UNIVERSITY OF ALBERTA
FACULTY OF ARTS
Department of Economics

Working Paper No. 2025-06

# Exchange Rate Predictability and Financial Conditions

**Sebastian Fossati**
**University of Alberta**

**Xiao Lu**
**University of Alberta**

September 2025

# Exchange Rate Predictability and Financial Conditions*

Sebastian Fossati† and Xiao Lu‡

University of Alberta

This version: September 12, 2025

**Abstract**

We model the conditional distribution of future exchange rate returns for nine currencies as a function of real-time financial conditions. We show that the lower and upper quantiles of the exchange rate return distribution exhibit significant in-sample co-movement with financial conditions. Similarly, the conditional moments of the out-of-sample forecast display time-varying patterns, with the variance and kurtosis showing the most pronounced changes during and after the 2008-09 financial crisis. Deteriorating financial conditions are associated with an increase in volatility, particularly for commodity currencies. Overall, we conclude that financial conditions capture tail dependencies in exchange rate returns and contain valuable information for out-of-sample prediction.

*Keywords*: exchange rates, financial conditions, NFCI, density forecasts
*JEL Codes*: C22, F31, G17

*We thank Chetan Dave for helpful comments and Maren Ulm for sharing data. We also thank seminar participants at University of Alberta, as well as attendees of the 59th Annual Meetings of the Canadian Economics Association (CEA), the 3rd Vienna Workshop on Economic Forecasting, and the 2025 meeting of the International Association for Applied Econometrics (IAAE) for their comments and suggestions.

†Corresponding author. Email: sfossati@ualberta.ca. Department of Economics, University of Alberta, Edmonton (AB), Canada. ORCID: https://orcid.org/0009-0007-8081-2197

‡Email: xlu15@ualberta.ca. Department of Economics, University of Alberta, Edmonton (AB), Canada.

# 1 Introduction

A well-established empirical fact in the macroeconomic forecasting literature is that exchange rates are hard to predict. In other words, it is hard to find macroeconomic variables that can help in forecasting exchange rate changes better than a simple random walk. This is known in the literature as the "exchange rate disconnect" puzzle and has been extensively documented in Meese and Rogoff (1983); Engel and West (2005); Rossi (2013); Itskhoki and Mukhin (2021); Engel and Wu (2023a), among many others. One of the reasons why forecasting exchange rates is difficult is that exchange rate predictability is unstable over time. For example, a model that works well in one period and using a certain performance metric may not work well in another period or using an alternative metric (Rossi, 2013; Byrne et al., 2018; Cheung et al., 2019; Colombo and Pelagatti, 2020; Engel and Wu, 2024).

Our paper contributes to an emerging literature that finds that financial variables may contain valuable information for predicting exchange rates. For example, Lilley et al. (2022) show that using measures of global risk demand and foreign bond purchases in the United States (U.S.) can help predict exchange rates. They note, however, that the predictive content of these variables is time-varying and appears to begin with the 2008-09 financial crisis. Similarly, Engel and Wu (2023b) find that using measures of the liquidity yield on government bonds can help outperform the random walk prediction of exchange rates in out-of-sample forecasting exercises if one focuses on the post-2008 sample period (see, also, Engel and Wu, 2024). A common feature of these papers is that their out-of-sample forecasting exercises focus on point forecasts of exchange rate returns, ignoring the distributional aspects of the forecasts.

In this paper, we model the conditional distribution of future exchange rate returns as a function of current (real-time) financial conditions. Our approach is similar to Adrian et al. (2019), who use a two-step procedure that combines quantile regressions (QR) and skewed-$t$ distributions to model the conditional distribution of U.S. real GDP growth using financial conditions. In our case, we follow this methodology to forecast the short-term conditional distribution of exchange rate returns computed from nine currencies relative to the U.S. dollar (USD). Namely, the Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), the euro (EUR), U.K. pound sterling (GBP), Japanese yen (JPY), Norwegian krone (NOK), New Zealand dollar (NZD), and Swedish krona (SEK).

Financial conditions refer to the overall state of financial markets and institutions, encompassing factors such as credit availability, risk levels, and leverage in the financial system. One measure of financial conditions that has gained popularity in recent years is the Chicago Fed's National Financial Conditions Index (NFCI). The NFCI measures financial conditions in the U.S. by aggregating indicators that capture aspects of risk, credit, and leverage within financial markets and institutions. The NFCI has been shown to capture historical periods of financial stress, such as the 2008-09 financial crisis, and provides a useful summary of financial conditions (Brave and Butters, 2011, 2012). A recent literature has shown that the NFCI can be used to forecast the lower quantiles of real GDP growth, known as "growth-at-risk" (Alessandri and Mumtaz, 2017; De Nicolò and Lucchetta, 2017; Adrian et al., 2019; Delle Monache et al., 2024). Amburgey and McCracken (2023) report notable gains when using real-time vintages to forecast U.S. real GDP growth over ex-post data, particularly in the periods leading to recessions.

Since the NFCI is updated weekly, we construct weekly exchange rate returns as 100 times the log difference of the exchange rate on the Wednesday of each week, aligned with each NFCI release. For our main analysis, our dataset consists of weekly data from January 1973 to July 2021, with real-time NFCI vintages available from January 1988 to July 2021. In the first step, we estimate predictive QR of the weekly exchange rate returns on current financial conditions to obtain the conditional quantiles of the future exchange rate returns. Our results show that the lower and upper quantiles of the exchange rate return distribution exhibit significant co-movement with current financial conditions. We find this result to be particularly pronounced for the AUD, CAD, and the NZD, which are commodity currencies that are more sensitive to changes in financial conditions (Chen and Rogoff, 2003; Bondatti and Rillo, 2022). Next, we use a skewed-$t$ distribution to characterize the predictive densities of the future exchange rate returns. The skewed-$t$ distribution is a generalization of the $t$-distribution that allows for skewness and heavy tails, making it a flexible distribution for modeling financial data (Azzalini and Capitanio, 2003; Adrian et al., 2019). We find that the conditional moments of the out-of-sample skewed-$t$ forecast exhibit time-varying patterns, with the variance and kurtosis showing the most pronounced changes during and after the 2008-09 financial crisis.

These results hold for all nine currencies and for different short-term forecast horizons (one-day, one-week, and four-weeks ahead). However, we find that the predictive content of the NFCI for exchange rate returns is time-varying and appears to begin in the early 2000s. We evaluate the robustness of our results to different forecast evaluation metrics, including the root mean squared error (RMSE), log Score, and Fluctuation tests (Giacomini and Rossi, 2010), and conclude that the NFCI provides valuable information for forecasting exchange rate returns, leading to more accurate density forecasts than the random walk benchmark.

Our paper also contributes to a recent literature on exchange rate tail risks. For example, Eguren-Martin and Sokol (2022) also use QR to show that the conditional distribution of monthly exchange rate returns responds to changes in global financial conditions. However, they do not evaluate the out-of-sample predictive content of financial conditions and, as a result, they cannot determine whether financial conditions are useful for monitoring tail risks in real time. Our results show that this is the case. Ferrara and Yapi (2022) use a similar approach to ours to show that Brexit-related uncertainty was associated with higher future depreciation risks of the GBP versus the euro. Smales (2022) who find that economic policy uncertainty indexes can help predict JPY exchange rates. Other works focusing on exchange rate tail risks include Bondatti and Rillo (2022) who find that exchange rates are exposed to downside tail risk with respect to some commodities and Cañon et al. (2024) who document exchange rate tail risks following unconventional monetary and liquidity policies.

The rest of the paper is organized as follows. Section 2 describes the measures of financial conditions and exchange rate returns used in the analysis. Section 3 presents the methods used to forecast exchange rate returns. Section 4 presents the empirical results of the analysis. Section 5 concludes.

# 2 Data

## 2.1 Financial conditions in real-time

Financial conditions refer to the overall state of financial markets and institutions, encompassing factors such as credit availability, risk levels, and leverage in the financial system. The Chicago Fed's NFCI measures financial conditions by aggregating 105 indicators that capture aspects of risk, credit, and leverage within financial markets and institutions. The index is constructed as the common factor in a dynamic factor model of the financial indicators, each expressed relative to its historical mean and standard deviation. A value of zero represents financial conditions at their historical average, while positive values indicate tighter-than-average conditions, and negative values suggest looser conditions. The adjusted-NFCI (ANFCI) further refines this measure by removing the influence of economic activity and inflation, isolating financial conditions independent of broader macroeconomic trends. Brave and Butters (2011, 2012) show that the NFCI captures historical periods of financial stress, such as the 2008-09 financial crisis, and provides a useful summary of financial conditions.[1]

The NFCI is updated weekly, with new data released every Wednesday at 8:30 a.m. ET, covering financial conditions through the previous Friday. Since the index incorporates a mix of weekly, monthly, and quarterly data, historical values can be revised as new information becomes available, particularly at the beginning of each month when many economic time series are updated. However, the NFCI is available in real-time only since May 2011, which limits its usefulness for out-of-sample forecasting exercises. To address this limitation, Amburgey and McCracken (2023) construct real-time vintages of the NFCI that allow for a more accurate assessment of the NFCI's predictive content in real-time. For example, they re-examine the effectiveness of the NFCI in forecasting U.S. real GDP growth and find notable gains in using real-time vintages over ex-post data as in Adrian et al. (2019), particularly in the periods leading up to recessions. This finding is particularly significant in our context, as exchange rate predictability is known to be time-varying and sensitive to changes in financial conditions that often precede economic downturns (Rossi, 2013; Cheung et al., 2019; Lilley et al., 2022; Engel and Wu, 2024).

In this paper, we use the unofficial NFCI vintages of Amburgey and McCracken

[1]The NFCI is available at https://www.chicagofed.org/research/data/nfci/current-data.

(2023) to examine the real-time predictive content of the NFCI for short-term exchange rate returns.[2] The NFCI vintages are available weekly from 1988/01/06 to 2021/07/28, with vintages starting in January 1973. Figure 1 plots the 2024/10/16 release of the ex-post NFCI and the real-time NFCI vintages for the period 2008 to 2009. As noted in Amburgey and McCracken (2023), the real-time vintages exhibit significant revisions and are more volatile than the ex-post NFCI, especially in the periods leading up to recessions.

[Figure 1 about here]

## 2.2 Exchange rate returns and the NFCI release dates

Our objective is to evaluate the real-time predictive content of the NFCI for exchange rate returns. To this end, we obtained daily nominal exchange rates for nine commonly analyzed currencies (see, e.g., Engel and Wu, 2023a) from the Federal Reserve Economic Data (FRED) database at the Federal Reserve Bank of St. Louis. Since the NFCI is updated weekly, we construct weekly exchange rate returns as 100 times the log difference of the exchange rate on the Wednesday of each week. More specifically, we align the weekly NFCI releases (Wednesdays at 8:30 a.m. ET) with the exchange rate returns obtained using noon buying rates in New York City to ensure that the NFCI is available at the time of the forecast. Note that while the NFCI is available at the time of the forecast (say, Wednesday at noon), the indicators used to construct the NFCI have been available since (at least) the previous Friday. As a result, our forecasting exercise evaluates the real-time predictive content of the NFCI for exchange rate returns, while establishing a lower bound for the predictability of exchange rates

[2]The real-time NFCI is available at https://www.stlouisfed.org/research/economists/mccracken.

using financial conditions.

Let $s_t$ denote the log of the nominal exchange rate of the foreign currency to one USD on the Wednesday of week $t$. The weekly $h$-step-ahead exchange rate return is defined as

$$r_{t+h} = 100 \times (s_{t+h} - s_t)\,, \tag{1}$$

with $t = 1, 2, \ldots, T$ denoting the weekly time index. Figure 2 plots the weekly exchange rate returns for the nine currencies in our sample from 1973/01/03 to 2024/10/16, while Figure A.1 of the online appendix plots the log exchange rates. Volatility in exchange rate returns is generally higher in recent years, particularly during periods of financial stress such as the 2008-09 financial crisis.

[Figure 2 about here]

# 3 Methods

## 3.1 Benchmark forecasts

The $h$-step-ahead exchange rate return is often modeled as a function of some variable or variables that are believed to contain predictive information about future exchange rates (see, e.g., Meese and Rogoff, 1983; Rossi, 2013; Colombo and Pelagatti, 2020; Engel and Wu, 2023a). For the $h$-step-ahead exchange rate return, $r_{t+h}$, we have the forecasting equation given by

$$r_{t+h} = \alpha_h + \boldsymbol{x}_t'\boldsymbol{\beta}_h + \varepsilon_{t+h}, \tag{2}$$

where $\boldsymbol{x}_t$ is the vector of current predictors. These return regressions are also used in studies on stock return predictability, such as those of Welch and Goyal (2008), Rapach et al. (2010), and Gargano and Timmermann (2014), among others.

To evaluate out-of-sample predictive content of $\boldsymbol{x}_t$, researchers typically estimate (2) over rolling or expanding samples to create forecasts that can be compared to the forecasts of some benchmark, usually the random walk. The $h$-step-ahead point forecast of the exchange rate return is given by

$$\hat{r}_{t+h} = \hat{\alpha}_h + \boldsymbol{x}_t' \hat{\boldsymbol{\beta}}_h, \tag{3}$$

while the $h$-step-ahead density forecast of the exchange rate return is usually obtained by assuming that the forecast error, $\varepsilon_{t+h}$, is normally distributed such that

$$\hat{f}(r_{t+h}) \sim \mathcal{N}(\hat{\alpha}_h + \boldsymbol{x}_t' \hat{\boldsymbol{\beta}}_h, \hat{\sigma}_h^2), \tag{4}$$

where $\hat{\sigma}_h^2$ is the estimated variance of the $h$-step-ahead prediction error.

We consider three forecasts for future exchange rate returns $r_{t+h}$ based on equations (2), (3), and (4). These are the following: a random walk benchmark forecast, a random walk with drift forecast, and a forecast obtained using the current value of the NFCI as the predictor.

**Random walk (RW)**. The RW forecast of no change in the (log) exchange rate is the most used benchmark in the exchange rate literature since Meese and Rogoff (1983). The RW with no drift assumes $\alpha_h = 0$ and $\boldsymbol{\beta}_h = \mathbf{0}$. As a result, the point forecast of the $h$-step-ahead exchange rate return is given by $\hat{r}_{t+h} = 0$, and a density forecast of the $h$-step-ahead exchange rate return is given by $\hat{f}(r_{t+h}) \sim \mathcal{N}(0, \hat{\sigma}_h^2)$.

**Random walk with drift (RW w/ drift)**. The RW with drift forecast of the (log) exchange rate allows for a constant trend in the exchange rate. The RW with drift assumes $\alpha_h \neq 0$ and $\boldsymbol{\beta}_h = \mathbf{0}$. As a result, the point forecast of the $h$-step-ahead exchange rate return is given by $\hat{r}_{t+h} = \hat{\alpha}_h$, and a density forecast of the $h$-step-ahead exchange rate return is given by $\hat{f}(r_{t+h}) \sim \mathcal{N}(\hat{\alpha}_h, \hat{\sigma}_h^2)$.

**OLS with NFCI (OLS)**. Finally, we consider a forecast based on the OLS estimation of equation (2) using the current NFCI as predictor of the future exchange rate return. In this case, the point forecast of the $h$-step-ahead exchange rate return is given by $\hat{r}_{t+h} = \hat{\alpha}_h + \boldsymbol{x}_t'\hat{\boldsymbol{\beta}}_h$, and a density forecast of the $h$-step-ahead exchange rate return is given by $\hat{f}(r_{t+h}) \sim \mathcal{N}(\hat{\alpha}_h + \boldsymbol{x}_t'\hat{\boldsymbol{\beta}}_h, \hat{\sigma}_h^2)$.

