Climate Risks and Realized Volatility of Major Commodity Currency Exchange Rates

Submission: February 2021

Matteo Bonato[*], Oğuzhan Çepni[†], Rangan Gupta[‡], Christian Pierdzioch[§]

Abstract

We find that climate-related risks forecast the intraday-data-based realized volatility of exchange-rate returns of eight major fossil fuel-exporters (Australia, Brazil, Canada, Malaysia, Mexico, Norway, Russia, and South Africa). We study a wide array of metrics capturing risks associated with climate change, derived from data directly on variables such as, for example, abnormal patterns of temperature. We control for various other moments (realized skewness, realized kurtosis, realized good and variance, upside and downside tail risk, and jumps) and estimate our forecasting models using random forests, a machine-learning technique tailored to analyze models with many predictors.

JEL Classifications: C22, C53, F31, Q54.

Keywords: Climate Risks; Commodity Currencies; Realized Variance; Forecasting

Conflicts of interest: The authors declare no conflict of interest.

[*]Department of Economics and Econometrics, University of Johannesburg, Auckland Park, South Africa; IPAG Business School, 184 Boulevard Saint-Germain, 75006 Paris, France. Email address: matteobonato@gmail.com.

[†]Copenhagen Business School, Department of Economics, Porcelænshaven 16A, Frederiksberg DK-2000, Denmark; Email address: oce.eco@cbs.dk

[‡]Department of Economics, University of Pretoria, Pretoria, 0002, South Africa; Email address: rangan.gupta@up.ac.za

[§]Department of Economics, Helmut Schmidt University, Holstenhofweg 85, P.O.B. 700822, 22008 Hamburg, Germany; Email address: macroeconomics@hsu-hh.de.

1 Introduction

Farhi and Gabaix (2016), based on significant earlier works by Rietz (1988) and Barro (2006), developed a new model of exchange rates, which highlights that the possibility of rare but extreme disasters is a key determinant of exchange-rate fluctuations (returns and volatility). In their model, a world disaster might occur at any point in time, and with disasters corresponding to bad times, such events cause an exchange-rate depreciation, where the extent depends on the relative riskiness of the two currencies being involved. Moreover, in their model, the probability of a world disaster and each country's exposure to such an event varies over time, resulting in large fluctuations (volatility) in exchange rates. Given these theoretical predictions, Gupta et al., (2019) empirically analyzed the in-sample predictability of international political crises acting as proxies for rare disaster risks[1] for the United States (U.S.) dollar-based exchange rates for Brazil, Russia, India, China, South Africa, and the United Kingdom (UK). Based on a non-parametric causality-in-quantiles test, and historical monthly data covering the period from 1918 to 2013, they showed that indeed rare disaster risks affect both returns and volatility over the majority of respective conditional distributions, with the predictability of volatility being relatively stronger than that of returns.[2]

Against this backdrop of in-sample evidence involving rare disaster risks and exchange-rate volatility, our objective is to conduct an elaborate out-of-sample forecasting analysis of the role of such risks for the predictability of exchange-rate volatility, since in-sample evidence does not necessarily translate into out-of-sample forecasting gains (Rapach and Zhou, 2013). In this regard, Campbell (2008), and more recently Rapach and Zhou (forthcoming), points out that the ultimate test of any predictive model (in terms of the econometric methodology(ies) and the predictor(s) being used) is in its out-of-sample forecasting performance.

The foreign exchange market is the largest and most liquid financial market in the world. As reported in the Triennial Survey of global foreign exchange market volumes of the Bank for International Settlement (BIS), the average daily turnover was 6.6 trillion US dollars in April of 2019 (up from 5.1 trillion U.S. dollars three years earlier). Naturally, in light of the size of this market, besides the statistical significance of a forecasting exercise, real-time forecasts are important to multinational firms, financial institutions, and traders aiming to hedge currency risks (Balcilar et al., 2016). Traders of foreign-currency options look to make profits by buying (selling) options if they expect volatility to rise above (fall below) of what is implied in currency option premiums. In addition, a large body of theoretical and empirical research has linked exchange-rate volatility to

[1]Park and Park (2020) used news articles reporting on North Korea to capture rare disaster risks, and found a strong depreciation of the South Korean won using non-parametric regressions- and event study-based approaches. Alternatively, identifying the number of confirmed cases associated with the ongoing COVID-19 pandemic as a rare disaster, Zhou et al., (2021), found evidence of a strong depreciation of the currencies of 27 advanced and emerging countries.

[2]In line with the theoretical claims, Gupta et al., (2019) also found that the effect is particularly strong on returns of emerging-market currencies, because such currencies are perceived to be more risky, when compared to that of the British pound.

trade and welfare (see, Clark et al., (2004), Asteriou et al., (2016), Senadza and Diaba (2017) for detailed reviews). Hence, accurate forecasting of exchange-rate volatility is of paramount importance to both investors and policymakers.

In this regard, the literature on forecasting exchange-rate volatility has primarily relied on univariate and multivariate versions of the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models, and also more recently multifractal models (see Pong et al., (2004), Hansen and Lunde (2005), Rapach and Strauss (2008), Christou et al., (2018) and Liu et al., (2020) for discussions of this literature). McAleer and Medeiros (2008), however, pointed out that the rich information contained in intraday data can lead to more accurate estimates and forecasts of daily volatility. Given this, we make use of variations of the Heterogeneous Autoregressive realized volatility (HAR-RV) model proposed by Corsi (2009) to forecast the realized volatility of the exchange-rate returns (i.e., the sum of non-overlapping squared high-frequency, 5-minute-interval intraday returns observed within a day, following Andersen and Bollerslev (1998)). In the process, we obtain an unconditional and observable metric of volatility. Moreover, the HAR-RV model has become increasingly popular because of its ability to capture well-established features of exchange-rate (and, in general, financial-returns) volatility such as long-memory and multi-scaling behaviour extensively documented in the empirical literature (e.g., Mei et al., (2017)). The key feature of the HAR-RV model is that it uses volatilities from different time resolutions to forecast the realized volatility of asset-price returns.

In order to capture the role of rare disaster risks in forecasting exchange-rate volatility, we augment the benchmark HAR-RV model to include proxies related to climate risks, as motivated by the burgeoning literature on "Climate Finance" (Giglio et al., (2021), Stroebel and Wurgler (2021)) over and above moments-based predictors (also derived from intraday data) namely, realized skewness, realized kurtosis, realized good and bad variance, upside and downside tail risk, and jumps. In other words, we, for the first time, go beyond the existing literature on forecasting exchange-rate realized volatility based on its moments (see for example, Andersen et al., (2003), Mcmillan and Speight (2012), Louzis et al., (2013), Barunik et al., (2016), Gkillas et al., (2019), Plíhal and Lyócsa (2021)) by considering daily metrics of climate-related risks.

The risks associated with climate change can be typically categorized into physical risks and transition risks. The former comprise risks arising due to, for example, rising temperatures, higher sea levels, destructive storms, and floods, and wildfires, while the latter arise in the wake of a gradual switch over to a low-carbon economy and comprise risk due to changes in the stance of climate policy, the development of disruptive green technologies, and changing patterns in consumer preferences. Understandably, as pointed out by Battiston et al., (2021) and Flori et al., (2021), every future scenario includes climate-related financial risks (though the level and form of the source of uncertainty may vary) due to the occurrence of rare disasters. Not surprisingly, climate risks, as captured by textual and narrative analyses of climate-change news (related to natural disasters, global warming, international summits, and climate policy) or via movements of temperature and precipiation, have provided a source of (high-frequency) proxies of possible forthcoming rare disaster risks (see, for example, Choi et al. (2020), Engle et al., (2020), Kapfhammer et al., (2020), Faccini et al., (2021)).

Motivated by this line of research, we use the information content of seasonal, predictable, and abnormal patterns of temperature, precipitation, number of heating degree days, number of cooling degree days, and wind speed to capture climate-related risks. In addition, as a check of the robustness of our results in terms of the proxies for rare disaster risks derived from climate risks, we use the daily Google Search Volume Index (SVI) of the topic "global warming" in a country, as well as the news trends (NT) function of the Bloomberg terminal to collect the news counts involving the term "climate change" for each country. The intensity of such keywords-based searches can also capture the transition risks component of climate change, over and above the physical risks.

The transition away from coal, gas, and oil, being bad news for commodity exporters, are likely to impact, just like physical risks, currency valuations of fossil-fuel exporters, and, hence, affect the variability of their exchange-rate returns. In light of this, as far as the exchange rates are concerned, we investigate the predictability of the realized volatility of the exchange rates of eight major fossil-fuel exporters (Australia, Brazil, Canada, Malaysia, Mexico, Norway, Russia, and South Africa), in line with the recent work by Kapfhammer et al., (2020). They, by proposing a novel news media-based measure of climate risk, show that when such risk is high, these major commodity currencies experience a persistent depreciation, and the relationship between commodity-price fluctuations and currencies tends to become weaker.[3]

Given that we extend the standard HAR-RV model to include several additional predictors, as our main econometric framework, we use a machine-learning technique to estimate our forecasting models, because overparameterization in a standard predictive regression framework results in poor out-of-sample performance. In this regard, we use random forests (Breiman, 2001) to estimate our forecasting models. Random forests are a machine-learning technique that is specifically tailored to operate in settings featuring a large array of predictors, while automatically recovering potential nonlinear links between the realized exchange-rate volatility and its various predictors, and any interaction effects between the predictors. As a matter of robustness, we also conduct our forecasting analysis by using the least absolute shrinkage and selection operator (Lasso) technique, proposed by Tibshirani (1996). The Lasso belongs to the spectrum of linear regression-analysis methods and performs, in a data-driven way, both model selection and regularization.

[3]For a related strand of recent research on commodity currencies, see, for example, Chen et a. (2010), who find that commodity-currency exchange rates predict subsequent commodity-price changes using quarterly data, and that evidence of the reverse predictability is not robust. Ferraro et. al. (2015) find that changes in commodity prices help to forecast exchange-rate returns at a daily frequency.

Our study contributes to the growing literature on reactions to climate risks and other external conditions. Duan et al. (2020) examined the pricing of carbon risk in the U.S. corporate bond market and find that carbon emissions intensity of firms provides predictability for future excess corporate bond returns. Atanasova and Schwartz (2019) showed that the expansion of fossil fuel reserves in commodity-exporter firms had a negative impact on company value, indicating that capital markets perceive fossil fuel as "stranded assets" in the transition to a lower-carbon economy. Our study, hence, differs from previous literature focusing on the pricing of carbon risk through the lens of the stock and bond market, because we examine how climate risk contributes to improve forecasts of exchange-rate volatility. Furthermore, using proxies for retail and institutional attention, our study adds to earlier literature linking investor attention and news media to asset prices (Andersson et al., 2016; Ben-Rephael et al., 2017; Choi et al., 2020; Engle et al., 2020). Finally, in contrast to earlier studies, which generally employ annual data and use local temperature measures (Colacito et al., 2019; Henseler and Schmuacher, 2019), our paper differs in that we directly investigate the predictability of exchange-rate volatility by constructing several high frequency measures of climate-related risks based on extreme weather conditions, which renders it possible to capture timely reaction of the foreign-exchange market to climate risk.

We organize the remainder of our paper as follows. In Section 2, we describe our machine-learning methodology, along with the derivation of realized volatility and its moments that we use for our empirical study. In Section 3, we present the data, while in Section 4 we discuss our empirical results both in terms of statistical and economic gains. In Section 5, we conclude.

2 Methodology

2.1 Empirical Models

The nucleus of our forecasting models is the familiar HAR-RV model (Corsi, 2009), which has been widely studied in the empirical-finance literature because it captures the long-memory and multi-scaling behavior typical of financial data. The HAR-RV model is given by

$$RV_{t+h} = \beta_0 + \beta_d RV_t + \beta_w RV_{w,t} + \beta_m RV_{m,t} + \epsilon_{t+h}, \tag{1}$$

where the index, h, denotes the forecast horizon, and (for $h > 1$) RV_{t+h} denotes the average realized variance over the h-days forecast horizon. We study a short (daily, $h = 1$), a medium (weekly, $h = 5$), and a long (monthly, $h = 22$) forecast horizon. Further, $RV_{w,t}$ is defined as the average RV from day $t-5$ to day $t-1$, while $RV_{m,t}$ is constructed as the average RV from day $t-22$ to day $t-1$. Hence, the HAR-RV model formalizes the heterogeneous market hypothesis (Müller et al., 1997) according to which financial markets are populated by market participants who trade at different time scales and who differ, therefore, with respect to their sensitivity to information flows at different time horizons.

