On the role of interest rate differentials in the dynamic asymmetry of exchange rates

J. Hambuckers[1,†] and M. Ulm[1]

Abstract

Motivated by the lack of empirical evidence in favor of the uncovered interest rate parity rule, we revisit the informational content of interest rate differentials (IRD) to explain daily exchange rates variations. Proposing a novel version of a GARCH model, we allow for the IRD to impact on the time-varying conditional asymmetry of the depreciation rate. We find IRD to be a significant factor for the Euro (EUR), the Swiss franc (CHF), the Swedish Krona, the Japanese Yen and the British Pound. These findings empirically support currency crash theories, suggesting that the larger the difference between interest rates, the more likely the high-yield currency appreciates on average but also exhibits greater risk of a large depreciation. Compared to random walk and buy-and-hold benchmarks, we document superior out-of-sample mean returns of a trading rule exploiting IRD information for EUR and CHF.

Keywords: Exchange rate, interest rate differential, GARCH, dynamic asymmetry.

JEL: C53, C58, C22, F31.

[1] University of Liège - HEC Liège, Belgium.

[†] Corresponding author: jhambuckers@uliege.be.

Acknowledgements: we are particularly grateful to the Editor (Prof. S. Mallick) and two anonymous referees for their comments and suggestions. J. Hambuckers acknowledges the financial support from the DFG within the framework of RTG 1644 *Scaling Problems in Statistics* and of the National Bank of Belgium (project REFEX). Parts of the research for this paper were conducted while the authors were affiliated with the University of Göttingen (Germany), Chairs of Statistics and Econometrics.

1 Introduction

The exact nature of the relationship between short-term interest rate differential (IRD) and currency depreciation remains an ongoing debate. Economic theory postulates that IRD and foreign exchange rates are linked over time via the uncovered interest rate parity (UIP) rule. In this framework, the currency of the high yield economy is expected to depreciate. Empirically, however, we often observe the opposite: several studies, among others Gabaix and Maggiori [2015], report an appreciation of the high yield currency over long periods of time, which contradicts UIP.

This apparent contradiction is a long-standing question in the finance literature. Early on, Meese and Rogoff [1983] noticed that models based on IRD cannot beat a simple random walk in predicting future exchange rates, raising the question of their predictive content. Despite considerable progress in terms of data availability and econometric techniques over the past 35 years, little has changed in this regard.

Several theoretical reasons have been suggested for this lack of performance: most notably the time-varying predictive content of the fundamentals such as IRD [Bacchetta and van Wincoop, 2013, Berge, 2014, Ismailov and Rossi, 2018], but also the misspecifications of the models traditionally employed to conduct these forecasts [Cheung et al., 2005, Rossi, 2013, Ismailov and Rossi, 2018, Amat et al., 2018]. Indeed, whereas most empirical studies focus on conditional mean forecasts [Husted et al., 2018], evidence suggests that exchange rates exhibit high-order dynamics and extremely weak mean dynamics [Chung and Hong, 2007, Brunnermeier et al., 2008, Ismailov and Rossi, 2018]. Concentrating on the second-order moment, Ulm and Hambuckers [2022] find a significant relationship between the volatility of exchange rate returns and changes in the IRD, transmitted via carry trading activities. Similarly, Ichiue and Koyama [2011] examine how exchange rate volatility is related to the failure of uncovered interest rate parity.

In Figure 1, we illustrate the relationship in higher orders by relating the empirical skewness of the residuals from a GARCH(1,1) model for the daily log-rate of change of the USD/EUR exchange rate to the corresponding IRD. The solid blue line in the two graphs displays the empirical skewness as a function of IRD for IRD being either negative (left graph) or positive (right graph). As a comparison, the black lines reflect the average (solid) and the 95% confidence interval (dashed) for the empirical skewness of the residuals of a GARCH(1,1) model, which was simulated independently of the IRD process. If the true underlying data generating process has its random part independent of IRD, then conditional skewness must be zero. The graphs, however, clearly suggest evidence of a relationship in higher orders between exchange rate returns and IRD. As a consequence, producing mean forecasts with tools such as vector autoregressive models might prove ineffective [Herwartz, 2017]. Instead, other strategies such as modelling conditional skewness should be considered [Chung and Hong, 2007, Brunnermeier et al., 2008, Anatolyev and Gospodinov, 2010, Liu, 2015].

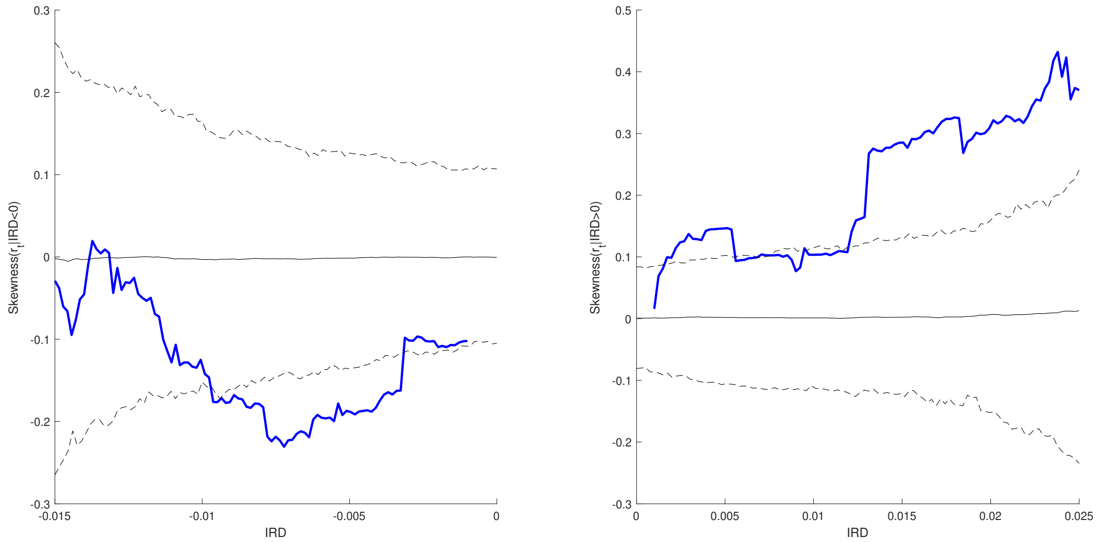

Figure 1: Solid blue line: Empirical skewness for residuals of a GARCH(1,1) computed on daily returns of USD/EUR over the period 1999M1-2019M3, conditional on observing IRDs being smaller(left panel) or larger (right panel) than zero. Dashed black line: 95% confidence intervals for the empirical skewness of GARCH(1,1) data simulated independently of the IRD process. Solid black line: average empirical skewness of GARCH(1,1) residuals independent from IRD. The x-axis is the IRD (in percent) expressed as a proportion (i.e. .01 = 1%).

Following the intuition depicted in Figure 1, we propose an econometric model in which we allow interest rates to convey information on the density of future depreciation rates. We focus in particular on the modelling of its asymmetry, since it is the major factor in determining the likelihood of an exchange rate depreciation. In our model, we refrain from establishing a link between the level of future (log) depreciation rates and interest rates. Consequently, our set-up is less restrictive in its assumptions, and allows us to test empirically general assertions such as "*is a currency more likely to appreciate when its interest rate is relatively high?*". In contrast to existing approaches that rely on "static" or empirical measures of skewness [Brunnermeier et al., 2008] or are employed in a two-step analysis [Iseringhausen, 2020], we propose a one-step time-series approach to investigate whether IRD carries predictive content for the direction of the change of exchange rates. Thus, we are able to hypothesize and test whether a large IRD is indicative of a higher risk of some reverting mechanism, i.e. of a future depreciation of the high yield currency, as suggested, e.g., by Farhi and Gabaix [2016][1]. We thereby assume the marginal effect of IRD to convey information on the *likelihood* of an appreciation or depreciation, instead of on the exchange rate movement itself.

To capture the link between IRD and exchange rates in higher orders, we combine a GARCH structure of the variance with a dynamic non-Gaussian distribution for the innovations, and let its asymmetry parameter evolve according to a time-series process augmented with exogenous predictors.

[1]In a theoretical exercise portrayed in this paper, the authors link the time-varying probability of rare economic disasters and the exposure of a country to such disasters to the risk of a depreciation. They argue that relatively risky countries feature high interest rates because investors need compensation for a potential depreciation in case of a disaster.

Our approach thus allows us to investigate the link between asymmetry and IRD in a single step. We outline it in more detail in Section 2.

We use the proposed model to study the depreciation rate of the US Dollar (USD) vis-a-vis two major currencies: the Euro (EUR) and the Swiss Franc (CHF). We investigate whether IRD is an important factor in the dynamic skewness of these currencies, and if economically meaningful forecasts can be derived from it. We also consider additional control factors such as past innovations, past skewness parameter, as well as the the CBOE volatility index (VIX), a measure of global uncertainty found to impact on the predictability of interest rates [Ismailov and Rossi, 2018]. To assess the economic significance of our findings, we look at the average return generated by a simple directional trading strategy based on our model, and compare it to several benchmarks. We conduct a persistence analysis to assess whether our results hold out-of-sample.

As highlighted by Hossfeld and MacDonald [2015], EUR and CHF are the only G10 currencies apart from USD exhibiting the features of safe haven and hedging currencies. It is therefore expected that, thanks to their economic fundamentals, the signal conveyed by IRD is magnified for these currencies in case of a carry trade reversal as described in Farhi and Gabaix [2016]. Moreover, Hsu et al. [2016] could only find mild or no evidence of profitable trading strategies at a daily frequency in earlier time periods for these currencies. As such, these currencies form an interesting but challenging laboratory to investigate the link between IRD and exchange rate skewness.

Based on our approach, we find that IRD is an important factor to model the dynamic skewness of the daily depreciation rate for EUR and CHF. The forecasting exercises reveal that the predictive performance of our model is statistically significant for both currencies. In contrast with previous studies limited to a monthly frequency, we demonstrate this effect at a daily frequency. Moreover, we find evidence of a trade-off between appreciation and the risk of a large depreciation: first, we observe that the larger the IRD, the more likely the high-yield currency is to appreciate. However, this comes at the cost of an increased likelihood of a currency crash, i.e. a large depreciation. In a robustness analysis presented in the supplementary material, we find evidence of a similar trade-off for three additional currencies (Japanese Yen, Swedish Krona and British Pound). Second, our results reveal that a depreciation shock to the home or the foreign currency on a given day makes an appreciation of the same currency more likely the next day. However, this shock is also associated with an increase in the likelihood of an extreme depreciation. This observation suggests that some self-fulfilling mechanism exists, in line with the findings of Habib and Stracca [2012]. Finally, for EUR, we find that an increase in global uncertainty or risk aversion is positively associated with a higher likelihood of appreciation. This also comes at the cost of an increasing risk of an extreme depreciation. These three effects combined underline the existence of a crash-risk trade-off: the larger the probability of an appreciation, the more likely we are to suffer an extreme depreciation.

We believe that our results provide useful insights for investors and policymakers. A better understanding of the exchange rate dynamics of a currency adds to the information available for risk management or investment analysis. Moreover, similarly to the volatility forecasts of an exchange rate [Poon and Granger, 2003], quantifying the risk of depreciation shocks may serve for monetary authorities and policymakers as a barometer for assessing the vulnerability of financial markets.

We have structured the paper as follows: in Section 2, we detail the features of the statistical model, the estimation approach, and the interpretation of the model. In Section 3, we perform the empirical analysis and investigate the economic significance of our results in Section 3.3. We conclude in Section 4.

2 Methodology

The fundamental feature of our econometric approach is a time-varying asymmetry of the depreciation rate that is allowed to depend on explanatory variables.

To set up our model, we build on a classic GARCH model [Engle, 1982, Bollerslev, 1986], and relax the usual assumptions of symmetric and time-constant distribution of the innovations (e.g. Gaussian or t-distributed). Instead, we combine a GARCH-type model with a sinh-arcsinh distribution for the innovations (SH, Jones and Pewsey [2009]), abbreviated GARCH-SH subsequently. Conveniently, the standardized SH distribution has two parameters (ϵ and δ) with interpretable meanings (asymmetry and shape), is centred on the Gaussian distribution (with $\epsilon = 0$ and $\delta = 1$), and has the single constraint[2] $\delta > 0$. Moreover, it accounts also for lighter tails than the normal distribution, a feature not possible with the skewed-t distribution. Finally, all its high-order moments exist without additional restrictions. This feature is particularly appealing in the analysis of exchange rates, as the existence of high-order moments is key to investigating the potential link between exchange rates and IRD. In our suggested GARCH-SH approach, we allow the asymmetry parameter ϵ to evolve according to an ARMAX structure, i.e. an autoregressive-moving average structure complemented by explanatory variables. Thereby, we can link the conditional distribution of exchange rate returns with the IRD as well as other relevant financial or economic factors, and account for dependence structures that go beyond the first and second moments. Furthermore, we allow the volatility level to enter the mean equation, defining a GARCH-in-mean model as in Glosten et al. [1993]. Empirically, the use of the contemporaneous volatility in the mean equation is motivated by Ranaldo and Söderlind [2010] and Menkhoff et al. [2012], who discovered a significant relation between the volatility and the expected depreciation of a currency. We detail the model and its essential features in the following subsections.

[2]Another constraint, although classical, is the finiteness of the parameters, which is needed to ensure that the distribution is proper.

