CENTRE FOR ECONOMETRIC ANALYSIS
CEA@BAYES

https://www.bayes.city.ac.uk/faculties-and-research/centres/cea

Bayes Business School (formerly Cass)
Faculty of Finance
106 Bunhill Row
London EC1Y 8TZ

Global Inflation: Implications for Forecasting and Monetary Policy

Marcelo C. Medeiros, Erik Christian Montes Schutte, Tobias Skipper Soussi

---

CEA@Bayes Working Paper Series

WP–CEA–08-2023

Global Inflation: Implications for forecasting and monetary policy

Marcelo C. Medeiros *
The University of Illinois at Urbana-Champaign
marcelom@illinois.edu

Erik Christian Montes Schütte
Aarhus University, CREATES, and DFI
christianms@econ.au.dk

Tobias Skipper Soussi
Aarhus University and CREATES
tss@econ.au.dk

April 14, 2023

*Corresponding author. Send correspondence to Marcelo C. Medeiros, Department of Economics at The University of Illinois at Urbana-Champaign,214 David Kinley Hall, 1407 West Gregory Drive, Urbana, IL 61801, USA; email: marcelom@illinois.edu. We are grateful to Dan Bernhardt, Eric Hillebrand, Stig Vinther Møller, Mikkel Bennedsen, Yunus Emre Ergemen, Niels Strange as well as conference and seminar participants at the second Workshop on Dimensionality Reduction and Inference in High-Dimensional Time Series, the 24th Brazilian Symposium on Probability and Statistics, the 19th School of Time Series and Econometrics, the 16th International Conference on Computational and Financial Econometrics, and the Department of Economics at the University of Illinois at Urbana-Champaign. Medeiros acknowledges partial financial support from CNPq and CAPES. Schütte acknowledges support from The Danish Council of Independent Research (DFF 7015-00017B). Soussi and Schütte are also grateful for research support from the Center for Research in Econometric Analysis of Time Series (CREATES) and Danish Finance Institute (DFI). The authors have no potential conflicts of interest to disclose.

Global Inflation
Implications for forecasting and monetary policy

Abstract

This paper considers inflation forecasting for a vast panel of countries. We combine the information from common factors driving global and country-specific inflation to build different models. We also rely on new advances in the Machine Learning literature. We show that random forests and neural networks are very competitive models, and their superiority, although stable across most of the time period considered, increases during recessions. We also show that it is easier to forecast countries with more developed economies. The forecasting gains seem to be partially explained by the degree of trade openness and inflation volatility within a year. Our results have two significant implications for monetary policy. First, our forecasts can serve as inflation expectations for countries where survey data are unavailable. Second, we shed some light on the links between inflation from different countries, facilitating the study of the transmission of monetary shocks.

JEL classifications: C53, C55, E31, E37, F15

Keywords: global inflation, inflation forecasting, machine learning, random forests, neural networks, shrinkage.

1. Introduction

Forecasting inflation is an essential and challenging task. Inflation forecasts can serve as a measure of inflation expectations and are essential for Central Banks to set interest rates and conduct monetary policy (Bernanke and Boivin 2003; Levin, Wieland, and Williams 2003; Mishkin 2007). While most papers in the literature focus on a single or a small set of countries, this paper considers the problem of simultaneously forecasting inflation from a large panel of countries. Our strategy explores a common factor structure among inflation series and potential (nonlinear) links across countries. Furthermore, we do not rely on additional variables apart from inflation and deterministic components, such as seasonal dummies. Therefore, this paper contributes to the literature on forecasting global inflation. It also sheds light on the drivers of the superior predictive performance of models based on cutting-edge machine-learning techniques.

There has been strong evidence in the literature that traditional benchmark inflation forecasting models can be outperformed by the combination of machine learning (ML) methods and a large set of economic predictors (big data); see, for example, Coulombe et al. (2020), Garcia, Medeiros, and Vasconcelos (2017), Medeiros and Mendes (2016), or Medeiros et al. (2021). On the other hand, Ciccarelli and Mojon (2010) show that the simple average of inflation data over OECD countries is a strong predictor of future inflation, and autoregressive models augmented with such an average also outperform well-established benchmarks. However, for several countries, the availability of economic data in real-time is more restricted. In some cases, real-time data are not available at all. Therefore, inspired by the evidence put forward in Ciccarelli and Mojon (2010), we take advantage of state-of-art ML methods to construct expanded indexes for global inflation and to explore the potential nonlinear relations among the dynamics of inflation in a large panel of countries using the recent *Global Database on Inflation* from Ha, Kose, and Ohnsorge (2021).

We start by constructing novel and more general measures of global inflation (factors) using a broader cross-section of countries compared to previous studies. Since our cross-section is much larger, we augment our models with regional factors. Besides the factors, we also include the high-dimensional panel of country-specific inflation series as predictors, thereby giving the models flexibility to choose between pre-computed factors and individual inflation series. To forecast inflation for all countries in our data set, we employ three machine learning methods: linear elastic net (EN), neural networks (NN), and random forests (RF). We also use an equal-weighted ensemble of these methods to account for model uncertainty and consider ensemble predictions. As benchmarks, we consider the random walk (RW) model, an autoregressive (AR) specification, and an AR model augmented with the global and regional factors (CM).

The purpose of this paper is not to determine the superiority of a comprehensive list of forecasting models. Rather, our goals are twofold. First, we aim to assess the generalizability of the findings in Ciccarelli and Mojon (2010) in a global inflation context and provide a straightforward yet systematic approach to constructing inflation forecasts for a large number of countries using the latest advancements in machine learning. Our forecasts serve as a useful alternative measure of inflation expectations for macroeconomic models utilized in policy analysis, and they have the added benefit of not requiring macroeconomic data at the national level, making them suitable for countries with limited or untimely data availability. Secondly, we seek to shed light on the drivers of potential performance gains over the benchmarks, furthering the discussion on the causes of global inflation.

1.1. Main Takeaways

Our results contribute to the literature in several directions. Corroborating previous findings in the literature, the RF model is generally the best predictive model, but the Ensemble model provides a similar, albeit slightly lower, performance. Models incorporating nonlinearities (RF, NN, and Ensemble) generally tend to outperform the linear EN model. Thus, considering nonlinearities in inflation dynamics is essential for most countries. The Ciccarelli and Mojon (2010) (CM) model works very well and provides nontrivial gains compared to the simpler AR model, irrespective of the type of factors used. Despite its simplicity, the CM model is also competitive with the EN model, which has access to all predictors besides the global and regional factors.

We dig deeper into the results of the best-performing model, the RF, and find that its performance gains are consistent throughout the entire out-of-sample period, with a sharper increase in relative performance during the 2008-09 recession, suggesting larger relative gains during recessions. When considering the performance of the RF model across countries, the highest accuracy is achieved for European countries. Overall, our results indicate the existence of a correlation between the degree of development of the economy and the performance of the RF model.

We also analyze how trade integration explains the relative performance of the RF model over the benchmark. Our empirical results support the hypothesis that trade integration is a critical determinant of the globalization of inflation and that these spillover effects in inflation happen in all regions.

Our findings build upon the work of Ciccarelli and Mojon (2010) by emphasizing the significance of global, regional, and country-specific factors in predicting future inflation. This underscores inflation as a global phenomenon and highlights the need for enhanced collaboration and coordination among central banks in formulating and executing monetary policy, especially for countries deeply integrated into the global economy. The global inflation database from Ha, Kose, and Ohnsorge (2021) proves valuable for forecasting inflation in less developed economies with limited macroeconomic data or reliable survey expectations, offering a more accurate alternative to random walk or autoregressive models for such countries. Accurate inflation forecasts are particularly crucial for these economies due to their higher inflation rate variability. Furthermore, as demonstrated in studies like Kamber and Wong (2020), global factors have a more pronounced impact on inflation in emerging market economies compared to advanced ones. Our methodology equips policymakers with a unified data source for establishing global and regional inflation benchmarks and provides a superior measure of inflation expectations for use in Phillips Curve specifications and other macroeconomic models.

1.2. Organization of the Paper

The remainder of the paper is organized as follows: Section 2 provides an overview of the data employed in the study. Section 3 details the forecasting methodology, briefly discussing the models used for inflation forecasting, evaluation measures, and interpretation tools. Section 4 presents the main findings, while Section 5 explores various channels of predictive performance. Section 6 showcases the empirical results through five case studies, and Section 7 offers concluding remarks. Additionally, we include a set of online appendices not intended for publication: Appendix A contains summary statistics for the data used in the study; Appendix B offers a detailed description of the forecasting models and the cross-validation scheme for tuning hyperparameters; Appendix C discusses alternative methods for constructing factors; and Appendix D presents supplementary empirical results excluded from the main paper due to space limitations.

2. Data

We source our data on inflation from the novel *Global Database on Inflation* from Ha, Kose, and Ohnsorge (2021). This database contains a wide range of inflation measures for many countries at monthly, quarterly, and yearly frequencies. Since we aim to utilize methods that thrive in high-dimensional settings, we focus on monthly measures. We focus on headline inflation since this measure has the broadest coverage in terms of the basket of goods and services it includes.[1] Our objective is to create a balanced panel of data with the largest

[1] Ha, Kose, and Ohnsorge (2021) define headline inflation as *changes in the prices of all goods and services in a basket of goods and services that is representative of consumer expenditures.*

possible cross-section of countries, and with minimal missing values. By having a balanced panel, we can make comparisons between countries without being influenced by differences in time periods. To achieve this, we limit our sample to the period from January 1980 to December 2019 (480 observations) and exclude countries with more than two consecutive missing values or a maximum of four missing values during this period. We use the Expectation-Maximization algorithm from Stock and Watson (2002a) to interpolate missing values, resulting in a total of only seven imputed values across all countries and periods in the data set.[2] This process results in inflation data for 91 different countries.