When we take into account for the moments of realized variance like, for example, realized skewness and realized kurtosis (described in details in Section 2.2) as additional predictors, we obtain the HAR-RV-M model:

$$RV_{t+h} = \beta_0 + \beta_d RV_t + \beta_w RV_{w,t} + \beta_m RV_{m,t} + \sum_{j=1}^{n} \theta_j M_{j,t} + \epsilon_{t+h}, \tag{2}$$

where $M_{j,t}$ denotes the realization of moment $j = 1, 2, , , ..., n$ on day t. When we extend the HAR-RV-M model to include the climate variables temperature, HDD, CDD, precipitation, wind speed, we obtain the HAR-RV-M-C model:

$$RV_{t+h} = \beta_0 + \beta_d RV_t + \beta_w RV_{w,t} + \beta_m RV_{m,t} + \sum_{j=1}^{n} \theta_j M_{j,t} + \sum_{j=1}^{5} \gamma_j C_{j,t} + \epsilon_{t+h}, \tag{3}$$

where $C_{jt}, j = 1, 2, ..., 5$ denotes the climate variables capturing extreme weather conditions. We obtain an alternative HAR-RV-M-climate model, when we use climate-related Google searches and U.S. news as our predictors. We call this alternative specification as HAR-RV-M-CN model. It is given by

$$RV_{t+h} = \beta_0 + \beta_d RV_t + \beta_w RV_{w,t} + \beta_m RV_{m,t} + \sum_{j=1}^{n} \theta_j M_{j,t} + \sum_{j=1}^{2} \gamma_j CN_{j,t} + \epsilon_{t+h}. \tag{4}$$

Since the extended HAR-RV models given in Equations (2) to (4) feature several predictors, we use a machine-learning techniques known as random forests as our main estimation tool. Random forests build forecasting models in a completely data-driven way, they account for any potential interaction effects between the predictors, and they automatically recover even complex potential nonlinear links between the dependent variable and its various predictors.

The basic idea underlying random forests is to replace the linear additive models specified in Equations (1) to (4) with a forecasting model assuming that the dependent variable can be expressed as a general (nonlinear) unknown function, f, of the predictors. In the case of the HAR-RV-M-C model, this idea can be formalized by writing

$$RV_{t+h} = f(RV_t, RV_{w,t}, RV_{m,t}, M_{1,t}, M_{2,t}, ..., M_{n,t}, C_{1,t}, C_{2,t}, ...). \tag{5}$$

Random forests, first proposed by Breiman (2001), are an ensemble machine-learning technique that recover the function, f, using a large number of individual random regression trees (for a textbook exposition, see Hastie et al., 2009). A regression tree, in turn, is made up of several nodes and branches that partition the predictor space systematically into non-overlapping regions in a hierarchical and binary manner. When a researcher grows a large regression tree that subdivides the predictor space into many such regions, it is possible to compute increasingly fine-grained predictions of realized exchange-rate volatility. At the same time, however, the structure of a large regression tree becomes increasingly complex and this complexity eventually deteriorates forecasting performance because of the necessarily arising over-fitting problem and sensitiveness to the data. A random forest overcomes this over-fitting problem by combining the forecasts from a large number of individual random regression trees, each fitted on bootstrapped sample of data. A random regression tree differs from a "simple" regression tree in that it uses a random subset of the predictors to grow the tree. These additional elements of randomness curb the impact of influential individual predictors (possibly due to past volatility clustering, jumps, or tail events) on predictions and, in addition, decorrelate the predictions from individual trees. In this way, averaging the predictions across the individual random regression trees leads to a stabilization of predictions and, thereby, superior forecasts of exchange-rate volatility.

In order to be more specific on how a regression tree, T, can be constructed in a completely data-driven manner, we recap that the basic idea is to form branches that partition the space of predictors, $x = (x_1, x_2, ...)$, into l non-overlapping regions, R_l. Starting at the the top level of a regression tree, a search-and-split algorithm is used to identify a first optimal partition of the data. This algorithm searches across the array of predictors and the corresponding candidates for split points, p, to form two half-planes, $R_1(s,p) = \{x_s | x_s \leq p\}$ and $R_2(s,p) = \{x_s | x_s > p\}$. The optimal combination of predictor and split point minimizes the following squared-error loss objective function:

$$\min_{s,p} \left\{ \min_{\bar{RV}_1} \sum_{x_s \in R_1(s,p)} (RV_i - \bar{RV}_1)^2 + \min_{\bar{RV}_2} \sum_{x_s \in R_2(s,p)} (RV_i - \bar{RV}_2)^2 \right\}, \tag{6}$$

where we have dropped the time index to keep the sub-indices simple. The subscript i denotes those observations of realized exchange-rate volatility, RV, that belong to a half-plane, and $\bar{RV}_k = \text{mean}\{RV_i | x_s \in R_k(s,p)\}, k = 1, 2$ denotes the half-plane-specific mean of RV.

At the second level of a regression tree, the objective function given in Equation (6) is used to find the optimal combination of predictors and split points for the respective half planes identified at the top level of the regression tree. Application of the search-and-split algorithm at the second tree level, thus, gives up to two second-level optimal splitting predictors and their corresponding optimal splitting points. Upon recursively applying the search-and-split algorithm in a top-down way to the next (third, fourth, ...) tree levels, we can build in a pyramid-like manner a hierarchical regression tree, where every new tree level adds an additional layer of complexity to the tree. Two simple approaches to terminate this tree-building process are to fix, in advance, a maximum number of terminal nodes for a regression tree or to require that the terminal nodes contain a minimum number of data points. Once the tree-building process has stopped, a researcher can use a regression tree with L regions to predict realized exchange-rate volatility by checking to which terminal region a realization of the predictors belongs. The corresponding region-specific mean of RV is then the optimal prediction of exchange-rate volatility:

$$T\left(\mathbf{x}_i, \{R_l\}_1^L\right) = \sum_{l=1}^{L} \bar{RV}_l \mathbf{1}(\mathbf{x}_i \in R_l), \tag{7}$$

where $\mathbf{1}$ denotes the indicator function. In the same way, an out-of-sample forecast is obtained by updating the predictor variables, trickling the new data down the in-sample hierarchical tree structure, and then using Equation (7) to predict exchange-rate volatility. Using to this end a large number of random regression trees (that is, regression trees that apply the search-and-split algorithm to a random subset of the predictors to build a tree) estimated on bootstrapped samples of the data gives random-forests out-of-sample predictions of exchange-rate volatility.

2.2 Realized Variance and Its Moments

Exchange-rate returns realized volatilities are estimated relying on the classical estimator of RV, i.e., as the sum of squared intraday returns (Andersen and Bollerslev, 1998):

$$RV_t = \sum_{i=1}^{M} r_{t,i}^2, \tag{8}$$

where $r_{t,i}$ denotes the intraday $M \times 1$ return vector, and $i = 1, ..., M$ is the number of intraday returns.

Upon using only positive or negative intraday returns, one can construct measures of semi-variance which can serve as measures of downside and upside risk, and capture the sign asymmetry of FX rates. These are called “bad’ (RVB) and “good’ (RVG) realized variance. We, thus, construct RVG and RVB by replacing RV ($= RVB + RVG$) in the above equation by RVB and RVG. In line with Barndorff-Nielsen et al. (2010), we compute the bad and good realized semi-variances as

$$RVB_t = \sum_{i=1}^{M} r_{t,i}^2 \ \mathbf{1}_{[(r_{t,i})<0]}, \tag{9}$$

$$RVG_t = \sum_{i=1}^{M} r_{t,i}^2 \ \mathbf{1}_{[(r_{t,i})>0]}. \tag{10}$$

Next, we compute the higher-moments (realized skewness, and realized kurtosis) of the daily FX returns distribution, that is, RSK and RKU. We consider RSK as a measure of the asymmetry of the daily exchange-rates returns distribution, and RKU as a measure that accounts for extremes.

We compute RSK on day t as

$$RSK_t = \frac{\sqrt{M}\sum_{i=1}^{M} r_{(i,t)^3}}{RV_t^{3/2}}, \tag{11}$$

and RKU on day t as

$$RKU_t = \frac{M\sum_{i=1}^{M} r_{(i,t)^4}}{RV_t^2}. \tag{12}$$

The scaling of RSK and RKU by $(M)^{1/2}$ and M ensures that magnitudes correspond to daily skewness and kurtosis.

In addition to RSK and RKU, we consider a jump component ($JUMPS$) in the FX price process. As shown by Andersen et al. (2005), jumps are both highly prevalent and distinctly less persistent than the continuous sample-path-variation process. Moreover, many jumps appear directly associated with specific macroeconomic news announcements. We use the result, derived by Barndorff-Nielsen and Shephard (2004), that the realized variance converges into permanent and discontinuous (jump) components as

$$\lim_{M\to\infty} RV_t = \int_{t-1}^{t} \sigma^2(s)ds + \sum_{j=1}^{N_t} k_{t,j}^2, \tag{13}$$

where N_t is the number of jumps within day t and $k_{t,j}$ is the jump size. This result implies that RV_t is a consistent estimator of the integrated variance $\int_{t-1}^{t}\sigma^2(s)ds$ plus the jump contribution. The asymptotic results derived by Barndorff-Nielsen and Shephard (2004, 2006) further show that

$$\lim_{M\to\infty} BV_t = \int_{t-1}^{t} \sigma^2(s)ds, \tag{14}$$

where BV_t is the realized bipolar variation defined as

$$BV_t = \mu_1^{-2}\left(\frac{M}{M-1}\right)\sum_{i=2}^{M} |r_{t,i-1}||r_{i,t}| = \frac{\pi}{2}\sum_{i=2}^{M} |r_{t,i-1}||r_{i,t}|, \tag{15}$$

where

$$\mu_a = E(|Z|^a), Z \sim N(0,1), a > 0. \tag{16}$$

A consistent estimator of the pure jump contribution can then be expressed as

$$J_t = RV_t - BV_t. \tag{17}$$

The significance of the jump component is tested relying on a formal test estimator proposed by

Brandorff-Nielsen and Shephard (2006) given by

$$JT_t = \frac{RV_t - BV_t}{(v_{bb} - v_{qq})\frac{1}{N}TP_t}, \tag{18}$$

where TP_t is the Tri-Power Quarticity:

$$TP_t = M\frac{M}{M-2}\left(\frac{\Gamma(0.5)}{2^{2/3}\Gamma(7/6)}\right)\sum_{i=3}^{M}|r_{t,i}|^{4/3}|r_{t,i-1}|^{4/3}|r_{t,i-2}|^{4/3}, \tag{19}$$

which converges to the Integrated Quarticity IQ_t

$$IQ_t \to \int_{t-1}^{t}\sigma^4(s)ds, \tag{20}$$

even in the presence of jumps. We use the notation $v_{bb} = \left(\frac{\pi}{2}\right) + \pi - 3$ and $v_{qq} = 2$. Note that for each t, $JT_t \sim N(0,1)$ as $M \to \infty$. As can be seen in Equation (17), the jump contribution to RV_t is either positive or null. Therefore, so as to avoid obtaining negative empirical contributions, we redefine, like Zhou and Zhu (2012), the jump measure as

$$RJ_t = max(RV_t - BV_t; 0). \tag{21}$$

Lastly, we consider the Hill tail-risk estimator (Hill, 1975), using $X_{t,i}$ the set of reordered intraday returns $r_{t,i}$, in such a way that

$$X_{t,i} \geq X_{t,j} \text{ for } i < j. \tag{22}$$

The Hill positive tail-risk estimator is then defined as

$$H_t^{up} = \frac{1}{k}\sum_{j=1}^{k}\ln(X_{t,i}) - \ln(X_{t,k}), \tag{23}$$

and the negative tail-risk estimator as

$$H_t^{down} = \frac{1}{k}\sum_{j=n-k}^{n}\ln(X_{t,i}) - \ln(X_{t,n-k}), \tag{24}$$

where k denotes the number of observation employed to achieve the chosen α tail interval.