## 3.2 QR-skewed-$t$ density forecasts

A recent literature has shown that the predictive content of financial conditions can be enhanced by using QR and a skewed-$t$ distribution to model the predictive distribution of real GDP growth (Adrian et al., 2019). For the $h$-step-ahead forecast of the exchange rate return, $r_{t+h}$, we implement the following two-step procedure. First, we estimate the QR of the future exchange rate return on the current NFCI to obtain the conditional predictive quantiles of the exchange rate return. The QR approach is similar to the OLS regression in equation (2), but differs in that it estimates coefficients for a specific percentile $\tau$. The predicted values from the QR are the quantiles of the exchange rate return conditional on the current NFCI, $\hat{Q}_{r_{t+h}|\boldsymbol{x}_t}(\tau|\boldsymbol{x}_t)$, and are given by

$$\hat{Q}_{r_{t+h}|\boldsymbol{x}_t}(\tau|\boldsymbol{x}_t) = \hat{\alpha}_h^\tau + \boldsymbol{x}_t'\hat{\boldsymbol{\beta}}_h^\tau. \tag{5}$$

Next, we use the conditional quantiles obtained from (5) to obtain the predictive density of the exchange rate return. We follow the procedure described in Adrian et al. (2019) and fit a skewed-$t$ distribution to obtain the predictive densities of the exchange rate return. The skewed-$t$ distribution is a generalization of the $t$-distribution that allows for skewness and fat tails (Azzalini and Capitanio, 2003). The skewed-$t$

distribution has four parameters: $\mu$ for location, $\sigma$ for scale, $\nu$ for fatness, and $\alpha$ for skewness. The density function of the skewed-$t$ distribution is given by

$$f(r;\mu,\sigma,\alpha,\nu) = \frac{2}{\sigma}\, t\left(\frac{r-\mu}{\sigma};\nu\right)\, T\left(\alpha\frac{r-\mu}{\sigma}\sqrt{\frac{\nu+1}{\nu+\frac{r-\mu}{\sigma}}};\nu+1\right) \tag{6}$$

where $f(\cdot)$ is the skewed-$t$ probability density function, $t(\cdot)$ is the probability density function of the $t$-distribution, and $T(\cdot)$ is the cumulative distribution function of the $t$-distribution. The skewed-$t$ distribution is a flexible distribution that can capture the skewness and fat tails often observed in financial data.

Following Adrian et al. (2019), for each week, we choose the four parameters of the skewed-$t$ distribution by minimizing the squared distance between the predicted 2.5, 5, 25, 75, 95, and 97.5 percent quantiles, $\hat{Q}_{r_{t+h}|\boldsymbol{x}_t}(\tau|\boldsymbol{x}_t)$, and the inverse cumulative distribution function of the skewed-$t$ distribution such that

$$\{\hat{\mu}_{t+h}, \hat{\sigma}_{t+h}, \hat{\alpha}_{t+h}, \hat{\nu}_{t+h}\} = \underset{\mu,\sigma,\alpha,\nu}{\operatorname{argmin}} \sum_{\tau}\left[\hat{Q}_{r_{t+h}|\boldsymbol{x}_t}(\tau|\boldsymbol{x}_t) - F^{-1}(\tau;\mu,\sigma,\alpha,\nu)\right]^2, \tag{7}$$

where $F^{-1}(\cdot)$ denotes the inverse cumulative distribution function of the skewed-$t$ distribution. In this case, the point forecast of the $h$-step-ahead exchange rate return is given by $\hat{r}_{t+h} = \hat{Q}_{r_{t+h}|\boldsymbol{x}_t}(0.5|\boldsymbol{x}_t)$, and the density forecast of the $h$-step-ahead exchange rate return is given by $\hat{f}(r_{t+h}|\boldsymbol{x}_t) = f(r_{t+h};\hat{\mu}_{t+h},\hat{\sigma}_{t+h},\hat{\alpha}_{t+h},\hat{\nu}_{t+h})$. For the remainder of the paper, we refer to this two-step procedure as the QR-skewed-$t$ approach.

## 3.3 Out-of-sample forecast evaluation

To evaluate the out-of-sample (OOS) predictive content of the NFCI for exchange rate returns, we conduct the following forecasting exercises. On the Wednesday of each week, we use the real-time (unofficial) NFCI vintages of Amburgey and McCracken (2023) to generate forecasts for the future exchange rate returns using the three benchmark forecasts (i.e., RW, RW with drift, and OLS), as well as the QR-skewed-$t$ forecast.

In each case, we generate a point forecast and a density forecast. Our main results are obtained using a rolling estimation window of 520 weeks (approximately 10 years) and a forecast horizon of one week (a one-step-ahead forecast). In addition, we consider a rolling estimation window of 1040 weeks (approximately 20 years) and an expanding window approach (recursive estimation), as well as forecast horizons of one day and four weeks. All our real-time OOS forecasting exercises run from 2000/01/05 to 2021/07/21, based on the real-time NFCI vintages of Amburgey and McCracken (2023) available at the time of writing.

To compare the sequences of $h$-step-ahead OOS forecasts of the exchange rate return, $r_{t+h}$, obtained using the current value of NFCI as a predictor with the forecasts obtained using a benchmark model (i.e., the RW forecast) we proceed as follows. First, we assume that the sample of size $T$ has been divided into an in-sample portion of size $R$ (e.g., 520 weeks) and an OOS portion of size $P$ (the remaining data). Next, we proceed to evaluate the accuracy of each set of $P$ OOS forecasts using an appropriate loss function. We evaluate the predictive accuracy of point forecasts using the RMSE. The RMSE is given by

$$\text{RMSE} = \sqrt{P^{-1} \sum_{t=R+h}^{T} (r_t - \hat{r}_t)^2}, \tag{8}$$

where $r_t$ is the realized value of the exchange rate return and $\hat{r}_t$ is the forecasted value of the exchange rate return. A lower RMSE indicates better predictive accuracy. Similarly, we evaluate the predictive accuracy of density forecasts using the log Score. The log score is a proper scoring rule (Gneiting and Raftery, 2007) and is given by

$$\log \text{Score} = P^{-1} \sum_{t=R+h}^{T} -\log \hat{f}(r_t), \tag{9}$$

where $r_t$ is the realized value of the exchange rate return and $\hat{f}(r_t)$ is the forecasted density of the exchange rate return. The log score is a measure of the log likelihood of the realized value of the exchange rate return under the forecasted density, and a lower log score indicates better predictive accuracy.

We evaluate the statistical significance of the OOS predictability results using the Diebold-Mariano (DM) test (Diebold and Mariano, 1995). For a given loss function $L(\cdot)$, we compute the sequence of $P$ OOS forecast loss differences $\{\Delta L_t\}_{t=R+h}^{T}$ and we are interested in testing the null hypothesis of equal predictive ability $H_0 : E\left[\Delta L_t\right] = 0$ against the alternative hypothesis that the model using the NFCI exhibits superior predictive ability. Diebold and Mariano (1995) propose an asymptotic test of equal (unconditional) predictive ability that can be applied to both nested and non-nested models (Giacomini and White, 2006). The DM statistic is given by

$$\text{DM} = \hat{\sigma}_P^{-1} P^{1/2} \overline{\Delta L}_P, \tag{10}$$

where $\overline{\Delta L}_P = P^{-1} \sum_{t=R+h}^{T} \Delta L_t$ and $\hat{\sigma}_P^2$ is the sample variance of $\Delta L_t$ if $h = 1$ or a HAC estimator of the long-run variance if $h > 1$. A common choice for the latter is the kernel-based estimator

$$\hat{\sigma}_P^2 = \sum_{j=-(q-1)}^{q-1} k(j/q) P^{-1} \sum_{t=R+h}^{T} \Delta L_t^* \Delta L_{t-j}^* \tag{11}$$

where $k(\cdot)$ is a kernel weight function, for example the Bartlett kernel of Newey and West (1987), $q$ is a bandwidth that grows with $P$, and $\Delta L_t^* = \Delta L_t - \overline{\Delta L}_P$ (see, for example, Andrews, 1991). Under the null hypothesis of equal predictive ability, $\text{DM} \xrightarrow{d} \mathcal{N}(0,1)$ as $P \to \infty$ and the null hypothesis of equal (unconditional) predictive ability is rejected at the 5% level if $|\text{DM}| > 1.96$. Giacomini and White (2006) note that since the null hypothesis is stated in terms of the estimated parameters, not the population parameters, this is a test of predictive ability in the finite sample (that is, given models, estimation windows, estimation procedures, etc.).

Since exchange rate predictability is known to be time-varying (Rossi, 2013; Engel and Wu, 2024), we are also interested in testing the null hypothesis of equal predictive ability against the alternative of local predictive ability. Giacomini and Rossi (2010) propose asymptotic tests of the joint hypothesis of equal and constant performance of the two models against the alternative of local predictive ability. In particular, the Fluctuation test is based on a rolling average of loss differences and the statistic is given by

$$\text{Fluct}_{t,m} = \hat{\sigma}_P^{-1} m^{-1/2} \sum_{j=t-m/2}^{t+m/2+1} \Delta L_j. \tag{12}$$

The statistic is computed for $t = R+h+m/2, \ldots, T-m/2+1$, with $m$ the window size and $\hat{\sigma}_P^2$ a HAC estimator of the long-run variance of $\Delta L_t$, for example the kernel-based estimator (11). The test rejects when $\max_t |\text{Fluct}_{t,m}| > k_\alpha$ with the critical value $k_\alpha$ obtained by simulation (Giacomini and Rossi, 2010).

# 4 Empirical results

## 4.1 In-sample return predictability

While the focus of this paper is on real-time out-of-sample predictability, we begin by briefly discussing the in-sample predictability of exchange rate returns using the ex-post NFCI (i.e., the 2024/10/16 release) as predictor. Figure 3 plots the OLS and QR in-sample estimates of the slope coefficients obtained from equations (3) and (5), respectively, for the one-week-ahead exchange rate returns and the sample period from 2000/01/01 to 2021/12/31. The OLS estimates (blue line) are found to be essentially zero for all currencies. This result is consistent with the literature that finds little evidence of in-sample predictability of exchange rates (see, e.g., Meese and Rogoff, 1983; Rossi, 2013). In contrast, the QR estimates (orange line) exhibit pronounced variation across the return distribution. In particular, the QR coefficients tend to be significantly negative at the lower quantiles and significantly positive at the upper quantiles. The shaded regions around the QR estimates denote the 95% confidence intervals, underscoring that these tail-dependent effects differ statistically from zero in most cases.

[Figure 3 about here]

Since exchange rates are quoted in units of the foreign currency per one USD, this result shows that when U.S. financial conditions tighten (i.e., $\Delta$NFCI $> 0$), such as higher interest rates and reduced credit supply, the distribution of the one-week-ahead exchange rate returns widens, with a higher probability of large negative returns and, also, a higher probability of large positive returns. This is particularly pronounced for the AUD, CAD, and the NZD, which are commodity currencies that tend to be more sensitive to changes in global financial conditions (Bondatti and Rillo, 2022).

Figure 4 plots the in-sample estimates of the slope coefficients but for the sample period from 1980/01/01 to 1999/12/31. In contrast to the results for the post-2000 period, both the OLS and QR coefficients are now essentially zero for all currencies. As a result, the predictive content of the NFCI for exchange rate returns is mostly non-existent in the pre-2000 period. This finding is consistent with a recent literature that finds that the predictive content of financial indicators for exchange rates is time-varying and appears to begin with the 2008-09 financial crisis (see, e.g., Lilley et al., 2022; Engel and Wu, 2023b, 2024).

[Figure 4 about here]

## 4.2 Out-of-sample return predictability in real-time

Next, we turn our attention to the OOS forecasting exercises, where we evaluate the predictive content of the NFCI for exchange rate returns using real-time vintages of the NFCI. Given our in-sample results reported in Figures 3 and 4 and the availability of real-time NFCI vintages at the time of writing (Amburgey and McCracken, 2023), the real-time OOS forecasting exercise runs from 2000/01/05 to 2021/07/21.

The top panel of Table 1 reports the RMSE of the one-week-ahead OOS forecasts of exchange rate returns. For the RW benchmark, the RMSE is reported for each currency. For the alternative models (RW with drift, OLS, and QR-skewed-$t$), the RMSE ratios of the alternative model to the RW benchmark are reported, with the one-sided DM test $p$-values in parentheses (Diebold and Mariano, 1995). A ratio below one indicates that the alternative model outperforms the RW benchmark. Our results show that none of the alternative forecasts outperforms the RW benchmark in terms of the OOS point forecasts.

[Table 1 about here]

The bottom panel of Table 1 reports the log Score of the one-week-ahead OOS forecasts of the exchange rate returns. For the RW benchmark, the log Score is reported for each currency. For the alternative models, the log Score ratios of the alternative model to the RW benchmark are reported, with the one-sided DM test $p$-values in parentheses. Our results show that the QR-skewed-$t$ forecast outperforms the RW benchmark in terms of log Score and the null hypothesis of equal predictive ability is rejected for all nine currencies (all $p$-values below 0.05). In contrast, the RW with drift and OLS forecasts do not exhibit improvements over the RW benchmark. Overall, the QR-skewed-$t$ forecast with the current NFCI yields the best density forecast in terms of log Score.

Figure 5 shows the realized exchange rate returns for each currency alongside the real-time one-week-ahead 95% prediction intervals from the RW benchmark and the QR-skewed-$t$ forecast. First, we note that the QR-skewed-$t$ forecasts exhibit notably wider bounds during episodes of heightened volatility, particularly during the 2008-09 financial crisis. Second, this widening of the prediction intervals is typically asymmetric. As a result, the QR-skewed-$t$ forecast adapts more effectively to sudden increases in risk, capturing both negative and positive tail movements in realized returns. In contrast, the RW prediction intervals remain relatively narrow, with substantial underestimation of the magnitude of realized returns during high stress periods. Overall, the QR-skewed-$t$ prediction intervals appear to provide better coverage of the large return swings observed during the 2008-09 financial crisis, underscoring the importance of accommodating tail risks when forecasting exchange rate movements under volatile market conditions. Furthermore, we observe that in the post-crisis period, the QR-skewed-$t$ prediction intervals are now narrower than the RW intervals, yet still capturing the bulk of the realized returns. As a result, by adapting to real-time financial conditions via the NFCI, the QR-skewed-$t$ density forecast captures changes in mean, volatility, and asymmetry, offering more accurate predictions than models that impose constant or symmetric conditional distributions.

[Figure 5 about here]

## 4.3 Sub-sample analysis

The results in Figure 5 show evidence of time variation in the conditional distribution of the one-week-ahead QR-skewed-$t$ density forecasts. As a result, we now examine the conditional moments of the QR-skewed-$t$ density forecasts. The conditional mean and variance of the one-week-ahead QR-skewed-$t$ density forecasts of the exchange rate returns are plotted in Figures 6 and 7, respectively. We note that while the conditional mean typically remains close to zero, a substantial spike appears during the 2008-09 financial crisis. The conditional variance also exhibits a time-varying pattern, with very large spikes during the 2008-09 financial crisis and, to a lesser extent, the COVID-19 pandemic. For instance, AUD, EUR, and NZD exhibit variance estimates that jump well above their usual levels in late 2008, reflecting the sudden and severe volatility in these currencies. In contrast, the CAD, CHF, JPY, and SEK show comparatively smaller upticks in conditional variance over the same period, though still higher than their pre-crisis baselines. The skewness and kurtosis of the QR-skewed-$t$ density forecasts (plotted in Figures A.6 and A.7 of the online appendix) also exhibit some time variation, but the patterns are less pronounced than those for the mean and variance. Outside the 2008-09 financial crisis, we find that the mean forecasts remain relatively stable even during periods of heightened financial stress, indicating that the NFCI does not provide a strong signal about the direction of exchange rate movements but rather about the level of uncertainty surrounding them.