Our sample includes countries with diverse inflation dynamics, featuring two particularly noteworthy issues. First, some countries experienced high inflation during the early part of our sample period. Second, there is significant heterogeneity in the degree of seasonality of inflation across countries, with some displaying strong seasonal patterns and others showing little to no seasonality. This heterogeneity is intrinsically interesting, as it suggests that some countries have a higher tendency towards periodic and predictable price adjustments, while others do not. To address the first issue, we take first differences in the monthly inflation series of countries that experienced hyperinflation during our sample period: Bolivia, Brazil, and Peru. We follow Cagan (1956) in defining hyperinflation as a monthly inflation rate of 50% or more in at least one month during the analyzed period.[3] To tackle the second issue, we use raw, non-seasonally adjusted data and incorporate monthly dummy variables as additional predictors in all forecasting models, allowing us to measure the significance of each country's unique seasonality concerning forecast accuracy.

Tables 7, 8 and 9 in the Appendix A show the minimum, maximum, average, standard deviation and first-order autocorrelation coefficient for monthly headline inflation, in percent, for all countries in our data set in the full sample period, 1980-2000 and 2000-2019, respec-

[2]Four observations were interpolated for Ethiopia, two for Suriname, and one for Myanmar.

[3]We take first differences in the inflation series of these countries when the series are used as predictors but leave the series in levels when forecasting them. This ensures a fair comparison across all inflation series. Our results are not sensitive to whether or not we first difference these inflation series in the panel of predictors.

tively. There is considerable heterogeneity across countries, with South American countries having much larger average monthly inflation across the whole sample period. Although our sample does not include the high inflationary period of the 1970s, the period from 1980 to 2000 has an average inflation level three times as large as the subsequent period.

3. Methodology and inference

This section describes the empirical methodology used in this paper and outlines the methods used to forecast inflation and draw inferences from the results.

3.1. The general framework

We consider the following general model:

$$\pi_{i,t+h}^{h} = G_{i,t,h}\big(\boldsymbol{\pi}_{i,t-p+1:t}^{AR}, \boldsymbol{\gamma}_{i,t}, g_t, r_t^R, \boldsymbol{x}_{i,t}\big) + u_{i,t+h}, \tag{3.1}$$

where $i = 1, ..., n$ denotes the country-specific index, $h = 1, ..., H$ represents the forecast horizon, $\pi_{i,t}$ is the log-difference of the headline CPI for country i at time t, and $\pi_{i,t+h}^{h} = \frac{1}{h}\sum_{j=1}^{h}\pi_{i,t+j}$ is the h month ahead inflation growth rate. The set of predictors may include p lags of the target variable to capture the auto-regressive structure in inflation, $\boldsymbol{\pi}_{i,t-p+1:t}^{AR} = (\pi_{i,t-p+1}, ..., \pi_{i,t})'$, a set of monthly dummy variables to capture the seasonal component, $\boldsymbol{\gamma}_{i,t}$, a global inflation factor, g_t, a relevant regional inflation factor, r_t^R, which depends on the region in which the country of interest is located, and a vector of international inflation series, $\boldsymbol{x}_{i,t} = (x_{i,1,t}, ..., x_{i,n,t})'$, which includes all $n-1$ countries in the panel except for the ith country. We do not include lags of the factors nor the inflation series for two reasons. First, given a large number of countries, forecast horizons, and estimation windows, we wish to simplify the estimated models as much as possible to make the interpretation hereof more straightforward. Second, preliminary results showed little to no forecast gains from including additional lags. Taking all into account, this means that our forecasting models consider 108 potential predictors. $G_{i,h,t}$ is the mapping between the predictors and future inflation, which may both be linear and nonlinear, and $u_{i,t+h}$ is a zero-mean random error. The direct forecast equation for country i at horizon h at time t is thus given by:

$$\widehat{\pi}_{i,t+h}^{h} = \widehat{G}_{i,t-R+1:t,h}\big(\boldsymbol{\pi}_{t-p+1:t}^{AR}, \boldsymbol{\gamma}_{i,t}, g_t, r_t^R, \boldsymbol{x}_{i,t}\big), \tag{3.2}$$

where $\widehat{G}_{i,t-R+1:t,h}$ is the estimated mapping function based on data from $t-R+1$ to t and R is the window size. More specifically, all forecasts are computed using a rolling window scheme with $R = 240 - p + 1$ observations. Our out-of-sample evaluation window starts in March 2000 and ends in December 2019. The primary motivation for a rolling window scheme is to lessen the effects of potential structural breaks in the data and avoid problems running superior predictive performance tests among nested models.[4] We set the number of lags, p, to four. We consider forecast horizons of $h = 1, 3, 6, 12$ months ahead. For the remainder of the paper, we omit country-specific indexing for simplicity of exposition, but everything presented from this point on is computed/estimated for each country-horizon pair.

The inclusion of a global inflation factor in our set of predictors is inspired by Ciccarelli and Mojon (2010), who find that a global inflation factor explains most of the variance in inflation across 22 OECD countries. Since our sample of countries is larger and more heterogeneous, we expand on their framework and consider five additional regional factors for Europe, North America, South America, Asia, and Africa.[5] The countries in each regional factor match the divisions given in Table 7 in Appendix A. We follow Ciccarelli and Mojon (2010) and use a simple cross-country average for all countries in the case of the global factor and constituent countries in each region in the case of regional factors. Section C in the Appendix considers two different methodologies to estimate the factors: PCA and the autoencoders based on neural networks.

[4]As Giacomini and White (2006) note, the unconditional Giacomini–White (GW) test is equivalent to the traditional Diebold–Mariano (DM) test when a rolling window scheme is used.

[5]Ciccarelli and Mojon (2010) only consider the 22 countries from the OECD.

3.2. Forecasting models

To exploit the high-dimensional nature of the predictor set, we consider three models from the machine learning literature that can handle many predictors. The first model is a penalized linear regression with the Elastic Net penalty of Zou and Hastie (2005), which recent papers have shown to work well for inflation forecasting in high-dimensional settings (e.g., Inoue and Kilian 2008, Medeiros and Mendes 2016, and Garcia, Medeiros, and Vasconcelos 2017). The second model we consider belongs to the class of Neural Networks, which have proven helpful for prediction in numerous domains, amongst them inflation (Coulombe 2022). The third model we consider is the Random Forest estimator of Breiman (2001), which is inspired by Medeiros et al. (2021) and Coulombe et al. (2020) who show that a Random Forest outperforms a vast array of other machine learning methods in forecasting US inflation in data-rich environments. The Neural Network and Random Forest models allow for the estimation of nonparametric nonlinear relations between predictors and the target variable. In this sense, we cover both a linear (Elastic Net) and nonlinear (Neural Network/Random Forest) mapping between the predictors and future inflation. These models are described in more detail in section B of the Appendix.

We also consider three benchmark models inspired by the literature. These are briefly described below.

3.2.1. Benchmark models

The first benchmark model is the Random Walk (RW) model in the vein of Atkeson and Ohanian (2001). We compute forecasts as follows:

$$\widehat{\pi}^h_{t+h} = \pi^h_t, \tag{3.3}$$

meaning that the RW forecast is nothing more than the time t value of the target variable.

The second benchmark model is an autoregressive (AR) model of order p, where p is selected using the Bayesian Information Criterion (BIC). Since some inflation series display a large seasonal component, we also include a set of monthly dummy variables:

$$\widehat{\pi}_{t+h}^{h} = \widehat{\phi}_0 + \widehat{\phi}_1 \pi_t + ... + \widehat{\phi}_p \pi_{t-p+1} + \sum_{s=1}^{11} \widehat{\beta}_s z_{s,t}, \tag{3.4}$$

where z_t represents the set of monthly dummy variables from January to November.[6]

Our final benchmark is inspired by Ciccarelli and Mojon (2010). We consider the model in Equation 3.4 but extend it by global and relevant regional inflation factors. We denote this model as "CM" going forward.

3.3. Evaluation measures and interpretation tools

We evaluate the quality of the forecasts using the root mean square error (RMSE) and the median absolute deviation from the median (MAD). For some model m and horizon h, we compute RMSE and MAD as:

$$\text{RMSE}_{h,m} = \sqrt{\frac{1}{\#OOS_h} \sum_{t \in OOS} \widehat{e}_{t,h,m}^2}, \tag{3.5}$$

$$\text{MAD}_{h,m} = \text{median}[|\widehat{e}_{t,h,m} - \text{median}(\widehat{e}_{t,h,m})|], \tag{3.6}$$

where $\widehat{e}_{t,h,m} = \pi_t^h - \widehat{\pi}_{t,m}^h$ and $\widehat{\pi}_{t,m}^h$ is the forecast of the h month ahead inflation growth at time t using model m and $\#OOS_h$ is the number of out-of-sample predictions made for horizon h. We use MAD, in addition to RMSE, to ensure that outliers do not drive the results.