2.1 Model specification and interpretation

We specify the exchange rate model according to the following set of equations: denoting the nominal exchange rate at time t by S_t, the log-rate of change $R_t = \log\left(S_t/S_{t-1}\right)$ follows a multiplicative heteroscedastic process of the form

$$
\begin{aligned}
R_t &= c + \lambda\sigma_t + r_t, & (1)\\
r_t &= \sigma_t z_t, & (2)\\
\sigma_t^2 &= \omega + \alpha\sigma_{t-1}^2 + \beta r_{t-1}^2, & (3)\\
z_t|\mathcal{I}_{t-1} &\overset{iid}{\sim} f(z_t;\epsilon_t,\delta|\mathcal{I}_{t-1}), & (4)\\
\epsilon_t &= g(\mathcal{I}_{t-1}), & (5)
\end{aligned}
$$

where c is a constant, σ_t^2 the conditional variance of r_t, and z_t the innovation at time t with mean zero and unit variance. $\mathcal{I}_t$ denotes the information set up to time t, comprising all values of z_t and vectors of covariates $\mathbf{x}_t$ up to time t. The probability density function (PDF) of the standardized sinh-arcsinh distribution with parameters ϵ_t and δ, conditional on $\mathcal{I}_{t-1}$ is denoted by $f(z;\epsilon_t,\delta|\mathcal{I}_{t-1})$. Moreover, $g(\cdot)$ is a parametric function linking the asymmetry parameter to past information. Expressions for the PDF, the value of the location and scale parameters in the standardized case, as well as a detailed discussion on the characteristics of the SH distribution can be found in the supplementary material. Note that we assume the shape parameter δ to be constant over time. Without loss of generality, we can relax this assumption to obtain a more flexible model, but prefer maintaining a more parsimonious one. Nevertheless, since the kurtosis depends on ϵ_t as well, dynamic kurtosis is automatically implied from our specification. Conditions stated in the supplementary material ensure that $\mathbb{E}(z_t) = 0$ and $\mathbb{E}(z_t^2) = 1$, so that σ_t^2 can be interpreted as the conditional variance of r_t.

More specifically, we define eq. (5) as a function of past innovations z_{t-1}, lagged values ϵ_{t-1}, as well as variables x_{t-1}, such as IRD_{t-1}, or other potential determinants of exchange rate asymmetry. Eq. (5) is therefore reformulated as:

$$
\epsilon_t = g(\mathcal{I}_{t-1}) = a_0 + a_1\epsilon_{t-1} + a_2 z_{t-1} + a_3 x_{t-1}. \qquad (6)
$$

This equation can be modified or restricted in several ways. For instance, assuming that all parameters in (6) take the value zero, we are back to the symmetric case. Setting $a_1 = a_2 = a_3 = 0$ leads to a model without dynamics but including asymmetry, whereas assuming a_1 and a_3 to be zero leads to a model in which only past innovations impact on the asymmetry. In the empirical analysis, a range of different specifications for eq. (6) is investigated. An overview is available in Appendix A.

2.2 Dynamic asymmetry and the currency crash risk

We employ (6) to study how the IRD, uncertainty (measured by the VIX), and past stochastic components influence the distribution of exchange rate returns. In particular, we suggest looking at three quantities: first, the probability of a positive shock (π_t), indicative of the likelihood of a depreciation of the home currency [3]. Furthermore, for economic interpretation, we are interested in two measures of *currency crash risk*, indicating the likelihood of a sudden large depreciation or large appreciation, of the home currency. These measures are denoted ρ_t^+ and ρ_t^-, respectively. Mathematically, the three quantities are defined as

$$\pi_t = \mathbb{P}(z_t > 0), \tag{7}$$

$$\rho_t^+ = \mathbb{P}(z_t > q^+), \tag{8}$$

$$\rho_t^- = \mathbb{P}(z_t < q^-), \tag{9}$$

where $q^+ > 0$ is large and $q^- < 0$ is small. In our empirical study, we use $q^+ = 2$ and $q^- = -2$, values corresponding to a two-standard deviation event and used as a reference point in Chernov et al. [2018] for currency crash risk. As such, ρ_t^+ and ρ_t^- indicate the probability that the exchange rate suffers from a random shock z_t at least as large as two standard deviations in absolute values.

The marginal effect of a change in one component in eq. (6) on these quantities is deduced from the sign of the regression coefficients. Table 1 summarizes the various scenarios for a variation in the level of IRD_{t-1} and highlights the important connections between the asymmetry parameter (ϵ_t), the likelihood of a depreciation (π_t), and the crash risks (ρ_t^+ and ρ_t^-): if the density is positively skewed (i.e. ϵ_t is positive), then an appreciation of the home currency is more likely than a depreciation (i.e. $\pi_t < 0.5$). If the asymmetry parameter becomes more positive, the density becomes more positively skewed, and an appreciation is even more likely (π_t becomes smaller). However, the risk of a large depreciation increases (ρ_t^+ increases). A similar reasoning holds for negative values of the asymmetry parameter ϵ_t and the risk ρ_t^- of a large appreciation. Figure 2 illustrates these movements in the PDF for a positive parameter ϵ_t. If the degree of skewness augments further (as indicated by the SH distribution in blue), the crash risk computed as the integral of the density beyond the blue vertical line increases.

2.3 Estimation procedure and inference

We estimate the model by means of maximum likelihood procedures. Denoting by Θ the vector of all parameters in equations 1 to 5, the time series of observed log-rate of change by $\mathbf{y_T} = \{R_t\}_{t=1,\ldots,T}$,

[3] Throughout the paper, we use the direct quotation for exchange rates, i.e. we express one unit of the foreign currency in terms of the home currency.

Table 1: Effects of a change in IRD_{t-1}

	$\text{IRD}_{t-1} \nearrow$		$\text{IRD}_{t-1} \searrow$	
	$a_3 > 0$	$a_3 < 0$	$a_3 > 0$	$a_3 < 0$
ϵ_t	+	-	-	+
π_t	-	+	+	-
ρ_t^+	+	-	-	+
ρ_t^-	-	+	+	-

Note: signs of the effect on the asymmetry parameter (ϵ_t), the likelihood of a depreciation (π_t), and the crash risks (ρ_t^+ and ρ_t^-) resulting from a change in IRD_{t-1} captured by a_3.

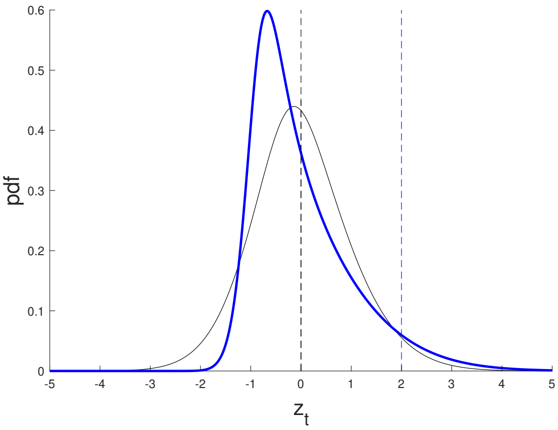

Figure 2: Example of standardized SH distributions with $\epsilon = 0.1$ (black) and $\epsilon = 0.7$ (blue) (we set $\delta = 0.85$). Those values imply skewness levels of 0.25 and 1.3, respectively. The vertical blue line indicates the threshold used for calculating the crash risk, i.e. the integral of the density beyond this threshold.

and assuming conditional independence, the conditional log-likelihood function $\mathcal{L}(\Theta; \mathbf{y}_T)$ is given by

$$\mathcal{L}(\Theta; \mathbf{y}_T) = \frac{1}{T} \sum_{t=1}^{T} \log \left(\frac{1}{\sigma_t} f((R_t - c - \lambda\sigma_t)/\sigma_t; \epsilon_t, \delta | \mathcal{I}_{t-1}) \right). \tag{10}$$

An estimator $\hat{\Theta}$ of Θ is obtained by maximizing numerically (10) with respect to Θ:

$$\hat{\Theta} = \arg\max_{\Theta} \mathcal{L}(\Theta; \mathbf{y}_T), \tag{11}$$

and subject to the constraints $\omega, \alpha, \beta > 0$, $\alpha + \beta < 1$ and $\delta > 0$. We do not set constraints on the other parameters[4]. A simulation study of the finite-sample performance of the proposed estimation method can be found in the supplementary material.

Under correct specification of the model and usual stationarity conditions, the Fisher information

[4] With regard to the choice of a starting value ϵ_0, we use $(a_0 + \sum_{j>2} a_j \bar{x}_j)/(1 - a_1)$. We also check *a posteriori* whether the estimated parameters ensure finite values of ϵ_t when $T \to +\infty$

matrix $H(\Theta)$ of (10) at $\hat{\Theta}$ can be used to test the following null hypothesis:

$$H_0 : \theta_i = 0, \tag{12}$$

where θ_i is the i^{th} element of Θ. To do so, we use the Wald-type test statistic

$$w_i = \hat{\theta}_i / \hat{\sigma}_{ii}, \tag{13}$$

where $\hat{\theta}_i$ is an estimator of θ_i and $\hat{\sigma}^2_{ii}$ is the i^{th} diagonal element of $H^{-1}(\hat{\Theta})$. Under H_0, $w_i \overset{as.}{\sim} N(0,1)$. In the supplementary material, we report a simulation study demonstrating that this approximation returns well-sized and respectably powerful tests. Similarly, restrictions in eq. (6) can be tested using likelihood ratio (LR) test statistics of the type

$$LR = -2(\mathcal{L}(\Theta_0; \mathbf{y}_T) - \mathcal{L}(\Theta_1; \mathbf{y}_T)), \tag{14}$$

with Θ_0 being a restricted version of Θ_1. Under the null hypothesis of the restricted model being the true one, we have the usual result $LR \overset{as.}{\sim} \chi^2_\nu$, ν being the number of restrictions.

2.4 Directional forecasting

As noted by Blaskowitz and Herwartz [2011], in the specific context of exchange rates, investors are particularly interested in the direction of the change of the market since this knowledge helps to hedge currency risk or devise an investment strategy. With the goal of assessing the economic significance of the proposed model with respect to exchange rates exhibiting high-order dynamics, we therefore focus on producing daily directional forecasts, and assess its directional accuracy. While in Section 3 we discuss in detail the results from an out-of-sample forecasting exercise, results from the in-sample analysis are provided in the supplementary material.

2.4.1 Directional forecasting strategy

To assess the proposed model, we consider the performance of a directional forecasting exercise, as found to be particularly appropriate to gauge the correctness of a time-varying asymmetry [Liu, 2015]. In a model without mean dynamics, a dynamically asymmetric distribution of the innovations implies that tomorrow's probability of a positive variation would be lower than a negative one. Similarly, the probability of a negative variation would be greater than a positive one. Consequently, knowing the degree and sign of asymmetry enables us to compute a probability of appreciation or depreciation, and to derive a forecasting strategy. An easy analogy can illustrate this: at any point in time, we bet on the tossing of a coin, heads or tails, in which the probabilities of each result are not equal and vary over time. Consequently, if we knew these probabilities, we could always bet on the most likely outcome.

With this model, we can compute the probability that the home currency depreciates (i.e. that $R_t > 0$) at each point in time, given the information set available at time $t-1$. This probability is denoted $p_{t|t-1}$ and is obtained from

$$\begin{aligned} p_{t|t-1} &= 1 - \mathbb{P}(R_t < 0|\mathcal{I}_{t-1}), &(15)\\ &= 1 - \mathbb{P}(c + \lambda\sigma_t + \sigma_t z_t < 0|\mathcal{I}_{t-1}), &(16)\\ &= 1 - \mathbb{P}(z_t < -c/\sigma_t - \lambda|\mathcal{I}_{t-1}), &(17)\\ &= 1 - F(-c/\sigma_t - \lambda; \epsilon_t, \delta|\mathcal{I}_{t-1}), &(18)\end{aligned}$$

where $F(\cdot)$ denotes the SH cumulative distribution function. Estimates of $\hat{p}_{t|t-1}$, for $t = 2, \ldots, T$, are obtained by plugging the vector of estimated parameters $\hat{\Theta}$ into (18). Then, the final directional forecast is obtained from the following indicator variable:

$$\hat{p}_t^* = \begin{cases} 1 \text{ if } \hat{p}_{t|t-1} > 0.5, \\ -1 \text{ otherwise.} \end{cases} \tag{19}$$

If $\hat{p}_{t|t-1}$ is larger than 0.5, we expect an increase in the exchange rate, and forecast an depreciation of the home currency. Otherwise, we expect a decrease in the exchange rate, and therefore forecast an appreciation of the home currency.

2.4.2 Trading strategy

To translate the directional forecasts into a trading strategy, we adopt the following rule: if the likelihood of a depreciation of the home currency is above 0.5 (i.e. if $\hat{p}_{t|t-1} > 0.5$), the investor takes a short position or owns the foreign currency. On the contrary, if the likelihood of an appreciation is above 0.5 (i.e. if $\hat{p}_{t|t-1} < 0.5$), the investor takes a long position in the home currency, i.e. owns USD. In the case of a constant asymmetry and negligible mean dynamics of the conditional density, the best trading strategy is either always to be in a short position for a negative asymmetry or always in a long position for a positive asymmetry.