[Figure 6 about here]

[Figure 7 about here]

An alternative approach to assessing the time-varying predictive performance of the density forecasts is to examine the cumulated sum of error differentials between the RW benchmark and a candidate prediction (e.g., QR-skewed-$t$), as in Gargano and Timmermann (2014) and Lyócsa et al. (2024). Figure 8 shows the one-week-ahead cumulated log Score differences for the three alternative density forecasts of the exchange rate returns—RW with drift, OLS, and QR-skewed-$t$—relative to the RW benchmark. A positive log Score difference indicates that the alternative forecast performs better than the benchmark. Across most currencies and time periods, the QR-skewed-$t$ forecasts shows a small but steady improvement over the RW benchmark. In addition, we observe that the cumulated log Score differences for the QR-skewed-$t$ forecasts exhibit pronounced jumps during the 2008-09 financial crisis and, to some extent, the COVID-19 pandemic, suggesting that the NFCI captures tail dependencies in exchange rate returns during turbulent economic phases. In contrast, the RW with drift and OLS forecasts generally track near or below zero, indicating that neither approach delivers superior performance relative to the RW benchmark.

[Figure 8 about here]

Another way to assess the time-varying predictive performance of the density forecasts is to report the Fluctuation statistic of Giacomini and Rossi (2010), as in Byrne et al. (2018) and Colombo and Pelagatti (2020). Figure 9 shows Fluctuation statistics for the log Score of the one-week-ahead OOS density forecasts of the exchange rate returns. Positive values indicate that the alternative forecast is more accurate than the RW benchmark. The dashed horizontal lines represent the 95% critical values for a one-sided test of the null hypothesis of equal predictive accuracy (Giacomini and Rossi, 2010). For all currencies, the null hypothesis is rejected at the 5% level for the QR-skewed-$t$ forecast, indicating that the forecasts outperform the RW benchmark. In contrast, the RW with drift and OLS Fluctuation statistics remain near or below the zero baseline for most of the sample, suggesting little improvement over the RW benchmark. The Fluctuation tests also highlight that the QR-skewed-$t$ advantage may be especially pronounced in episodes of heightened financial instability, lending further support to the notion that heavy-tailed or distributionally flexible models can capture the underlying risks in exchange rate movements more effectively than simpler alternatives.

[Figure 9 about here]

## 4.4 Calibration

Diebold et al. (1997) propose the use of the probability integral transform (PIT) as a tool for assessing the calibration of density forecasts. Figure 10 shows the histograms of the PIT of the one-week-ahead density forecasts of the exchange rate returns from the RW benchmark (grey line) and the QR-skewed-$t$ alternative (orange line). The horizontal axis corresponds to selected PIT quantiles, while the vertical axis shows the ratio of the observed PIT frequencies to those implied by a perfectly calibrated forecast distribution. The dashed horizontal line at unity thus denotes perfect calibration. For most currencies, the RW forecast densities exhibit deviations from the uniform distribution with the PIT histograms showing a lower frequency of observed returns in the tails than implied by the uniform distribution (i.e., values below 1 at the tails) and a higher frequency of returns near the center (i.e., values above 1 at the center). As a result, the RW forecasts tend to underpredict the likelihood of large exchange rate returns (i.e., the predicted intervals are too narrow, on average). In contrast, the QR-skewed-$t$ forecasts often exhibit better calibration of the predictive densities. This is particularly evident for the AUD, EUR, GBP, JPY, NZD, and SEK.

[Figure 10 about here]

## 4.5 Additional results

### 4.5.1 Other estimation windows

Table 1 showed the results for the one-week-ahead forecast horizon, using a 10 year rolling window (i.e., 520 weeks). Overall, we concluded that the QR-skewed-$t$ forecasts consistently outperform the RW benchmark across all currencies. We find very similar results when using a 20 year rolling window (Table A.1 of the online appendix) and a small loss in predictive ability when using an expanding window (Table A.2 of the online appendix). Overall, stronger predictive ability is observed using rolling windows compared to an expanding window.

### 4.5.2 Other forecast horizons

Exchange rate predictability can vary significantly across different forecast horizons (Rossi, 2013). Table 1 showed an average improvement in predictive accuracy of (about) four percent for the QR-skewed-$t$ forecasts over the RW benchmark for the one-week-ahead forecast horizon. The results for the one-day-ahead forecast horizon (Table A.3 of the online appendix) show an average improvement in predictive accuracy of (about) eight percent for the QR-skewed-$t$ forecasts over the RW benchmark. As a result, we find stronger evidence of predictability at even shorter horizons. In contrast, OOS predictability is weaker at the four-week-ahead forecast horizon (Table A.4 of the online appendix). The average improvement in predictive accuracy is (about) two percent for the QR-skewed-$t$ forecasts over the RW benchmark. Overall, our results point to the robustness of the QR-skewed-$t$ forecasts at short forecast horizons.

## 4.6 Controlling for interest rate differentials

So far we have analyzed the predictive content of the NFCI as a single predictor of exchange rate returns. The literature suggests that other predictors (e.g., interest rate differentials, price differentials, output differentials) may also play a significant role (Rossi, 2013). However, most of these predictors are not available at high frequency, which limits their usefulness for real-time short-term forecasting. In this section, we evaluate the robustness of our results to the inclusion of interest rate differentials (IRD) as an additional predictor. IRD have been extensively used to forecast exchange rates (see, e.g., Meese and Rogoff, 1983; Molodtsova and Papell, 2009; Giacomini and Rossi, 2010; Byrne et al., 2018; Ulm and Hambuckers, 2022; Ostry, 2023). In particular, Ulm and Hambuckers (2022) find that daily exchange rate volatility is driven by IRD, a result consistent with our finding that the financial indicators may contain information on higher-order moments of exchange rate returns.

To compute the IRD, we obtained daily three-month Libor rates for CHF, EUR, GBP, and JPY, and three-month Euro Libor rates for CAD and AUD ending in 2017/12/20, as in Ulm and Hambuckers (2022).[3] To control for IRD, we augment the predictor vector $\boldsymbol{x}_t$ in equations (3) and (5) to include the current value of the NFCI and the IRD. As a result, in this section we focus on these six currencies against the USD and focus our attention on the sample period from 2000/01/01 to 2017/12/20. The in-sample estimates of the slope coefficients for the one-week-ahead exchange rate returns show that the NFCI coefficients remain unchanged when controlling for IRD (Figure A.8 of the online appendix). The top panel of Table 2 reports the log Score of the one-week-ahead OOS forecasts of the exchange rate returns for the period 2008/01/02 to 2017/12/20 without controlling for IRD, while the bottom panel reports the log

[3]We thank Maren Ulm for providing the data on interest rate differentials.

Score when controlling for IRD. Overall, we find that the log Score of QR-skewed-$t$ forecasts are not affected in a substantial way by the inclusion of IRD. As a result, controlling for IRD does not change the finding that the QR-skewed-$t$ outperforms the RW benchmark in terms of density forecast accuracy.

[Table 2 about here]

# 5 Conclusion

We use real-time vintages of the NFCI to model the short-term conditional distribution of exchange rate returns. Our in-sample results show that the lower and upper quantiles of the exchange rate return distribution exhibit significant co-movement with current financial conditions. This result is stronger for commodity currencies (i.e., the AUD, CAD, and the NZD) and is only present in the post-2000 period. We also find evidence of out-of-sample predictability of exchange rate returns for all nine currencies and for the three short-term forecast horizons considered (one-day, one-week, and four-weeks ahead). Our results show that the conditional moments of the density forecast exhibit time-varying patterns, with the variance and kurtosis showing the most pronounced changes during and after the 2008-09 financial crisis. We conclude that the NFCI provides valuable information for forecasting exchange rate returns, leading to more accurate density forecasts than the random walk benchmark.

# References

Adrian, T., Boyarchenko, N., and Giannone, D. (2019). Vulnerable Growth. *American Economic Review*, 109(4):1263–1289.

Alessandri, P. and Mumtaz, H. (2017). Financial conditions and density forecasts for us output and inflation. *Review of Economic Dynamics*, 24:66–78.

Amburgey, A. J. and McCracken, M. W. (2023). On the real-time predictive content of financial condition indices for growth. *Journal of Applied Econometrics*, 38(2):137–163.

Andrews, D. W. K. (1991). Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation. *Econometrica*, 59(3):817.

Azzalini, A. and Capitanio, A. (2003). Distributions Generated by Perturbation of Symmetry with Emphasis on a Multivariate Skew $t$ -Distribution. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 65(2):367–389.

Bondatti, M. and Rillo, G. (2022). Commodity tail-risk and exchange rates. *Finance Research Letters*, 47:102937.

Brave, S. and Butters, R. A. (2011). Monitoring financial stability: A financial conditions index approach. *Economic Perspectives*, 35(First Quarter):22–43.

Brave, S. and Butters, R. A. (2012). Diagnosing the financial system: financial conditions and financial stress. *International Journal of Central Banking*, 8(2):191–239.

Byrne, J. P., Korobilis, D., and Ribeiro, P. J. (2018). On the sources of uncertainty in exchange rate predictability. *International Economic Review*, 59(1):329–357.

Cañon, C., Gerba, E., Pambira, A., and Stoja, E. (2024). An unconventional fx tail risk story. *Journal of International Money and Finance*, 148:103152.

Chen, Y.-c. and Rogoff, K. (2003). Commodity currencies. *Journal of international Economics*, 60(1):133–160.

Cheung, Y.-W., Chinn, M. D., Pascual, A. G., and Zhang, Y. (2019). Exchange rate prediction redux: New models, new data, new currencies. *Journal of International Money and Finance*, 95:332–362.

Colombo, E. and Pelagatti, M. (2020). Statistical learning and exchange rate forecasting. *International Journal of Forecasting*, 36(4):1260–1289.

De Nicolò, G. and Lucchetta, M. (2017). Forecasting Tail Risks. *Journal of Applied Econometrics*, 32(1):159–170.

Delle Monache, D., De Polis, A., and Petrella, I. (2024). Modeling and Forecasting Macroeconomic Downside Risk. *Journal of Business & Economic Statistics*, 42(3):1010–1025.

Diebold, F. X., Gunther, T. A., and Tay, A. (1997). Evaluating density forecasts.

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

Eguren-Martin, F. and Sokol, A. (2022). Attention to the tail (s): Global financial conditions and exchange rate risks. *IMF Economic Review*, 70(3):487.

Engel, C. and West, K. (2005). Exchange Rates and Fundamentals. *Journal of Political Economy*, 113(3):485–517.

Engel, C. and Wu, S. P. Y. (2023a). Forecasting the U.S. Dollar in the 21st Century. *Journal of International Economics*, 141:103715.

Engel, C. and Wu, S. P. Y. (2023b). Liquidity and Exchange Rates: An Empirical Investigation. *The Review of Economic Studies*, 90(5):2395–2438.

Engel, C. M. and Wu, S. P. Y. (2024). Exchange Rate Models are Better than You Think, and Why They Didn't Work in the Old Days. *SSRN Electronic Journal*.

Ferrara, L. and Yapi, J. (2022). Measuring exchange rate risks during periods of uncertainty. *International Economics*, 170:202–212.

Gargano, A. and Timmermann, A. (2014). Forecasting commodity price indexes using macroeconomic and financial predictors. *International Journal of Forecasting*, 30(3):825–843.

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

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

Gneiting, T. and Raftery, A. E. (2007). Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102(477):359–378.

Itskhoki, O. and Mukhin, D. (2021). Exchange Rate Disconnect in General Equilibrium. *Journal of Political Economy*, 129(8):2183–2232.

Lilley, A., Maggiori, M., Neiman, B., and Schreger, J. (2022). Exchange Rate Reconnect. *The Review of Economics and Statistics*, 104(4):845–855.

Lyócsa, Š., Plíhal, T., and Vỳrost, T. (2024). Forecasting day-ahead expected shortfall on the eur/usd exchange rate: The (i) relevance of implied volatility. *International Journal of Forecasting*, 40(4):1275–1301.

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

Molodtsova, T. and Papell, D. H. (2009). Out-of-sample exchange rate predictability with taylor rule fundamentals. *Journal of international economics*, 77(2):167–180.

Newey, W. K. and West, K. D. (1987). A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3):703.

Ostry, D. A. (2023). Tails of foreign exchange-at-risk (fear).

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

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

Smales, L. A. (2022). The influence of policy uncertainty on exchange rate forecasting. *Journal of Forecasting*, 41(5):997–1016.

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

Welch, I. and Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies*, 21(4):1455–1508.

**Tables and Figures**

Table 1: OOS one-week ahead forecast performance with real-time NFCI.

| | AUD | CAD | CHF | EUR | GBP | JPY | NOK | NZD | SEK |
|---|---|---|---|---|---|---|---|---|---|
| **A: RMSE** | | | | | | | | | |
| RW | 1.69 | 1.21 | 1.47 | 1.33 | 1.33 | 1.31 | 1.66 | 1.75 | 1.53 |
| | – | – | – | – | – | – | – | – | – |
| RW w/ drift | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.89) | (0.91) | (0.67) | (0.92) | (0.86) | (0.96) | (0.87) | (0.81) | (0.88) |
| OLS | 1.01 | 1.01 | 1.00 | 1.01 | 1.00 | 1.00 | 1.01 | 1.00 | 1.00 |
| | (0.82) | (0.78) | (0.90) | (0.82) | (0.60) | (0.85) | (0.77) | (0.73) | (0.74) |
| QR-skewed-$t$ | 1.00 | 1.01 | 1.01 | 1.01 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.35) | (0.86) | (0.97) | (0.76) | (0.40) | (0.92) | (0.49) | (0.67) | (0.70) |
| **B: log Score** | | | | | | | | | |
| RW | 1.97 | 1.65 | 1.82 | 1.69 | 1.73 | 1.70 | 1.95 | 2.01 | 1.85 |
| | – | – | – | – | – | – | – | – | – |
| RW w/ drift | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.86) | (0.87) | (0.59) | (0.92) | (0.87) | (0.95) | (0.87) | (0.77) | (0.91) |
| OLS | 1.00 | 1.01 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.78) | (0.80) | (0.88) | (0.82) | (0.58) | (0.80) | (0.76) | (0.69) | (0.73) |
| QR-skewed-$t$ | **0.93** | **0.97** | **0.95** | **0.97** | **0.96** | **0.98** | **0.95** | **0.96** | **0.97** |
| | (0.00) | (0.03) | (0.01) | (0.01) | (0.01) | (0.02) | (0.01) | (0.00) | (0.00) |

Notes: This table presents the forecast performance of the one-week ahead out-of-sample forecasts of exchange rate returns with the real-time NFCI and a 10 year rolling window. Performance is evaluated using the RMSE and the log score in the 2000/01/05 – 2021/07/21 test set. For the random walk (RW) benchmark, the RMSE and log score are reported. For the alternative forecasts (RW with drift, OLS, and QR-skewed-$t$), the RMSE and log score are reported relative to the RW benchmark. A value less than 1 indicates that the alternative model has better predictive accuracy than the RW benchmark. The numbers in parentheses are the p-values of the one-sided robust Diebold-Mariano test for the null hypothesis of equal predictive accuracy between the RW benchmark and the alternative model. The alternative hypothesis is that the alternative model has better predictive accuracy than the RW benchmark. Newey-West standard errors with $h-1$ lags are applied.