To test if the forecasts from distinct models differ, we consider the multi-horizon average superior predictive ability (SPA) test by Quaedvlieg (2021). The test considers all horizons simultaneously, and the null hypothesis is that the model in question does not beat the benchmark model on average. At the same time, the alternative entails superior average performance across the horizons considered. We follow Quaedvlieg (2021) and set the block

[6] We omit one dummy variable to avoid perfect multicollinearity.

length in the moving block bootstrap to three and the number of bootstrap resamples to 999. We give the individual forecast horizons equal weight and compute the test statistic based on squared error loss differentials.

To analyze the stability of predictive accuracy through time, we plot the cumulative sum of squared error difference (CSSED) between the model of interest and the benchmark model (Welch and Goyal 2008). We compute the CSSED for model m and horizon h at time t as:

$$\text{CSSED}_{m,t}^{h} = \sum_{i=t_0}^{t} \left[(\pi_t^h - \tilde{\pi}_t^h)^2 - (\pi_t^h - \widehat{\pi}_{m,t}^h)^2 \right], \tag{3.7}$$

where t_0 and t are the beginning and the end of the evaluation period and $\tilde{\pi}_t^h$ is the forecast produced by the benchmark model. A positive $\text{CSSED}_{m,t}^{h}$ at time t implies that the model of interest is outperforming the benchmark up to time t, and vice versa for a negative value.

Machine learning models, such as the RF and NN models, can be black boxes for which it is difficult to calculate the marginal effect of a variable and thus obtain a measure that allows interpreting the model. To surmount this challenge, we use SHAP values (Lundberg and Lee 2017). SHAP values are currently considered the state-of-the-art method for interpreting machine learning models, as they have many desirable properties.[7] They stem from game theory as a general solution to the problem of attributing a payoff obtained in a cooperative game to the individual players based on their marginal contribution to the game (Shapley 1953). In this setting, we can think of the variables as the players and the prediction as the payoff. We compute SHAP values in `Python` using the `shap` package.

We denote the SHAP value of model m corresponding to the kth predictor for the tth observation in the training sample by $\phi_{k,t}^{(m)}$. This SHAP value gives us the contribution in model m of the kth predictor to the prediction of the target variable for the tth observation (measured in terms of the deviation from the mean of the target over the training sample).

[7] See Molnar (2020) for a detailed exposition of SHAP values.

An important property of SHAP values is efficiency:

$$\hat{y}_t^{(m)} = \sum_{i=1}^{n} \phi_{k,t}^{(m)}, \tag{3.8}$$

meaning that the prediction of model m at time t, $\hat{y}_t^{(m)}$, can be decomposed into the sum of SHAP values for all the predictors of the model. From this result, we have that a variable importance measure for predictor k in model m over the training sample $t = 1, ..., T$ can be constructed as:

$$VI_k^{(m)} = \frac{1}{T} \sum_{t=1}^{T} |\phi_{k,t}^{(m)}|. \tag{3.9}$$

$VI_k^{(m)}$ provides a measure of the overall importance of predictor k to the fitted model. Since we use a rolling training window estimation, we have a sequence of variable importance measures for each predictor (one for each training window). Therefore, to obtain an overall measure of importance across the whole sample period, we calculate the $VI_k^{(m)}$ for all predictors in each estimation window and then take an average of these.

4. Results

This section presents the main empirical results for the models we consider: the Random Walk (RW), the seasonal autoregressive model (AR), the Ciccarelli and Mojon (2010) model (CM), the Elastic Net (EN), the Neural Network (NN), the Random Forest (RF), as well as an equal-weighted Ensemble model consisting of the EN, NN, and RF models. We focus on out-of-sample forecasting accuracy to avoid overfitting issues in machine learning models. Our results demonstrate the superior predictive power of the global inflation panel compared to univariate models. To validate our predictions, we show their strong correlation with forecasts using a wider range of predictors. Our results also support using the global panel as a substitute for lacking national data in developing countries. We delve deeper into the results by examining the overall forecasting performance and accuracy over time.

Figure 1: Forecasting outperformance of an Ensemble of RF, NN, and EN models vis-à-vis the RW model for $h = 1$

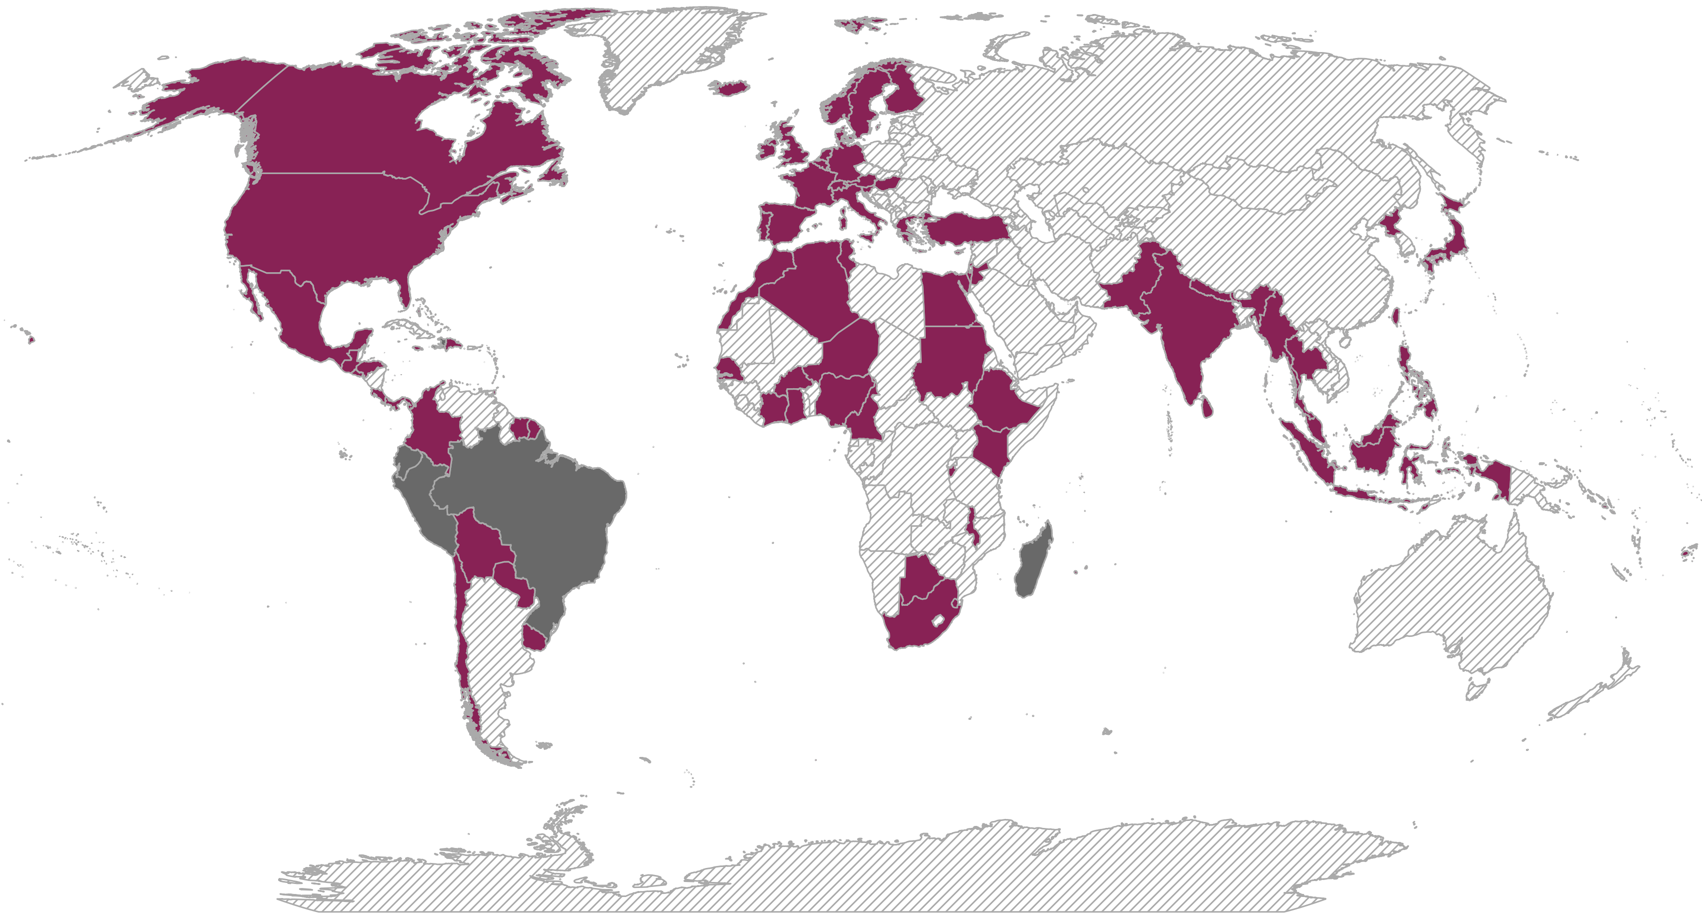

The map shows the outperformance of an equal-weighted Ensemble of RF, NN, and EN models vis-à-vis the RW model for $h = 1$. Countries colored in dark red have a RMSE ratio for the equal-weighted Ensemble vis-à-vis the RW model below 1, implying better forecasting performance for the one-month ahead forecast horizon. Countries colored in grey have a RMSE ratio above 1, implying worse forecasting performance. A total of 86 out of 91 countries have a ratio below 1.

