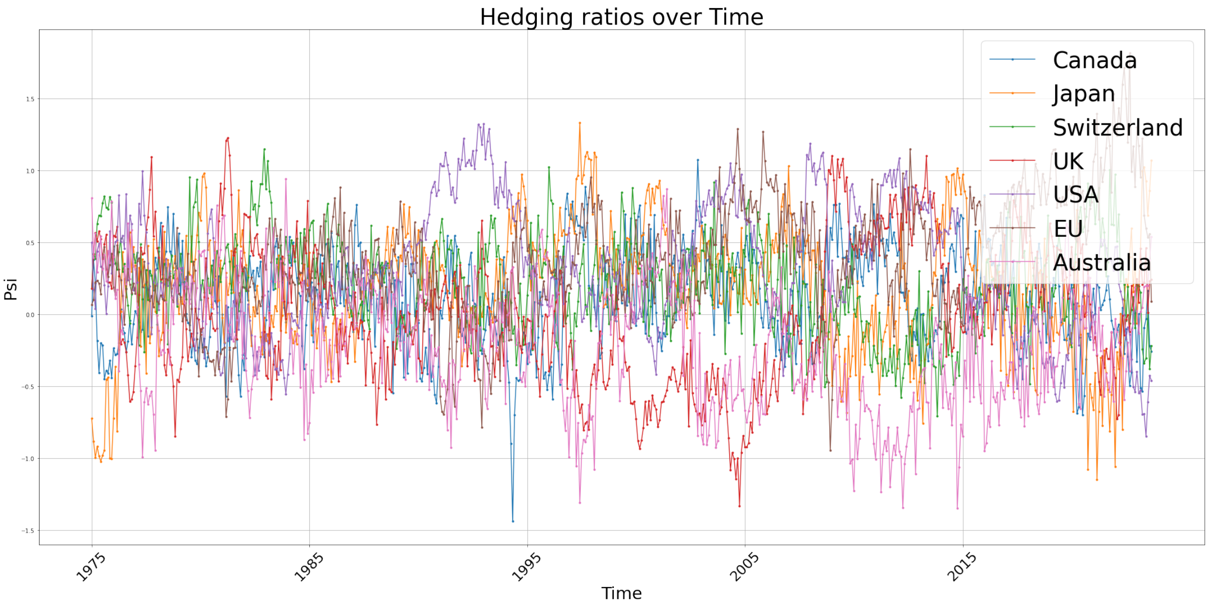

Figure 4: Economies' hedging weights over time, in the period from January 1975 until September 2023

A.3 OOS results for different training windows

Table 16: Out-of-sample performance using an expanding window

Training Window	15 M	30 M	60 M	90 M	120 M
mean (%)	22.8	21.6	19.4	15.6	12.6
std (%)	34.7	32.7	28.9	27.5	26.2
SR	0.66	0.65	0.67	0.57	0.48

Investor has a MV utility function with $\gamma = 5$

Table 17: Out-of-sample performance using a sliding window

Training Window	15 M	30 M	60 M	90 M	120 M
mean (%)	58.7	33.6	18.4	16.3	13.6
std (%)	152	69.6	40.1	32.1	28.9
SR	0.39	0.48	0.46	0.51	0.47

Investor has a MV utility function with $\gamma = 5$