2.4.3 Performance criteria

To assess the quality of these forecasts, we use a set of measures that account for different aspects of the performance. First, we use the correct classification rate over h time periods (starting in $t+1$), given by

$$\text{CR} = \frac{1}{h} \sum_{j=t+1}^{t+h} \mathbb{1}(sign(R_j) = \hat{p}_j^*), \tag{20}$$

where $\hat{p}_j^*$ is given by equation (19), $\mathbb{1}(\cdot)$ denotes an indicator function taking value 1 if the condition in parentheses is met, and $sign(.)$ denotes the sign function. This criterion measures the raw performance of a model in a pure classification exercise. To assess the performance in terms of CR, we use the independence test from Pesaran and Timmermann [2009], which accounts for serial correlation. Second, we use the mean return obtained with our directional forecasts over the same period, given by

$$\hat{m} = \frac{1}{h} \sum_{j=t+1}^{t+h} \hat{p}_j^* R_j. \tag{21}$$

Diebold and Mariano [1995], Blaskowitz and Herwartz [2011], and Elliott and Timmermann [2016] argue that employing a realized economic value is often more sensible than a statistical value in evaluating the usefulness of a forecast. In particular, $\hat{m}$ measures the economic usefulness of "being right", i.e. it combines the correct prediction with the timing of a success [Blaskowitz and Herwartz, 2011, 2014]. Hence, if we predict the correct direction of change, we make a gross profit of R_j, whereas a loss of the same amount is suffered if the prediction is wrong. Such a criterion is used throughout the finance literature to assess trading rules; see Bajgrowicz and Scaillet [2012] for an example.

We compare the performance of the GARCH-SH model with the performance of random walk, inverted random walk, always-short, and buy-and-hold approaches, as well as with the most relevant sub-specifications of the dynamic skewness model. Random walk (RW^+) directional predictions must be understood as predicting tomorrow's direction of change using today's sign of the return. In contrast, directional predictions by the inverted random walk (RW^-) for tomorrow are based on the opposite of today's sign of the return. Always-short (AS) and buy-and-hold (BH) strategies always predict an appreciation or a depreciation of the home currency, respectively.

The significance of the out-of-sample performance is stressed by means of the Diebold and Mariano [1995] (DM) test, the conditional predictive ability test of Giacomini and White [2006], and the fluctuation test of Giacomini and Rossi [2010]. The last two tests have the advantage of explicitly covering loss functions that are based on direction of change, account for estimated parameters, and enable both the comparison of nested and non-nested models. This allows their application to out-of-sample versions of (20) and (21). These tests are denoted by GW and GR, respectively.

3 Empirical results

We present the results of our empirical study on the asymmetry of exchange rates, as outlined in the previous sections. We first describe the data in Section 3.1. In Section 3.2, we briefly present the final model specification and discuss the empirical findings of including the IRD into the exchange rate asymmetry. In Section 3.3, we discuss the detailed results from the out-of-sample forecasting exercise.

3.1 Data

We use daily foreign exchange rates expressed as USD per unit of foreign currency, for the EUR and CHF. The rates correspond to noon buying rates in New York for cable transfers payable and are available from the Board of Governors of the Federal Reserve System. We compute the log-rate of change $R_t = \log(S_t/S_{t-1})$, where S_t is the nominal exchange rate at time t. The considered time period ranges from 6 January 1999 to 25 March 2019. The final samples comprise 5,074 observations for each currency. Dates for which exchange rate data were missing were removed from the analysis.

Interest rates are the three-month London Inter-Bank Offered Rate (LIBOR) for the respective currencies[5], as employed by Jordà and Taylor [2012], Ismailov and Rossi [2018] and Du et al. [2018]. All interest rate data were retrieved from the website of the Federal Reserve of St. Louis. Missing LIBOR data are replaced by the previously observed rate. The interest rate differential IRD_t is defined as $\text{LIBOR}_t^{Home} - \text{LIBOR}_t^{Foreign}$, so that positive values correspond to situations in which the home currency is the investment currency. For negative values of IRD, the home currency becomes the funding currency. Throughout this article, we define USD as the home currency. Furthermore, VIX data are daily closing prices and are provided by the CBOE. Missing data have been replaced by the first prior price available (54 occurrences).

Figures 3 and 4 display the exchange rate time series and the corresponding log-rate of change. Figure 5 portrays interest rates and the VIX. Several events, such as negative LIBOR rates, the soar of the VIX, the financial crisis, or the removal of the CHF ceiling rate [6], might indicate instabilities in the relationship between exchange rates and IRD (please refer to the discussions in Giacomini and Rossi [2010], Bacchetta and van Wincoop [2013], and Ismailov and Rossi [2018] for examples). This observation motivates us to study the performance of the model using rolling-window estimates in Section 3.3, as well as to repeat the main analysis on subsamples in the supplementary material. We also report the result of a similar analysis for the Swedish Krona (SEK), the Japanese Yen (JPY),

[5]Although LIBOR has been involved in a manipulation scandal (see, e.g., Fouquau and Spieser [2015]), we use this rate for consistency with previous studies. An avenue for further research in that respect would be to consider other reference rates such as swap rates instead.

[6]In response to a continuously strong appreciation of CHF, a cap of 1.20 CHF/EUR was introduced by the Swiss National Bank (SNB) on September 6, 2011. In light of the increasing difficulty for the SNB to defend the ceiling, it was abandoned on January 15, 2015. These two events are responsible for the outliers visibles on Figure 4.

and the British Pound (GBP) in the supplementary material.

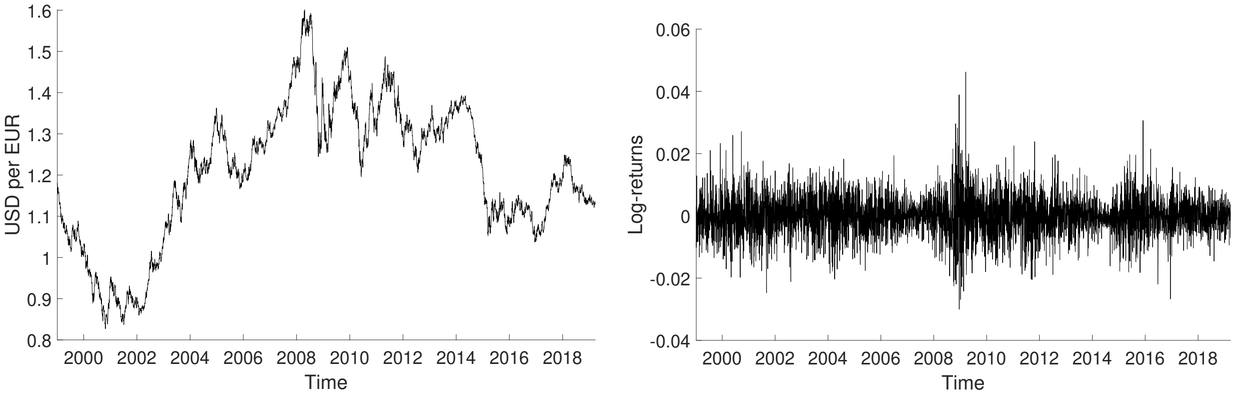

(i) (ii)

Figure 3: (i) Daily exchange rate of EUR against USD, and (ii) log-returns (right) over the period 6 January 1999 - 25 March 2019.

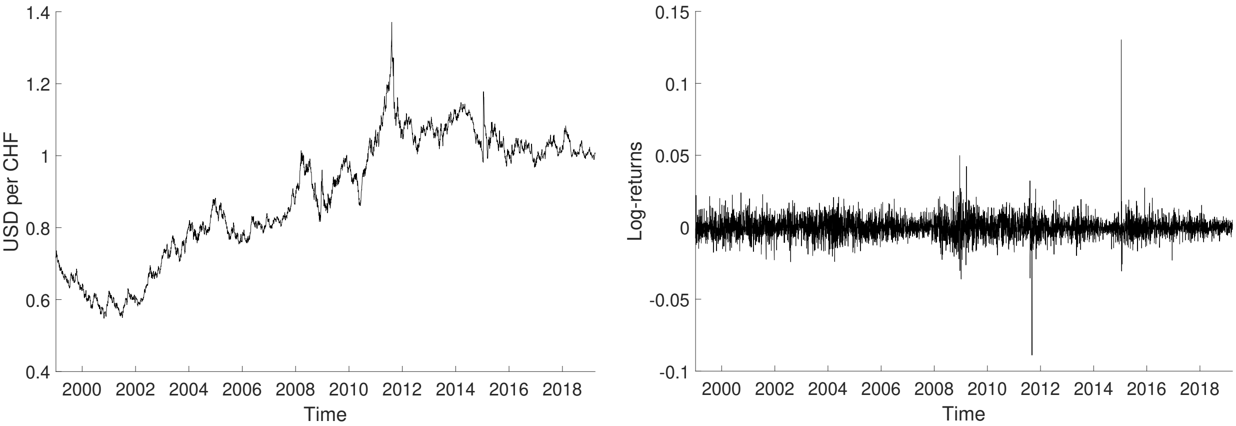

(i) (ii)

Figure 4: (i) Daily exchange rate of CHF against USD, and (ii) log-returns over the period 6 January 1999 - 25 March 2019.

3.2 Model specification and economic interpretation

In the following, we briefly discuss the results of a detailed investigation into different specifications for the general skewness equation (6) estimated on the full sample. In the rest of this section, we concentrate on the results implied by the most relevant model specifications. Additional results for a broader set of specifications can be found in the supplementary material. A detailed overview of the different skewness specifications and their abbreviations henceforth used is provided in Table A1 in Appendix A. Throughout this empirical analysis, we assume a GARCH-in-Mean process such that

$$\mathbb{E}_t(R_t) = c + \lambda_t \sigma_t \tag{22}$$

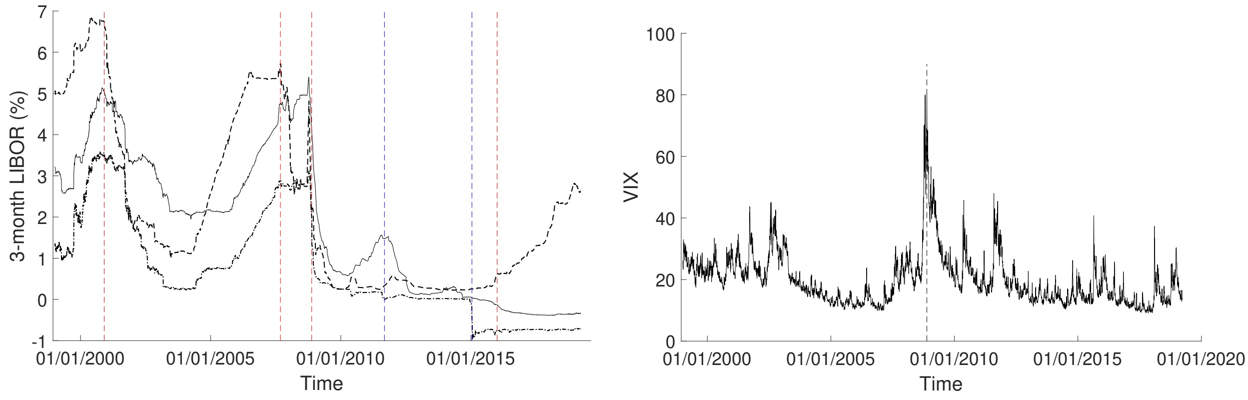

(i) (ii)

Figure 5: (i) Three-month LIBOR rates (EUR: solid, CHF: dashed-dotted, USD: dashed) and (ii) VIX. The dashed vertical lines indicate remarkable events: drop in the Dow Jones index by 445 basis points, dotcom bubble crash, liquidity crisis of 2007, banks bailout of 2008, and the hike of federal fund rate in 2016 (red), establishment and removal of CHF ceiling (blue).

3.2.1 Model specification for EUR and CHF

Table 2 contains information on the model selection for the three most relevant specifications of the skewness equation[7]. ARMAX(2), ARMAX(IRD), and MAX(IRD) display the lowest values for the Bayesian information criterion (BIC) among all considered specifications for EUR. Using likelihood ratio tests against simpler nested alternatives, we find the ARMAX(2)[8] specification to be superior to all considered alternatives for EUR. Hence, the presence of IRD in the skewness specification appears necessary to provide a good fit, and it is well complemented by the informational content of the VIX. Additional test results supporting this specification can be found in the Appendix, Table B3. Moreover, we repeat the analysis conducted in the introduction for the residuals from the proposed model. Figure 6 displays the empirical skewness of subgroups of residuals obtained from ARMAX(2) and pooled by IRD levels. The skewness levels are close to zero for every IRD level, suggesting that our dynamic asymmetry component captures well the high-order dependence.

Turning to CHF, we find the MAX(IRD) specification to be superior to all other considered models in terms of BIC. Results from likelihood ratio tests comparing the more complex specification ARMAX(2) against the simpler alternatives ARMAX(IRD) and MAX(IRD) do not provide evidence that complementing the IRD information by the VIX and an autoregressive component would yield a superior model. Repeating the preliminary analysis on the empirical skewness of the residuals for MAX(IRD) in Figure 7, we find skewness levels close to zero for every IRD level and much smaller than those obtained with a GARCH(1,1) model. In the next section, we provide an economic interpretation of the ARMAX(2) model for EUR and the MAX(IRD) model for CHF, which are the best performing

[7] A more comprehensive overview of model selection criteria for the skewness specifications including interest rates or the VIX are reported in Appendix B, Table B3.