4.1. The predictive power of global inflation

One of our central propositions is that the global inflation panel can serve as an effective alternative to univariate models in countries where comprehensive and current monthly macroeconomic data is scarce. This implies that our predictive models should surpass univariate benchmark models and, if feasible, generate forecasts approximately on par with the most accurate predictions made by policymakers. Additionally, we contend that our forecasts can function as a valuable supplementary measures for inflation expectations, which requires the predictions to be superior to those generated by univariate models. In order to validate these claims, we demonstrate that the global inflation panel is capable of producing forecasts that outperform univariate models and exhibit a notable correlation with the most dependable forecasts created by policymakers. It is important to clarify that we do not assert that the forecasts generated by the global inflation panel can surpass the best predictions made by policymakers, but rather that they can serve as a trustworthy alternative in the absence of such predictions.

The map in Figure 1 illustrates the superior performance of the Ensemble model, which takes advantage of global inflation information compared to the RW benchmark model. In 86 out of 91 countries in our sample, the RMSE ratio of the Ensemble model was smaller than 1, indicating a significant improvement in forecasting accuracy.[8] Further, comparing the Ensemble with the AR model showed similar results, with 80 countries having a ratio below 1. Maps showing the magnitude of the country-specific RMSE ratios can be found in Appendix D. Similar maps for other models are available upon request.

We evaluate the correlation between the inflation forecasts produced by the Organisation for Economic Co-operation and Development (OECD) and those generated by our equal-weighted Ensemble model, which incorporates the EN, NN, and RF models. Given that the OECD's forecasts are only available for seasonally adjusted inflation series with a three-

[8]For 3, 6, and 12 months ahead, we find that the Ensemble model outperforms the RW benchmark model for 86, 81, and 66 countries, respectively. These maps can be found in Appendix D.

Table 1: Correlation between OECD forecasts and Ensemble forecasts

Cou	BEL	CAN	CHE	CHL	CRI	DEU	DNK	ESP	FIN	FRA	GBR	GRC
ρ	0.56	0.57	0.71	0.59	0.89	0.61	0.68	0.81	0.71	0.65	0.63	0.74
p-val	0.00	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Cou	HUN	IRL	ISR	ITA	JPN	KOR	LUX	MEX	NLD	PRT	SWE	USA
ρ	0.79	0.74	0.70	0.81	0.29	0.78	0.61	0.54	0.65	0.74	0.68	0.58
p-val	0.00	0.00	0.00	0.00	0.02	0.00	0.00	0.00	0.00	0.00	0.00	0.00

The table reports the correlation (ρ) between the OECD forecasts and the out-of-sample forecast made by an equal-weighted Ensemble of RF, NN, and EN models at a horizon of three months ahead ($h = 3$) as well as the p-value (p-val) of the F-statistic of overall regression significance of a regression of the OECD forecasts on the Ensemble forecasts and constant. The countries in the table are: Belgium (BEL), Canada (CAN), Switzerland (CHE), Chile (CHL), Costa Rica (CRI), Germany (DEU), Denmark (DNK), Spain (ESP), Finland (FIN), France (FRA), United Kingdom (GBR), Greece (GRC), Hungary (HUN), Ireland (IRL), Israel (ISR), Italy (ITA), Japan (JPN), South Korea (KOR), Luxemburg (LUX), Mexico (MEX), Netherlands (NLD), Portugal (PRT), Sweden (SWE), and the United States (USA).

month horizon, we adjusted our target variables and predictors accordingly. We focus on the Ensemble forecasts for a horizon of three months. As a widely recognized and established benchmark, the OECD's forecasts are derived from a combination of expert judgment and model-based analysis, considering the economic conditions of individual countries and the global economy. By comparing our forecasts to those of the OECD, we aim to demonstrate the reliability and validity of our predictions. The results of this analysis, conducted on a sample of 24 countries for which forecasts overlap with our sample, are presented in Table 6. On average, the correlation is 0.67, indicating that the global inflation panel could potentially substitute for countries without comprehensive forecasts and limited macroeconomic predictors for model-based forecasts. An exception is Japan, where the correlation is only 0.24, implying that relying solely on global inflation is insufficient to forecast inflation accurately in this country. The overall high correlation suggests that the global inflation database may be a good way to forecast inflation for developing countries without a large panel of timely macroeconomic indicators.

4.2. Aggregate results

Table 2 reports the main forecasting results as summary statistics across all models and horizons under consideration. Due to space constraints, we cannot show the forecasting results for each model and each country at each horizon. Therefore, we opt to focus on summary statistics and averages across countries. In Section 6 we show results for a selected number of countries. Columns (1)-(4) report the ratio of average RMSE for the model of interest against the average RMSE for the RW model for horizons $h = 1, 3, 6, 12$, respectively. Columns (5)-(8) report the corresponding ratios of average MAD for the model of interest against the average MAD for the RW model.[9] Columns (9) and (10) report the share of times each model achieved the lowest RMSE and MAD across all country/horizon pairs. Column (11) shows the average of the p-values from the average multi-horizon SPA test proposed by Quaedvlieg (2021). Finally, column (12) reports the proportion of times each model achieved a p-value below the conventional 5% level. We use the squared error loss differential concerning the RW model in the test. The first panel reports the results for the benchmark models, and the second panel reports the results for the machine learning models that use the whole panel (EN, NN, RF, Ensemble)

Several facts emerge from the table. First, the RF model tends to be the best-performing model, closely followed by the Ensemble model. The RF model has the lowest RMSE ratio across all horizons, but the Ensemble model does achieve a lower MAD ratio at horizons of 3 and 6 months ahead. Both models show the greatest relative improvement at three and six months horizons with RMSE/MAD ratios that are 3-6% lower than the next best model, which are usually the NN and EN models. The RF and Ensemble models also have the highest frequency of outperformance, with the RF model having a slight edge over the Ensemble when using the RMSE ratio and vice-versa when using the MAD ratio. The only model with an average SPA p-value below the conventional 5% level is the RF model, and

[9] An alternative approach to evaluate the average performance of the models would be to compute ratios first and then take an average afterward. We take this approach in Table 11 in Appendix D. This approach does not alter the overall conclusions but is subject to the inherent asymmetry of ratios.

this significance level is achieved individually for 86 out of 91 countries. When looking at the machine learning models, non-linear models (NN and RF) tend to outperform their linear counterparts (EN), highlighting the potential significance of nonlinearities in global inflation. Despite its simplicity, the CM model is often competitive with the EN model, and it has a higher percentage of outperformance than the latter even though the EN model has access to all predictors besides the global and regional factors. Lastly, it is worth noting that the RW benchmark model is difficult to outperform at the $h = 12$ horizon. With respect to RMSE ratios, only the RF model achieves a lower ratio than the RW benchmark model at the one-year forecast horizon. This decrease in relative performance at $h = 12$ suggests that the predictive information embedded in other countries' inflation rates is most beneficial at short and medium horizons. This finding also implies that inflation transmission generally occurs at horizons shorter than one year.

Overall, we find all models considered in this paper to improve on the RW benchmark model at the short to medium-horizon forecasts ($h = 1, 3, 6$). An increase in model complexity is generally accompanied by improved forecast accuracy. The RF model emerges as the top-performing model, closely followed by the Ensemble model, which provides the advantage of forecast diversification. We also find that the relatively simple CM model is competitive and often outperforms the EN model, which has access to all inflation series and regional and global factors. Thus, when nonlinearities are excluded from the functional form, the global and regional factors appear to embed most of the relevant information driving the comovement in inflation across the globe.

Table 2: Forecasting results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	RMSE				MAD				Frequency of outperformance		Multi-horizon SPA	
Model	$h = 1$	$h = 3$	$h = 6$	$h = 12$	$h = 1$	$h = 3$	$h = 6$	$h = 12$	% min. RMSE	% min. MAD	Average	Share < 0.05
RW RMSE/MAD	0.009	0.007	0.005	0.003	0.005	0.004	0.003	0.002	8%	10%		
AR	0.814	0.777	0.866	1.153	0.807	0.680	0.740	1.051	9%	15%	0.131	82%
CM	0.801	0.754	0.829	1.111	0.802	0.676	0.742	1.047	16%	11%	0.103	87%
EN	0.786	0.756	0.840	1.128	0.768	0.661	0.721	1.063	5%	4%	0.117	85%
NN	0.778	0.744	0.809	1.063	0.765	0.639	0.657	0.938	8%	13%	0.147	82%
RF	**0.737**	**0.697**	**0.750**	**0.970**	**0.710**	0.607	0.646	**0.920**	**28%**	22%	**0.042**	**95%**
Ensemble	0.740	0.702	0.765	1.013	0.716	**0.603**	**0.643**	0.927	26%	**24**%	0.071	90%

The table reports several summary statistics for out-of-sample forecast accuracy for each model, computed across all 91 countries. The out-of-sample evaluation window runs from March 2000 to December 2019. Columns (1), (2), (3), and (4) report the average RMSE for the RW model and the ratios of the average RMSE across all countries for $h = 1, 3, 6, 12$ where the ratios are computed between the model of interest and the RW model as a benchmark. Columns (5), (6), (7), and (8) report the corresponding average MAD for the RW model and the ratios of average MAD for the remaining models. Columns (9) and (10), respectively, report the number of times each model achieved the lowest RMSE and MAD across all country/horizon pairs as a share. Finally, columns (11) and (12) report the mean p-value and the share of p-values below the conventional 5% level across all countries of the average SPA test of Quaedvlieg (2021). The test is based on squared errors. The best results for each horizon are highlighted in bold.