Table 2: OOS one-week ahead forecast performance with real-time NFCI and IRDs.

| | AUD | CAD | CHF | EUR | GBP | JPY |
|---|---|---|---|---|---|---|
| **A: log Score without IRD** | | | | | | |
| RW | 2.13 | 1.87 | 2.00 | 1.80 | 1.89 | 1.83 |
| | – | – | – | – | – | – |
| RW w/ drift | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.89) | (0.95) | (0.58) | (0.92) | (0.87) | (0.91) |
| OLS | 1.01 | 1.01 | 1.00 | 1.01 | 1.00 | 1.00 |
| | (0.81) | (0.72) | (0.85) | (0.82) | (0.57) | (0.70) |
| QR-skewed-$t$ | **0.91** | **0.92** | **0.93** | **0.98** | **0.94** | 1.01 |
| | (0.02) | (0.00) | (0.07) | (0.07) | (0.05) | (0.63) |
| **B: log Score with IRD** | | | | | | |
| RW | 2.13 | 1.87 | 2.00 | 1.80 | 1.89 | 1.83 |
| | – | – | – | – | – | – |
| RW w/ drift | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.89) | (0.95) | (0.58) | (0.92) | (0.87) | (0.91) |
| OLS | 1.01 | 1.01 | 1.00 | 1.01 | 1.00 | 1.00 |
| | (0.88) | (0.70) | (0.87) | (0.83) | (0.61) | (0.73) |
| QR-skewed-$t$ | **0.93** | **0.92** | **0.92** | **0.98** | **0.95** | 1.02 |
| | (0.04) | (0.00) | (0.05) | (0.08) | (0.07) | (0.70) |

Notes: This table presents the forecast performance of the one-week ahead out-of-sample forecasts of exchange rate returns with the real-time NFCI, interest rate differentials (IRD), and a 10 year rolling window. Performance is evaluated using the log score in the 2008/01/02 – 2017/12/20 test set. For the random walk (RW) benchmark, the log score is reported. For the alternative forecasts (RW with drift, OLS, and QR-skewed-$t$), the log score is reported relative to the RW benchmark. A value less than 1 indicates that the alternative model has better predictive accuracy than the RW benchmark. The numbers in parentheses are the p-values of the one-sided robust Diebold-Mariano test for the null hypothesis of equal predictive accuracy between the RW benchmark and the alternative model. The alternative hypothesis is that the alternative model has better predictive accuracy than the RW benchmark. Newey-West standard errors with $h-1$ lags are applied.

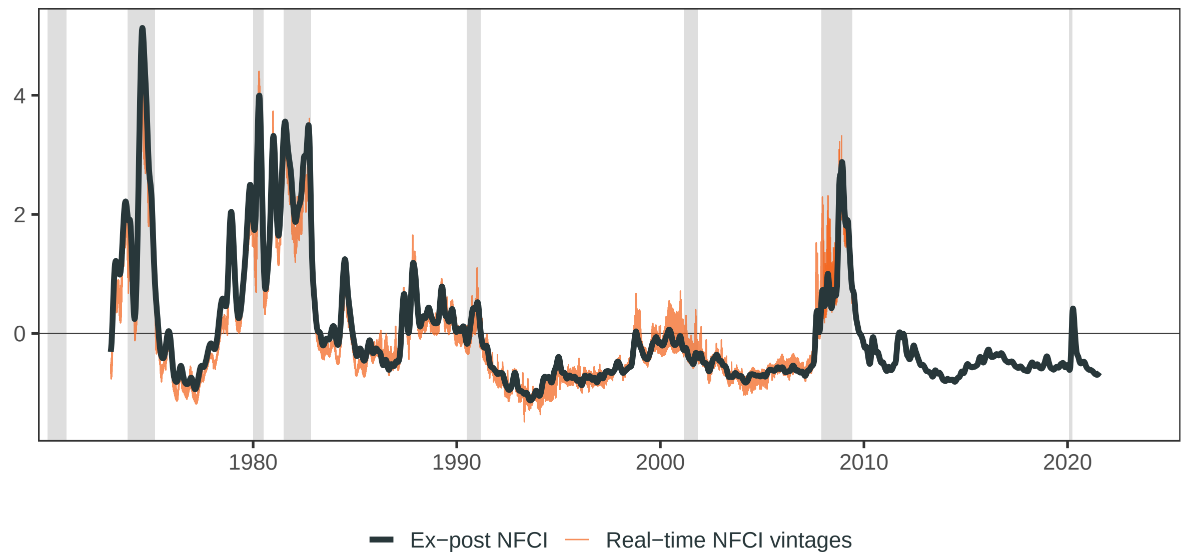


Figure 1: Time series plot of the 2024/10/16 release of the ex-post NFCI and the real-time NFCI vintages from Amburgey and McCracken (2023) corresponding to the years 2008 to 2009. Shaded areas denote NBER recessions.

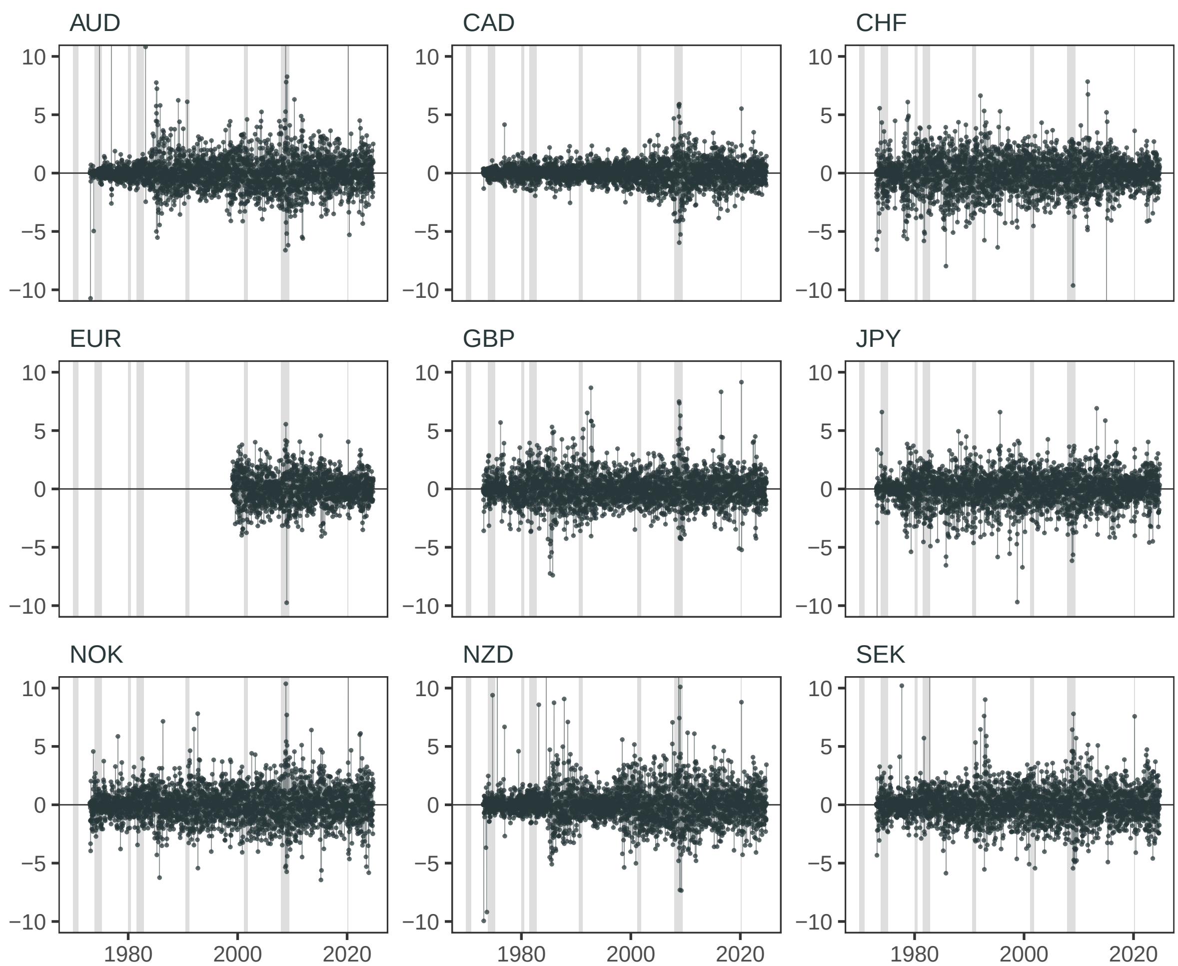


Figure 2: Time series plot of the weekly exchange rate returns (100 times the log difference of the exchange rate on the Wednesday of each week). Shaded areas denote NBER recessions.

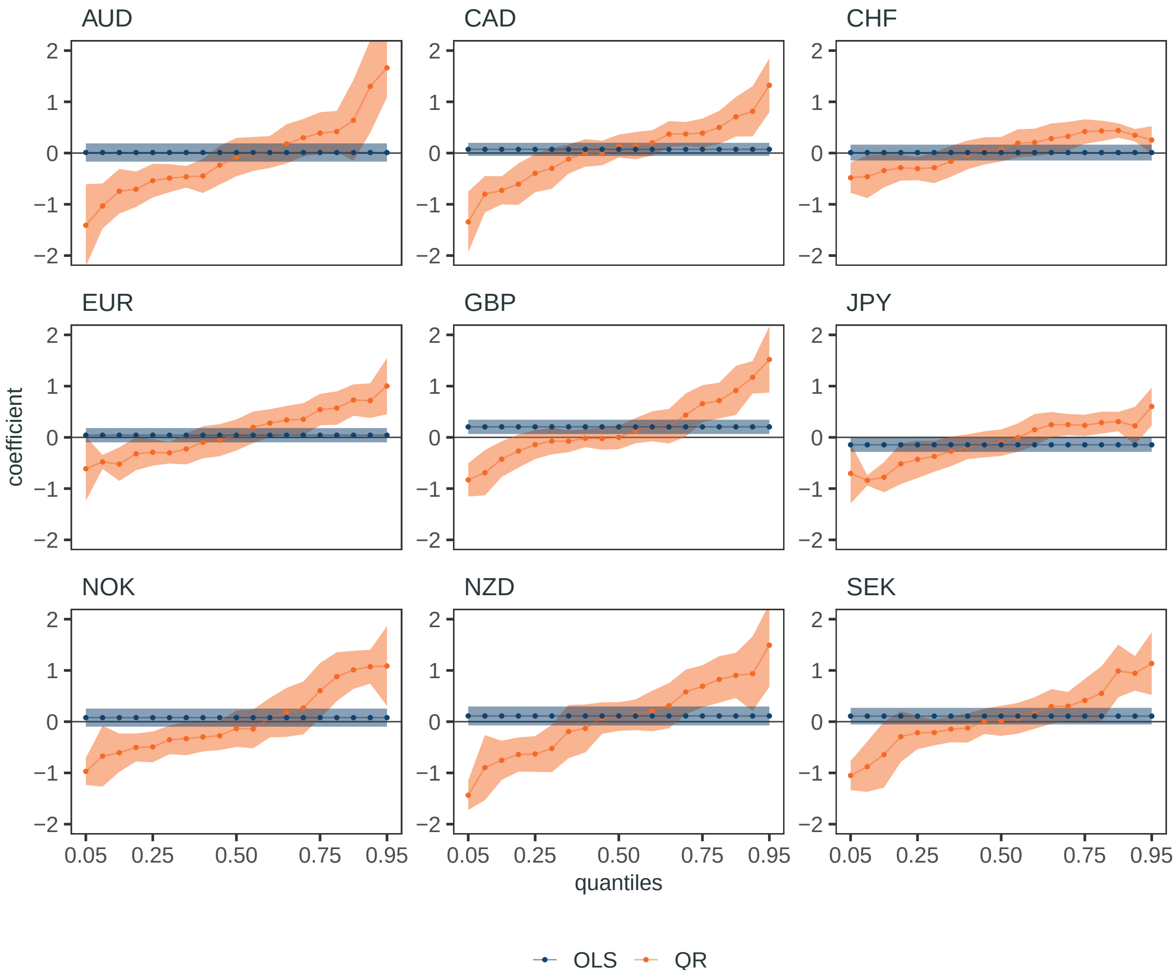


Figure 3: In-sample OLS and QR coefficients for the one week ahead exchange rate returns on the ex-post NFCI. Shaded areas denote 95% confidence intervals. Sample period: 2000/01/01 – 2021/12/31.

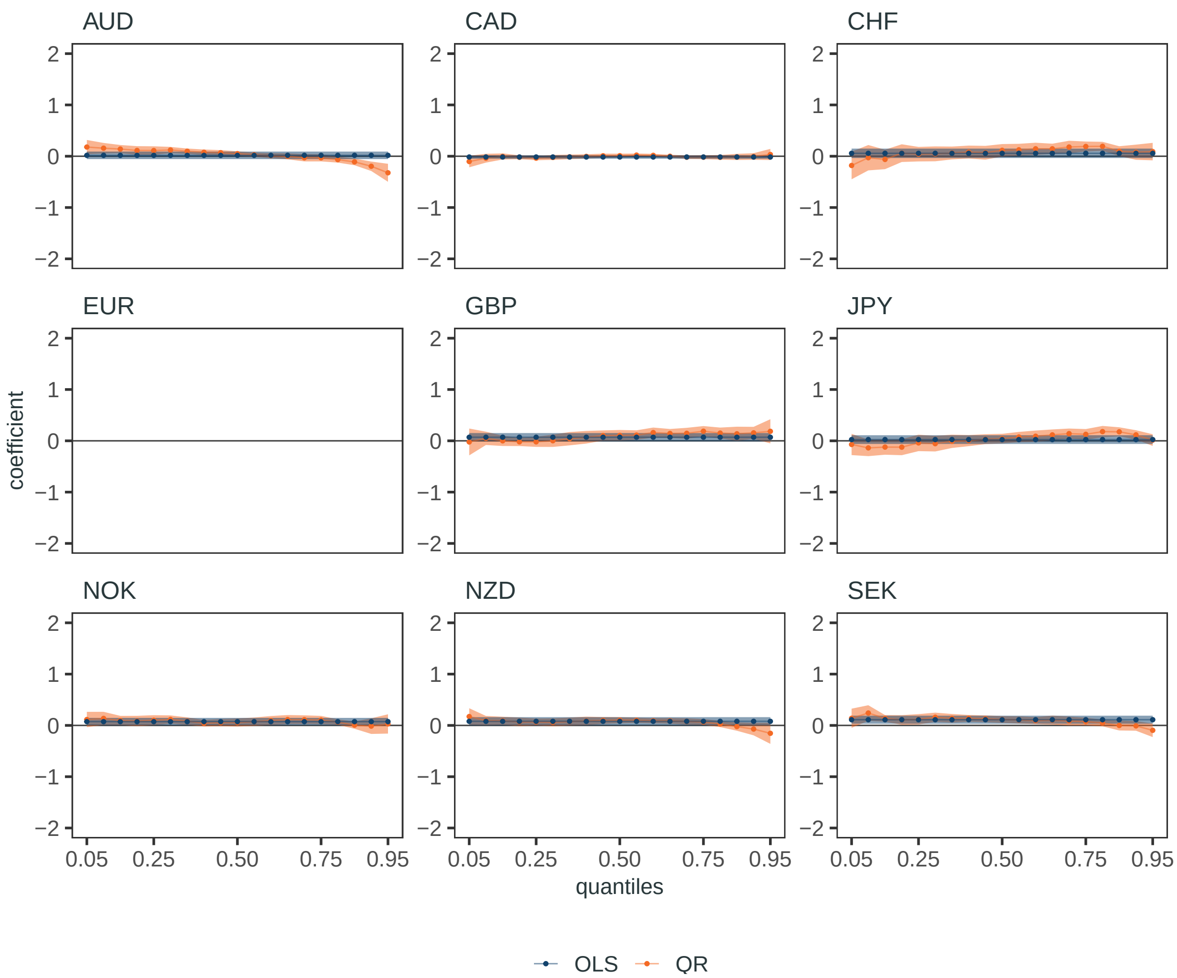


Figure 4: In-sample OLS and QR coefficients for the one week ahead exchange rate returns on the ex-post NFCI. Shaded areas denote 95% confidence intervals. Sample period: 1980/01/01 – 1999/12/31.