3 Data

We collect 5-minute interval intraday data on the currencies of major commodity-exporter economies in terms of fossil fuel against the US dollar over a 24 hour trading day. The commodity-exporter economies are Australia (USDAUD), Brazil (USDBRL), Canada (USDCAD), Malaysia (USDMYR), Mexico (USDMXN), Norway (USDNOK), Russia (USDRUB), and South Africa (USDUZAR). Using the intraday data and formulas introduced in Section 2, we compute daily measures of realized variance, the corresponding good and bad variants, and the other covariates, i.e., realized jumps, realized skewness, realized kurtosis, upside tail risk, downside tail risk, and jumps. Because the realized variance is known to exhibit occasional large peaks, we use the natural log of the square root (that is, volatility) of realized variance as our dependent variable, and, when computing forecasts (and, thus, reverting back to anti-logs), account for this transformation by adding a Jensen-Ito term to forecasts. Table 1 reports the number of observations and the start and end dates of our sample period for every exchange rate based on the data availability. We also report summary statistics of our dependent variable along with its first five coefficients of auto-correlation. The slowly decaying coefficients of auto-correlation underline that the HAR-RV model is a useful starting point for our forecasting analysis.

– Please include Table 1 about here. –

We obtain daily weather data from the Bloomberg terminal, produced by the National Climatic Data Center (NCDC), for our sample countries and the US[4]. The weather data captures several factors, including temperature, precipitation, number of heating degree days (HDD), number of cooling degree days (CDD), and wind speed as described below:

- **Temperature**: The average temperature (usually of the high and low) that was observed between 7am and 7pm local time, expressed in Fahreinheit.
- **HDD**: The number of degrees below 65 degrees Fahrenheit of the mean temperature used to estimate the energy needed to heat a building.
- **CDD**: The number of degrees above 65 degrees Fahrenheit of the mean temperature used to estimate the energy needed to cool a building.
- **Precipitation**: The amount of rain, snow, sleet or hail that falls in a specific location.
- **Wind speed**: Average of sustained winds which does not include wind gust, expressed in knots.

Following the approach described by Choi et al., (2020), we decompose the weather-related variables into three components that account for seasonal, predictable, and abnormal patterns. In particular, for each country, i, and day, t, we compute the daily $WeatherMeasure_{it}$ using the

[4]Since we collect our currency data against the US dollar, we include the climate variables of the US to capture the relative riskiness of the two currencies being involved.

following formula:

$$WeatherMeasure_{it} = Aver\_WeatherMeasure_{it} + Mon\_WeatherMeasure_{it} + Ab\_WeatherMeasure_{it}, \quad (25)$$

where $WeatherMeasure_{it}$ denotes temperature, HDD, CDD, precipitation, or wind speed. Moreover, $Aver\_WeatherMeasure_{it}$ is the average of $WeatherMeasure_{it}$ in country i over the 120 months prior to t, $Mon\_WeatherMeasure_{it}$ is the average deviation of the $WeatherMeasure_{it}$ from the monthly average temperature in country i in the same calendar month over the last 10 years minus $Aver\_WeatherMeasure_{it}$, and $Ab\_WeatherMeasure_{it}$ is the remainder, which represents unusual deviations from local weather conditions and, thus, is the focus of our analysis. Finally, we standardize these abnormal deviations, usually known as the standardized anomaly (Kim et al., 2021).

To capture how people react to abnormal local weather conditions, we use the daily Google Search Volume Index (SVI) of the topic “global warming” in a country, representing the retail investor attention. Considering that the attention to abnormal weather conditions also can be amplified through communication channels and the media, we use the news trends (NT) function of the Bloomberg terminal to collect the news counts, including the term “climate change” for each country. Based on a vast archive of news stories and social media posts from over 150,000 sources, the NT function makes it possible to search specific keywords and obtain the historical volume of relevant news. Given that institutional investors who work in asset management, banking, and institutional financial services primarily use Bloomberg terminals (Ben-Rephael et al., 2017), it is reasonable to assume that institutional investors who have more resources and incentives to pay attention to news quickly will follow the news that appears on Bloomberg terminal. Hence, we use this measure to proxy for institutional investors’ attention.

4 Empirical Results

4.1 Statistical Gains

In order to set the stage for our forecasting analysis, we lay out in Table 2 for recursive-estimation window.[5] We use a training period of one year (250 observations) to initialize the estimations. The table depicts the ratios of the root-mean-squared-forecasting errors (RMSFEs) calculated by dividing the RMSFE statistic for a benchmark model by the RMSFE statistic of the rival model. Whenever the ratio exceeds unity, the rival model outperforms the benchmark model in terms of the RMSFE

[5] We use the R language and environment for statistical computing (R Core Team 2021) to conduct our forecasting anaqlysis. We use the R add-on package “grf” (Tibshirani et al. 2021) to estimate random forests. A random forest features 1,000 individual random regression trees, where 10 is the minimum size of a terminal node, and floor(# number of predictors/3) predictors are used for random splitting. We construct the data matrix such that the number of forecasts is the same for all forecast horizons.

statistic. The benchmark and rival models are reported in the first column of the table, and we estimate both models by random forests. Using random forests for both models, even when the core HAR-RV model is the benchmark model, guarantees that variations in the reported RMSFE ratios reflect a variation in the vector of predictors that define the benchmark and rival models, not a variation in the estimation technique. In addition to the core HAR-RV model, we study the following three models:

1. The HAR-RV-M model extends the HAR-RV model to include the following "moments": realized skewness, realized kurtosis, realized good variance, realized bad variance, upside tail risk, downside tail risk, and jumps.
2. The HAR-RV-M-C model extends the HAR-RV-M model to include the following domestic and U.S. climate variables: Temperature, HDD, CDD, precipitation, wind speed.
3. The HAR-RV-M-CN model extends the HAR-RV-M model to include the climate-related Google searches and U.S. news.

The results we summarize in Table 2 show that, for all three forecast horizons under scrutiny, the overwhelming majority of the RMSFE ratios are larger than unity. Hence, the HAR-RV-M model performs better than the core HAR-RV model in terms of the RMSFE statistic, while the HAR-RV-M-C and HAR-RV-M-CN models further improve upon the HAR-RV-M model. The only exceptions arise in the case of the Russian ruble. In sum, the climate-risk variables go beyond the battery of widely-studied standard predictors like realized higher-order moments, jumps, and tail risks in improving the accuracy of forecasts of the subsequent exchange-rate volatility. In economic terms, a possible explanation for the predictive power of climate-related risks could rely on the notion that climate risks cause economic harm and uncertainty, which are quickly priced by economic agents in the foreign-exchange market, thereby, affecting the exchange-rate volatility.

– Please include Table 2 about here. –

In Table 3, we compare random forests with the ordinary-least-squared (OLS) estimator and the least absolute shrinkage and selection operator (Lasso) estimator (Tibshirani, 1996). The Lasso estimator is a model shrinkage and predictor-selection technique that is an attractive alternative to the standard OLS estimator given that our HAR-RV-M-C and HAR-RV-M-CN models feature a relatively large array of predictor variables. The Lasso estimator achieves model shrinkage and predictor selection by adding to the standard quadratic OLS loss function a penalty term that consists of a weighting parameter (that is, the shrinkage parameter) multiplied by the sum of the absolute values of the estimated coefficients (that is, an L1 penalty term). Depending on the value of the weighting parameter, the Lasso estimator shrinks the magnitude of the estimated coefficients of the forecasting model, or even sets to zero some of the estimated coefficients.[6] Importantly, the Lasso estimator resembles the OLS estimator insofar as it retains the assumption of a linear forecasting model. Hence, the results we report in Table 3 shed light on whether a nonlinear forecasting model, that is random forests, is superior relative to two popular linear forecasting models.

– Please include Table 3 about here. –

We use the RMSFE statistic to compare in Table 3 random forests with the OLS and the Lasso estimators, where we focus on the HAR-RV-M-C model. The results for the HAR-RV-M-CN model are similar and are not reported (but available upon request). The results for the RMSFE ratios clearly demonstrate that random forests outperform the two alternative linear estimators, where the additional forecasting gains from using random forests tend to increase for several of the studied currency pairs in the length of the forecast horizon.

In Table A1 at the end of the paper (Appendix), we compare random forests with two popular variants of the Lasso estimator: The Ridge regression estimator and an elastic net. The Ridge regression estimator replaces the L1 penalty terms of the Lasso with an L2 penalty term. The elastic net, in turn, is in our analysis a simple equally-weighted combination of the Ridge regression estimator and the Lasso. Again we find that random forests perform well relative to its two competitors, especially at the intermediate and long forecast horizons.

– Please include Figure 2 about here. –

Figure 2 illustrates why random forests tend to perform better than the estimators that retain the assumption of a linear forecasting model. Focusing on Google SVI of the topic "global warming" and four selected currency pairs, the figure presents partial-dependence functions that visualize how our dependent variable (that is, $\ln(\sqrt{RV})$) responds to a variation in the Google SVI..[7] While not all partial-dependence functions suggest that the link between realized exchange-rate volatility and the climate-risk predictors is nonlinear, the selected partial-dependence functions that we plot in Figure 2 make clear that such nonlinearities are present in the data. Random forests are tailored to detect such nonlinearities, and to exploit them to improve forecast accuracy.

– Please include Table 4 about here. –

Table 4 supplements our analysis so far by summarizing the results (p-values) of comparisons of forecast accuracy by means of formal statistical tests. Because comparing forecasts from different

[6]We use the R add-on package "glmnet" (Friedman et al., 2010) to implement the Lasso estimator, where we use 10-fold cross validation to select the weighting parameter that minimizes the mean cross-validated error.

[7]We use the full sample of data to estimate the partial dependence functions. We use the R add-on package randomForestSRC (Ishwaran and Kogalur 2021) to trace out the partial-dependence functions. Sampling is with replacement, the minimum node size is five, and one third of the predictors are used for splitting. Details are available from the authors upon request.

random forests is complicated due to their complex structure, we present results for both the Clark and West (2007) test and the Diebold and Mariano (1995) test, as modified by Harvey et al. (1997). The latter test should be considered to be approximate because we use a recursive-estimation window. Moreover, the interpretation of the test results should acknowledge that different random forests cannot be regarded as nested models in a strict sense. Notwithstanding, it is reassuring that the test results largely are in line with the results of the analysis of the RMSFE ratios. We observe that the tests yield highly significant results in the overwhelming majority of cases. Insignificant test results mainly obtain for the Russian rubel and for the long forecast horizon when we use the modified Diebold-Mariano test compare the forecasts from the HAR-RV model with the forecasts that we compute by means of the HAR-RV-M model.

In order to shed further light on the statistical significance of our results, we report in Table A2 the results of encompassing regressions (Fair and Shiller, 1990). For the encompassing regressions, we regress the actual realized variance on a constant and the forecasts derived from the benchmark and the rival model. The estimation results demonstrate that the coefficient we estimate for the rival model is significantly different from zero in the overwhelming majority of cases, whereas the coefficient that we estimate for the benchmark model is significantly different from zero occasionally only. Hence, we conclude that in the overwhelming majority of cases the rival model encompasses the predictive value of the benchmark model and, in addition, contains additional predictive value not already captured by the benchmark model.

– Please include Table 5 about here. –

As an alternative setting, we consider a fixed estimation window. To this end, we use the first half of the sample period for estimation of the models and the second half for forecasting. The RMSFE ratios we report in Table 5 corroborate that, with few exceptions, the climate-risk predictors help to improve the accuracy of forecasts relative to the HAR-RV-M benchmark model, which, in turn, produces better forecasts than the HAR-RV benchmark model.

The fixed estimation window lends itself to study slightly modified forecasting settings. First, we use the fixed estimation window to study the predictive value of the climate-risk predictors for the anti-log of exchange-rate volatility. Second, we consider exchange-rate volatility $t + h$ periods ahead rather than its average over the forecast horizon as our predictand. Results (not reported for the sake of brevity, but available upon request from the authors) are in line with the results we report in Table 5.