Figure 2: RMSE distributions

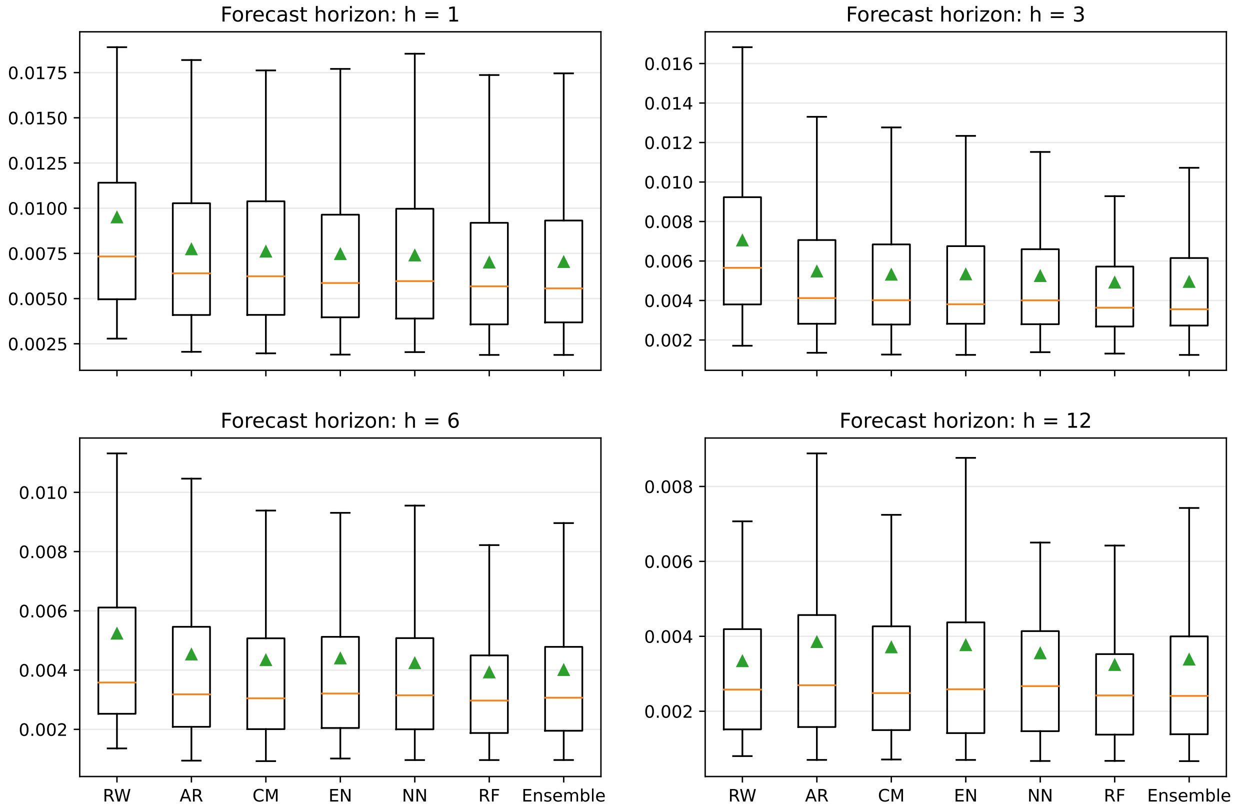

The figure shows RMSE box plots for the Random Walk (RW), Autoregressive model (AR), Ciccarelli and Mojon (2010) model (CM), Elastic Net (EN), Neural Network (NN), Random Forest (RF) and Ensemble for $h = 1, 3, 6, 12$. The extremes of the box correspond to the 1st, Q1, and 3rd, Q3, quartile implying that 50% of the RMSEs are located within these bounds. The upper whisker is computed as Q3 + 1.5 × IQR, while the lower is computed as Q1 - 1.5 × IQR, where IQR is the interquartile range. The green triangle reports the average RMSE, while the orange horizontal line shows the median RMSE.

The previous results refer to averages, telling us little about the distribution of performance accuracy across countries in the panel. Figure 2 shows a box plot of RMSE statistics for the RW, AR, CM, EN, NN, RF, and Ensemble models. We focus on RMSE, a measure of absolute predictive accuracy, to complement the picture provided by Table 2, which focuses on relative performance (RMSE ratios). Figure 17 in Appendix D shows RMSE ratio box plots. In general, the distribution of RMSEs aligns with the findings in Table 2. The RF model is distinguished by a consistently narrower interquartile range, particularly at three and six-month horizons, suggesting that the RF model tends to exhibit significantly lower downside risk in terms of forecast accuracy. For short ($h = 1$) and long ($h = 12$) forecast horizons, the distributions across models appear relatively similar, indicating that simpler models like RW and AR can be highly competitive for short and long forecast horizons. This supports the notion that most global inflation spillover occurs in the medium term. The distributions of the CM, EN, and NN models are fairly comparable, although the NN model tends to exhibit lower variance. Interestingly, the CM model, despite its simplicity, delivers relatively robust performance. Overall, the plot verifies that the RF model surpasses all other models in performance, which is why much of the following analysis will center on the RF model.

4.3. Forecast accuracy across time

In this section, we look at the distribution of forecast accuracy across time. Figure 3 plots the 80th, 60th, 50th, 40th, and 20th percentiles of the CSSED distribution at each time t for the RF model across all forecast horizons. Note that we show the distribution of CSSED lines across all countries in the sample, which implies that any two points on any given line do not necessarily represent the same country. While this approach does not allow us to make conclusions about the CSSED for any specific country, it provides a much more general picture of the predictive ability of the RF model relative to the RW benchmark throughout the evaluation period.[10] We generally observe that the relative advantage of the RF model over the RW benchmark is almost monotonically increasing as a function of time, though at different rates for different percentiles. This is especially true for the shorter forecast horizons ($h = 1, 3, 6$), where all the percentile-specific-CSSED lines steadily increase throughout the entire out-of-sample period. This indicates that the CSSED distribution is steadily shifting upwards and, thus, that the RF model provides consistent forecast gains over the benchmark model throughout the entire out-of-sample period. Nonetheless, there

[10]Figures 18-23 in Appendix D plots the CSSED for specific countries for the AR, CM, EN, NN, RF, and Ensemble models.

was a sharper increase in relative performance during the 2008-09 recession, implying larger relative gains during recessions.

The conclusion is less optimistic, though unsurprising, given the results in Table 2 for the one-year ahead forecasts, where the lower percentiles of the CSSED distribution are mostly stable or even declining during the early part of the forecast evaluation period. Following the recession, all percentiles of the CSSED distribution increase steeply before leveling off for the remainder of the out-of-sample period. In Appendix D, we show in figures 24-28 that the above conclusions are largely unchanged for the AR, CM, EN, NN, and Ensemble models. As such, other models considered in this paper, which use varying degrees of additional information relative to the RW model, show an ability to consistently beat the benchmark model throughout the entire evaluation period, at least for the shorter horizons.

5. Where is predictive performance coming from?

The previous sections show that nonlinear models, and in particular, the RF model, outperform linear models. In this section, we dig deeper into the sources of outperformance. We begin by looking at the role of global trade integration. We then analyze what role the global and regional factors play and whether these are complements or substitutes for the high-dimensional panel that includes all individual inflation series. Finally, we analyze the role of nonlinearities.

5.1. The role of trade integration and global inflation

As a first step at understanding the sources of forecasting performance, we analyze the role of trade integration. Since the RF model successfully predicts inflation using other countries' inflation, our results can be seen as supportive of inflation as a global phenomenon (Kose, Otrok, and Whiteman 2003; Ciccarelli and Mojon 2010; Neely and Rapach 2011; Mumtaz and Surico 2012). This global comovement of inflation, often called the *"global inflation*