[8] That is, the specification of eq. (6) including an autoregressive component, lagged innovation, and two exogeneous predictors, namely IRD_{t-1} and VIX_{t-1}. See Appendix A for more details.

models based on model selection criteria. Estimated conditional volatility, skewness, and kurtosis for these models are displayed in Figure 8.

Table 2: Model selection and specification criteria

EUR	ARMAX(2)	ARMAX(IRD)	MAX(IRD)
BIC	-37,960.02	-37,960.32	-37,964.14
LR	-	8.24***	12.94***
CHF			
BIC	-37,246.26	-37,254.44	-37,261.96
LR	-	0.35	1.37

Notes: The line *LR* displays the likelihood ratio test statistics between ARMAX(2) and the competing models. *** denotes tests significant at the 1% level.

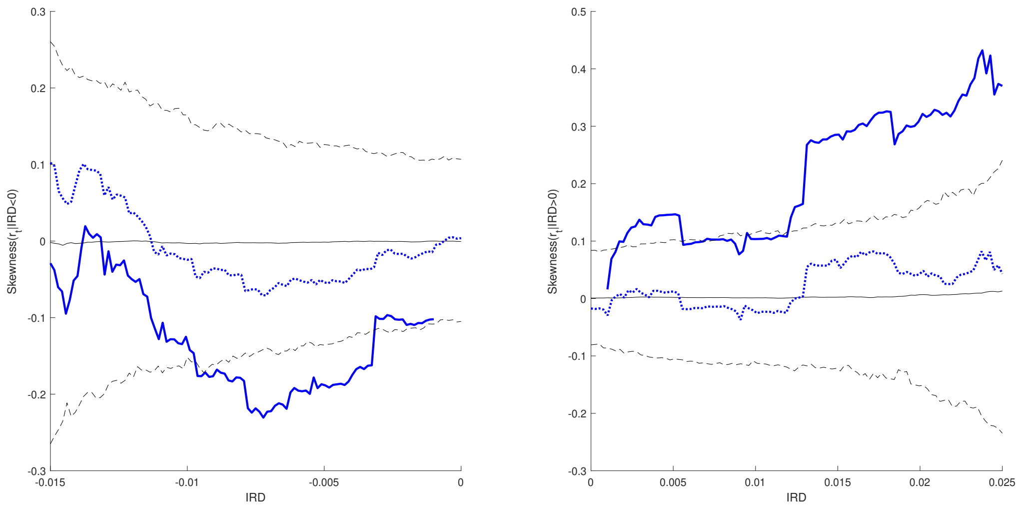

Figure 6: Empirical skewness of USD/EUR residuals, conditional on observing IRD smaller (resp. larger) than zero. Solid blue: residuals of a GARCH(1,1) model. Dashed blue: residuals of ARMAX(2). Solid black: average empirical skewness of GARCH(1,1) residuals independent of IRD. The x-axis is the IRD expressed as a proportion (i.e. .01 = 1%).

3.2.2 Effect of IRD on depreciation and currency crash risk

In Table 3, we report the estimated coefficients of the final specifications for the two currencies. We start with an interpretation of the parameter a_3, which captures the link between IRD and the conditional distribution of the depreciation rate. In particular, we look at the marginal effect of a change in IRD on the probability of a depreciation of USD, measured by π_t, and on the crash risks[9], measured by ρ_t^+ and ρ_t^-. We find a positive and significant regression effect of IRD for both

[9]As a reminder, ρ_t^+ measures a sudden depreciation of USD, and ρ_t^- a sudden appreciation.

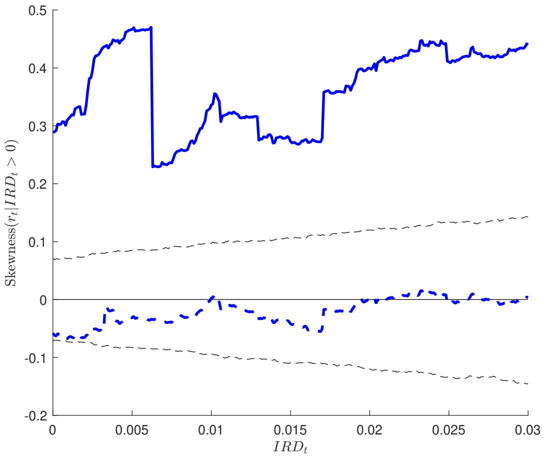

Figure 7: Empirical skewness of USD/CHF residuals, conditional on observing IRD larger than zero (dashed blue). Solid blue: residuals of a GARCH(1,1) model. See Figure 6 for more explanations.

currencies[10]. As outlined in Section 2.2, this result implies that an increase in USD interest rates is associated with an increase in the probability of an appreciation of USD. Analogously, an increase in the foreign interest rate leads to a higher likelihood of an USD depreciation. It suggests also that a large and positive IRD opens up the possibility of profitable carry trades, whereby USD is the investment currency. In this setup, market participants could borrow the foreign currency at a low rate, buy USD, invest them at a higher rate and still expect an appreciation of USD. These results fit into the theoretical framework of Farhi and Gabaix [2016], in which the currency of the country with high interest rates appreciates, conditional on no disaster occurring.

Simultaneously, though, an increase in positive IRD has an opposite effect on ρ_t^+: it becomes more likely to observe an extremely positive shock, synonymous with a large depreciation of USD[11]. This observation is in line with Farhi and Gabaix [2016] and Jurek [2014], who associate IRD with the risk of a currency crash: the larger the IRD, the greater the realignment pressures. Hence, we are more likely to observe a reverting move or a crash on the exchange rate market. A potential explanation for this effect is the increasing share of market participants involved in carry trades when IRDs are large [Brunnermeier et al., 2008]: the larger the IRD, the more carry trade investors fear realignments of the exchange rate. As a consequence, they may unwind their positions in the investment currency, leading to the appreciations of the funding currency. This result also highlights the potential endogeneity of the reverting mechanism, as suggested in Farhi and Gabaix [2016]: out of fear, investors turn themselves into a force of realignment that leads to a sudden depreciation.

Figure 9 depicts the empirical connection between IRD and the probabilities π_t, ρ_t^+, and ρ_t^- for

[10] A similar result is also obtained for CHF when using only the pre-capping period, i.e. up to 6 September 2011, which removes the two outliers displayed on Figure 4. For this period, we obtain a coefficient of 6.03, significant at the 1% level. Full results are available upon demand.

[11] A similar reasoning holds if the foreign currency is the funding currency, leading to a increase in ρ_t^- when IRD becomes more negative.

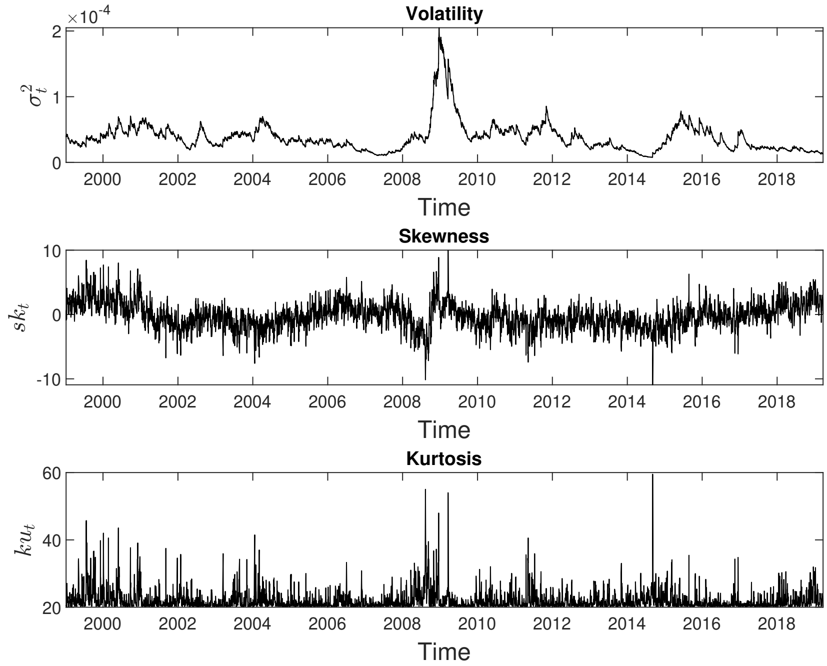

(i) USD/EUR (ARMAX(2))

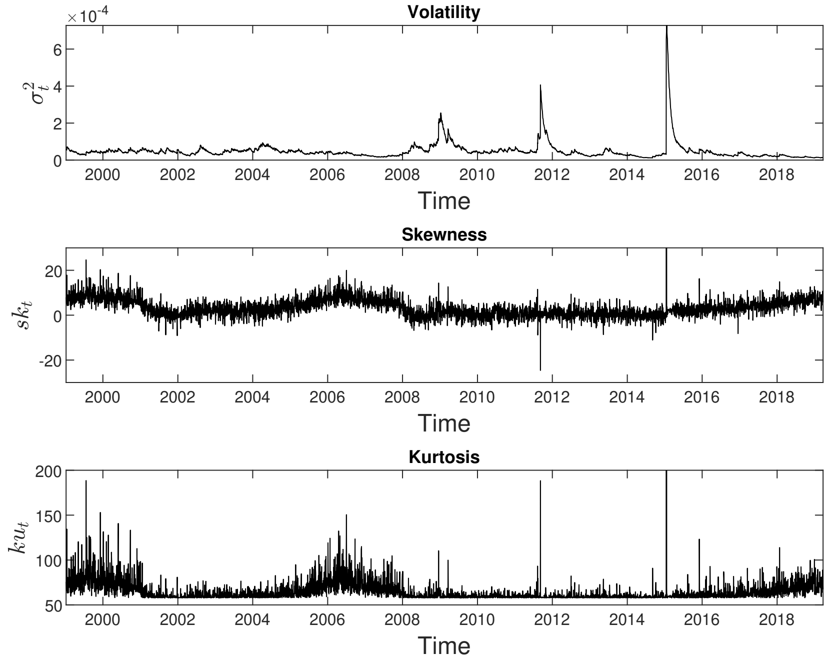

(ii) USD/CHF (MAX(IRD))

Figure 8: Conditional variance, skewness and kurtosis for EUR (top) and CHF (bottom).

EUR and CHF, as unraveled by the models: the larger the IRD, the lower the probability of an overall depreciation of the home currency, but the higher the probability of a large depreciation. Thereby, a change in the IRD implies a more than proportional increase in the risk of a large depreciation compared with the expected change in the total probability of a depreciation. The estimates suggest that while an increasing IRD favors the investment currency on average, it also increases more than proportionally the crash risk of this currency.

Table 3: Parameter estimates for final model specifications

Time period 1999M1 - 2019M3						
Currency	Specification	a_0	a_1	a_2	IRD (a_3)	VIX (a_4)
USD/EUR	ARMAX(2)	-0.043 *	0.631 ***	0.042 ***	1.520 *	0.171 *
		(0.025)	(0.184)	(0.014)	(0.813)	(0.101)
USD/CHF	MAX(IRD)	-0.000	-	0.040***	3.448***	-
		(0.018)	-	(0.013)	(0.890)	-

Estimated coefficients for the constant, AR, and MA parameters (a_0, a_1, a_2) in the skewness equation (6). IRD and VIX refer to the estimated parameters of the corresponding predictors. Standard errors are placed in parentheses. *, **, and *** denote Wald tests significant at the 10%, 5%, and 1% test levels.

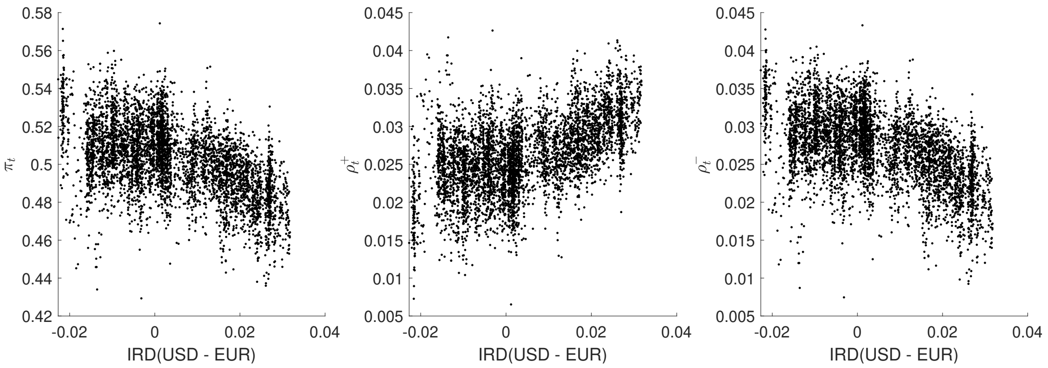

(i) USD/EUR (ARMAX(2))

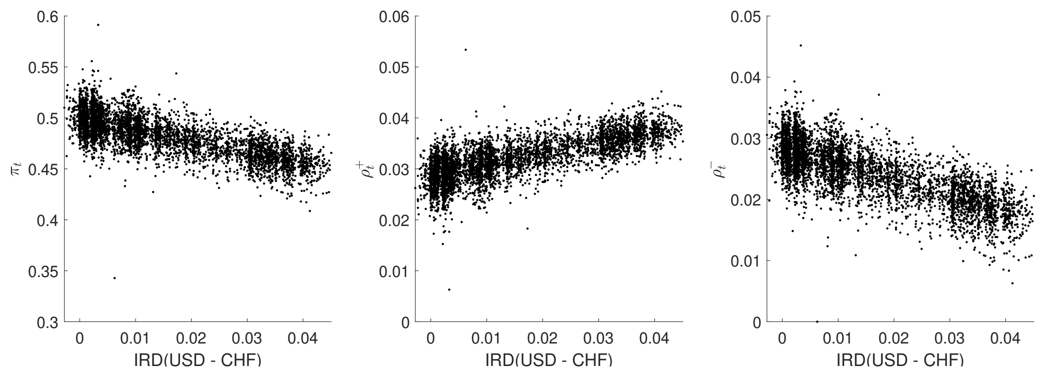

(ii) USD/CHF (MAX(IRD))

Figure 9: Standardized effects of IRD on depreciation and currency crash risks, measured by π_t, ρ_t^+ and ρ_t^- with $q^+ = 2$ and $q^- = -2$ (from left to right). Estimates are obtained from the ARMAX(2) model for EUR (top) and MAX(IRD) for CHF (bottom).