– Please include Table 6 about here. –

Table 6 reports results for a rolling-estimation window (of length 250 observations). The results of the rolling-estimation window are weaker than those for the recursive and the fixed estimation window. The RMSFE ratios exceed unity often by a small margin only, and mainly for the short forecast horizon, for the long forecast horiozn when we compare the HAR-RV-M and the HAR-RV-M-C models, and for the intermediate and long forecast horizons when we compare the HAR-RV-M and the HAR-RV-M-CN models. Hence, the climate-risk predictors also have predictive value in the rolling-estimation-window setting, but, from a forecasting perspective, the recursive and fixed estimation windows that we study in our forecasting analysis clearly appear to be the preferred schemes.

– Please include Table 7 about here. –

Finally, we present, for a recursive-estimation window, in Table 7 results for the mean-absolute-forecast-error (MAFE) ratios, which attaches a smaller weight to very large (positive or negative) forecast errors than the RMSFE statistic and, thereby, somewhat mitigates the impact of peaks in exchange-rate volatility and volatility clustering on forecast performance. The results corroborate the results for the RMSFE ratios insofar as the HAR-RV-M dominates the HAR-RV model (at the short and intermediate forecast horizons), while the HAR-RV-M-C and HAR-RV-M-CN models clearly outperform the HAR-RV-M model. Hence, we conclude that accounting for the climate-risk predictors helps to improve the accuracy of forecasts of exchange-rate volatility, a result that apparently is of paramount importance for international investors and risk managers.

4.2 Utility Gains

We next shed light on the economic benefits a forecast consumer obtains from using the the climate-risk predictors. To this end, we adapt the method proposed by West et al., (1993). We consider a forecast consumer who has a utility function of the format $U_{t+1} = W_{t+1} - \frac{1}{2}\gamma W_{t+1}^2$, where $\gamma > 0$ is a parameter. For simplicity, we focus on the case $h = 1$. The budget constraint is of the format $W_{t+1} = W_t\left[f_t(R^* + e_{t+1}) + (1 - f_t)R\right]$, where $0 < f_t < 1$ denotes a portfolio share (the choice variable of the forecast consumer), e_{t+1} denotes the exchange rate (assumed to follow a random work without drift), and R^* an R denote the foreign and domestic interest rates (assumed to be constant). We assume $\mu = R^* - R > 0$ for the interest-rate differential. Choosing the portfolio share to maximize expected utility, $E_tU_{t+1} = E_t\left[W_{t+1} - \frac{1}{2}\gamma W_{t+1}^2\right]$, where E_t denotes the conditional-expectations operator, and assuming non-satiation (that is, $1 - \gamma WR > 0$), yields an optimal portfolio share given by

$$f_{t+1} = \frac{\mu}{\gamma W}\left[1 - \gamma WR\right]\frac{1}{\mu^2 + \hat{RV}_{t+1}}, \tag{26}$$

where a hat denotes a forecast, $\hat{RV}_{t+1}$, of the realized exchange-rate variance. It should be noted that $\gamma W = \frac{\delta}{1+\delta}$, where δ is the coefficient of relative risk aversion (CRRA). Upon using the result for

f_{t+1} in the formula for expected utility, we obtain

$$E_t U_{t+1} = \left[c + d\, u(RV_{t+1}, \hat{RV}_{t+1})\right] W, \tag{27}$$

where $c = R - \frac{1}{2}\gamma W R^2$, and $d = \frac{\mu^2}{\gamma W}\left[1 - \gamma W R\right]^2$, and

$$u(RV_{t+1}, \hat{RV}_{t+1}) = \left[\frac{1}{\mu^2 + \hat{RV}_{t+1}} - \frac{1}{2}\frac{\mu^2 + RV_{t+1}}{(\mu^2 + \hat{RV}_{t+1})^2}\right]. \tag{28}$$

Next, we use the ex-post actual values of the realized exchange-rate variance as a proxy of the conditional expectation of RV_{t+1}, and approximate $E_t U_{t+1}$ by computing average utility, $\bar{U} = (c + d\bar{u})W$, across the out-of-sample forecasting periods for our forecasting models. We can then easily compare two models, let us denote them as Models 1 and 2, in terms of $\bar{U}_1$ and $\bar{U}_2$. Suppose Model 1 is the better of the two models. Then, in order to make it equivalent to Model 2 in terms of average utility, we apparently have to reduce the wealth associated with Model 1 by some amount, ΔW, such that $(c + d\bar{u}_1)(W - \Delta W) = (c + d\bar{u}_2)W$, which implies $\frac{\Delta W}{W} = 1 - \frac{\bar{U}_2}{\bar{U}_1}$. Like West et al. (1993), we express this ratio as an average "fee" converted to annual basis points by computing $252 \times (1 - \frac{\bar{U}_2}{\bar{U}_1}) \times 100 \times 100$, assuming that a year has 252 trading days, and where multiplication by the first 100 turns the expression into percentage, and the second 100 gives basis points. This fee summarizes how much a forecast consumer would be willing to pay to use the forecasts from the superior forecasting model rather than the sub-optimal model.

– Please include Table 8 about here. –

Table 8 summarizes the results for three numerical values of δ, where we use in this numerical example $R = 0.05$ and $\mu = 0.05$. Corroborating the results of the statistical forecast criteria, we find that a forecast consumer benefits from using the climate-risk predictors to compute forecasts of exchange-rate volatility in the majority of cases, with exceptions arising for few model configurations only for the exchange-rates of the dollar vis-à-vis the rubel and the ringgit. Overall, our findings indicate that better predictability of realized volatility can translate into noticeable utility gains for a risk-averse investor who would be ready to pay a portfolio-management fee to have access to the forecast computed by means of the model augmented to include climate-related risks.

5 Concluding Remarks

Forecasting exchange-rate volatility is of apparent importance for currency trading, optimal cross-border asset allocation, identifying and managing currency exposure, and pricing and trading of derivative securities like foreign-currency options. Because exchange-rate fluctuations can have substantial implications for trade and welfare, forecasting exchange-rate volatility is also a key concern for policymakers, especially when major commodity currencies exhibit large and sudden swings. It is, therefore, not surprising that a large and significant literature on modeling and forecasting exchange-rate volatility exists. Our contribution to this literature is that we argue that accounting for climate-related risks has the potential to improve the accuracy of forecasts of the realized volatility of major exchange rates vis-à-vis the dollar.

In our forecasting analysis, we analyze the role of climate risks in forecasting the intraday-data-based daily realized volatility of exchange-rate returns of eight major fossil fuel-exporters, namely Australia, Brazil, Canada, Malaysia, Mexico, Norway, Russia, and South Africa. We study a wide array of metrics capturing risks associated with climate change, derived from data directly on variables such as abnormal patterns of temperature, precipitation, number of heating degree days (HDD), number of cooling degree days (CDD), and wind speed, as well as Google search volume and media coverage on the topic. Because we also control for various other moments (realized skewness, realized kurtosis, realized good and variance, upside and downside tail risk, and jumps) in our forecasting models, we end up with many predictors and, hence, use a machine-learning approach, i.e., random forests, to study the predictive value of climate-related risks for the subsequent realized exchange-rate volatility. Our findings clearly show that climate-related risks do have incremental predictive value, both in statistical and economic terms, for exchange-rate volatility. Inspection of partial dependence plots suggests that the link between climate-related risks and subsequent exchange-rate volatility tends to be nonlinear, which explains why random forests work well in our forecasting analysis.

As part of future research, it is be interesting to extend our analyses to other asset markets like, equities and bonds, especially in a regional context, given that climate change, while typically interpreted as an aggregate risk factor, often has a significant local risk component.

References

Andersen, T.G., and Bollerslev, T. (1998). Answering the skeptics: yes, standard volatility models do provide accurate forecasts. International Economic Review, 39(4), 885–905.

Andersen T.G., Bollerslev T., Diebold F.X. (2007). Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility. Review of Economics and Statistics, 89(4), 701–720.

Andersen, T.G., Bollerslev, T., Diebold, F.X., and Labys, P. (2003). Modeling and Forecasting Realized Volatility. Econometrica, 71(2), 579–625.

Andersson, M., Bolton, P., and Samama, F. (2016). Hedging climate risk. Financial Analysts Journal, 72(3), 13-32.

Asteriou, D., Masatci, K., and Pilbeam, K. (2016). Exchange rate volatility and international trade: International evidence from the MINT countries. Economic Modelling, 58, 133–140.

Atanasova, C., and Schwartz, E. S. (2019). Stranded fossil fuel reserves and firm value (No. w26497). National Bureau of Economic Research.

Balcilar, M., Gupta, R., Kyei, C., and Wohar, M.E. (2016). Does Economic Policy Uncertainty Predict Exchange Rate Returns and Volatility? Evidence from a Nonparametric Causality-in-Quantiles Test. Open Economies Review, 27(2), 229–250.

Barndorff-Nielsen, O.E. and Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. Journal of Financial Econometrics, 2, 1–37.

Barndorff-Nielsen, O.E. and Shephard, N. (2006). Econometrics of testing for jumps in financial economics using bipower variation. Journal of Financial Econometrics, 4, 1–30.

Barndorff-Nielsen, O.E., Kinnebrouk, S., and Shephard, N. (2010). Measuring downside risk: realised semivariance. Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle, Edited by T. Bollerslev, J. Russell, and M. Watson, 117-136. Oxford University Press.

Barro, R.J. (2006). Rare Disasters and Asset Markets in the Twentieth Century. Quarterly Journal of Economics, 121, 823–866.

Barunik, J., Krehlik, T., and Vacha, L. (2016). Modeling and forecasting exchange rate volatility in time-frequency domain. European Journal of Operational Research, 251(1), 329–340.

Battiston, S., Dafermos, Y., and Monasterolo, I. (2021). Climate risks and financial stability. Journal of Financial Stability, 54, 100867.

Ben-Rephael, A., Da, Z., and Israelsen, R. D. (2017). It depends on where you search: Institutional investor attention and underreaction to news. Review of Financial Studies, 30(9), 3009−3047.

Breiman, L. (2001). Random forests. Machine Learning, 45, 5−32.

Campbell, J.Y., (2008) Viewpoint: estimating the equity premium, Canadian Journal of Economics, 41, 1−21.

Chen, Y.-C., Rogoff, K.S., and Rossi, B. (2010). Can exchange rate forecast commodity prices? Quarterly Journal of Economics, 125(3), 1145−1194.

Choi, D., Gao, Z., and Jiang, W. (2020). Attention to global warming. Review of Financial Studies, 33(3), 1112−1145.

Clark, P.B., Wei, S-J., Tamirisa, N.T., Sadikov, A.M., and Zeng, L. (2004). A new look at exchange rate volatility and trade Flows. International Monetary Fund, Washington, D.C.

Clark, T.D., and West, K.D. (2007). Approximately normal tests for equal predictive accuracy in nested Benchmark/rival model. Journal of Econometrics, 138, 291−311.

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

Christou, C., Gupta, R., Hassapis, C., and Suleman, M.T. (2018). The role of economic uncertainty in forecasting exchange-rate returns and realized volatility: Evidence from quantile predictive regressions. Journal of Forecasting, 37(7), 705−719.

Colacito, R., Hoffmann, B., and Phan, T. (2019). Temperature and growth: A panel analysis of the United States. Journal of Money, Credit and Banking, 51(2-3), 313-368.

Diebold, F.X., and Mariano, R.S. (1995). Comparing predictive accuracy. Journal of Business and Economic Statistics, 13, 253−263.

Duan, T., Li, F. W., and Wen, Q. (2021). Is Carbon Risk Priced in the Cross-Section of Corporate Bond Returns?. Available at SSRN 3709572.

Engle, R.F., Giglio, S., Kelly, B., Lee, H., and Stroebel, J. (2020). Hedging climate change news. Review of Financial Studies, 33(3), 1184−1216.

Faccini, R., Matin, R., and Skiadopoulos, G. (2021). Dissecting climate risks: Are they reflected in stock prices? Available at SSRN: https://ssrn.com/abstract=3795964.

Fair, R.C. and Shiller, R. J. (1990). Comparing information in forecasts from econometric models. American Economic Review, 80(3), 375−389.

Farhi, E., and Gabaix, X. (2016). Rare disasters and exchange rates. Quarterly Journal of Economics, 131(1), 1−52.

Ferraro, D., Rogoff, K. and Rossi, B. (2015). Can oil prices forecast exchange rates? An empirical analysis of the relationship between commodity prices and exchange rates, Journal of International Money and Finance 54, 116−141.