Figure 3: CSSED for the RF model

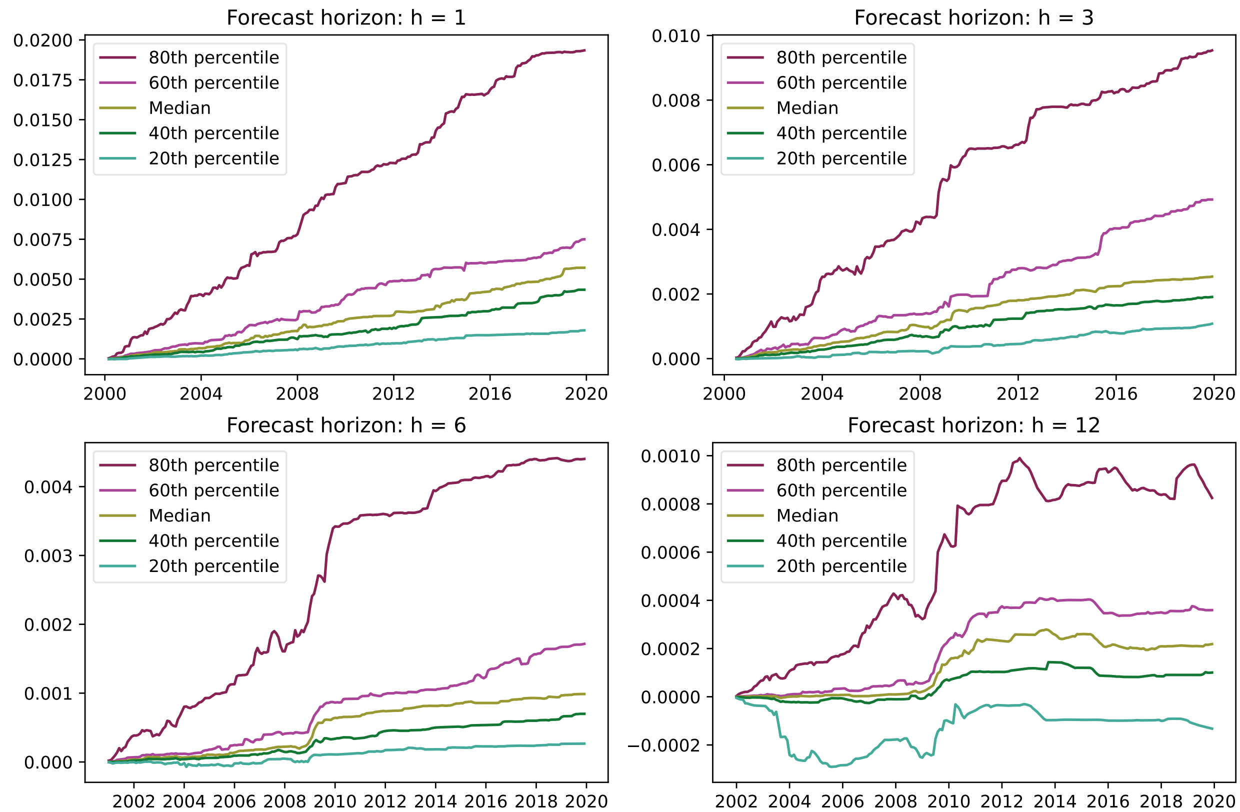

The figure shows the cumulative sum of squared differences (CSSED) for the RF model at $h = 1, 3, 6, 12$. We plot various percentiles of the CSSED distribution through time. A single line in the figure thus does not portray a specific country but a specific point in the CSSED distribution through time. We plot the 80th (dark red), 60th (purple), median (yellow), 40th (green) and 20th (turquoise) percentiles for each horizon.

hypothesis", is assumed to be the result of, among other things, the growing trade integration between countries (Attinasi, Balatti, et al. 2021).[11] Thus, a natural hypothesis to test is whether forecasting performance for a given country is positively related to the degree to which that country is integrated into global trade. To test this hypothesis, we run a fixed-effects panel regression that links the forecasting outperformance of the RF model (against the RW model) to trade openness. Specifically, we run the following fixed effects panel regression:

$$D_{i,t} = \alpha_i + \gamma_t + \beta Openness_{i,t} + \boldsymbol{\delta}' \boldsymbol{Z}_{i,t} + \varepsilon_{i,t} \tag{5.1}$$

[11] That inflation is affected by trade is not a new concept in economics. Indeed, the link between the two variables has long been incorporated into economic models (Shinkai 1973; Kingston and Turnovsky 1978).

where $D_{i,t} = \left(\frac{1}{12}\sum_{s=1}^{12}(\hat{e}_{i,s}^{(h),RW})^2 - \frac{1}{12}\sum_{s=1}^{12}(\hat{e}_{i,s}^{(h),RF})^2\right)$, is the annualized forecast loss differential between the RW and RF models for country i in year t, with $\hat{e}_{i,s}^{(h),RW}$ and $\hat{e}_{i,s}^{(h),RF}$ denoting the forecast error for country i in month $s = 1, ..., 12$ at horizon h, for the RW and RF models, respectively. The loss differential is aggregated to the yearly frequency since most of the predictors are only available for all countries at the yearly frequency. Trade openness is defined as the sum of official trade in imports and exports of goods and services, scaled by domestic output, $Openness_{i,t} = \frac{Imports_{i,t}+Exports_{i,t}}{GDP_{i,t}}$. $\boldsymbol{Z}_{i,t}$ is a vector of control variables that includes GDP growth, GDP per capita growth, the unemployment rate, the average inflation rate in year t, and the realized variance of inflation in year t, calculated as $RV_{i,t} = \frac{1}{12}\sum_{s=1}^{12}(\pi_{i,s} - \bar{\pi}_i)^2$, where $\pi_{i,s}$ is the inflation rate for country i in month s, and $\bar{\pi}_i$ denoting the average monthly inflation rate for country i in year t. All variables are standardized to have a mean of zero and a variance of one for ease of interpretation. We compute standard errors using the panel corrected standard errors (PCSE) of Beck and Katz (1995) with cross-sectional clustering. The results are robust to other ways of computing standard errors.

Table 3 shows the results of these panel regressions for all horizons considered. Openness is positively related to forecasting outperformance at all horizons except $h = 12$. The result for $h = 12$ is not surprising since the RF model only marginally outperforms the RW model (on average) at the yearly horizon. However, for the horizons at which there is outperformance, the variable is generally statistically significant at the 1% and 5% levels. In terms of control variables, the ones that stand out are the average inflation rate and the realized variance of inflation, both of which are also positively related to forecast outperformance, implying that the RF model also tends to outperform in countries and periods in which there are high inflation rates and high volatility. On the other hand, GDP per capita growth is only marginally significant at horizons of three and six months ahead. Finally, it is worth noting that GDP growth and unemployment are generally negatively related to outperformance, although these results are only marginally significant for some horizons.

Table 3: Panel regression: determinants of forecasting outperformance

Openness	ΔGDP	ΔGDP per cap.	Unemp. rate	Avg. infl.	RV of infl.	R^2_{within}
			$h=1$			
0.129						[0.037]
(4.696)						
0.108	-2.075	1.637	0.016	0.135		[0.074]
(3.855)	(-3.862)	(3.424)	(0.520)	(4.407)		
0.057	-0.904	0.823	-0.007	0.0378	0.565	[0.335]
(2.371)	(-1.729)	(1.719)	(-0.320)	(1.426)	(23.315)	
			$h=3$			
0.099						[0.032]
(3.621)						
0.066	-2.025	1.559	-0.050	0.146		[0.070]
(2.365)	(-3.362)	(3.163)	(-1.184)	(4.799)		
0.012	-0.801	0.709	-0.075	0.044	0.590	[0.185]
(0.513)	(-1.902)	(1.877)	(-3.325)	(1.748)	(25.181)	
			$h=6$			
0.129						[0.140]
(4.843)						
0.112	-0.644	0.674	-0.034	0.088		[0.147]
(3.938)	(-1.289)	(1.528)	(-1.212)	(2.836)		
0.081	-0.087	0.301	-0.045	0.036	0.307	[0.222]
(3.000)	(-0.190)	(0.749)	(-1.736)	(1.211)	(11.004)	
			$h=12$			
0.003						[0.228]
(0.142)						
-0.012	0.229	-0.152	-0.014	0.028		[0.242]
(-0.422)	(0.498)	(-0.370)	(-0.553)	(0.912)		
-0.014	0.263	-0.175	-0.015	0.026	0.020	[0.248]
(-0.489)	(0.570)	(-0.427)	(-0.570)	(0.813)	(0.705)	

The table reports the slope coefficients, t-statistics (in parenthesis) and within R^2 in brackets, for a fixed effects panel regression of the form, $D_{i,t} = \alpha_i + \gamma_t + \beta Openness_{i,t} + \boldsymbol{\delta}' \boldsymbol{Z}_{i,t} + \varepsilon_{i,t}$, where $D_{i,t} = \left(\frac{1}{12}\sum_{s=1}^{12}(\hat{e}_{i,s}^{(h),RW})^2 - \frac{1}{12}\sum_{s=1}^{12}(\hat{e}_{i,s}^{(h),RF})^2\right)$, is the annualized forecast loss differential between the RW and the RF models for country i in year t. Trade openness is the sum of imports and exports divided by GDP, and $\boldsymbol{Z}_{i,t}$ is a vector of control variables that includes GDP growth, GDP per capita growth, the unemployment rate, the average inflation rate and the realized variance of inflation. All variables are standardized to have a mean of zero and a variance of one for ease of interpretation. We compute standard errors using the panel corrected standard errors (PCSE) of Beck and Katz (1995) using cross-sectional clustering. The sample period is 2000 to 2019.

As an additional step in understanding the role of trade integration and global inflation for our results, we examine the overall importance the RF model attributes to its predictors. We calculate variable importance using Equation 3.9 and, for simplicity of exposition, we categorize the predictors into groups. As a result, we have five geographical clusters of inflation series, a group of seasonal dummies, a group with the AR components, and finally, regional and global factors. Figure 4 displays the aggregate Variable Importance (VI) measures for each group across all forecast horizons under consideration. The distribution of VI measures for regional groups as predictors appears relatively even, which is consistent for targets in all regions. European inflation series receive more weight within the model for shorter horizons, while South American countries receive more weight for longer horizons. Factors generally receive a lower weight in the model; however, it's important to note that each factor is a single variable, while the geographical inflation groups consist of an average of 18 variables.[12]

Notably, the seasonal dummies and AR components, which are the only variables unique to each target, are assigned relatively low importance by the RF model, often representing just 20-30% of the total importance. This further emphasizes that it is primarily global inflation, captured by both the factors and the high-dimensional panel of inflation series, that drives the robust forecasting performance of the RF model.

Overall, our results support the idea that trade integration is an important determinant of the globalization of inflation and that these spillover effects in inflation happen in all regions.

5.2. The role of global and regional factors

The general model stated in Equation 3.1 includes both the high-dimensional panel of inflation, *and* a global and regional factor. However, this leads to a natural question: *is the*

[12]Figure 29 in Appendix D presents the equivalent figure with each group divided by the number of predictors in that group. The figure reveals that global and regional factors typically have more weight than any individual country.