3.2.3 Effect of past unexpected shocks

We now look at the effect of past innovations on the asymmetry. We find a_2 to be positive and significant for all specifications and both currencies, with similar magnitudes. Hence, past positive shocks have a positive effect on the likelihood of an appreciation of USD, but also a positive effect on large depreciation. In other words, the larger an unexpected depreciation on one day, the more likely the appreciation on the next day on average, but also the higher the likelihood of a very large depreciation. We suggest that this effect is connected to the existence of self-fulfilling mechanisms, as found by Habib and Stracca [2012]. According to them, exchange rates fluctuate around some equilibrium value. As a result, unexpected depreciation is followed by appreciation periods. However, large unexpected depreciation may lead more and more economic agents to believe in a future depreciation and to short USD, thus increasing the risk of a sudden USD crash. If this phenomenon takes place at a time of high volatility, shocks will be amplified. Such a mechanism is consistent not only with clustered volatility, but also with clustered signs.

3.2.4 Effect of uncertainty on USD/EUR

Besides IRD, the VIX was included in the model for USD/EUR as an additional predictor. The corresponding coefficient a_4 is found to be positive. Thus, for high values of the VIX, an appreciation of USD against EUR becomes more likely. Simultaneously, though, the likelihood of a currency crash increases as well. A higher likelihood of appreciation of USD over other currencies in times of financial stress (as captured by the VIX) is in line with several findings related to safe haven currencies and funding liquidity constraints: Habib and Stracca [2012] find that the larger the size of the economy, relative to world GDP, the higher the currency excess returns in times of financial stress.

3.2.5 Crash-risk trade-off

Summarizing the analysis, we draw the following conclusions. First, the larger the difference between interest rates, the more likely the high-yield currency is to appreciate. However, this "local" higher appreciation of the high-yield currency comes with the cost of a greater likelihood of a currency crash, i.e. of a large depreciation. Figure 10 illustrates this result, displaying the risk of an extreme depreciation shock (ρ_t^-) suffered by the foreign currency as a function of the probability of appreciation of the foreign currency (π_t) vis-a-vis the US Dollar. We observe a clear, upward relationship for both EUR and CHF. Secondly, we observe that a large unexpected depreciation makes an appreciation more likely the next day. The reciprocal situation applies for a large unexpected appreciation. Moreover, the former is also associated with an increase in the likelihood of an extreme depreciation, suggesting the existence of some self-fulfilling mechanism. Finally, an increase in global uncertainty or risk aversion, as measured by the VIX, is positively associated with a higher likelihood of appreciation for the USD against the EUR, albeit at the cost of an increasing risk of an extreme depreciation. These three effects combined lead to a crash-risk trade-off: the larger the probability of an appreciation, the more likely the currency will suffer an extreme depreciation (ceteris paribus). Additional results regarding the estimated crash-risk measures, especially along the time dimension, can be found in the supplementary material, Section 2.4.

In a robustness check challenging the generality of our findings and presented in details in the supplement (Section 3), we repeat our analysis for three additional currencies, namely SEK, JPY and GBP. Using the same time period, we find for these currencies that conditional skewness is well described by the proposed model. In Table 9 of the supplement, we report BIC values and results of likelihood-ratio tests suggesting that the MAX(IRD) and ARMAX(IRD) specifications should be favored for all currencies, and are superior to the ARMA specification. Moreover, similarly to what is depicted in Figures 6, 7 and 10, we observe a positive association between π_t and ρ_t^- for these three currencies. In Table 10 of the supplement, we report coefficient estimates for the MAX(IRD) specification. All coefficients associated to IRD are found to be significantly positive at the 1% test level and of similar magnitude than what is observed for EUR and CHF. This brings further evidence of a crash risk trade-off also for currencies not considered as safe heaven or hedging currencies.

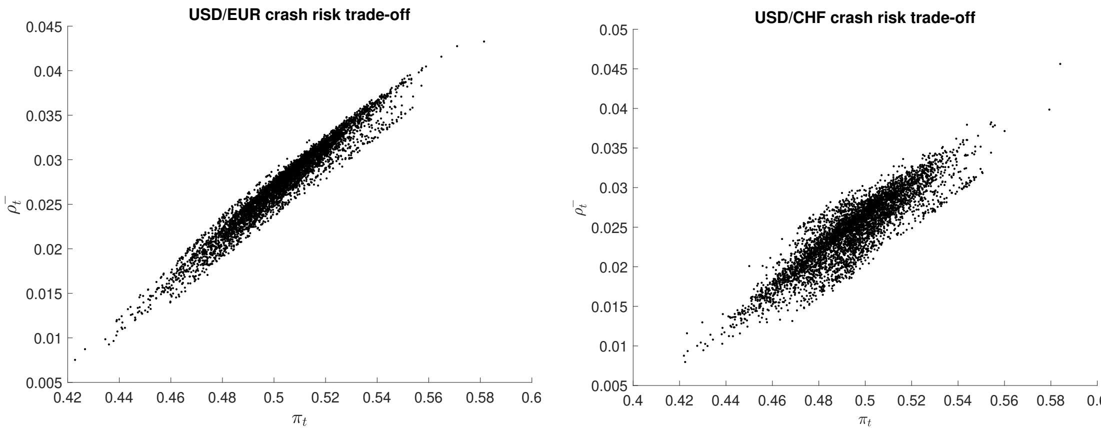

(i) USD/EUR (ARMAX(2))

(ii) USD/CHF (MAX(IRD))

Figure 10: Crash-risk trade-off between the probability of an appreciation (π_t) and the probability of an extreme depreciation (ρ_t^-) for the foreign currency.

3.3 Forecasting with dynamic asymmetry

In this section, we investigate the extent to which the proposed modeling approach translates into economically and statistically significant forecasts. The questions to be answered are, hence: could an investor have correctly guessed the direction of change of exchange rates using the suggested framework? Is the dynamic skewness, and in particular the effect of IRD, sufficiently strong to be exploited and yield an investment strategy with positive returns? In the following, we present the detailed results of an out-of-sample assessment. Results on the in-sample performance can be found in the supplementary material.

3.3.1 Out-of-sample performance

Out-of-sample directional forecasts are obtained using rolling-window estimates of the parameters. The initial training period ranges from 6 January 1999 to 1 December 2014 (3999 observations), covering, hence, both non-crisis and crisis data, as well as information from the more recent period with low interest rates. We re-estimate the parameters of the model every five days, and predict the direction of change up to March 2019, resulting in one-step-ahead predictions for 1,075 days.

We compare these results to a series of naive benchmarks: the BH and AS strategies are similar to forecasting exchange rates with a momentum strategy using the training data to determine the direction of the momentum. The RW^+/RW^-strategies can also be regarded as a one-day momentum strategy. To account for the effect of choosing the side of the benchmark (i.e. if we use the long or short side of the BH/AS and RW^+/RW^- strategies), we use the best performing benchmark during the in-sample period to conduct out-of-sample forecasts.

The forecasting strategies are assessed based on the performance criteria defined in Section 2.4. Moreover, to test for significant differences in forecasting abilities, we use the DM, GW, and GR tests discussed in Section 3.2. In particular, the GR test allows to control for variations in forecasting performance over time, contrary to other tests that only take the average performance into account.

EUR out-of-sample performance

Table 4 contains the performance criteria for selected skewness specifications. Using the ARMAX(2) specification for the skewness parameter yields an average performance $\hat{m}$ of 5.20% in equivalent yearly rate over the forecast horizon. This is the best result across all tested models, with ARMAX(IRD) being the second best performing model with a profit of 2.76%.[12] Figure 11 illustrates the compounded value over time of investing 1 USD at the beginning of the forecasting period according to trading rules defined by either ARMAX(2) or ARMAX(IRD). Although there are losses at the onset of the period, we rapidly make a profit between 2017 and 2018. On the contrary, both benchmark models return a negative performance with -1.26% for the random walk strategy and -2.26% for the momentum strategy, as indicated by column 3 in Table 4. Looking at a more complex benchmark, such as the ARMA specification which does not include IRD (Table B4 in Appendix B.3), we find the ARMAX(IRD) specification to be still superior, delivering a 20% increase in profit. Nevertheless, this result suggest the importance of the ARMA structure in the dynamic skewness. It also potentially points out to a fluctuating predictability of IRD, consistent with the pocket-of-predictability phenomenon [Farmer et al., 2023].

[12]Detailed out-of-sample results for all tested specifications can be found in Appendix B.3.

In terms of classification performance displayed in columns 1 and 2 in Table 4, ARMAX(IRD) obtains a correct classification rate of 52.09%, and is the only model with a sign forecasting ability found significant by the test of Pesaran and Timmermann [2009]. Moreover, we test whether the out-of-sample profit $\hat{m}$ obtained with ARMAX(2) and ARMAX(IRD) is significantly different from 0. To do so, we use the GR test for several $\tau \in [.1, .85]$. The loss function used for the test is given by eq. (21). We reject the null hypothesis of no profit at the 5% test level for ARMAX(2). These results are displayed in Figure 12, upper panels. [13] The results in columns 4 and 5 in Table 4 indicate that the GW and DM tests are inconclusive.

Table 4: Out-of-sample forecasting performance of selected skewness specifications

USD/EUR (2014M12 - 2019M3)								
	Classification		average profit			average excess performance		
Specification	CR	PT09	$\hat{m}$	DM	GW	$\Delta\hat{m}$	DM	GW
ARMAX(2)	51.16%	0.30	5.20%	1.21	4.73	6.46%	0.94	2.59
ARMAX(IRD)	52.09%	4.08**	2.76%	0.64	0.92	4.02%	0.58	0.57
RW^+/RW^-	48.00%	0.00	-1.26%	-0.34	0.16	-	-	-
BH/AS	47.81%	-	-2.26%	-0.55	0.31	-	-	-

Table entries denote correct classification rate (CR), average profit ($\hat{m}$), and average excess performance over the random walk benchmark ($\Delta\hat{m}$) for selected specifications of the skewness parameter. Entries for PT09, DM, GW display the respective test statistics. ** indicates tests significant at the 5% level.

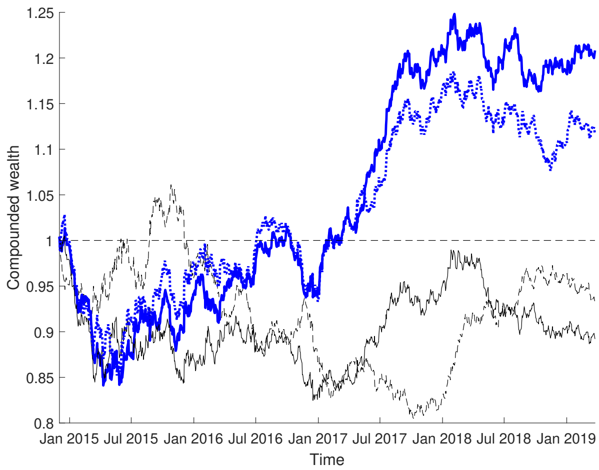

Figure 11: For EUR, compounded value of an initial investment of 1 USD in the trading rules derived from ARMAX(2) (solid blue), ARMAX(IRD) (dotted blue), random walk (dashed black), and momentum (solid black).

Furthermore, it is of interest whether there are significant differences with respect to the momentum

[13] Similar figures for the ARMAX(IRD) specification can be found in the supplementary material.

benchmark. Therefore, we repeat the procedure using the following loss function:

$$L_{\Delta}^{(1)} = \frac{1}{h} \sum_{j=t+1}^{t+h} (\hat{p}_j^* - \hat{p}_j^{*,RW}) R_j, \tag{23}$$

where $\hat{p}_j^{*,RW}$ is the sign forecast obtained from the random walk benchmark. The results displayed in Figure 12 (lower panel) indicate that the performance of the ARMAX(2) model is significantly better than the benchmark for some periods of time. Comparable results are obtained for the ARMAX(IRD) model and are displayed in the supplementary material [14].

We conclude that the IRD and the VIX prove useful to predict the future direction of change of the USD/EUR exchange rate. In particular, a changing IRD intensity as observed over the period 2016-2018, carries economically significant predictive power.