Flori, A., Pammolli, F., and Spelta, A. (2021). Commodity prices co-movements and financial stability: a multidimensional visibility nexus with climate conditions. Journal of Financial Stability, 54, 100876.

Friedman, J. Hastie, T., and Tibshirani, R. (2010). Regularization paths for Generalized Linear Models via coordinate descent. Journal of Statistical Software, 33(1), 1−22.

Giglio, S., Kelly, B., and Stroebel, J. (2021). Climate finance. Annual Review of Financial Economics, 13, 15−36.

Gkillas, K., Gupta, R., and Pierdzioch, C. (2019). Forecasting (downside and upside) realized exchange-rate volatility: Is there a role for realized skewness and kurtosis? Physica A: Physica A: Statistical Mechanics and its Applications, 532(2), 121867.

Gupta, R., Suleman, M.T., and Wohar, M.E. (2019). Exchange rate returns and volatility: The role of time-varying rare disaster risks. European Journal of Finance, 25(2), 190−203.

Hansen, P.R., and Lunde, A. (2005). A forecast comparison of volatility models: does anything beat a GARCH (1, 1)? Journal of Applied Econometrics, 20(7), 873−889.

Harvey, D., Leybourne, S., and Newbold, P. (1997). Testing the equality of prediction mean squared errors. International Journal of Forecasting, 13(2), 281−291.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). The elements of statistical learning: Data mining, inference, and prediction. 2nd ed., Springer: New York, NY, USA.

Henseler, M., and Schumacher, I. (2019). The impact of weather on economic growth and its production factors. Climatic Change, 154(3), 417-433.

Hill, B. (1975). A simple general approach to inference about the tail of a distribution. Annals of Statistics, 3, 1163−1173.

Ishwaran, H., and Kogalur, U.B. (2021). Fast unified random forests for survival, regression, and classification (RF-SRC), R package version 2.12.1.

Kapfhammer, F., Larsen, V.H., and Thorsrud, L.A. (2020). Climate risk and commodity currencies. Centre for Applied Macro and Petroleum economics (CAMP), BI Norwegian Business School, Working Papers No. 10/2020.

Kim, H.S., Matthes, C., and Phan, T. (2021). Extreme weather and the macroeconomy. Federal reserve Bank of Richmond, Working Paper No. 21−14.

Liu, R., Demirer, R., Gupta, R., and Wohar, M.E. (2020). Volatility forecasting with bivariate multifractal models. Journal of Forecasting, 39(2), 155−167.

Louzis, D.P., Xanthopoulos-Sisinis, S., and Refenes, A.P. (2013). The role of high-frequency intra-daily data, daily range and implied volatility in multi-period value-at-risk forecasting. Journal of Forecasting, 32(6), 561−576.

McAleer, M., and Medeiros, M.C. (2008). Realized volatility: A review. Econometric Reviews 27, 10−45.

Mcmillan, D.G., and Speight, A.E. (2012). Daily FX volatility forecasts: can the GARCH (1, 1) model be beaten using high-frequency data? Journal of Forecasting, 31(4), 330−343.

Mei, D., Liu, J., Ma, F., and Chen, W. (2017). Forecasting stock market volatility: Do realized skewness and kurtosis help?, Physica A: Statistical Mechanics and its Applications, 481, 153−159.

Müller, U.A., M.M. Dacorogna, Davé, R.D., Olsen, R.B., and Pictet, O.V. (1997). Volatilities of different time resolutions − aAnalyzing the dynamics of market components. Journal of Empirical Finance, 4(2−3), 213−239.

Park, C., and Park, S. (2020). Rare disaster risk and exchange rates: An empirical investigation of South Korean exchange rates under tension between the two Koreas. Finance Research letters, 36, 101314.

Plíhal, T., and Lyócsa, Š. (2021). Modeling realized volatility of the EUR/USD exchange rate: Does implied volatility really matter? International Review of Economics & Finance, 71, 811−829.

Pong, S., Shackleton, M.B., Taylor, S.J., and Xu, X. (2004). Forecasting currency volatility: A comparison of implied volatilities and AR(FI)MA models. Journal of Banking & Finance, 28(10), 2541−2563.

R Core Team (2021). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL: https://www.R-project.org/.

Rapach, D.E., and Strauss, J.K. (2008). Structural breaks and GARCH models of exchange rate volatility. Journal of Applied Econometrics, 23(1), 65−90.

Rapach, D.E., and Zhou, G. (2013). Forecasting stock returns. Handbook of Economic Forecasting, 2 (Part A), Graham Elliott and Allan Timmermann (Eds.), Amsterdam: Elsevier, 328−383.

Rapach, D.E., and Zhou, G. (Forthcoming). Asset pricing: Time-series predictability. The Oxford Research Encyclopedia of Economics and Finance.

Rietz, T. (1988). The equity risk premium: A solution. Journal of Monetary Economics, 22, 117−131.

Senadza, B., and Diaba, D.D. (2017). Effect of exchange rate volatility on trade in Sub-Saharan Africa? Journal of African Trade, 4(1-2), 20−36.

Stroebel, J., and Wurgler, J. (2021). What do you think about climate finance? Journal of Financial Economics, 142, 487−498.

Tibshirani, R. (1996) Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society, Series B, 58, 267−288.

West, K.D., Edison, H.J, and Cho, D. (1993). A utility-based comparison of some models of exchange rate volatility. Journal of International Economics,35, 23−45.

Zhou, H., Yu, M., Li, J., and Qin, Q. (2021). Rare disasters, exchange rates, and macroeconomic policy: Evidence from COVID-19? Economics Letters, 209, 110099.

Figure 1: Realized variance ($\ln(\sqrt{RV})$)