Figure 4: Variable importance for the RF model

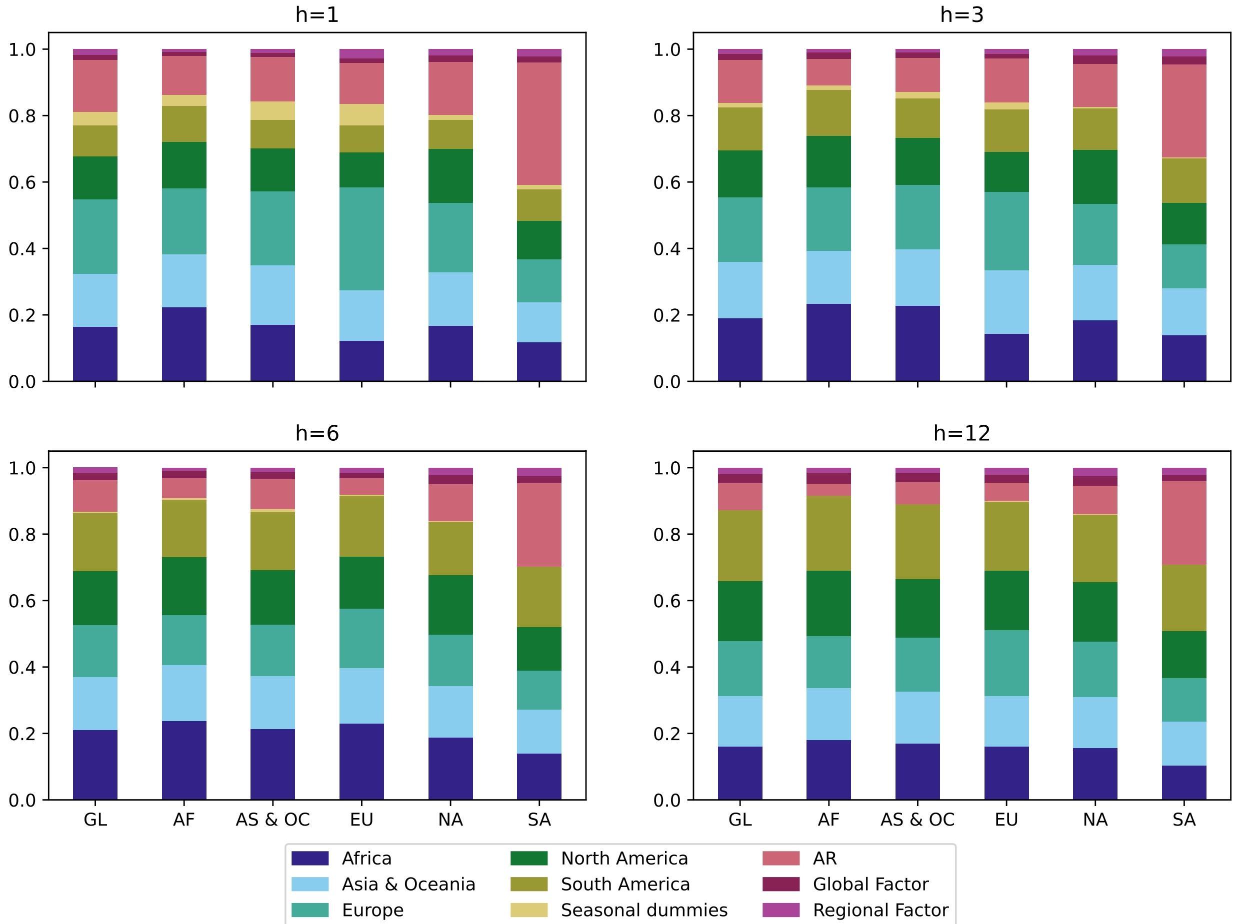

The figure shows the aggregate variable importance (VI) measure of each group of predictors for the RF model at forecast horizons of $h = 1, 2, 6, 12$. The group to which the target variable belongs is shown on the x-axis. The groups are Global (GL), Africa (AF), Asia & Oceania (AS & OC), Europe (EU), North America (NA), and South America (SA). The VI measure is computed using Equation 3.9. The values are scaled to sum to one.

information in the high-dimensional panel of international inflation already embedded in the global and regional factors, or do they act as compliments? This question is closely related to the forecasting methodology since we are indirectly also asking whether the high-dimensional machine learning methods we employ can successfully recover a plausible factor structure in the data or whether these factors need to be explicitly constructed as inputs for the forecasting model.

To answer this question, we use the Harvey, Leybourne, and Newbold (1998) tests for forecast encompassing. The tests consider a convex combination of the RF forecast made using the high-dimensional panel, $\hat{\pi}^h_{I,t+h}$, and the RF forecast made with the factors $\hat{\pi}_{F,t+h}$:

$$\pi^h_{t+h} = \lambda^h_I \hat{\pi}^h_{I,t+h} + \lambda^h_F \hat{\pi}^h_{F,t+h}, \tag{5.2}$$

where $\lambda^h_I + \lambda^h_F = 1$. If $\lambda^h_I = 0$, the factors encompass the high-dimensional forecasts in the sense that the latter do not contribute any relevant information in addition to that contained in the factors. The alternative, $\lambda^h_I > 0$, suggests that the factors do not encompass the information in the high-dimensional panel. Table 4 shows the average λ^h_I and λ^h_F and the share of countries for which certain scenarios based on the statistical significance of λ^h_I and λ^h_F apply. The first scenario is λ^h_I statistically significant while λ^h_F is insignificant, with the second scenario being the reversed case (the two first rows in the second panel). In these scenarios, one would conclude that the inflation panel or the factors encompass the other. The two final scenarios (the two final rows in the second panel) are those in which either both λ^h_I and λ^h_F are statistically significant, implying that both the inflation panel and the factors provide significant information, or that both are insignificant.

Overall, these encompassing tests imply that the high-dimensional panel of inflation series and factors act more as complements than substitutes since the average weight assigned to each is close to 0.5, except for $h = 12$, where the high-dimensional panel gets a higher weight. The p-values on the first two scenarios have an almost equal share of around one-third, which

Table 4: Forecast encompassing tests

Summary statistics	Forecast horizon $h=1$	$h=3$	$h=6$	$h=12$
Average λ_I^h	0.466	0.479	0.568	0.684
Average λ_F^h	0.534	0.521	0.432	0.316
$P_I \leq \alpha \wedge P_F > \alpha$, share in %	30%	31%	42%	33%
$P_I > \alpha \wedge P_F \leq \alpha$, share in %	34%	29%	24%	16%
$P_I \leq \alpha \wedge P_F \leq \alpha$, share in %	34%	30%	12%	0%
$P_I > \alpha \wedge P_F > \alpha$, share in %	2%	11%	22%	51%

The first panel reports the weights from the Harvey, Leybourne, and Newbold (1998) encompassing test as an average computed across all 91 countries. Using the corresponding country-specific p-values, denoted P_I and P_F, and $\alpha = 0.05$, the second panel report the share of countries for which certain scenarios apply. Columns might not sum to 1 due to rounding.

suggests that either of them can encompass the other; however, it is not clear that either of them dominates, with the exception again of $h = 12$, where the high-dimensional panel is significant more often. The third case, in which both sets of predictors are significant, also happens in approximately one-third of the cases for short horizons but decreases as the horizon increases, reaching zero at $h = 12$. Finally, the case where neither is significant happens very seldom for short horizons but reaches more than 50% for the yearly horizon.

As an additional robustness check, we also estimate the factors using two alternative methodologies: PCA and autoencoders. The results of this analysis are shown in Appendix C, and they demonstrate that using alternative techniques for factor estimation has a negligible impact on forecasting performance.

5.3. The role of nonlinearities

In order to assess the role of nonlinearities, we follow Medeiros et al. (2021) and estimate a hybrid RF/OLS model. The procedure follows: For each bootstrap sample $b = 1, ..., B$, grow a single tree with $k = 20$ nodes and save the splitting variables chosen by the tree.[13]

[13] We choose $k = 20$ in order to prevent overfitting the linear model and because this typically matches the number of splitting variables chosen by the random forest.

Next, compute a forecast of the target variable using OLS on the chosen splitting variables, denoted by $\widehat{\pi}^{b,h}_{t+h}$. Finally, we take an average of all OLS forecasts to obtain the final forecast: $\widehat{\pi}^{h}_{t+h} = B^{-1}\sum_{b=1}^{B}\widehat{\pi}^{b,h}_{t+h}$. As such, RF/OLS is a linearized version of the RF model. If the RF/OLS model performs similarly to the RF model, it would indicate that nonlinearities are unimportant. Additionally, if the RF/OLS model improves on other linear models while underperforming the RF model, it would indicate that the gains observed for the RF model are both due to superior variable selection and nonlinearities in inflation.

Table 5: Forecast results: RF/OLS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	RMSE				MAD			
Model	$h = 1$	$h = 3$	$h = 6$	$h = 12$	$h = 1$	$h = 3$	$h = 6$	$h = 12$
RF/OLS RMSE/MAD	0.007	0.005	0.004	0.004	0.004	0.003	0.002	0.002
EN	1.021	1.028	1.032	1.056	0.998	1.016	1.023	1.030
RF	0.957	0.948	0.921	0.908	0.923	0.934	0.917	0.891

The table reports the forecast results for the RF/OLS model with respect to both the RMSE and MAD evaluation metrics and the EN and RF models relative to the RF/OLS model. Columns (1), (2), (3), and (4) report the average RMSE computed across all 91 countries for the RF/OLS model and the ratio of the average RMSE of the EN and RF models relative to the RF/OLS model for $h = 1, 3, 6, 12$. Columns (5), (6), (7), and (8) report the corresponding results for the MAD evaluation metric.