CHF out-of-sample performance

We produce similar out-of-sample forecasts for CHF. The various indicators are reported in Table 5 for the main specifications. The skewness specifications based on IRD or mixing both IRD and VIX prove to perform better than the others. In particular, MAX(2), MAX(IRD) and ARX(IRD) exhibit an average profit $\hat{m}$ ranging between 6.49% and 7.64% in equivalent yearly rates (Table 5). To illustrate these results, we display the compounded value over time of investing 1 USD in our trading rules in December 2014 in Figure 13. Most of the performance seems to be concentrated at the beginning of the period, i.e. between 2015 and mid 2016. All models including the IRD were successful in correctly betting on the surge in CHF value on 15 January 2015, partially explaining the relatively high profit compared to EUR. The random walk benchmark was also successful in that respect but, contrary to dynamic skewness models, its performance decreased substantially later on, and the final profit is close to zero. A more stringent comparison with the ARMA specification of the skewness parameter reveals a profit of 2.75%, which is still 56% to 64% smaller than the ones obtained with the models including IRD. As for EUR, this model exhibits also less desirable features, such as a larger maximum drawdown (see Appendix B.3, Table B7). The positive performance of the ARMA model confirms the important of the autoregressive and moving average components in the skewness dynamics.

We report the results of the GR test for the MAX(IRD) specification but comparable results are obtained for the alternative specifications including IRD information in the supplementary material. The upper panels in Figure 14 display the GR statistics used to test the null hypothesis $H_0 : |\hat{m}| \leq 0$. We reject this hypothesis at the 5% test level for various values of τ. As suggested by Figure

[14] Similar results are achieved when the loss function is adapted to test for superior sign predictability. We also conducted these tests using the best (ex-post) BS/AS benchmark (in this case, this is the AS benchmark, with a correct classification rate of 51.35%), i.e. without accounting for in-sample benchmark selection. The results stay unchanged for ARMAX(2). Results are provided in the supplementary material.

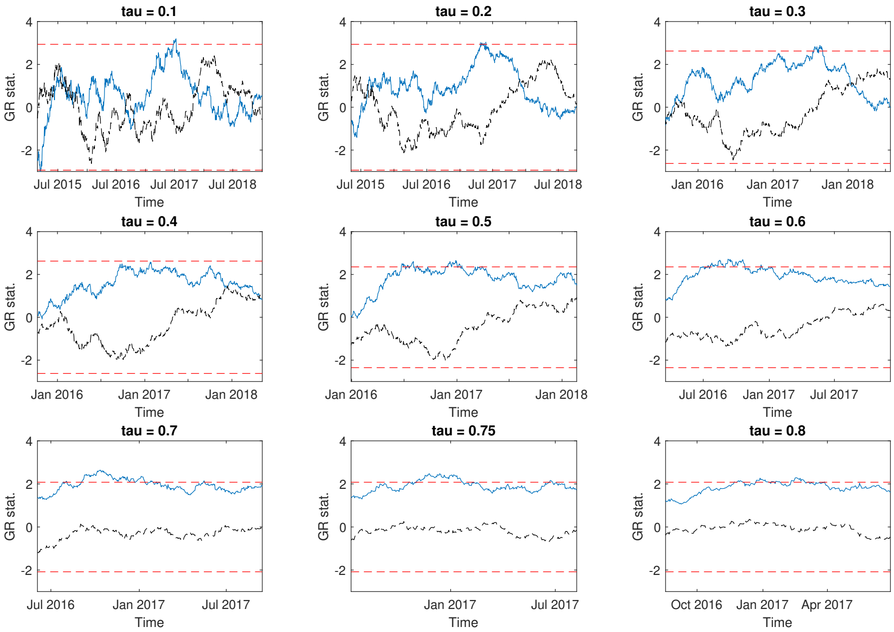

(i) GR test for ARMAX(2) ($H_0 : |\hat{m}| \leq 0$)

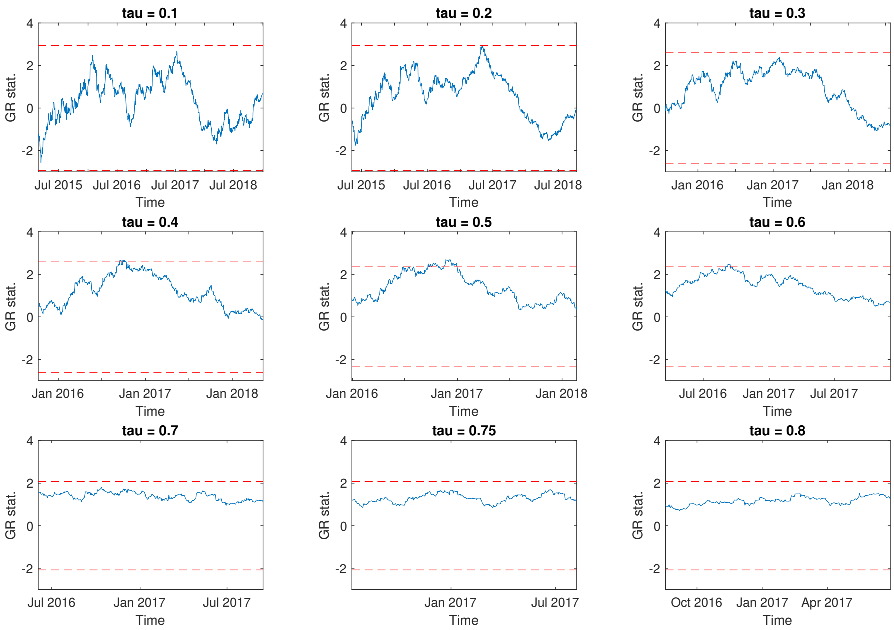

(ii) GR test for ARMAX(2) ($H_0 : |\Delta\hat{m}| \leq 0$)

Figure 12: For EUR, GR test statistic (blue) with $\tau \in [.1, .85]$ using (i) $\hat{m}$ and (ii) $L_{\Delta}^{(1)}$ as loss functions. If the statistic is above the rejection threshold (dashed red), we reject the null hypothesis (i) $H_0 : |\hat{m}| \leq 0$ (upper panel) or (ii) $H_0 : |\Delta\hat{m}| \leq 0$ (lower panel). $\Delta\hat{m}$ is the average difference in profit with respect to the random walk benchmark. Dashed black: test statistic for the RW strategy.

13, the significant performance always occurs in 2015-2016. Using $L_{\Delta}^{(1)}$ as loss function and testing $H_0 : |\Delta\hat{m}| \leq 0$, we obtain similar results as displayed in the lower panels of Figure 14. Overall, these results highlight that the out-of-sample predictive ability of the models augmented with IRD benefits from the sharp appreciation of the CHF at the beginning of 2015.

Robustness analysis: out-of-sample performance for SEK, JPY and GBP

Table 5: Out-of-sample forecasting performance for selected specifications

USD/CHF (2014M12 - 2019M3)								
	Classification		average profit			average excess performance		
Specification	CR	PT09	$\hat{m}$	DM	GW	$\Delta\hat{m}$	DM	GW
MAX(2)	50.79%	0.31	7.64%	1.33	2.97	6.69%	0.86	1.18
MAX(IRD)	50.60%	0.12	6.49%	1.22	2.70	5.54%	0.71	0.65
ARX(IRD)	50.79%	0.58	7.63%	1.46	4.62*	6.68%	1.09	1.45
RW^+/RW^-	47.72%	0.00	0.95%	0.18	1.22	-	-	-
BH/AS	47.07%	-	-0.69%	-0.13	0.80	-	-	-

Please refer to the notes of Table 4 for explanations on the table entries.

In Section 3, Table 11 of the supplement, we report detailed results of an out-of-sample performance analysis for SEK, JPY and GBP. In contrast to EUR and CHF, we do not produce a significantly positive out-of-sample profit for these currencies. For SEK, however, out-of-sample profit is large (up to 4.31%) and substantially higher than for benchmark models (1.28% for ARMA, -5.17% for RW). For JPY and GBP, the out-of-sample performance is negligible or even negative. These results suggest that, although IRD explains well the dynamic skewness in the regression analysis, its effect is not strong enough to generate economically useful forecasts with our model.

Following Hossfeld and MacDonald [2015], these discrepancies might stem from the nature of the currency: CHF and EUR are considered safe heaven and hedge currencies, that are expected to appreciate even in times of global market downturn. On the contrary, SEK was found to be a speculative currency, i.e. a currency depreciating during stock market downturn. Neither JPY nor GBP exhibited any specific features. Our results may, hence, reflect that the predictive content of IRD is modulated by factors such as market state or market stress level. We let this question be answered in future research.

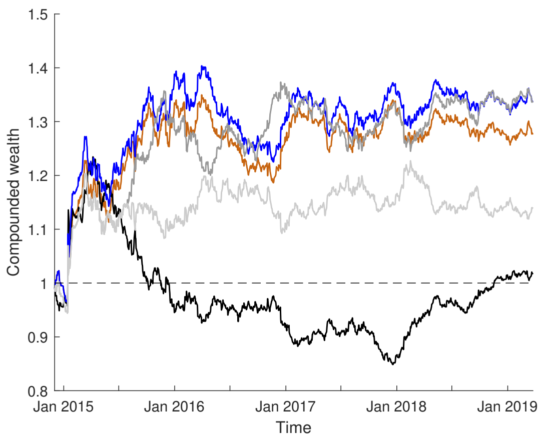

Figure 13: For CHF, compounded value of an initial investment of 1 USD in the trading rules derived from the models MAX(2) (blue), MAX(IRD) (brown), ARX(IRD) (dark grey), random walk (solid black) and constant asymmetry (light grey).

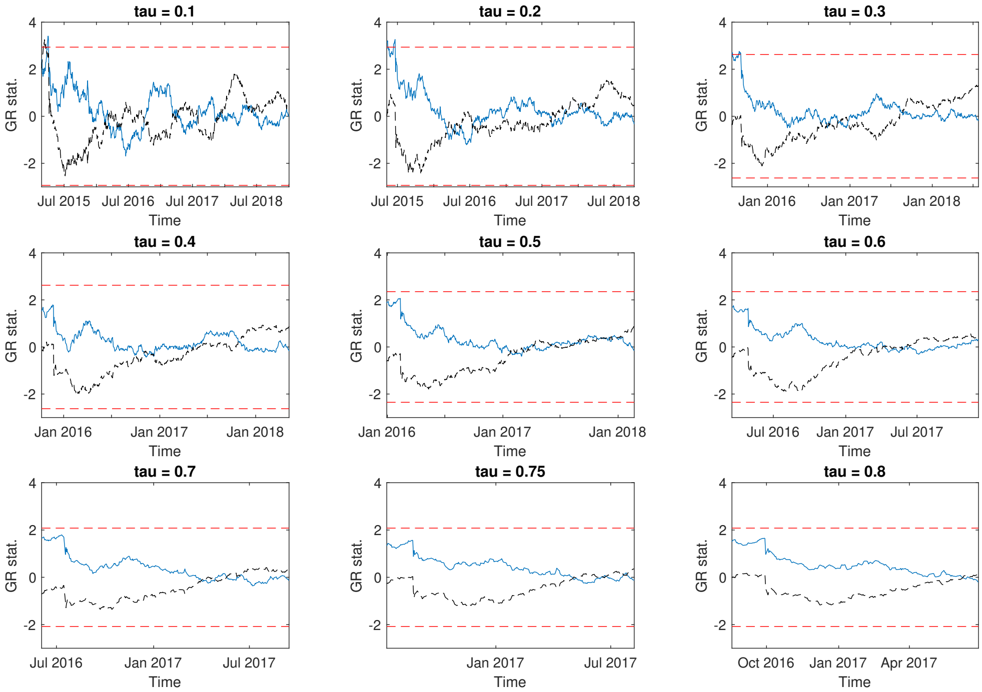

(i) GR test for MAX(IRD) ($H_0 : |\hat{m}| \leq 0$)

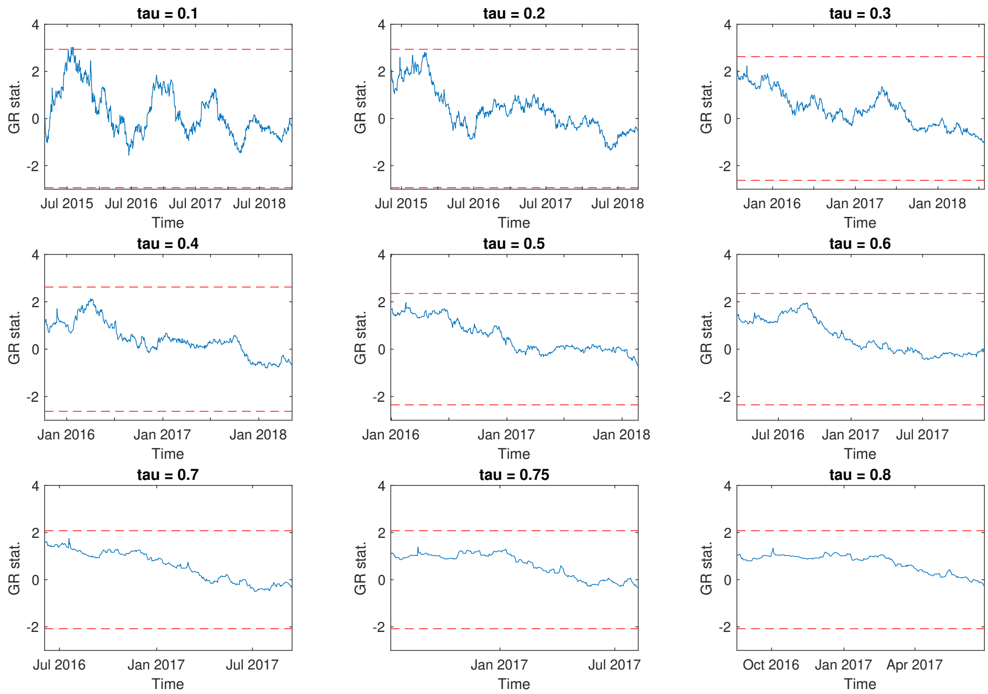

(ii) GR test for MAX(IRD) ($H_0 : |\Delta\hat{m}| \leq 0$)

Figure 14: For CHF, GR test statistic (blue) with $\tau \in [.1, .85]$ using (i) $\hat{m}$ and (ii) $L_{\Delta}^{(1)}$ as loss functions. If the statistic is above the rejection threshold (dashed red), we reject the null hypothesis at the 5% test level. Dashed black: test statistics for the RW benchmark.