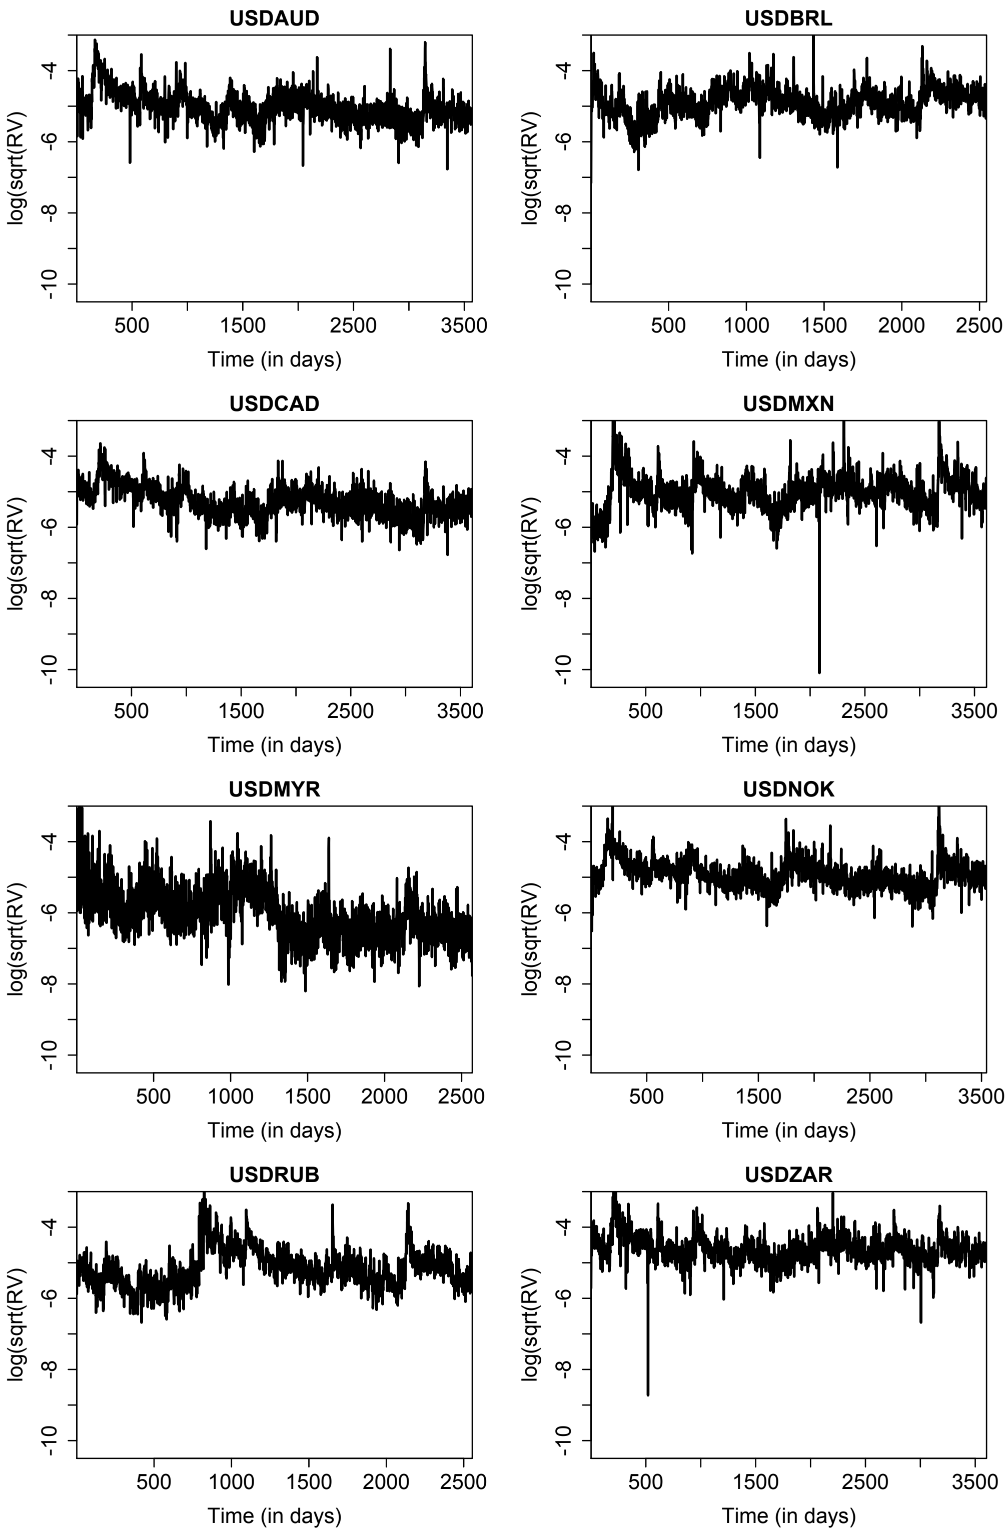

Figure 2: Selected partial dependence functions

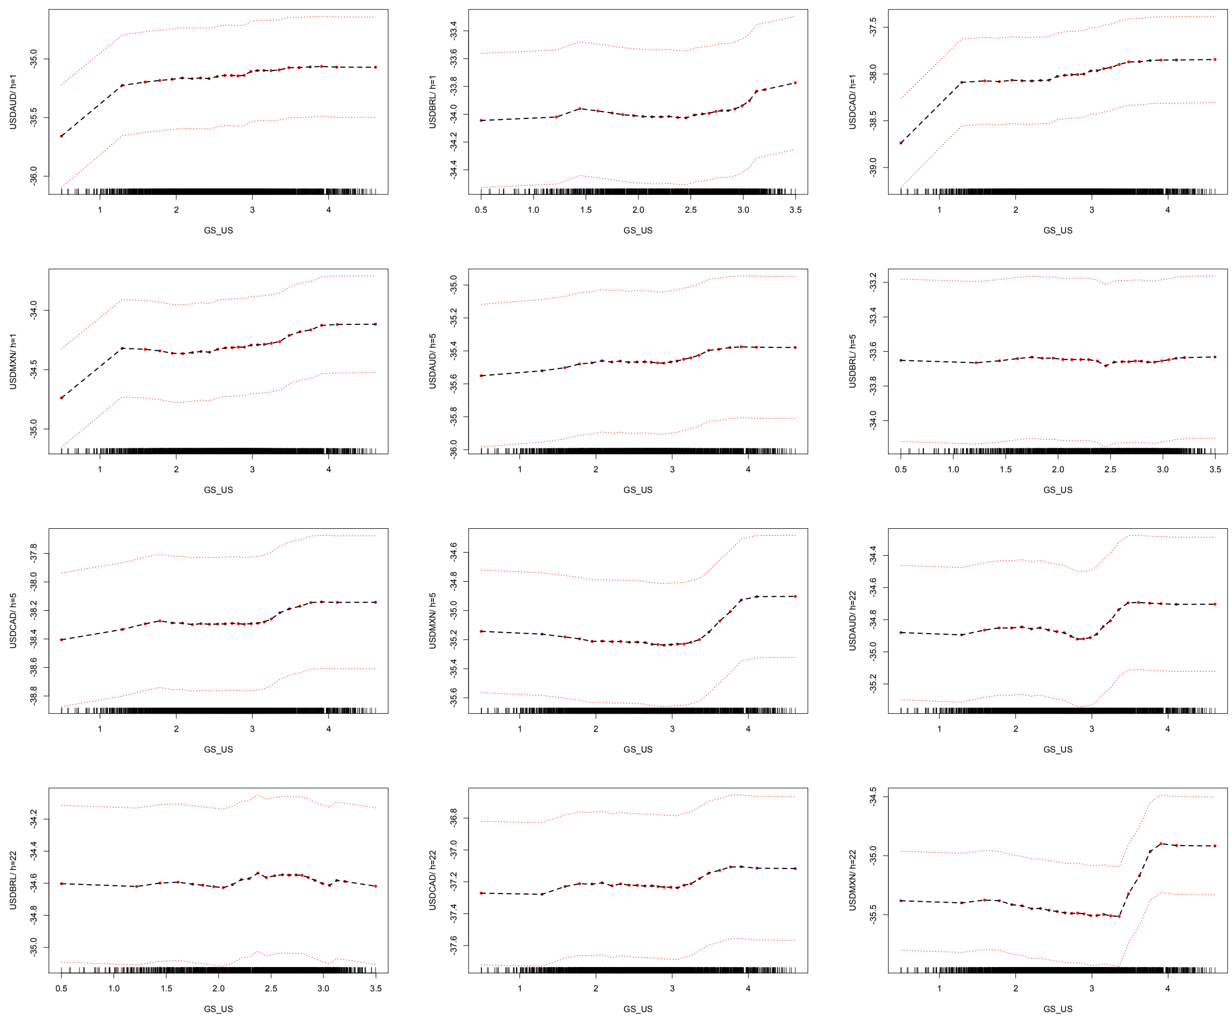

Red points/black dashed lines: partial values. Dashed red lines: smoothed error band (plus/minus two standard errors).

Table 1: Summary statistics of realized exchange-rate volatility

Exchange rate	Obs.	Start	End	Min.	1st Qua.	Median	Mean	3rd Qua.	Max.	AC(1)	AC(2)	AC(3)	AC(4)	AC(5)
USDAUD	3572	01/07/2008	11/8/2021	-6.7689	-5.3357	-5.0798	-5.0485	-4.8120	-3.1343	0.8041	0.7337	0.7126	0.7308	0.7350
USDBRL	2545	08/29/2011	11/8/2021	-7.1565	-5.1610	-4.8871	-4.9029	-4.6277	-2.5910	0.7023	0.6441	0.6083	0.6068	0.6002
USDCAD	3608	01/02/2008	11/8/2021	-6.7727	-5.5847	-5.3219	-5.2912	-5.0165	-3.6415	0.7849	0.7489	0.7337	0.7307	0.7466
USDMXN	3610	01/02/2008	11/8/2021	-10.0934	-5.2692	-5.0055	-5.0057	-4.7472	-2.3062	0.7891	0.7411	0.7050	0.6952	0.6888
USDMYR	2569	09/1/2011	11/8/2021	-8.2009	-6.4848	-6.0323	-5.9817	-5.4795	-2.2344	0.6185	0.5483	0.5342	0.5708	0.6819
USDNOK	3545	03/31/2008	11/8/2021	-6.5090	-5.1844	-4.9527	-4.9223	-4.6995	-2.4357	0.7779	0.7139	0.6835	0.6980	0.6988
USDRUB	2555	08/31/2011	11/8/2021	-6.6790	-5.5390	-5.2224	-5.1745	-4.8665	-2.2180	0.8537	0.8006	0.7613	0.7437	0.7333
USDZAR	3603	01/3/2008	11/8/2021	-8.7250	-4.8394	-4.6490	-4.6185	-4.4263	-2.6363	0.7021	0.6405	0.6060	0.5983	0.5998

Summary statistics are for $\ln(\sqrt{RV})$. Obs= number of observations. AC(x)= Coefficient of autocorrelation at lag x.

Table 2: RMSFE ratios

Exchange rate / benchmark vs. rival model	h=1	h=5	h=22
USDAUD / HAR-RV vs. HAR-RV-M	1.0333	1.0170	1.0311
USDBRL / HAR-RV vs. HAR-RV-M	1.0224	1.0245	1.0188
USDCAD / HAR-RV vs. HAR-RV-M	1.0310	1.0229	1.0016
USDMXN / HAR-RV vs. HAR-RV-M	1.0200	1.0152	1.0142
USDMYR / HAR-RV vs. HAR-RV-M	1.0137	1.0222	1.0097
USDNOK / HAR-RV vs. HAR-RV-M	1.0175	1.0092	1.0057
USDRUB / HAR-RV vs. HAR-RV-M	1.0107	-	-
USDZAR / HAR-RV vs. HAR-RV-M	1.0272	1.0195	1.0133
USDAUD / HAR-RV-M vs. HAR-RV-M-C	1.0337	1.0240	1.0614
USDBRL / HAR-RV-M vs. HAR-RV-M-C	1.0202	1.0312	1.0358
USDCAD / HAR-RV-M vs. HAR-RV-M-C	1.0296	1.0304	1.0270
USDMXN / HAR-RV-M vs. HAR-RV-M-C	1.0217	1.0291	1.0347
USDMYR / HAR-RV-M vs. HAR-RV-M-C	1.0108	1.0338	1.0329
USDNOK / HAR-RV-M vs. HAR-RV-M-C	1.0170	1.0164	1.0129
USDRUB / HAR-RV-M vs. HAR-RV-M-C	1.0051	-	-
USDZAR / HAR-RV-M vs. HAR-RV-M-C	1.0226	1.0224	1.0322
USDAUD / HAR-RV-M vs. HAR-RV-M-CN	1.0300	1.0263	1.1384
USDBRL / HAR-RV-M vs. HAR-RV-M-CN	1.0237	1.0316	1.0519
USDCAD / HAR-RV-M vs. HAR-RV-M-CN	1.0317	1.0463	1.0832
USDMXN / HAR-RV-M vs. HAR-RV-M-CN	1.0240	1.0275	1.0534
USDMYR / HAR-RV-M vs. HAR-RV-M-CN	1.0114	1.0303	1.0347
USDNOK / HAR-RV-M vs. HAR-RV-M-CN	1.0230	1.0241	1.0418
USDRUB / HAR-RV-M vs. HAR-RV-M-CN	1.0149	-	-
USDZAR / HAR-RV-M vs. HAR-RV-M-CN	1.0238	1.0360	1.0809

Both the benchmark and the rival model are estimated by random forests. For both models the root-mean-squared-forecasting error (RMSFE) is computed. The RMSFE ratio is constructed by dividing the RMSFE of the benchmark model by the RMSFE of the rival model. A RMSFE ratio larger than unity indicates that the rival model yields more accurate forecasts than the benchmark model. Results are for a recursive estimation window. Training period = 250 observations. For better readability of the estimation results, the table depicts the RMSFE ratios that exceed unity only. The parameter h denotes the forecast horizon. HAR-RV-M = The HAR-RV model is extended to include the following "moments": realized skewness, realized kurtosis, realized good variance, realized bad variance, upside tail risk, downside tail risk, jumps. HAR-RV-M-C = The HAR-RV-M model is extended to include the following domestic and U.S. climate variables: Temperature, HDD, CDD, precipitation, wind. HAR-RV-M-CM = The HAR-RV-M model is extended to include the climate-related Google searches and U.S. news.

Table 3: Comparing random forests with competing methods (RMSFE ratios)

Exchange rate / benchmark vs. rival model	h=1	h=5	h=22
USDAUD / HAR-RV-M-C-ols vs. HAR-RV-M-C-RF	-	1.0258	1.0324
USDBRL / HAR-RV-M-C-ols vs. HAR-RV-M-C-RF	1.3290	1.0148	1.0210
USDCAD / HAR-RV-M-C-ols vs. HAR-RV-M-C-RF	-	1.0108	1.0415
USDMXN / HAR-RV-M-C-ols vs. HAR-RV-M-C-RF	1.0444	1.0292	1.0489
USDMYR / HAR-RV-M-C-ols vs. HAR-RV-M-C-RF	1.0006	1.0214	1.0629
USDNOK / HAR-RV-M-C-ols vs. HAR-RV-M-C-RF	-	1.0299	1.0651
USDRUB / HAR-RV-M-C-ols vs. HAR-RV-M-C-RF	1.1024	1.5260	4.6521
USDZAR / HAR-RV-M-C-ols vs. HAR-RV-M-C-RF	1.0211	1.0218	1.0259
USDAUD / HAR-RV-M-C-Lasso vs. HAR-RV-M-C-RF	1.0375	1.0706	1.0666
USDBRL / HAR-RV-M-C-Lasso vs. HAR-RV-M-C-RF	-	1.0074	1.0247
USDCAD / HAR-RV-M-C-Lasso vs. HAR-RV-M-C-RF	-	1.0030	1.0308
USDMXN / HAR-RV-M-C-Lasso vs. HAR-RV-M-C-RF	-	1.0150	1.0182
USDMYR / HAR-RV-M-C-Lasso vs. HAR-RV-M-C-RF	1.0018	1.0201	1.0604
USDNOK / HAR-RV-M-C-Lasso vs. HAR-RV-M-C-RF	-	1.0620	1.0637
USDRUB / HAR-RV-M-C-Lasso vs. HAR-RV-M-C-RF	1.4198	2.1212	4.9172
USDZAR / HAR-RV-M-C-Lasso vs. HAR-RV-M-C-RF	1.0163	1.0207	1.0216

Both the benchmark and the rival model include the following domestic and U.S. climate variables: Temperature, HDD, CDD, precipitation, wind. The benchmark model is estimated either by the OLS technique or the Lasso estimator. The rival model is estimated by random forests. For both models the root-mean-squared-forecasting error (RMSFE) is computed. The RMSFE ratio is constructed by dividing the RMSFE of the benchmark model by the RMSFE of the rival model. A RMSFE ratio larger than unity indicates that the rival model yields more accurate forecasts than the benchmark model. Results are for a recursive estimation window. Training period = 250 observations. For better readability of the estimation results, the table depicts the RMSFE ratios that exceed unity only. The parameter h denotes the forecast horizon.