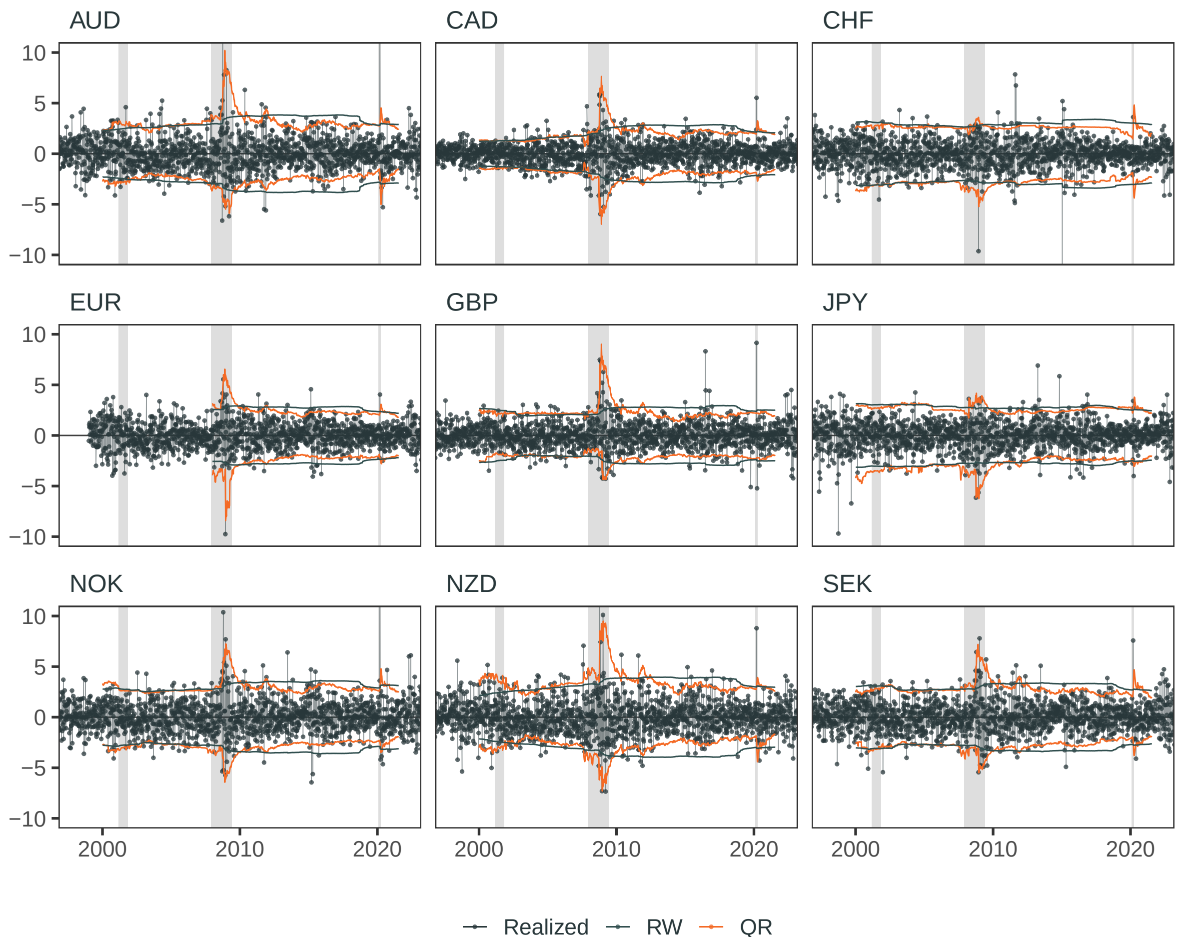


Figure 5: Time series plot of the realized returns and the one week ahead out-of-sample 95% prediction intervals of exchange rate returns with the real-time NFCI and a 10 year rolling window. Results shown for the random walk (RW) benchmark and the QR-skewed-$t$ (QR) density forecasts. Shaded areas denote NBER recessions.

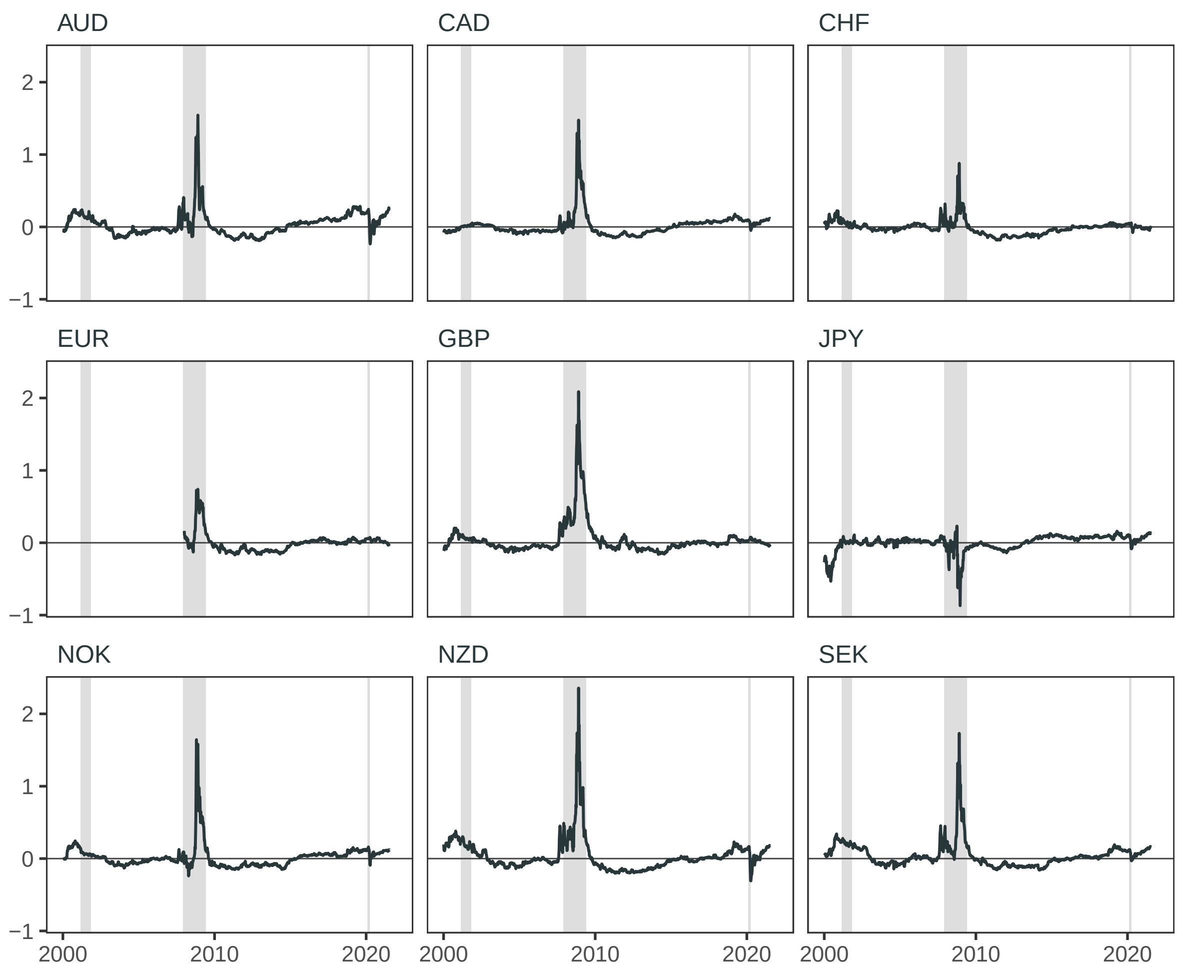


Figure 6: Time series plot of the conditional mean of the one week ahead out-of-sample density forecasts of exchange rate returns with the real-time NFCI and a 10 year rolling window. Shaded areas denote NBER recessions.

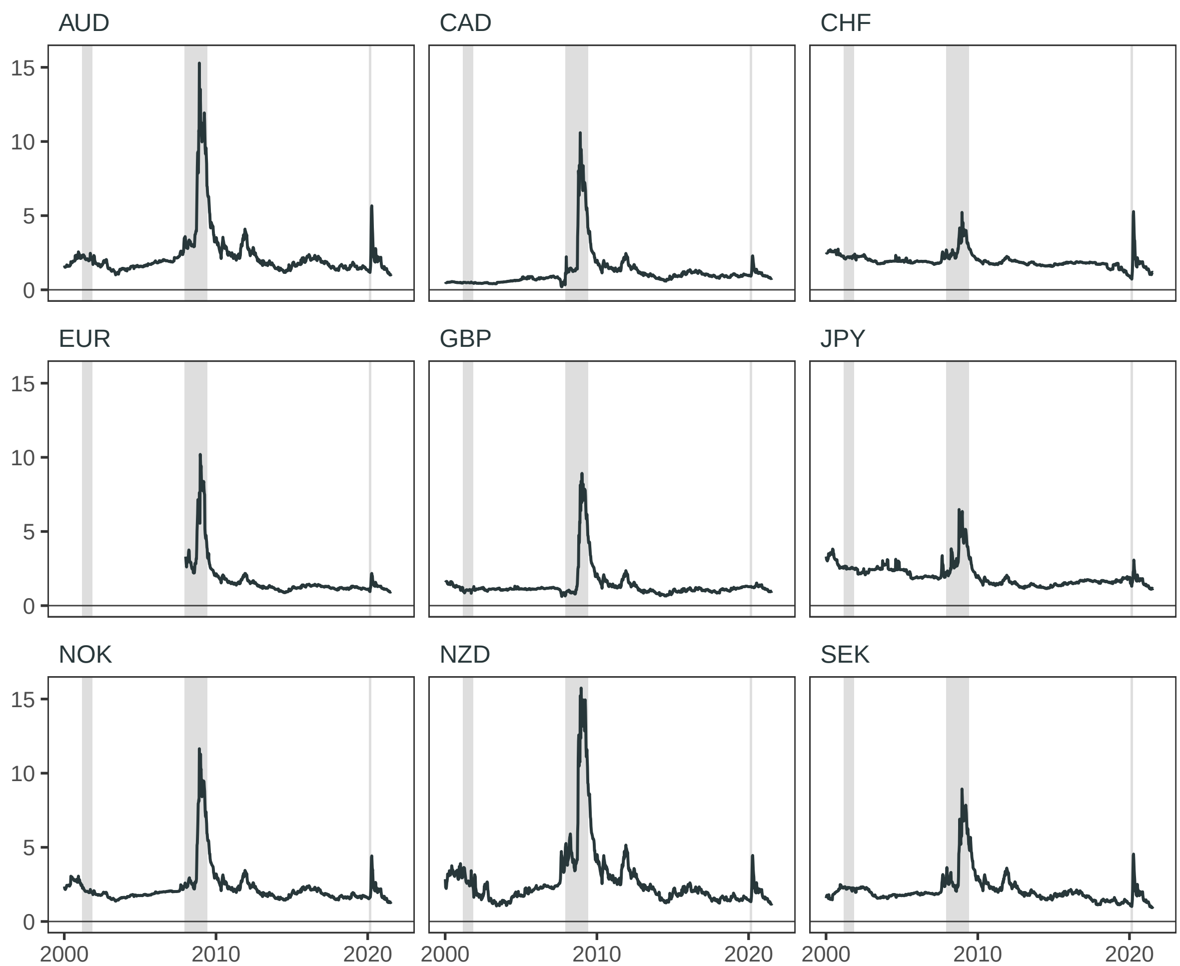


Figure 7: Time series plot of the conditional variance of the one week ahead out-of-sample density forecasts of exchange rate returns with the real-time NFCI and a 10 year rolling window. Shaded areas denote NBER recessions.


Figure 8: Cumulative log Score differences (relative to the RW benchmark) of the one week ahead out-of-sample density forecasts of exchange rate returns with the real-time NFCI and a 10 year rolling window. Shaded areas denote NBER recessions.

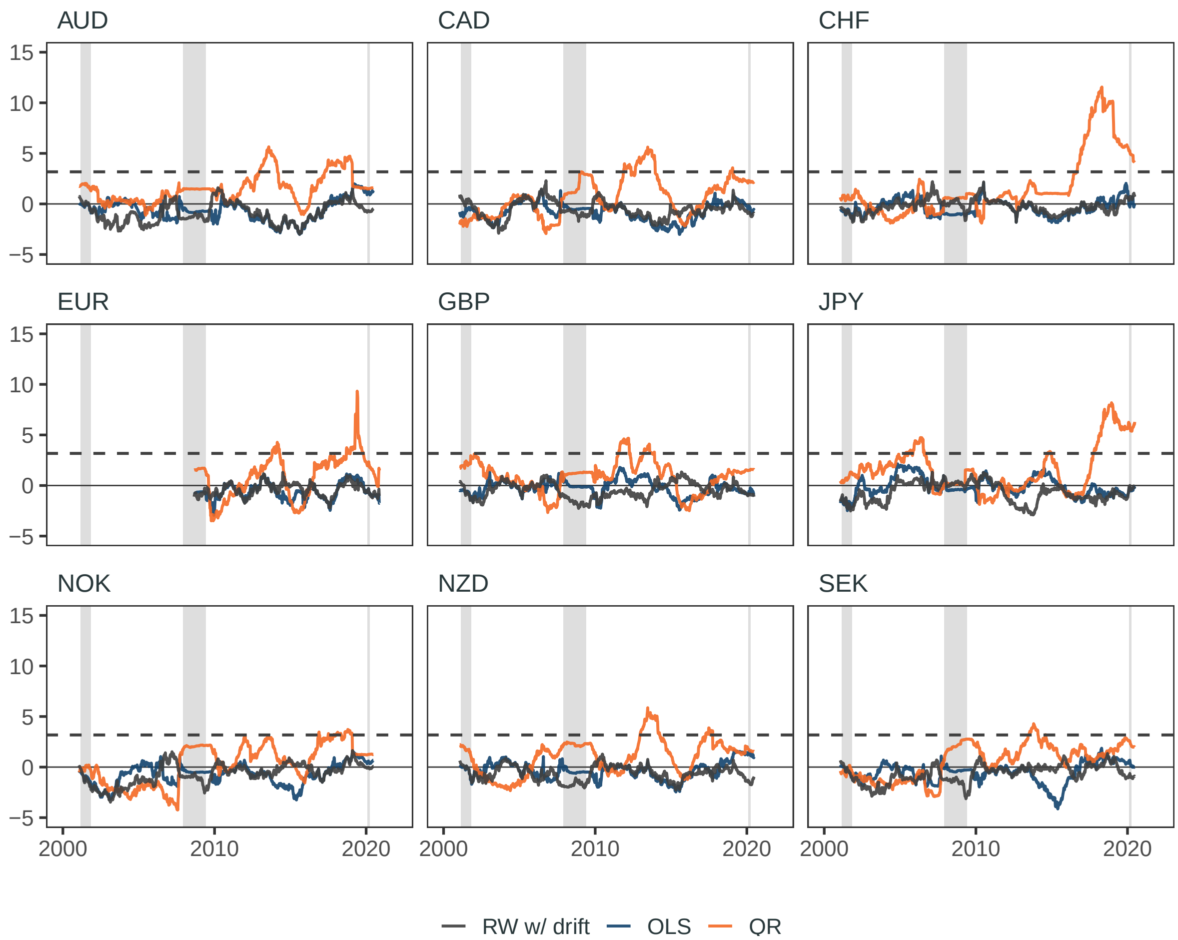


Figure 9: Fluctuation tests for the log Score of the one week ahead out-of-sample density forecasts of exchange rate returns with the real-time NFCI and a 10 year rolling window. Positive values of the Fluctuation statistic imply that the alternative model is better than the random walk (RW) benchmark. The dashed line represents the 95% critical values for a one-sided test of the null hypothesis of equal predictive accuracy between the RW benchmark and the alternative model. Shaded areas denote NBER recessions. Test set: 2000/01/05 – 2021/07/21.