4 Conclusion

Using a model that allows for conditional dynamic asymmetry, we revisit the link between interest rate differentials and daily exchange rate returns. Applying this approach to the study of EUR and CHF daily exchange rates vis-a-vis the US Dollar, we discover that the larger the difference between interest rates, the more likely the high-yield currency is to appreciate. However, it occurs at the cost of an increase in the likelihood of a large depreciation. In robustness checks, we find similar dynamics for SEK, JPY and GBP. This result is consistent with the theoretical framework of Farhi and Gabaix [2016] and Brunnermeier et al. [2008], in which a brutal unwinding of carry trade positions generates a currency crash.

Furthermore, when accounting for the effect of financial uncertainty, we find that USD is more likely to appreciate with respect to EUR when the VIX increases, but also that it is then exposed to a higher risk of a currency crash. These results are in line with Menkhoff et al. [2012], Bekaert et al. [2013] and Habib and Stracca [2012], who suggest that liquidity shortage and increasing risk aversion lead investors towards buying USD, but also that this increasing uncertainty leads to an increase in the likelihood of a large USD depreciation.

We use our model to predict the direction of change in exchange rates and use these forecasts to build a transparent and theoretically motivated trading strategy. The detected effects are sufficiently large to generate statistically significant economic gains for EUR and CHF, both in- and out-of-sample. A limitation of the present analysis is that we do not account for the selection of the model itself. Therefore nothing guarantees that one could have obtained a profit *ex ante*, as assessed, e.g. in Bajgrowicz and Scaillet [2012]. However, as discussed in Inoue and Kilian [2005], in-sample results typically exhibit a higher power in performance tests. Hence, the consistency between in-sample and out-of-sample tests limits the risk of our findings being spurious.

From a policy standpoint, our results suggest that favoring an increase in IRD correlates with systemic issues such as a high likelihood of a large depreciation. They also highlight the importance of self-fulfilling mechanisms and volatility in currency crashes, suggesting that preventing depreciation or appreciation shocks in periods of high uncertainty would reduce crash risk. Future research could extend this approach to the currencies of emerging markets in which monetary policy authorities are faced with both large IRD with the USD and substantial exchange rate fluctuations. Other potential mitigating factors, such as TED spreads or corporate bond spreads, could be investigated to challenge the robustness of our findings.

References

C. Amat, T. Michalski, and G. Stoltz. Fundamentals and exchange rate forecastability with simple machine learning methods. *Journal of International Money and Finance*, 88:1–24, 2018.

S. Anatolyev and N. Gospodinov. Modeling financial return dynamics via decomposition. *Journal of Business & Economic Statistics*, 28(2):232–245, 2010.

P. Bacchetta and E. van Wincoop. On the unstable relationship between exchange rates and macroeconomic fundamentals. *Journal of International Economics*, 91(1):18 – 26, 2013.

P. Bajgrowicz and O. Scaillet. Technical trading revisited: False discoveries, persistence tests, and transaction costs. *Journal of Financial Economics*, 106(3):473–491, 2012.

G. Bekaert, M. Hoerova, and Duca M.L. Risk, uncertainty and monetary policy. *Journal of Monetary Economics*, 60(7):771–788, 2013.

T.J. Berge. Forecasting disconnected exchange rates. *Journal of Applied Econometrics*, 29(5):713–735, 2014.

J. Berkowitz. Testing density forecasts, with applications to risk management. *Journal of Business & Economic Statistics*, 19(4): 465–474, 2001.

A.E. Bernardo and O. Ledoit. Gain, loss, and asset pricing. *Journal of Political Economy*, 108(1):144–172, 2000.

O. Blaskowitz and H. Herwartz. On economic evaluation of directional forecasts. *International Journal of Forecasting*, 27(4): 1058–1065, 2011.

O. Blaskowitz and H. Herwartz. Testing the value of directional forecasts in the presence of serial correlation. *International Journal of Forecasting*, 30(1):30–42, 2014.

T. Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986.

M. Brunnermeier, S. Nagel, and L. Pedersen. Carry trades and currency crashes. *NBER Macroeconomics Annual*, 23:313–347, 2008.

M. Chernov, J. Graveline, and I. Zviadadze. Crash risk in currency returns. *Journal of Financial and Quantitative Analysis*, 53 (1):137–170, 2018.

Y-W. Cheung, M.D. Chinn, and A.G. Pascual. Empirical exchange rate models of the nineties: Are any fit to survive? *Journal of International Money and Finance*, 24(7):1150–1175, 2005.

J. Chung and Y. Hong. Model-free evaluation of directional predictability in foreign exchange markets. *Journal of Applied Econometrics*, 22(5):855–889, 2007.

F.X. Diebold and R.S. Mariano. Comparing Predictive Accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263, 1995.

J.A. Doornik and H. Hansen. An omnibus test for univariate and multivariate normality. *Oxford Bulletin of Economics and Statistics*, 70(s1):927–939, 2008.

W. Du, A. Tepper, and A. Verdelhan. Deviations from covered interest rate parity. *The Journal of Finance*, 73(3):915–957, 2018.

G. Elliott and A. Timmermann. *Economic Forecasting.* Princeton University Press, 2016. ISBN 978-0-691-14013-1.

R.F. Engle. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, 50(4):987–1007, 1982.

M. Farhi and X. Gabaix. Rare disasters and exchange rates. *Quarterly Journal of Economics*, 131(1):1–52, 2016.

L. Farmer, L. Schmidt, and A. Timmermann. Pockets of predictability. *Journal of Finance*, 78(3):1279–1341, 2023.

J. Fouquau and P.K. Spieser. Statistical evidence about LIBOR manipulation: A "Sherlock Holmes" investigation. *Journal of Banking & Finance*, 50:632–643, 2015.

X. Gabaix and M. Maggiori. International liquidity and exchange rate dynamics. *Quarterly Journal of Economics*, 130(3):1369–1420, 2015.

R. Giacomini and B. Rossi. Forecast comparisons in unstable environments. *Journal of Applied Econometrics*, 25(4):595–620, 2010.

R. Giacomini and H. White. Tests of Conditional Predictive Ability. *Econometrica*, 74(6):1545–1578, 2006.

L. Glosten, R. Jagannathan, and D. Runkle. On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. *Journal of Finance*, 48(5):1779–1801, 1993.

M. Habib and L. Stracca. Getting beyond carry trade: What makes a safe haven currency? *Journal of International Economics*, 87(1):50–64, 2012.

H. Herwartz. Stock return prediction under GARCH - an empirical assessment. *International Journal of Forecasting*, 33(3):569–580, 2017.

O. Hossfeld and R. MacDonald. Carry funding and safe haven currencies: A threshold regression approach. *Journal of International Money and Finance*, 59:185–202, 2015.

P.-H. Hsu, Y. Hsu, and C. Kuan. Testing the predictive ability of technical analysis using a new stepwise test without data snooping bias. *Journal of Empirical Finance*, 17(3):471–484, 2010.

P.-H. Hsu, M. Taylor, and Z. Wang. Technical trading: Is it still beating the foreign exchange market? *Journal of International Economics*, 102:188–208, 2016.

L. Husted, J. Rogers, and B. Sun. Uncertainty, currency excess returns, and risk reversals. *Journal of International Money and Finance*, 88:228–241, 2018.

H. Ichiue and K. Koyama. Regime switches in exchange rate volatility and uncovered interest parity. *Journal of International Money and Finance*, 30(7):1436–1450, 2011.

A. Inoue and L. Kilian. In-sample or out-of-sample tests of predictability: Which one should we use? *Econometric Reviews*, 23(4): 371–402, 2005.

M. Iseringhausen. The time-varying asymmetry of exchange rate returns: A stochastic volatility – stochastic skewness model. *Journal of Empirical Finance*, 58:275–292, 2020.

A Ismailov and B. Rossi. Uncertainty and deviations from uncovered interest rate parity. *Journal of International Money and Finance*, 88:242–259, 2018.

M.C. Jones and A. Pewsey. Sinh-arcsinh distributions. *Biometrika*, 96(4):761–780, 2009.

O. Jordà and A.M. Taylor. The carry trade and fundamentals: Nothing to fear but FEER itself. *Journal of International Economics*, 88(1):74 – 90, 2012.

J.W. Jurek. Crash-neutral currency carry trades. *Journal of Financial Economics*, 113(3):325 – 347, 2014.

X. Liu. Modeling time-varying skewness via decomposition for out-of-sample forecast. *International Journal of Forecasting*, 31(1): 30–42, 2015.

R. Meese and K. Rogoff. Empirical exchange rate models of the seventies. Do they fit out of sample? *Journal of International Economics*, 14(1-2):3–24, 1983.

L. Menkhoff, L. Sarno, M. Schmeling, and A. Schrimpf. Carry Trades and Global Foreign Exchange Volatility. *Journal of Finance*, 67(2):681–717, 2012.

H. Pesaran and A. Timmermann. Testing dependence among serially correlated multicategory variables. *Journal of the American Statistical Association*, 104(485):325–337, 2009.

S.H. Poon and Clive W. J. Granger. Forecasting Volatility in Financial Markets: A Review. *Journal of Economic Literature*, XLI: 478–539, 2003.

A. Ranaldo and P. Söderlind. Safe haven currencies. *Review of Finance*, 14:385–407, 2010.

B. Rossi. Exchange rate predictability. *Journal of Economic Literature*, 51(4):1063–1119, 2013.

M. Ulm and J. Hambuckers. Do interest rate differentials drive the volatility of exchange rates? evidence from an extended stochastic volatility model. *Journal of Empirical Finance*, 65:125–148, 2022.

A Overview of skewness specifications

In our empirical analysis, we consider up to 12 specifications of the general skewness equation given by (6). In Table A1, we detail these equations.

Table A1: Tested specifications of the skewness equation.

Specification number	Name	Equation
(1)	**CST**	$\epsilon_t = a_0.$
(2)	**ARX(VIX)**	$\epsilon_t = a_0 + a_1\epsilon_{t-1} + a_3\text{VIX}_{t-1}.$
(3)	**ARX(IRD)**	$\epsilon_t = a_0 + a_1\epsilon_{t-1} + a_4\text{IRD}_{t-1}.$
(4)	**ARX(2)**	$\epsilon_t = a_0 + a_1\epsilon_{t-1} + a_3\text{IRD}_{t-1} + a_4\text{VIX}_{t-1}.$
(5)	**MA**	$\epsilon_t = a_0 + a_2 z_{t-1}.$
(6)	**ARMA**	$\epsilon_t = a_0 + a_1\epsilon_{t-1} + a_2 z_{t-1}.$
(7)	**MAX(VIX)**	$\epsilon_t = a_0 + a_2 z_{t-1} + a_3\text{VIX}_{t-1}.$
(8)	**MAX(IRD)**	$\epsilon_t = a_0 + a_2 z_{t-1} + a_4\text{IRD}_{t-1}.$
(9)	**MAX(2)**	$\epsilon_t = a_0 + a_2 z_{t-1} + a_3\text{IRD}_{t-1} + a_4\text{VIX}_{t-1}.$
(10)	**ARMAX(VIX)**	$\epsilon_t = a_0 + a_1\epsilon_{t-1} + a_2 z_{t-1} + a_3\text{VIX}_{t-1}.$
(11)	**ARMAX(IRD)**	$\epsilon_t = a_0 + a_1\epsilon_{t-1} + a_2 z_{t-1} + a_4\text{IRD}_{t-1}.$
(12)	**ARMAX(2)**	$\epsilon_t = a_0 + a_1\epsilon_{t-1} + a_2 z_{t-1} + a_3\text{IRD}_{t-1} + a_4\text{VIX}_{t-1}.$

Rows indicate the different versions of the skewness equations that are considered in the empirical analysis.

B Additional results

B.1 In-sample results

In this section, we display the performance measures of the in-sample analysis, for EUR/USD (Table B1) and CHF/USD (Table B2). Additional performance measures and computational details of the in-sample analysis can be found in the supplementary material. The columns *Full. SSPA* and *Int. SSPA* give the p-values of the stepwise superior predictability test of Hsu et al. [2010], including either all models, or only the most sophisticated ones. The column *p-val.* contains the p-value of a bootstrap test of no superior performance for a single model.