Table 4: Formal tests of forecast accuracy

Exchange rate / benchmark vs. rival model	h=1	h=5	h=22
Clark-West test			
USDAUD / HAR-RV vs. HAR-RV-M	0.0000	0.0000	0.0000
USDBRL / HAR-RV vs. HAR-RV-M	0.0000	0.0000	0.0000
USDCAD / HAR-RV vs. HAR-RV-M	0.0000	0.0000	0.0001
USDMXN / HAR-RV vs. HAR-RV-M	0.0000	0.0001	0.0025
USDMYR / HAR-RV vs. HAR-RV-M	0.0000	0.0000	0.0003
USDNOK / HAR-RV vs. HAR-RV-M	0.0000	0.0011	0.0000
USDRUB / HAR-RV vs. HAR-RV-M	0.0008	-	-
USDZAR / HAR-RV vs. HAR-RV-M	0.0000	0.0000	0.0000
USDAUD / HAR-RV-M vs. HAR-RV-M-C	0.0000	0.0000	0.0000
USDBRL / HAR-RV-M vs. HAR-RV-M-C	0.0000	0.0000	0.0000
USDCAD / HAR-RV-M vs. HAR-RV-M-C	0.0000	0.0000	0.0000
USDMXN / HAR-RV-M vs. HAR-RV-M-C	0.0000	0.0000	0.0003
USDMYR / HAR-RV-M vs. HAR-RV-M-C	0.0000	0.0000	0.0000
USDNOK / HAR-RV-M vs. HAR-RV-M-C	0.0000	0.0008	0.0000
USDRUB / HAR-RV-M vs. HAR-RV-M-C	0.0186	-	0.0603
USDZAR / HAR-RV-M vs. HAR-RV-M-C	0.0000	0.0000	0.0000
USDAUD / HAR-RV-M vs. HAR-RV-M-CN	0.0000	0.0004	0.0000
USDBRL / HAR-RV-M vs. HAR-RV-M-CN	0.0000	0.0000	0.0000
USDCAD / HAR-RV-M vs. HAR-RV-M-CN	0.0000	0.0000	0.0000
USDMXN / HAR-RV-M vs. HAR-RV-M-CN	0.0000	0.0000	0.0000
USDMYR / HAR-RV-M vs. HAR-RV-M-CN	0.0000	0.0000	0.0000
USDNOK / HAR-RV-M vs. HAR-RV-M-CN	0.0000	0.0000	0.0000
USDRUB / HAR-RV-M vs. HAR-RV-M-CN	0.0011	-	-
USDZAR / HAR-RV-M vs. HAR-RV-M-CN	0.0000	0.0000	0.0000
Diebold-Mariano test			
USDAUD / HAR-RV vs. HAR-RV-M	0.0000	0.0416	0.0255
USDBRL / HAR-RV vs. HAR-RV-M	0.0001	0.0018	0.0327
USDCAD / HAR-RV vs. HAR-RV-M	0.0000	0.0045	-
USDMXN / HAR-RV vs. HAR-RV-M	0.0034	-	-
USDMYR / HAR-RV vs. HAR-RV-M	0.0116	0.0001	-
USDNOK / HAR-RV vs. HAR-RV-M	0.0093	0.0686	-
USDRUB / HAR-RV vs. HAR-RV-M	0.0743	-	-
USDZAR / HAR-RV vs. HAR-RV-M	0.0000	0.0051	0.0531
USDAUD / HAR-RV-M vs. HAR-RV-M-C	0.0000	0.0346	0.0203
USDBRL / HAR-RV-M vs. HAR-RV-M-C	0.0001	0.008	0.0023
USDCAD / HAR-RV-M vs. HAR-RV-M-C	0.0000	0.0004	0.0324
USDMXN / HAR-RV-M vs. HAR-RV-M-C	0.0011	0.0171	0.0448
USDMYR / HAR-RV-M vs. HAR-RV-M-C	0.0371	0.0001	0.0741
USDNOK / HAR-RV-M vs. HAR-RV-M-C	0.0001	0.0741	0.0109
USDRUB / HAR-RV-M vs. HAR-RV-M-C	-	-	-
USDZAR / HAR-RV-M vs. HAR-RV-M-C	0.0001	0.0021	0.0025
USDAUD / HAR-RV-M vs. HAR-RV-M-CN	0.0000	-	0.0039
USDBRL / HAR-RV-M vs. HAR-RV-M-CN	0.0000	0.0066	0.0026
USDCAD / HAR-RV-M vs. HAR-RV-M-CN	0.0000	0.0000	0.0015
USDMXN / HAR-RV-M vs. HAR-RV-M-CN	0.0004	0.0129	0.0823
USDMYR / HAR-RV-M vs. HAR-RV-M-CN	0.0369	0.0002	0.0804
USDNOK / HAR-RV-M vs. HAR-RV-M-CN	0.0001	0.0292	0.0001
USDRUB / HAR-RV-M vs. HAR-RV-M-CN	0.0422	-	-
USDZAR / HAR-RV-M vs. HAR-RV-M-CN	0.0000	0.0002	0.0000

Both the benchmark and the rival model are estimated by random forests. For both models Clark-West and modified Diebold-Mariano tests are computed (p-values; robust standard errors). Results are for a recursive estimation window. Training period = 250 observations. For better readability of the estimation results, the table depicts the p-values that are smaller than 10% only. The parameter h denotes the forecast horizon. HAR-RV-M = The HAR-RV model is extended to include the following “moments”: realized skewness, realized kurtosis, realized good variance, realized bad variance, upside tail risk, downside tail risk, jumps. HAR-RV-M-C = The HAR-RV-M model is extended to include the following domestic and U.S. climate variables: Temperature, HDD, CDD, precipitation, wind. HAR-RV-M-CM = The HAR-RV-M model is extended to include the climate-related Google searches and U.S. news.

Table 5: RMSFE ratios for a fixed estimation window

Exchange rate / benchmark vs. rival model	h=1	h=5	h=22
USDAUD / HAR-RV vs. HAR-RV-M	1.0283	1.0431	1.0640
USDBRL / HAR-RV vs. HAR-RV-M	1.0291	1.0502	1.0691
USDCAD / HAR-RV vs. HAR-RV-M	1.0328	1.0602	1.0542
USDMXN / HAR-RV vs. HAR-RV-M	1.0151	1.0357	1.0129
USDMYR / HAR-RV vs. HAR-RV-M	1.0777	1.0900	1.1432
USDNOK / HAR-RV vs. HAR-RV-M	1.0363	1.0352	1.0467
USDRUB / HAR-RV vs. HAR-RV-M	1.0365	1.0265	-
USDZAR / HAR-RV vs. HAR-RV-M	1.0323	1.0443	1.0296
USDAUD / HAR-RV-M vs. HAR-RV-M-C	1.0289	1.032	1.0907
USDBRL / HAR-RV-M vs. HAR-RV-M-C	1.0284	1.0548	1.0604
USDCAD / HAR-RV-M vs. HAR-RV-M-C	1.0275	1.0498	1.0557
USDMXN / HAR-RV-M vs. HAR-RV-M-C	1.0222	1.0434	1.0028
USDMYR / HAR-RV-M vs. HAR-RV-M-C	1.0352	-	1.0485
USDNOK / HAR-RV-M vs. HAR-RV-M-C	1.0257	1.0344	1.0421
USDRUB / HAR-RV-M vs. HAR-RV-M-C	1.0354	1.0125	-
USDZAR / HAR-RV-M vs. HAR-RV-M-C	1.0297	1.0406	1.0179
USDAUD / HAR-RV-M vs. HAR-RV-M-CN	1.0209	1.0192	1.1703
USDBRL / HAR-RV-M vs. HAR-RV-M-CN	1.0268	1.0331	1.0431
USDCAD / HAR-RV-M vs. HAR-RV-M-CN	1.0198	1.0498	1.0601
USDMXN / HAR-RV-M vs. HAR-RV-M-CN	1.0205	1.0142	-
USDMYR / HAR-RV-M vs. HAR-RV-M-CN	1.0722	1.100	1.0880
USDNOK / HAR-RV-M vs. HAR-RV-M-CN	1.0175	1.0147	1.0465
USDRUB / HAR-RV-M vs. HAR-RV-M-CN	1.0417	1.0200	-
USDZAR / HAR-RV-M vs. HAR-RV-M-CN	1.0239	1.0321	-

Both the benchmark and the rival model are estimated by random forests. For both models the root-mean-squared-forecasting error (RMSFE) is computed. The RMSFE ratio is constructed by dividing the RMSFE of the benchmark model by the RMSFE of the rival model. A RMSFE ratio larger than unity indicates that the rival model yields more accurate forecasts than the benchmark model. Results are for a fixed estimation window, where 50% of the sample are used for estimating the models. For better readability of the estimation results, the table depicts the RMSFE ratios that exceed unity only. The parameter h denotes the forecast horizon. HAR-RV-M = The HAR-RV model is extended to include the following “moments”: realized skewness, realized kurtosis, realized good variance, realized bad variance, upside tail risk, downside tail risk, jumps. HAR-RV-M-C = The HAR-RV-M model is extended to include the following domestic and U.S. climate variables: Temperature, HDD, CDD, precipitation, wind. HAR-RV-M-CM = The HAR-RV-M model is extended to include the climate-related Google searches and U.S. news.

Table 6: RMSFE ratios for a rolling estimation window

Exchange rate / benchmark vs. rival model	h=1	h=5	h=22
USDAUD / HAR-RV vs. HAR-RV-M	1.0187	-	-
USDBRL / HAR-RV vs. HAR-RV-M	1.0131	-	-
USDCAD / HAR-RV vs. HAR-RV-M	1.0219	-	-
USDMXN / HAR-RV vs. HAR-RV-M	1.0131	-	-
USDMYR / HAR-RV vs. HAR-RV-M	1.0047	-	-
USDNOK / HAR-RV vs. HAR-RV-M	-	-	-
USDRUB / HAR-RV vs. HAR-RV-M	1.0059	-	-
USDZAR / HAR-RV vs. HAR-RV-M	1.0079	-	1.0046
USDAUD / HAR-RV-M vs. HAR-RV-M-C	1.0118	-	-
USDBRL / HAR-RV-M vs. HAR-RV-M-C	1.0047	-	-
USDCAD / HAR-RV-M vs. HAR-RV-M-C	1.0157	-	1.0121
USDMXN / HAR-RV-M vs. HAR-RV-M-C	1.0047	-	-
USDMYR / HAR-RV-M vs. HAR-RV-M-C	-	-	1.0039
USDNOK / HAR-RV-M vs. HAR-RV-M-C	-	-	-
USDRUB / HAR-RV-M vs. HAR-RV-M-C	1.0063	-	1.0046
USDZAR / HAR-RV-M vs. HAR-RV-M-C	1.0048	-	1.0462
USDAUD / HAR-RV-M vs. HAR-RV-M-CN	1.0193	1.0014	1.0065
USDBRL / HAR-RV-M vs. HAR-RV-M-CN	1.0133	1.0066	1.0262
USDCAD / HAR-RV-M vs. HAR-RV-M-CN	1.0244	1.0045	1.0237
USDMXN / HAR-RV-M vs. HAR-RV-M-CN	1.0191	-	-
USDMYR / HAR-RV-M vs. HAR-RV-M-CN	1.0030	1.0037	1.0052
USDNOK / HAR-RV-M vs. HAR-RV-M-CN	-	-	-
USDRUB / HAR-RV-M vs. HAR-RV-M-CN	1.0074	-	-
USDZAR / HAR-RV-M vs. HAR-RV-M-CN	1.0090	1.0110	1.0544

Both the benchmark and the rival model are estimated by random forests. For both models the root-mean-squared-forecasting error (RMSFE) is computed. The RMSFE ratio is constructed by dividing the RMSFE of the benchmark model by the RMSFE of the rival model. A RMSFE ratio larger than unity indicates that the rival model yields more accurate forecasts than the benchmark model. Results are for a rolling-estimation window. Length of the rolling-estimation window = 250 observations. For better readability of the estimation results, the table depicts the RMSFE ratios that exceed unity only. The parameter h denotes the forecast horizon. HAR-RV-M = The HAR-RV model is extended to include the following “moments”: realized skewness, realized kurtosis, realized good variance, realized bad variance, upside tail risk, downside tail risk, jumps. HAR-RV-M-C = The HAR-RV-M model is extended to include the following domestic and U.S. climate variables: Temperature, HDD, CDD, precipitation, wind. HAR-RV-M-CM = The HAR-RV-M model is extended to include the climate-related Google searches and U.S. news.