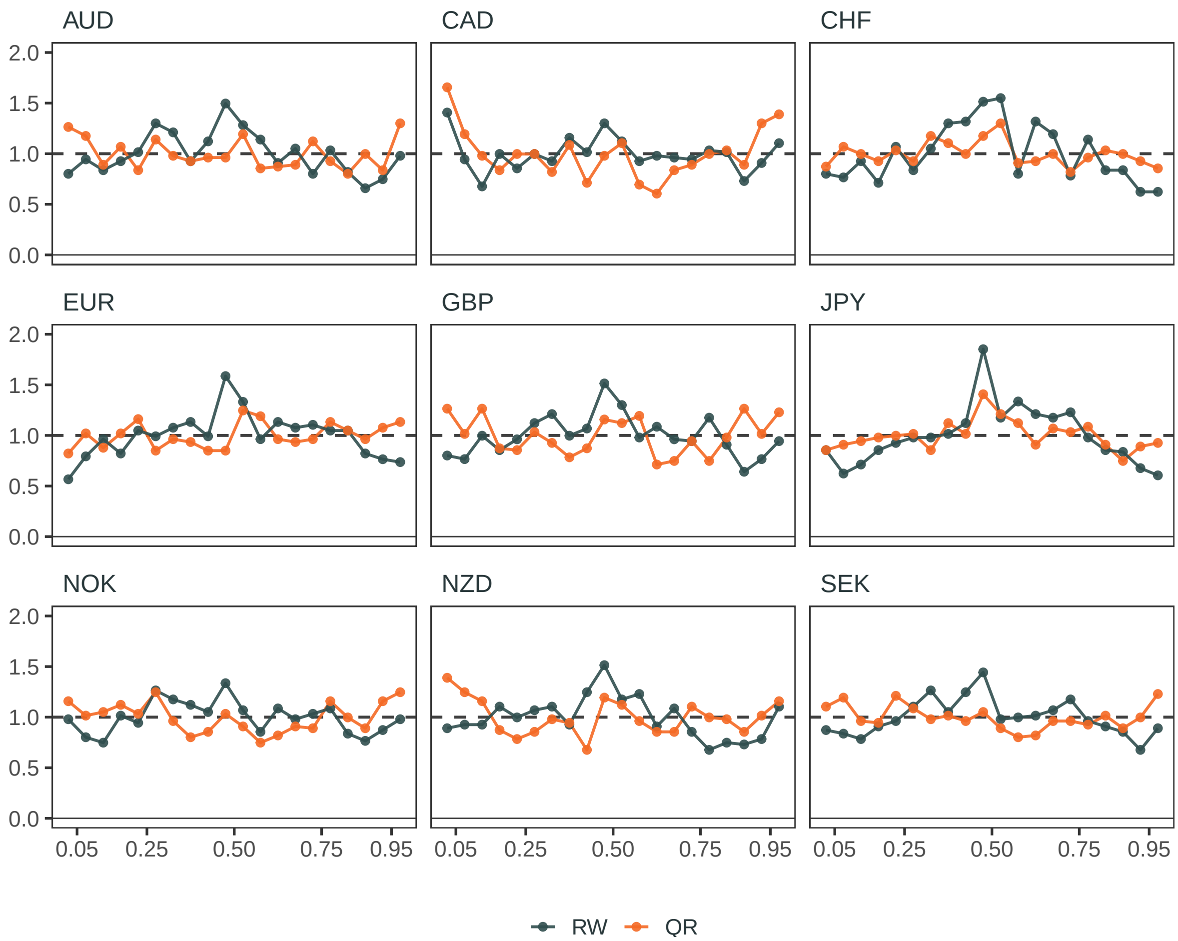


Figure 10: Probability integral transform (PIT) histograms of the one week ahead out-of-sample density forecasts of exchange rate returns with the real-time NFCI and a 10 year rolling window. Results shown for the random walk (RW) benchmark and the QR-skewed-$t$ (QR) density forecasts. The dashed line represents a perfectly calibrated density forecast (i.e., the histogram of a uniform distribution). Test set: 2000/01/05 – 2021/07/21.

# Online Appendix

Table A.1: OOS one-week ahead forecast performance with real-time NFCI and a 20 year rolling window.

| | AUD | CAD | CHF | EUR | GBP | JPY | NOK | NZD | SEK |
|---|---|---|---|---|---|---|---|---|---|
| **A: RMSE** | | | | | | | | | |
| RW | 1.69 | 1.21 | 1.47 | 1.33 | 1.33 | 1.31 | 1.66 | 1.75 | 1.53 |
| | – | – | – | – | – | – | – | – | – |
| RW w/ drift | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.88) | (0.97) | (0.46) | (0.95) | (0.62) | (0.85) | (0.80) | (0.77) | (0.70) |
| OLS | 1.00 | 1.00 | 1.00 | 1.01 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.80) | (0.88) | (0.76) | (0.80) | (0.53) | (0.66) | (0.78) | (0.70) | (0.69) |
| QR-skewed-$t$ | 1.00 | 1.00 | 1.01 | 1.01 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.14) | (0.94) | (0.93) | (0.74) | (0.36) | (0.86) | (0.19) | (0.62) | (0.54) |
| **B: log Score** | | | | | | | | | |
| RW | 1.98 | 1.76 | 1.81 | 1.70 | 1.71 | 1.70 | 1.94 | 2.03 | 1.84 |
| | – | – | – | – | – | – | – | – | – |
| RW w/ drift | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.82) | (0.88) | (0.45) | (0.95) | (0.66) | (0.85) | (0.77) | (0.72) | (0.66) |
| OLS | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.78) | (0.89) | (0.76) | (0.80) | (0.55) | (0.61) | (0.77) | (0.69) | (0.68) |
| QR-skewed-$t$ | **0.93** | **0.92** | **0.95** | **0.97** | **0.95** | **0.97** | **0.96** | **0.95** | **0.98** |
| | (0.01) | (0.00) | (0.01) | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) |

Notes: This table presents the forecast performance of the one-week ahead out-of-sample forecasts of exchange rate returns with the real-time NFCI and a 20 year rolling window. Performance is evaluated using the RMSE and the log score in the 2000/01/05 – 2021/07/21 test set. For the random walk (RW) benchmark, the RMSE and log score are reported. For the alternative forecasts (RW with drift, OLS, and QR-skewed-$t$), the RMSE and log score are reported relative to the RW benchmark. A value less than 1 indicates that the alternative model has better predictive accuracy than the RW benchmark. The numbers in parentheses are the p-values of the one-sided robust Diebold-Mariano test for the null hypothesis of equal predictive accuracy between the RW benchmark and the alternative model. The alternative hypothesis is that the alternative model has better predictive accuracy than the RW benchmark. Newey-West standard errors with $h-1$ lags are applied.

Table A.2: OOS one-week ahead forecast performance with real-time NFCI and an expanding estimation window.

| | AUD | CAD | CHF | EUR | GBP | JPY | NOK | NZD | SEK |
|---|---|---|---|---|---|---|---|---|---|
| **A: RMSE** | | | | | | | | | |
| RW | 1.69 | 1.21 | 1.47 | 1.33 | 1.33 | 1.31 | 1.66 | 1.75 | 1.53 |
| | – | – | – | – | – | – | – | – | – |
| RW w/ drift | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.88) | (0.98) | (0.53) | (0.96) | (0.70) | (0.86) | (0.80) | (0.92) | (0.80) |
| OLS | 1.00 | 1.00 | 1.00 | 1.01 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.79) | (0.92) | (0.62) | (0.80) | (0.31) | (0.81) | (0.63) | (0.54) | (0.47) |
| QR-skewed-$t$ | 1.00 | 1.00 | 1.00 | 1.01 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.68) | (0.79) | (0.83) | (0.74) | (0.23) | (0.88) | (0.46) | (0.51) | (0.45) |
| **B: log Score** | | | | | | | | | |
| RW | 1.97 | 1.82 | 1.81 | 1.70 | 1.71 | 1.70 | 1.95 | 1.99 | 1.85 |
| | – | – | – | – | – | – | – | – | – |
| RW w/ drift | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| (0.86) | (0.96) | (0.51) | (0.96) | (0.70) | (0.86) | (0.81) | (0.91) | (0.80) | |
| OLS | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.78) | (0.95) | (0.59) | (0.80) | (0.28) | (0.80) | (0.67) | (0.56) | (0.49) |
| QR-skewed-$t$ | 0.96 | **0.92** | **0.96** | **0.97** | **0.95** | **0.97** | **0.97** | 0.98 | **0.99** |
| | (0.12) | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.04) | (0.13) | (0.05) |

Notes: This table presents the forecast performance of the one-week ahead out-of-sample forecasts of exchange rate returns with the real-time NFCI and an expanding estimation window. Performance is evaluated using the RMSE and the log score in the 2000/01/05 – 2021/07/21 test set. For the random walk (RW) benchmark, the RMSE and log score are reported. For the alternative forecasts (RW with drift, OLS, and QR-skewed-$t$), the RMSE and log score are reported relative to the RW benchmark. A value less than 1 indicates that the alternative model has better predictive accuracy than the RW benchmark. The numbers in parentheses are the p-values of the one-sided robust Diebold-Mariano test for the null hypothesis of equal predictive accuracy between the RW benchmark and the alternative model. The alternative hypothesis is that the alternative model has better predictive accuracy than the RW benchmark. Newey-West standard errors with $h-1$ lags are applied.

Table A.3: OOS one-day ahead forecast performance with real-time NFCI and a 10 year rolling window.

| | AUD | CAD | CHF | EUR | GBP | JPY | NOK | NZD | SEK |
|---|---|---|---|---|---|---|---|---|---|
| **A: RMSE** | | | | | | | | | |
| RW | 0.81 | 0.58 | 0.76 | 0.67 | 0.58 | 0.64 | 0.83 | 0.90 | 0.75 |
| | – | – | – | – | – | – | – | – | – |
| RW w/ drift | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.66) | (1.00) | (0.68) | (0.85) | (0.75) | (0.70) | (0.94) | (0.87) | (0.87) |
| OLS | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.01 | 1.00 | 1.00 | 1.00 |
| | (0.92) | (0.73) | (0.92) | (0.83) | (0.83) | (0.88) | (0.80) | (0.82) | (0.92) |
| QR-skewed-$t$ | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.58) | (0.92) | (0.80) | (0.93) | (0.48) | (0.60) | (0.25) | (0.69) | (0.60) |
| **B: log Score** | | | | | | | | | |
| RW | 1.24 | 0.91 | 1.20 | 1.03 | 0.89 | 0.98 | 1.25 | 1.36 | 1.14 |
| | – | – | – | – | – | – | – | – | – |
| RW w/ drift | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.63) | (0.99) | (0.68) | (0.87) | (0.71) | (0.68) | (0.94) | (0.81) | (0.90) |
| OLS | 1.01 | 1.01 | 1.00 | 1.00 | 1.01 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.93) | (0.75) | (0.91) | (0.83) | (0.81) | (0.83) | (0.84) | (0.84) | (0.94) |
| QR-skewed-$t$ | 0.96 | **0.90** | **0.82** | **0.93** | **0.94** | **0.93** | **0.94** | **0.93** | **0.96** |
| | (0.13) | (0.00) | (0.07) | (0.01) | (0.01) | (0.00) | (0.02) | (0.00) | (0.02) |

Notes: This table presents the forecast performance of the one-day ahead out-of-sample forecasts of exchange rate returns with the real-time NFCI and a 10 year rolling window. Performance is evaluated using the RMSE and the log score in the 2000/01/05 – 2021/07/21 test set. For the random walk (RW) benchmark, the RMSE and log score are reported. For the alternative forecasts (RW with drift, OLS, and QR-skewed-$t$), the RMSE and log score are reported relative to the RW benchmark. A value less than 1 indicates that the alternative model has better predictive accuracy than the RW benchmark. The numbers in parentheses are the p-values of the one-sided robust Diebold-Mariano test for the null hypothesis of equal predictive accuracy between the RW benchmark and the alternative model. The alternative hypothesis is that the alternative model has better predictive accuracy than the RW benchmark. Newey-West standard errors with $h-1$ lags are applied.

Table A.4: OOS four-week ahead forecast performance with real-time NFCI and a 10 year rolling window.

| | AUD | CAD | CHF | EUR | GBP | JPY | NOK | NZD | SEK |
|---|---|---|---|---|---|---|---|---|---|
| **A: RMSE** | | | | | | | | | |
| RW | 3.35 | 2.33 | 2.87 | 2.67 | 2.55 | 2.59 | 3.20 | 3.45 | 3.04 |
| | – | – | – | – | – | – | – | – | – |
| RW w/ drift | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.66) | (0.74) | (0.48) | (0.71) | (0.53) | (0.71) | (0.63) | (0.53) | (0.66) |
| OLS | 1.00 | 1.01 | 1.01 | 1.01 | 1.00 | 1.00 | 1.00 | 1.00 | 1.01 |
| | (0.59) | (0.74) | (0.64) | (0.63) | (0.48) | (0.58) | (0.60) | (0.55) | (0.58) |
| QR-skewed-$t$ | 1.00 | 1.01 | 1.03 | 1.02 | 0.99 | 1.01 | 1.00 | 1.01 | 1.01 |
| | (0.56) | (0.71) | (0.78) | (0.73) | (0.39) | (0.66) | (0.52) | (0.64) | (0.61) |
| **B: log Score** | | | | | | | | | |
| RW | 2.66 | 2.32 | 2.48 | 2.39 | 2.38 | 2.38 | 2.60 | 2.70 | 2.54 |
| | – | – | – | – | – | – | – | – | – |
| RW w/ drift | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.61) | (0.68) | (0.43) | (0.72) | (0.54) | (0.70) | (0.64) | (0.49) | (0.71) |
| OLS | 1.00 | 1.01 | 1.00 | 1.01 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | (0.54) | (0.76) | (0.63) | (0.64) | (0.47) | (0.52) | (0.59) | (0.50) | (0.59) |
| QR-skewed-$t$ | **0.96** | 1.00 | **0.98** | **0.97** | 0.98 | **0.99** | 0.99 | **0.97** | 0.99 |
| | (0.03) | (0.52) | (0.03) | (0.01) | (0.14) | (0.02) | (0.14) | (0.01) | (0.11) |

Notes: This table presents the forecast performance of the four-week ahead out-of-sample forecasts of exchange rate returns with the real-time NFCI and a 10 year rolling window. Performance is evaluated using the RMSE and the log score in the 2000/01/05 – 2021/07/21 test set. For the random walk (RW) benchmark, the RMSE and log score are reported. For the alternative forecasts (RW with drift, OLS, and QR-skewed-$t$), the RMSE and log score are reported relative to the RW benchmark. A value less than 1 indicates that the alternative model has better predictive accuracy than the RW benchmark. The numbers in parentheses are the p-values of the one-sided robust Diebold-Mariano test for the null hypothesis of equal predictive accuracy between the RW benchmark and the alternative model. The alternative hypothesis is that the alternative model has better predictive accuracy than the RW benchmark. Newey-West standard errors with $h-1$ lags are applied.