Table B1: In-sample forecasting performance for EUR

	USD/EUR (1999M1 - 2019M3)									
Specification	CR	PT09	$\hat{m}$	*Full. SSPA*	*Int. SSPA*	*p-val.*	$\Delta\hat{m}$	*Full. SSPA*	*Int. SSPA*	*p-val.*
ARMAX(2)	52.83%	6.52**	5.45%	0.083	0.039	0.023	5.22%	0.124	0.074	0.022
ARMAX(IRD)	52.49%	3.02*	3.49%	0.401	0.257	0.074	3.27%	0.347	0.252	0.094
ARMAX(VIX)	51.57%	2.10	1.36%	0.898	0.739	0.295	1.14%	0.746	0.615	0.287
MAX(2)	52.04%	4.15**	3.40%	0.424	0.275	0.099	3.18%	0.391	0.286	0.110
MAX(IRD)	52.3%	7.20***	2.17%	0.752	0.551	0.181	1.95%	0.581	0.455	0.193
MAX(VIX)	51.47%	3.21*	2.69%	0.614	0.419	0.093	2.47%	0.34	0.248	0.087
ARMA	50.5%	1.92	-1.30%	1.000	-	0.295	-1.51%	0.973	-	0.287
MA	51.68%	5.01**	2.52%	0.668	-	0.068	2.29%	0.294	-	0.06
ARX(2)	51.31%	0.18	2.12%	0.763	-	0.261	1.9%	0.700	-	0.259
ARX(IRD)	51%	0.06	1.39%	0.892	-	0.348	1.17%	0.789	-	0.363
ARX(VIX)	50.25%	0.01	0.57%	0.973	-	0.459	0.35%	0.881	-	0.464
CST	49.79%	0.02	0.42%	0.978	-	0.458	0.2%	0.898	-	0.473
RW^+/RW^-	50.7%	0.00	0.22%	0.983	-	0.467	-	-	-	-
BH/AS	50.1%	-	0.19%	0.985	-	0.470	-	-	-	-

$*$, $**$ and $***$ indicate tests significant at the 10%, 5% or 1% levels, respectively. $\Delta\hat{m}$ refers to the average *excess performance* over the random walk benchmark.

Table B2: In-sample forecasting performance for CHF

	USD/CHF (1999M1 - 2019M3)									
Specification	CR	PT09	$\hat{m}$	*Full. SSPA*	*Int. SSPA*	*p-val.*	$\Delta\hat{m}$	*Full. SSPA*	*Int. SSPA*	*p-val.*
ARMAX(2)	52.52%	11.52***	3.91%	0.314	0.175	0.059	-0.01%	0.859	0.736	1.000
ARMAX(IRD)	52.19%	7.70***	3.27%	0.444	0.264	0.085	-0.64%	0.912	0.815	1.000
ARMAX(VIX)	51.16%	1.26	0.84%	0.939	0.713	0.355	-3.07%	0.994	0.973	1.000
MAX(2)	52.33%	9.72***	6.20%	0.043	0.023	0.005	2.28%	0.555	0.421	0.223
MAX(IRD)	52.21%	7.37***	5.89%	0.060	0.032	0.007	1.97%	0.606	0.456	0.256
MAX(VIX)	50.99%	0.79	-0.23%	1.000	1.000	1.000	-4.14%	0.999	0.990	1.000
ARMA	51.28%	1.81	1.11%	0.913	-	0.302	-2.81%	0.992	-	1.000
MA	50.87%	0.38	-0.45%	1.000	-	1.000	-4.36%	1.000	-	1.000
ARX(2)	50.91%	0.09	0.19%	0.985	-	0.458	-3.73%	0.993	-	1.000
ARX(IRD)	51.58%	0.32	4.64%	0.177	-	0.037	0.73%	0.804	-	0.428
ARX(VIX)	50.06%	1.12	-1.67%	1.000	-	1.000	-5.59%	1.000	-	1.000
CST	50.12%	1.27	-1.58%	1.000	-	1.000	-5.49%	1.000	-	1.000
RW^{+}/RW^{-}	50.87%	0.00	3.92%	0.309	-	0.055	-	-	-	-
BH/AS	49.96%	-	3.72%	0.855	-	0.239	-	-	-	-

*, ** and *** indicate tests significant at the 10%, 5% and 1% levels, respectively. $\Delta\hat{m}$ refers to the average *excess performance* over the random walk benchmark.

B.2 Robustness and additional test results

In this section, we display additional results related to the specification tests of Berkowitz [2001], Doornik and Hansen [2008], and Anderson-Darling in Table B3 for both EUR and CHF.

Table B3: Model selection and specification criteria.

EUR	ARMAX(2)	ARMAX(IRD)	ARMAX(VIX)	MAX(2)	MAX(IRD)	MAX(VIX)
LLF	-19,026.94	-19,022.82	-19,019.48	-19,024.26	-19,020.47	-19,017.16
AIC	-38,031.88	-38,025.64	-38,018.96	-38,028.52	-38,022.94	-38,016.32
BIC	-37,960.02	-37,960.32	-37,953.64	-37,963.20	-37,964.14	-37,957.53
LR	-	8.24***	14.92***	5.36**	12.94***	19.56***
BK	0.702	0.758	0.840	0.252	0.735	0.770
DH	2.885	3.030	3.217	2.758	2.895	3.019
AD	0.353	0.371	0.328	0.369	0.339	0.357
CHF	**ARMAX(2)**	**ARMAX(IRD)**	**ARMAX(VIX)**	**MAX(2)**	**MAX(IRD)**	**MAX(VIX)**
LLF	-18,670.06	-18,669.88	-18,662.59	-18,669.52	-18,669.38	-18,661.98
AIC	-37,318.11	-37,319.77	-37,305.18	-37,319.05	-37,320.75	-37,305.97
BIC	-37,246.26	-37,254.44	-37,239.86	-37,253.73	-37,261.96	-37,247.18
LR	-	0.35	14.94***	1.07	1.37	16.15***
BK	3.32	3.36	3.06	3.39	3.41	3.13
DH	18.63***	18.74***	16.57***	18.95***	19.04***	16.90***
AD	1.08	1.09	1.07	1.12	1.12	1.11

LLF denotes the value of the negative log-likelihood function. The line *LR* displays the likelihood ratio test statistics between ARMAX(2) and the competing models. The lines labelled BK, DH and AD report the test statistics for Berkowitz [2001], Doornik and Hansen [2008] and Anderson-Darling tests, respectively. *** denote tests significant at the 1% level.

B.3 Out-of-sample forecasting performance

In this section, we provide additional results for the out-of-sample analysis discussed in Section 3. In Tables B4 to B7, we display the detailed out-of-sample performance results for each tested skewness specification.

Table B4: Out-of-sample forecasting performance for EUR

USD/EUR (2014M12 - 2019M3)								
	Classification		average profit			average excess performance		
Specification	CR	PT09	$\hat{m}$	DM	GW	$\Delta\hat{m}$	DM	GW
ARMAX(2)	51.16%	0.30	5.20%	1.21	4.73	6.46%	0.94	2.59
ARMAX(IRD)	52.09%	4.08**	2.76%	0.64	0.92	4.02%	0.58	0.57
ARMAX(VIX)	50.23%	1.99	1.62%	0.37	0.62	2.88%	0.44	0.97
MAX(2)	49.40%	0.15	-0.60%	-0.14	0.08	0.66%	0.09	0.14
MAX(IRD)	49.30%	0.02	-2.77%	-0.64	0.53	-1.51%	-0.21	0.05
MAX(VIX)	47.72%	1.00	-2.89%	-0.67	0.73	-1.63%	-0.24	0.07
ARMA	51.16%	2.32	2.27%	0.50	0.25	3.53%	0.49	0.23
MA	48.47%	0.22	-2.70%	-0.61	0.63	-1.44%	-0.20	0.05
ARX(2)	49.21%	0.18	-1.59%	-0.36	0.16	-0.33%	-0.05	1.43
ARX(IRD)	48.28%	0.29	-4.18%	-0.99	1.05	-2.92%	-0.50	0.38
ARX(VIX)	47.72%	0.11	0.37%	0.08	0.00	-1.63%	0.29	0.09
CST	47.44%	1.63	-3.05%	-0.74	0.85	-1.79%	-0.31	0.23
RW^+/RW^-	48.00%	0.00	-1.26%	-0.34	0.16	-	-	-
BH/AS	47.81%	-	-2.26%	-0.55	0.31	-	-	-

Table entries denote correct classification rate (CR), average profit ($\hat{m}$), and average excess performance over the random walk benchmark ($\Delta\hat{m}$). Entries for PT09, DM, GW display the respective test statistics. ** indicates tests significant at the 5% levels.

Table B5: Additional out-of-sample measures for USD/EUR

USD/EUR (2014M12 - 2019M3)						
Specification	Sharpe	Skew	AUC	AUC*	G/L	Max. Draw.
ARMAX(2)	0.59	0.03	0.53	0.53	1.10	0.17
ARMAX(IRD)	0.31	-0.04	0.53	0.51	1.05	0.17
ARMAX(VIX)	0.18	0.10	0.53	0.52	1.03	0.16
MAX(2)	-0.07	0.05	0.51	0.52	0.99	0.20
MAX(IRD)	-0.31	0.08	0.52	0.51	0.95	0.23
MAX(VIX)	-0.33	0.06	0.51	0.51	0.95	0.22
ARMA	0.26	0.04	0.53	0.51	1.04	0.20
MA	-0.31	0.05	0.52	0.51	0.95	0.23
ARX(2)	-0.18	-0.10	0.50	0.50	0.97	0.29
ARX(IRD)	-0.47	-0.10	0.50	0.50	0.92	0.24
ARX(VIX)	0.04	0.20	0.49	0.50	1.01	0.13
CST	-0.35	0.09	0.50	0.51	0.94	0.21
RW^+/RW^-	-0.14	-0.28	0.48	0.49	0.98	0.24
BH/AS	-0.26	0.09	-	-	0.96	0.18

Additional out-of-sample performance measures for the profit of the trading rules derived from the different models (USD/EUR). The column entries denote the Sharpe ratio, the skewness of the daily profit, the area under the correct classification frontier (AUC), its return-weighted version (AUC*) of Jordà and Taylor [2012], the gain-loss (G/L) ratio of Bernardo and Ledoit [2000] and the maximum drawdown on the compounded profits.

Table B6: Out-of-sample forecasting performance for CHF

USD/CHF (2014M12 - 2019M3)								
	Classification		average profit			average excess performance		
Specification	CR	PT09	$\hat{m}$	DM	GW	$\Delta\hat{m}$	DM	GW
ARMAX(2)	51.16%	1.27	5.16%	1.07	1.21	4.21%	0.54	4.95*
ARMAX(IRD)	51.53%	0.92	5.10%	1.06	1.12	4.15%	0.52	5.70*
ARMAX(VIX)	51.26%	2.26	3.32%	0.59	1.87	2.37%	0.24	1.18
MAX(2)	50.79%	0.31	7.64%	1.33	2.97	6.69%	0.86	1.18
MAX(IRD)	50.60%	0.12	6.49%	1.22	2.70	5.54%	0.71	0.65
MAX(VIX)	51.26%	3.55*	1.71%	0.32	1.18	0.76%	0.08	0.91
ARMA	52.47%	3.20*	3.21%	0.55	1.77	2.26%	0.23	1.18
MA	51.72%	4.71**	2.60%	0.48	1.51	1.65%	0.17	1.03
ARX(2)	51.16%	1.47	6.85%	1.31	3.76	5.90%	1.01	1.26
ARX(IRD)	50.79%	0.58	7.63%	1.46	4.62*	6.68%	1.09	1.45
ARX(VIX)	50.70%	0.15	3.19%	0.62	2.22	2.24%	0.39	0.57
CST	48.56%	0.32	3.72%	0.75	2.00	2.77%	0.48	0.43
RW^+/RW^-	47.72%	0.00	0.95%	0.18	1.22	-	-	-
BH/AS	47.07%	-	-0.69%	-0.13	0.80	-	-	-

Please refer to the notes in Table B4 for explanations on the table entries. * and ** indicate tests significant at the 10% and 5% levels.

Table B7: Additional out-of-sample forecasting performance for CHF

USD/CHF (2014M12 - 2019M3)						
Specification	Sharpe	Skew	AUC	AUC*	G/L	Max. Draw.
ARMAX(2)	0.48	6.02	0.50	0.51	1.10	0.11
ARMAX(IRD)	0.47	6.04	0.51	0.52	1.10	0.11
ARMAX(VIX)	0.31	-6.79	0.50	0.50	1.07	0.18
MAX(2)	0.71	6.20	0.51	0.52	1.16	0.13
MAX(IRD)	0.60	6.17	0.51	0.53	1.13	0.12
MAX(VIX)	0.16	-6.82	0.51	0.50	1.03	0.18
ARMA	0.30	-6.89	0.53	0.50	1.06	0.18
MA	0.24	-6.84	0.51	0.50	1.05	0.18
ARX(2)	0.63	6.25	0.52	0.55	1.14	0.16
ARX(IRD)	0.71	6.15	0.51	0.55	1.16	0.12
ARX(VIX)	0.30	6.21	0.53	0.52	1.06	0.19
CST	0.34	6.30	0.52	0.51	1.08	0.10
RW^+/RW^-	0.09	6.87	0.49	0.51	1.02	0.31
BH/AS	-0.06	6.86	-	-	0.99	0.19

Please refer to the notes in Table B5 for explanations on the table entries.