Table 7: MAFE ratios

Exchange rate / benchmark vs. rival model	h=1	h=5	h=22
USDAUD / HAR-RV vs. HAR-RV-M	1.0413	1.0140	1.0191
USDBRL / HAR-RV vs. HAR-RV-M	1.0351	1.0172	1.0063
USDCAD / HAR-RV vs. HAR-RV-M	1.0293	1.0127	-
USDMXN / HAR-RV vs. HAR-RV-M	1.0391	1.0142	-
USDMYR / HAR-RV vs. HAR-RV-M	1.0265	1.0179	1.0119
USDNOK / HAR-RV vs. HAR-RV-M	1.0313	1.0030	-
USDRUB / HAR-RV vs. HAR-RV-M	1.0239	1.0013	-
USDZAR / HAR-RV vs. HAR-RV-M	1.0337	1.0239	1.0063
USDAUD / HAR-RV-M vs. HAR-RV-M-C	1.0407	1.0158	1.0341
USDBRL / HAR-RV-M vs. HAR-RV-M-C	1.0327	1.0273	1.0289
USDCAD / HAR-RV-M vs. HAR-RV-M-C	1.0269	1.0228	1.0215
USDMXN / HAR-RV-M vs. HAR-RV-M-C	1.0406	1.0309	1.0187
USDMYR / HAR-RV-M vs. HAR-RV-M-C	1.0272	1.0342	1.0348
USDNOK / HAR-RV-M vs. HAR-RV-M-C	1.0284	1.0092	1.0225
USDRUB / HAR-RV-M vs. HAR-RV-M-C	1.0186	1.0195	1.0344
USDZAR / HAR-RV-M vs. HAR-RV-M-C	1.0306	1.0307	1.0285
USDAUD / HAR-RV-M vs. HAR-RV-M-CN	1.0419	1.0432	1.1339
USDBRL / HAR-RV-M vs. HAR-RV-M-CN	1.0335	1.0224	1.0376
USDCAD / HAR-RV-M vs. HAR-RV-M-CN	1.0364	1.0505	1.0960
USDMXN / HAR-RV-M vs. HAR-RV-M-CN	1.0513	1.0480	1.0798
USDMYR / HAR-RV-M vs. HAR-RV-M-CN	1.0287	1.0356	1.0433
USDNOK / HAR-RV-M vs. HAR-RV-M-CN	1.0314	1.0303	1.0887
USDRUB / HAR-RV-M vs. HAR-RV-M-CN	1.0209	1.0125	1.0101
USDZAR / HAR-RV-M vs. HAR-RV-M-CN	1.0413	1.0582	1.0972

Both the benchmark and the rival model are estimated by random forests. For both models the mean-absolute-forecasting error (MAFE) is computed. The MAFE ratio is constructed by dividing the MAFE of the benchmark model by the MAFE of the rival model. A MAFE ratio larger than unity indicates that the rival model yields more accurate forecasts than the benchmark model. Results are for a recursive-estimation window. Training period = 250 observations. For better readability of the estimation results, the table depicts the RMSFE ratios that exceed unity only. The parameter h denotes the forecast horizon. HAR-RV-M = The HAR-RV model is extended to include the following “moments”: realized skewness, realized kurtosis, realized good variance, realized bad variance, upside tail risk, downside tail risk, jumps. HAR-RV-M-C = The HAR-RV-M model is extended to include the following domestic and U.S. climate variables: Temperature, HDD, CDD, precipitation, wind. HAR-RV-M-CM = The HAR-RV-M model is extended to include the climate-related Google searches and U.S. news.

Table 8: Economic benefits of using the climate-risk predictors

Exchange rate / benchmark vs. rival model	$\delta = 1$	$\delta = 3$	$\delta = 5$
USDAUD / HAR-RV vs. HAR-RV-M	22.8081	22.216	22.0205
USDBRL / HAR-RV vs. HAR-RV-M	63.9011	62.2258	61.6727
USDCAD / HAR-RV vs. HAR-RV-M	8.9087	8.6790	8.6032
USDMXN / HAR-RV vs. HAR-RV-M	35.4924	34.5656	34.2596
USDMYR / HAR-RV vs. HAR-RV-M	0.8254	0.8043	0.7973
USDNOK / HAR-RV vs. HAR-RV-M	28.1385	27.4038	27.1612
USDRUB / HAR-RV vs. HAR-RV-M	100.2394	97.6172	96.7515
USDZAR / HAR-RV vs. HAR-RV-M	82.7249	80.5287	79.804
USDAUD / HAR-RV-M vs. HAR-RV-M-C	26.5456	25.8565	25.6289
USDBRL / HAR-RV-M vs. HAR-RV-M-C	56.4785	54.9978	54.5089
USDCAD / HAR-RV-M vs. HAR-RV-M-C	8.5492	8.3288	8.2560
USDMXN / HAR-RV-M vs. HAR-RV-M-C	39.9869	38.9427	38.5979
USDMYR / HAR-RV-M vs. HAR-RV-M-C	-	-	-
USDNOK / HAR-RV-M vs. HAR-RV-M-C	40.8051	39.7396	39.3879
USDRUB / HAR-RV-M vs. HAR-RV-M-C	-	-	-
USDZAR / HAR-RV-M vs. HAR-RV-M-C	65.8739	64.125	63.5479
USDAUD / HAR-RV-M vs. HAR-RV-M-CN	7.536	7.3404	7.2758
USDBRL / HAR-RV-M vs. HAR-RV-M-CN	73.6711	71.7396	71.102
USDCAD / HAR-RV-M vs. HAR-RV-M-CN	7.9529	7.7479	7.6802
USDMXN / HAR-RV-M vs. HAR-RV-M-CN	28.2678	27.5296	27.2859
USDMYR / HAR-RV-M vs. HAR-RV-M-CN	-	-	-
USDNOK / HAR-RV-M vs. HAR-RV-M-CN	94.0533	91.5975	90.7867
USDRUB / HAR-RV-M vs. HAR-RV-M-CN	233.6994	227.5864	225.5681
USDZAR / HAR-RV-M vs. HAR-RV-M-CN	47.4598	46.1999	45.7841

Both the benchmark and the rival model are estimated by random forests. Results are for a recursive-estimation window. Training period = 250 observations. For better readability of the results, the table only depicts positive fees (expressed in annual basis points) a forecast consumer would be willing to pay for using the forecasts implied by the rival model. The parameter h denotes the forecast horizon. HAR-RV-M = The HAR-RV model is extended to include the following "moments": realized skewness, realized kurtosis, realized good variance, realized bad variance, upside tail risk, downside tail risk, jumps. HAR-RV-M-C = The HAR-RV-M model is extended to include the following domestic and U.S. climate variables: Temperature, HDD, CDD, precipitation, wind. HAR-RV-M-CM = The HAR-RV-M model is extended to include the climate-related Google searches and U.S. news.

Appendix

Table A1: Comparing random forests with Ridge estimator and an elastic net (RMSFE ratios)

Exchange rate / benchmark vs. rival model	h=1	h=5	h=22
USDAUD / HAR-RV-M-C-Ridge vs. HAR-RV-M-C-RF	1.0052	1.0190	1.0299
USDBRL / HAR-RV-M-C-Ridge vs. HAR-RV-M-C-RF	1.8521	1.0383	1.0257
USDCAD / HAR-RV-M-C-Ridge vs. HAR-RV-M-C-RF	-	1.0016	1.0302
USDMXN / HAR-RV-M-C-Ridge vs. HAR-RV-M-C-RF	1.0066	1.0195	1.0079
USDMYR / HAR-RV-M-C-Ridge vs. HAR-RV-M-C-RF	1.0009	1.0223	1.0640
USDNOK / HAR-RV-M-C-Ridge vs. HAR-RV-M-C-RF	-	1.0116	1.0503
USDRUB / HAR-RV-M-C-Ridge vs. HAR-RV-M-C-RF	2.7294	2.1413	6.4866
USDZAR / HAR-RV-M-C-Ridge vs. HAR-RV-M-C-RF	1.0174	1.0186	1.0190
USDAUD / HAR-RV-M-C-Elastic net vs. HAR-RV-M-C-RF	1.0032	1.0695	1.0508
USDBRL / HAR-RV-M-C-Elastic net vs. HAR-RV-M-C-RF	-	1.0126	1.0217
USDCAD / HAR-RV-M-C-Elastic net vs. HAR-RV-M-C-RF	-	1.0014	1.0294
USDMXN / HAR-RV-M-C-Elastic net vs. HAR-RV-M-C-RF	1.0189	1.0129	1.0135
USDMYR / HAR-RV-M-C-Elastic net vs. HAR-RV-M-C-RF	1.0017	1.0199	1.0607
USDNOK / HAR-RV-M-C-Elastic net vs. HAR-RV-M-C-RF	-	1.0416	1.0574
USDRUB / HAR-RV-M-C-Elastic net vs. HAR-RV-M-C-RF	1.4351	2.1256	4.9324
USDZAR / HAR-RV-M-C-Elastic net vs. HAR-RV-M-C-RF	1.0151	1.0198	1.0188

Both the benchmark and the rival model include the following domestic and U.S. climate variables: Temperature, HDD, CDD, precipitation, wind. The benchmark model is estimated either by the Ridge estimator or an elastic net. The latter is computed as the equally weighted combination of the Lasso and the Ridge estimator. The rival model is estimated by random forests. For both models the root-mean-squared-forecasting error (RMSFE) is computed. The RMSFE ratio is constructed by dividing the RMSFE of the benchmark model by the RMSFE of the rival model. A RMSFE ratio larger than unity indicates that the rival model yields more accurate forecasts than the benchmark model. Results are for a recursive estimation window. Training period = 250 observations. For better readability of the estimation results, the table depicts the RMSFE ratios that exceed unity only. The parameter h denotes the forecast horizon.

Table A2: Results of encompassing regressions

Exchange rate / benchmark vs. rival model	benchmark/h=1	rival/h=1	benchmnark/h=5	rival/h=5	benchmark/h=22	rival/h=22
USDAUD / HAR-RV vs. HAR-RV-M	-	0.0000	-	0.0000	-	0.0000
USDBRL / HAR-RV vs. HAR-RV-M	-	0.0000	-	0.0000	-	0.0000
USDCAD / HAR-RV vs. HAR-RV-M	-	0.0000	-	0.0000	0.0096	0.0000
USDMXN / HAR-RV vs. HAR-RV-M	-	0.0000	-	0.0004	-	0.0007
USDMYR / HAR-RV vs. HAR-RV-M	-	0.0000	-	0.0000	-	0.0007
USDNOK / HAR-RV vs. HAR-RV-M	-	0.0000	-	0.0159	0.0778	0.0150
USDRUB / HAR-RV vs. HAR-RV-M	-	0.0010	0.0856	-	0.0078	-
USDZAR / HAR-RV vs. HAR-RV-M	-	0.0000	-	0.0000	0.0094	0.0000
USDAUD / HAR-RV-M vs. HAR-RV-M-C	-	0.0000	-	0.0000	-	0.0000
USDBRL / HAR-RV-M vs. HAR-RV-M-C	-	0.0000	-	0.0000	-	0.0000
USDCAD / HAR-RV-M vs. HAR-RV-M-C	-	0.0000	-	0.0000	-	0.0000
USDMXN / HAR-RV-M vs. HAR-RV-M-C	-	0.0000	-	0.0001	-	0.0000
USDMYR / HAR-RV-M vs. HAR-RV-M-C	-	0.0000	-	0.0000	-	0.0000
USDNOK / HAR-RV-M vs. HAR-RV-M-C	-	0.0000	-	-	0.0342	0.0000
USDRUB / HAR-RV-M vs. HAR-RV-M-C	-	0.0034	0.0656	-	0.0606	-
USDZAR / HAR-RV-M vs. HAR-RV-M-C	-	0.0000	-	0.0000	-	0.0000
USDAUD / HAR-RV-M vs. HAR-RV-M-CN	-	0.0000	-	0.0039	-	0.0000
USDBRL / HAR-RV-M vs. HAR-RV-M-CN	-	0.0000	-	0.0000	-	0.0000
USDCAD / HAR-RV-M vs. HAR-RV-M-CN	-	0.0000	-	0.0000	-	0.0000
USDMXN / HAR-RV-M vs. HAR-RV-M-CN	-	0.0000	0.0377	0.0000	-	0.0016
USDMYR / HAR-RV-M vs. HAR-RV-M-CN	-	0.0000	-	0.0000	-	0.0000
USDNOK / HAR-RV-M vs. HAR-RV-M-CN	-	0.0006	-	0.0142	-	0.0000
USDRUB / HAR-RV-M vs. HAR-RV-M-CN	-	0.0095	-	-	0.0173	-
USDZAR / HAR-RV-M vs. HAR-RV-M-CN	-	0.0000	-	0.0000	-	0.0000

Both the benchmark and the rival model are estimated by random forests. Results are for a recursive-estimation window. Training period = 250 observations. The encompassing regressions are estimated using the ordinary least squared (OLS) technique. For better readability of the estimation results (results for the intercept term are not shown), the table depicts the p-values (based on robust standard errors) that are smaller than 10% only. The parameter h denotes the forecast horizon. HAR-RV-M = The HAR-RV model is extended to include the following "moments": realized skewness, realized kurtosis, realized good variance, realized bad variance, upside tail risk, downside tail risk, jumps. HAR-RV-M-C = The HAR-RV-M model is extended to include the following domestic and U.S. climate variables: Temperature, HDD, CDD, precipitation, wind. HAR-RV-M-CM = The HAR-RV-M model is extended to include the climate-related Google searches and U.S. news.