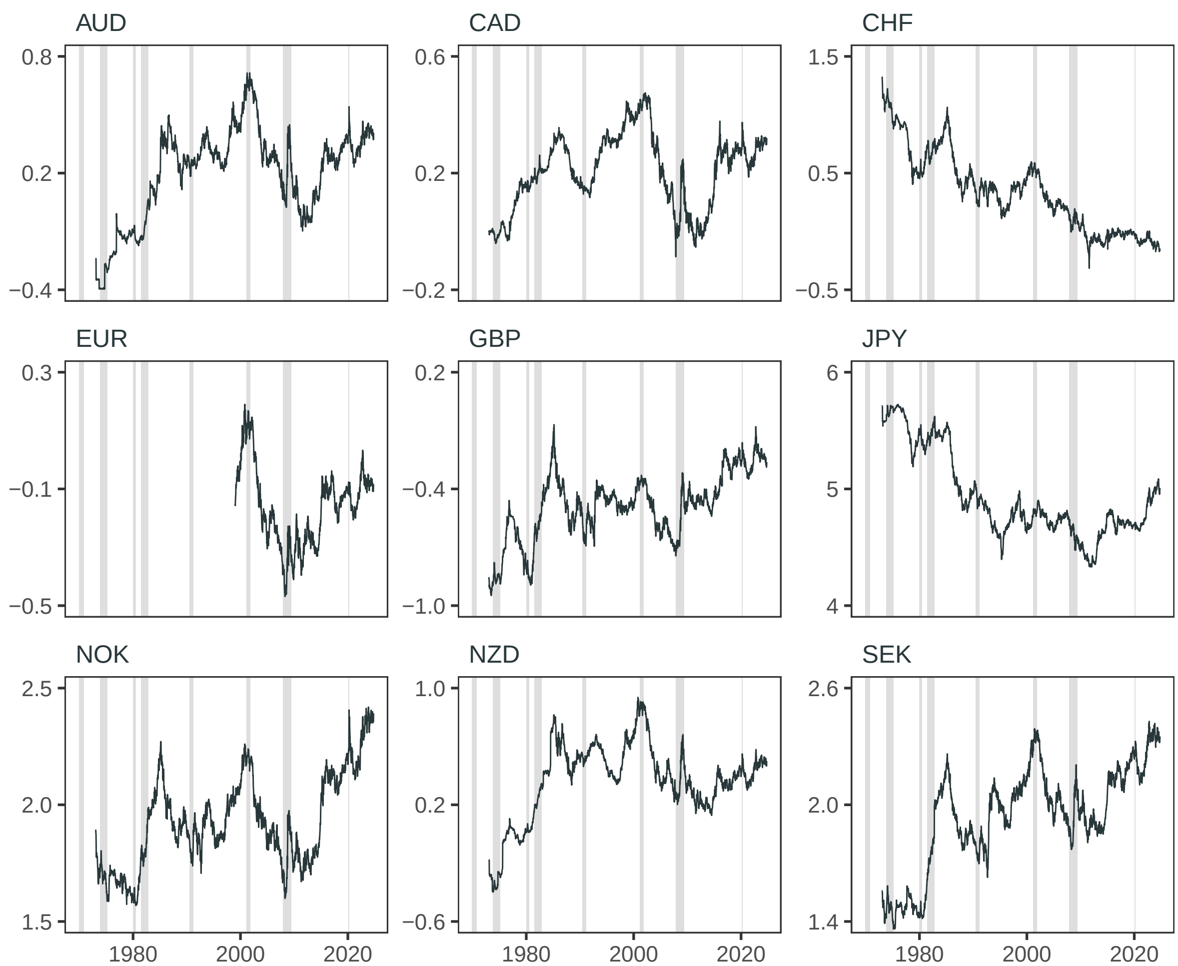


Figure A.1: Time series plot of the exchange rates (log of foreign currency to one U.S. Dollar on the Wednesday of each week). Shaded areas denote NBER recessions. Source: FRED, Federal Reserve Bank of St. Louis.

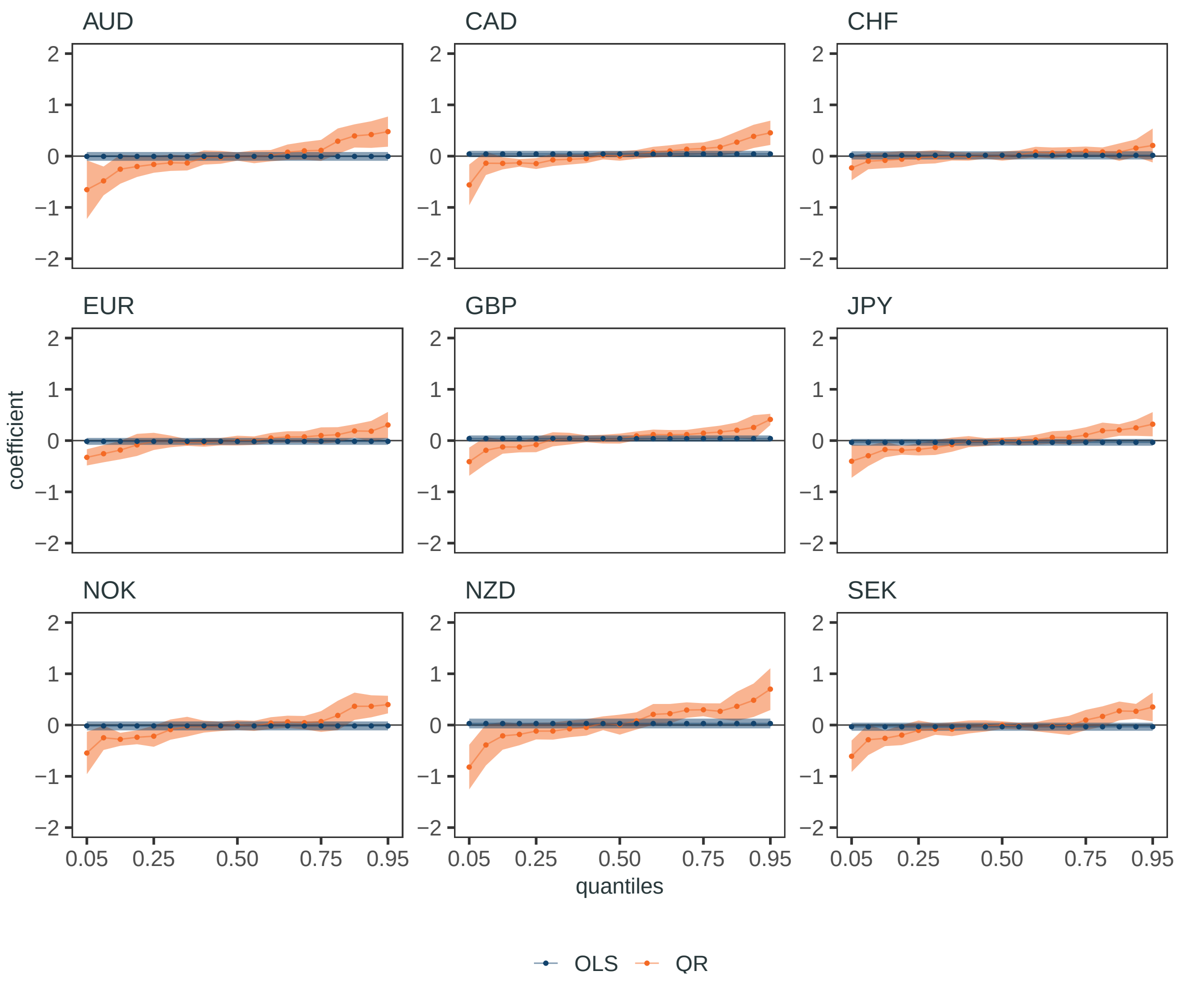


Figure A.2: In-sample OLS and QR coefficients for the one day ahead exchange rate returns on the ex-post NFCI. Shaded areas denote 95% confidence intervals. Sample period: 2000/01/01 – 2021/12/31.

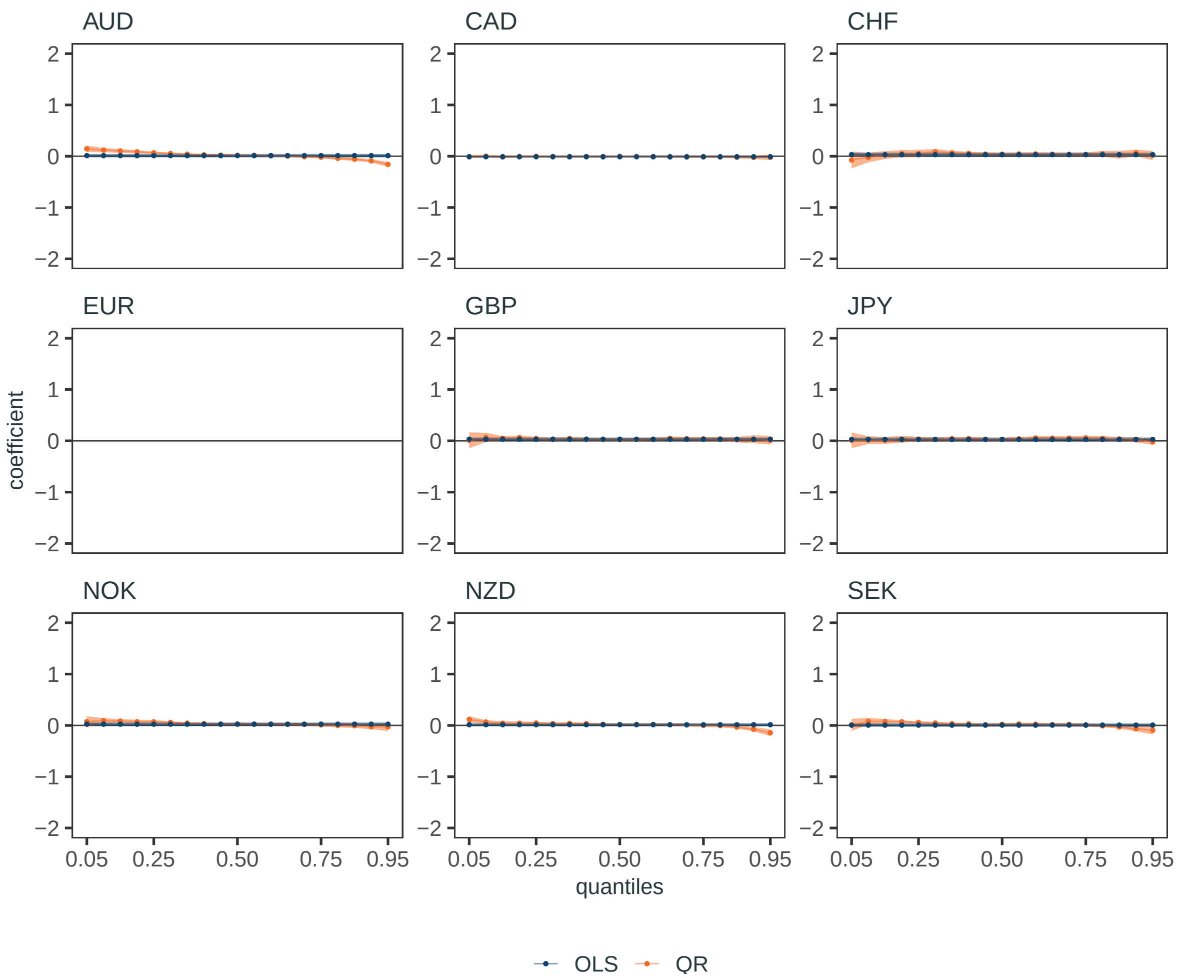


Figure A.3: In-sample OLS and QR coefficients for the one day ahead exchange rate returns on the ex-post NFCI. Shaded areas denote 95% confidence intervals. Sample period: 1980/01/01 – 1999/12/31.

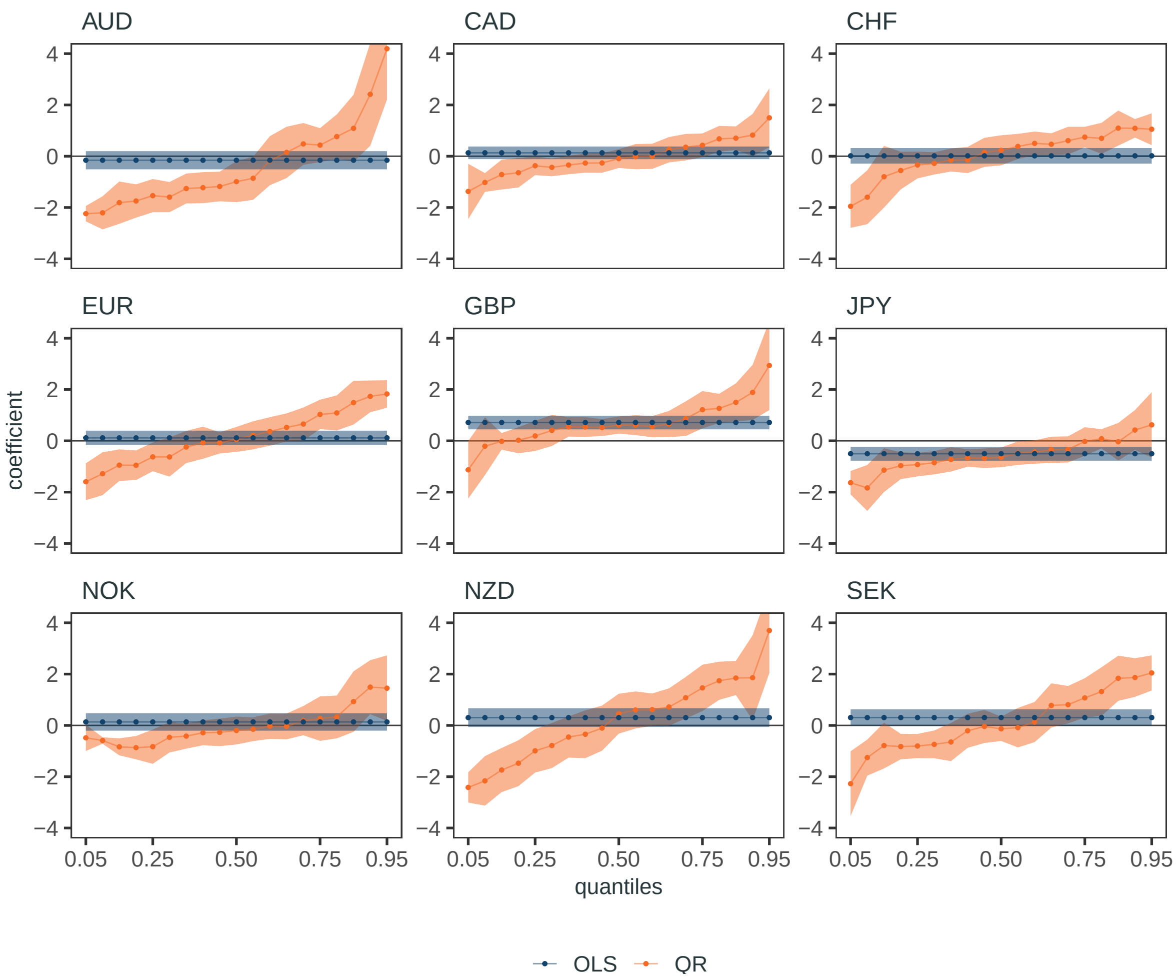


Figure A.4: In-sample OLS and QR coefficients for the four weeks ahead exchange rate returns on the ex-post NFCI. Shaded areas denote 95% confidence intervals. Sample period: 2000/01/01 – 2021/12/31.

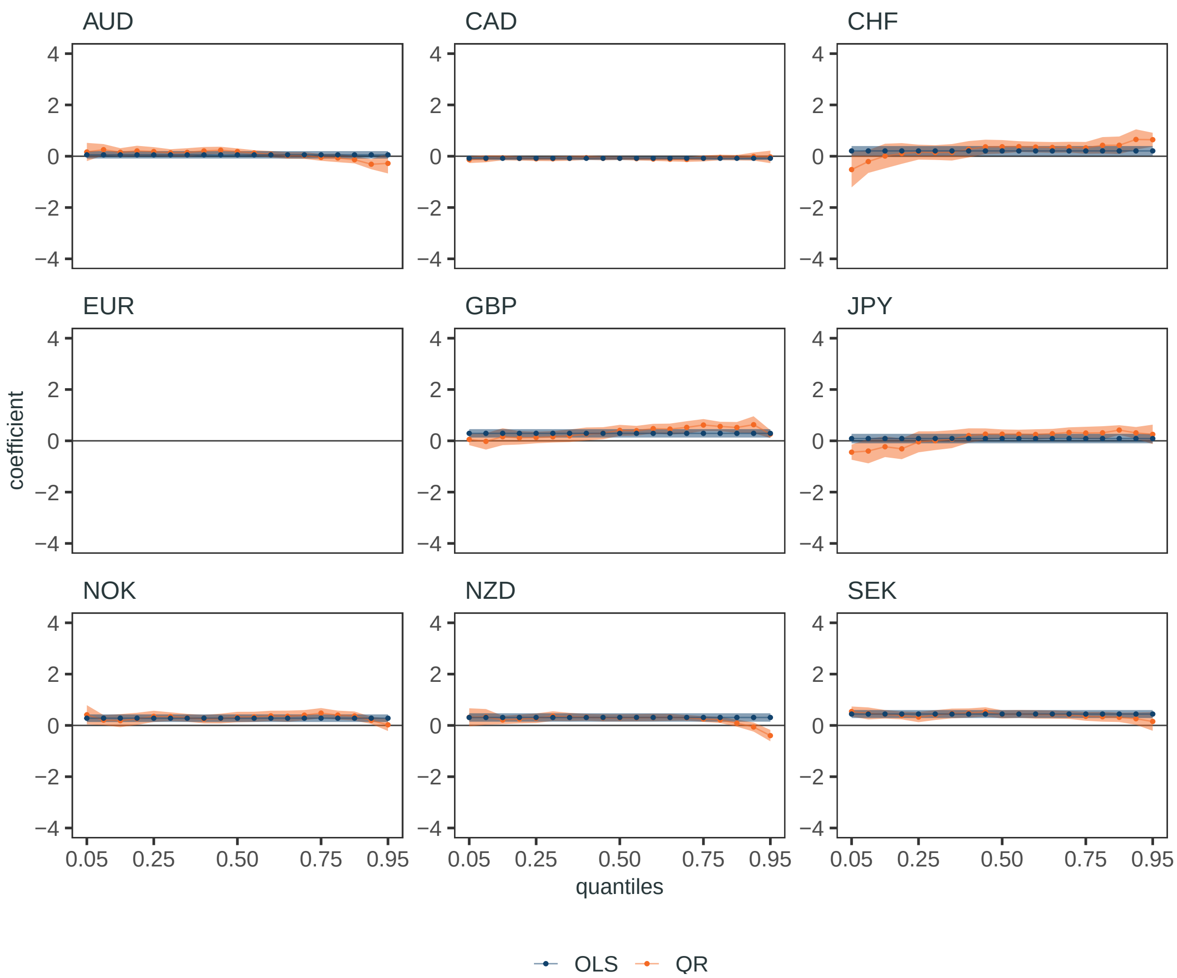


Figure A.5: In-sample OLS and QR coefficients for the four weeks ahead exchange rate returns on the ex-post NFCI. Shaded areas denote 95% confidence intervals. Sample period: 1980/01/01 – 1999/12/31.

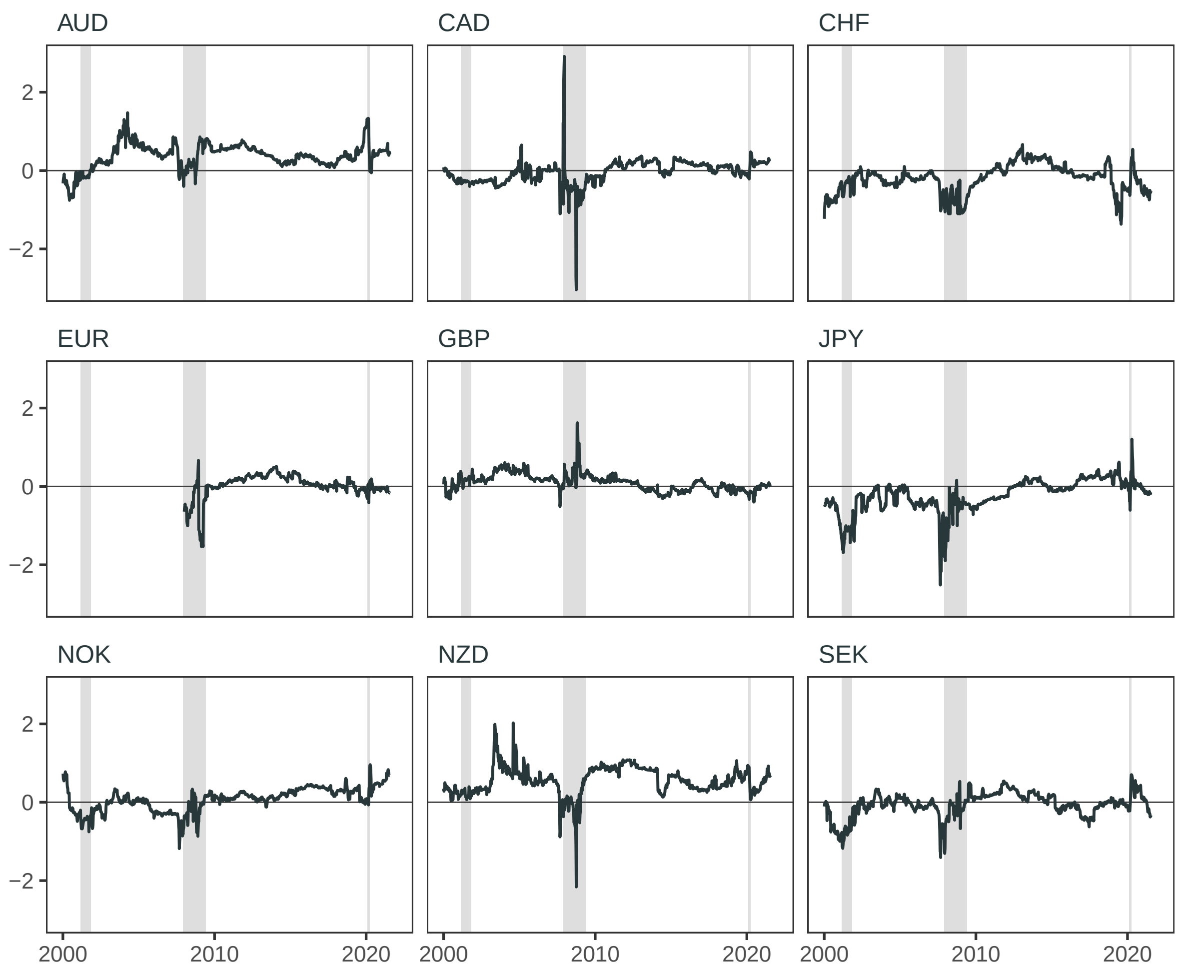


Figure A.6: Time series plot of the conditional skewness of the one week ahead out-of-sample density forecasts of exchange rate returns with the real-time NFCI and a 10 year rolling window. Shaded areas denote NBER recessions.

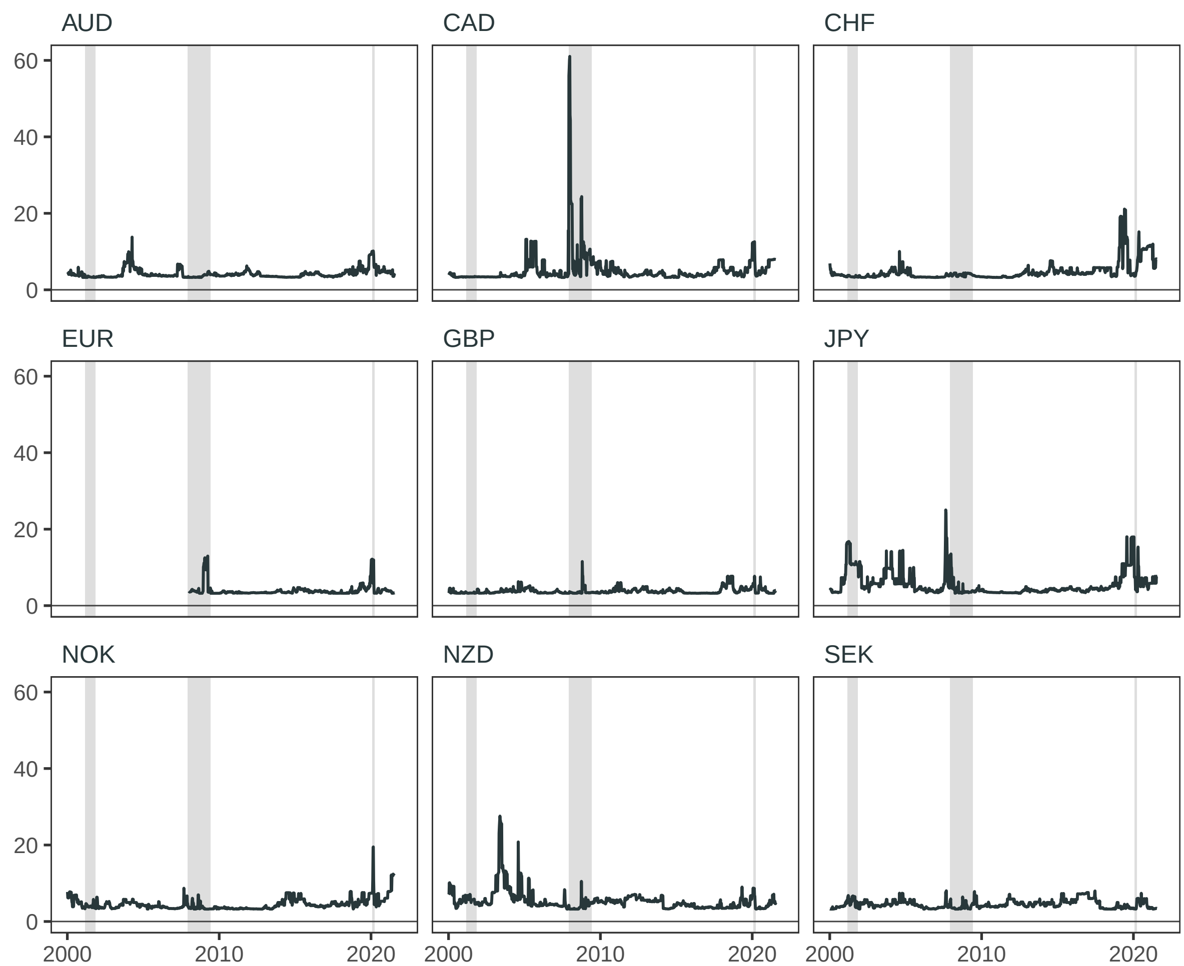


Figure A.7: Time series plot of the conditional kurtosis of the one week ahead out-of-sample density forecasts of exchange rate returns with the real-time NFCI and a 10 year rolling window. Shaded areas denote NBER recessions.

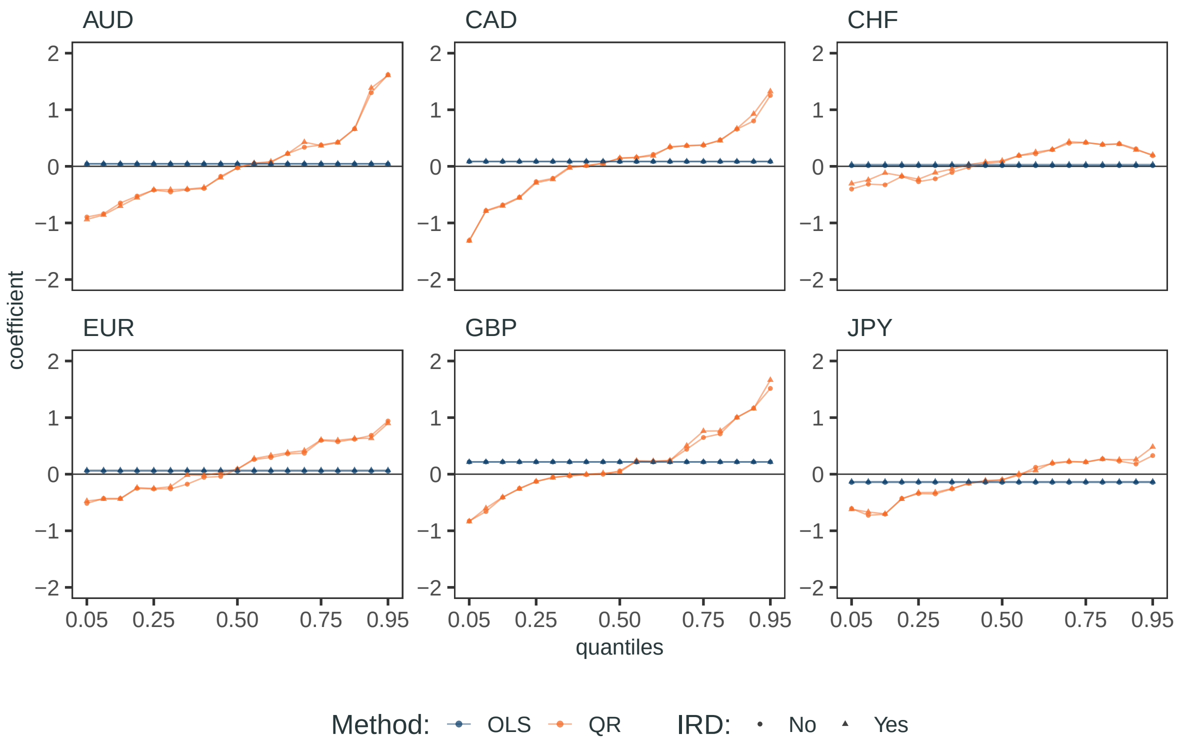


Figure A.8: In-sample OLS and QR coefficients for the one week ahead exchange rate returns on the ex-post NFCI with and without interest rate differentials (IRD). Sample period: 2000/01/01 – 2017/12/20.

# Department of Economics, University of Alberta
# Working Paper Series

| |
|---|
| **2025-05:** The Happiness Crisis for Younger Generations in Canada and the United States: What is Different and What is Not – **Huang, H.,** Helliwell, J., Norton, M. |
| **2025-04:** Take the Load Off: Time and Technology as Determinants of Electricity Demand Response – Bailey, M., **Brown, D.**, Shaffer, B., Wolak, F. |
| **2025-03:** Is There a Happiness Crisis Among Young Canadians? – **Huang, H.,** Helliwell, J., Norton, M. |
| **2025-02:** Employment Adjustments to Increased Imports: Evidence from a Developing Country – **Ural Marchand, B.** |
| **2025-01:** Estimating Discrete Choice Demand Models with Sparse Market-Product Shocks – Lu, Z., **Shimizu, K.** |
| **2024-11:** Canada's "COVID-19 Referendum": Voting in the Early Federal Election of 2021 – **Marchand, J.,** Wang, Y. |
| **2024-10:** On Prior Confidence and Belief Updating – Chan, K., Charness, G., **Dave, C.,** Reddinger, J. |
| **2024-09:** Merger Review under Asymmetric Information – **Langinier, C.,** Ray Chaudhuri, A. |
| **2024-08:** Environmental Regulations and Green Innovation: The Role of Trade and Technology Transfer – **Langinier, C.,** Martinez-Zarzoso, I., Ray Chaudhuri, A. |
| **2024-07:** Green Patents in an Oligopolistic Market with Green Consumers– **Langinier, C.,** Ray Chaudhuri, A. |
| **2024-06:** The Effect of Patent Continuation on the Patenting Process – **Langinier, C.** |
| **2024-05:** Applying Indigenous Approaches to Economics Instruction – Chavis, L., **Wheeler, L.** |
| **2024-04:** Electric Vehicles and the Energy Transition: Unintended Consequences of a Common Retail Rate Design – Bailey, M., **Brown, D.,** Myers, E., Shaffer, B., Wolak, F. |
| **2024-03:** Electricity Market Design with Increasing Renewable Generation: Lessons From Alberta – **Brown, D.,** Olmstead, D., Shaffer, B. |
| **2024-02:** Evaluating the Role of Information Disclosure on Bidding Behavior in Wholesale Electricity Markets – **Brown, D.,** Cajueiro, D., **Eckert, A., Silveira, D.** |
| **2024-01:** Medical Ethics and Physician Motivations – **Andrews, B.** |
| **2023-14:** Employers' Demand for Personality Traits and Provision of Incentives – **Brencic, V., McGee, A.** |
| **2023-13:** Demand for Personality Traits, Tasks, and Sorting – **Brencic, V., McGee, A.** |
| **2023-12:** Whoever You Want Me to Be: Personality and Incentives – **McGee, A.,** McGee, P. |
| **2023-11:** Gender Differences in Reservation Wages in Search Experiments – **McGee, A.,** McGee, P. |
| **2023-10:** Designing Incentive Regulation in the Electricity Sector – **Brown, D.,** Sappington, D. |
| **2023-09:** Economic Evaluation under Ambiguity and Structural Uncertainties – **Andrews, B.** |
| **2023-08:** Show Me the Money! Incentives and Nudges to Shift Electric Vehicle Charge Timing – Bailey, M., **Brown, D.,** Shaffer, B., Wolak, F. |
| **2023-07:** Screening for Collusion in Wholesale Electricity Markets: A Review of the Literature – **Brown, D., Eckert, A., Silveira, D.** |
| **2023-06:** Information and Transparency: Using Machine Learning to Detect Communication – **Brown, D.,** Cajueiro, D., **Eckert, A., Silveira, D.** |