Table 5 reports the forecasting results for the RF/OLS model and the forecast results of the EN and RF models relative to the RF/OLS model for comparison. Similarly to Table 2, we report averages across countries. Columns (1)-(4) report the average RMSE for the RF/OLS model and the average RMSE relative to the RF/OLS model for the EN and RF models for horizons $h = 1, 3, 6, 12$, respectively. Columns (5)-(8) report the corresponding results for the MAD evaluation metric. Overall, the RF/OLS model performs well and follows the patterns observed in Table 2. Compared to the EN and RF models, we find that the RF/OLS model achieves lower average errors than the EN model, except for $h = 1$ using the MAD evaluation metric, and strictly higher average errors compared to the RF model for both evaluation metrics. These results imply that the dominance of the RF model is both due to its ability to approximate nonlinear functions *and* superior variable selection.

The marginal improvement of the RF model over the RF/OLS model is about four to eight percentage points at $h = 1$, depending on the evaluation metric, but grows almost linearly to about ten percentage points at $h = 12$, implying that the importance of allowing for nonlinearities increases with the forecast horizon.

6. Case studies

To complement our analysis of aggregate results, we zoom into five countries as case studies to illustrate empirical results for individual inflation series. We chose the United States, Nigeria, Brazil, Japan, and the United Kingdom. These countries are chosen to represent large economies within each of the regions considered, except for Oceania, which we leave out to conserve space. Table 6 shows the RMSE and MAD of the RW model and ratios relative to this benchmark for all models considered for each of these countries. The results in the table stress the point that, although the RF is the best model on average, there is heterogeneity in model performance across countries. For example, in the US, the best model tends to be the AR model, which implies that international inflation bears little predictive power for US inflation. It is possible that the US, as the largest economy in the world, acts more like a driver of global inflation, and therefore predictive power is unidirectional, going from the US towards other countries but not the other way around. For Brazil and Nigeria, the RF model tends to produce sizeable gains with respect to the RW benchmark, especially at horizons between three and six months, which is also where predictability tends to peak across countries. At short and long horizons, the RW model is competitive. This pattern tends to be common in countries with volatile inflation. The United Kingdom and Japan display a similar pattern in which the Ensemble model tends to perform best but within close range to other models. Interestingly, the gains in relation to the AR model are relatively modest, implying that although international inflation series contain useful predictive information, this informational content is limited in the case of these two countries.

If we look across models, we note that the EN model is likely to underperform the much sparser CM model, which again shows that if the model is linear, most informational content in global and regional inflation is already embedded in the factors. The RF model showcases its stability across countries since it tends to perform close to the best model in cases where it is not the best model.

To further analyze the role of nonlinearities, we analyze the functional form of predictive relationships. To do this, we plot the SHAP values for each individual observation across the span of the predictive variable. We select a representative candidate with important nonlinearities amongst the top five most important predictors (according to the measure in Equation 3.9). Since we use a recursive estimation scheme, we look at the fitted model for the period in which the RF model has the largest difference in squared loss differential relative to the RW model. The plots are shown in Figure 5. The SHAP plots generally indicate strong nonlinearities in the predictive relationships. Some of these predictors display an "S" shape, implying that the effect is small for low values of the predictor. It then increases as the value of the predictor increases but then flattens out at high values. Other predictors, for example, inflation in Mexico when predicting inflation in the US at a horizon of $h = 6$, display a steep function for low values of Mexican inflation, which then flattens out once Mexican inflation approaches 1%. Overlapping SHAP values (y-axis) for the same value of the predictor (x-axis) imply that the predictor in question also exhibits interactions with other predictors. This pattern is evident for many of the predictors, for example, Germany as a predictor for US inflation at $h = 1$. The highly non-linear nature of the RF model, in combination with the high-dimensional nature of our panel, implies that these interactions come from many different predictors.

As a final exercise, we assess the combined significance of the predictors in the RF model for the five countries of interest. To calculate the variable importance, we utilize Equation 3.9. Following the methodology used in Figure 4, we categorize the predictors into five geographical groups of inflation series, a group of seasonal dummies, a group of AR components, and regional and global factors. The aggregate VI measures for each group are displayed in Figure 6 for horizons $h = 1$ and $h = 6$ with values normalized to sum to one. The results reveal that inflation levels from various regions and global factors are given more weight than unique predictors, with the exception of Brazil, where past inflation rates receive the highest weight from the RF model.

At the shorter horizon ($h = 1$), the RF model assigns greater importance to the geographical group of countries that includes the country of interest or to a group of countries with a strong trade relationship. For instance, European and Asian countries are considered most important in predicting inflation in the United States, while other African countries play a significant role in predicting inflation in Nigeria, and other European countries are critical in predicting inflation in the United Kingdom. At the longer horizon ($h = 6$), we observe a more balanced weight assigned to the five geographical groups of inflation series, suggesting that the importance of considering global inflation levels may increase with the forecast horizon.

Table 6: Forecasting results: case studies

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	RMSE				MAD			
Model	$h = 1$	$h = 3$	$h = 6$	$h = 12$	$h = 1$	$h = 3$	$h = 6$	$h = 12$
United States								
RW RMSE/MAD	0.004	0.004	0.003	0.002	0.002	0.002	0.002	0.001
AR	**0.754**	**0.572**	**0.539**	0.783	**0.672**	**0.528**	**0.405**	0.691
CM	0.760	0.577	0.541	0.788	0.690	0.547	0.429	0.731
EN	0.811	0.622	0.596	0.764	0.754	0.638	0.529	0.701
NN	0.865	0.634	0.575	**0.758**	0.838	0.599	0.501	**0.656**
RF	0.848	0.615	0.557	0.777	0.865	0.548	0.455	0.718
Ensemble	0.814	0.611	0.559	0.761	0.850	0.582	0.463	0.692
Nigeria								
RW RMSE/MAD	0.017	0.012	0.009	0.005	**0.004**	0.005	0.004	0.002
AR	0.896	0.862	0.950	1.340	1.541	1.294	1.550	1.722
CM	0.882	0.820	0.859	1.148	1.696	1.267	1.439	1.527
EN	0.786	0.705	0.741	1.081	1.116	0.760	0.816	1.128
NN	0.825	0.769	0.731	1.006	1.316	0.844	0.942	1.077
RF	0.779	**0.676**	**0.620**	**0.889**	1.013	0.710	0.767	1.014
Ensemble	**0.774**	0.685	0.644	0.906	1.115	**0.704**	**0.709**	**0.921**
Brazil								
RW RMSE/MAD	**0.003**	0.003	0.003	**0.003**	0.002	0.002	0.002	**0.001**
AR	3.245	3.920	5.136	7.784	2.609	2.749	4.143	7.669
CM	3.469	4.357	5.074	7.372	2.414	2.713	3.786	4.488
EN	2.132	3.387	4.564	7.120	1.741	2.671	3.626	7.036
NN	2.476	3.396	3.876	4.620	1.465	1.159	1.186	1.350
RF	1.026	**0.901**	**0.812**	1.132	**0.900**	**0.871**	**0.869**	1.160
Ensemble	1.585	2.160	2.619	3.757	1.294	1.349	1.711	2.852
Japan								
RW RMSE/MAD	0.004	0.003	0.002	0.003	0.002	0.002	0.001	0.001
AR	0.752	0.661	0.711	0.879	**0.753**	0.627	0.643	0.882
CM	0.731	0.612	**0.672**	0.853	0.756	0.546	**0.595**	**0.681**
EN	0.749	0.632	0.727	0.833	0.759	0.569	0.699	0.709
NN	0.709	0.632	0.708	0.876	0.821	0.549	0.727	0.807
RF	0.712	0.616	0.680	0.834	0.802	**0.471**	0.649	0.808
Ensemble	**0.698**	**0.610**	0.684	**0.825**	0.783	0.514	0.662	0.742
United Kingdom								
RW RMSE/MAD	0.004	0.003	0.001	0.001	0.002	0.002	0.001	**0.000**
AR	0.496	0.537	0.809	1.089	0.523	0.491	0.842	1.070
CM	0.486	0.501	0.755	1.044	**0.483**	**0.433**	0.755	1.056
EN	0.487	0.495	0.749	0.944	0.555	0.449	0.798	1.075
NN	0.517	0.549	0.767	0.994	0.536	0.504	0.721	1.163
RF	0.491	0.521	0.719	0.975	0.493	0.498	0.666	1.052
Ensemble	**0.468**	**0.494**	**0.711**	**0.923**	0.537	0.469	**0.665**	1.102

The table reports out-of-sample forecast summary statistics for each model in consideration for the countries of the United States, Nigeria, Brazil, Japan, and the United Kingdom. The out-of-sample evaluation window runs from March 2000 to December 2019. Columns (1), (2), (3), and (4) report the RMSE for the RW model and RMSE ratios for the remaining models at $h = 1, 3, 6, 12$, where the ratios are computed between the model of interest and the RW model as a benchmark. Columns (5), (6), (7), and (8) report the corresponding MAD for the RW model and the MAD ratios for the remaining models. The best results for each horizon are highlighted in bold.

Figure 5: SHAP partial dependence plots

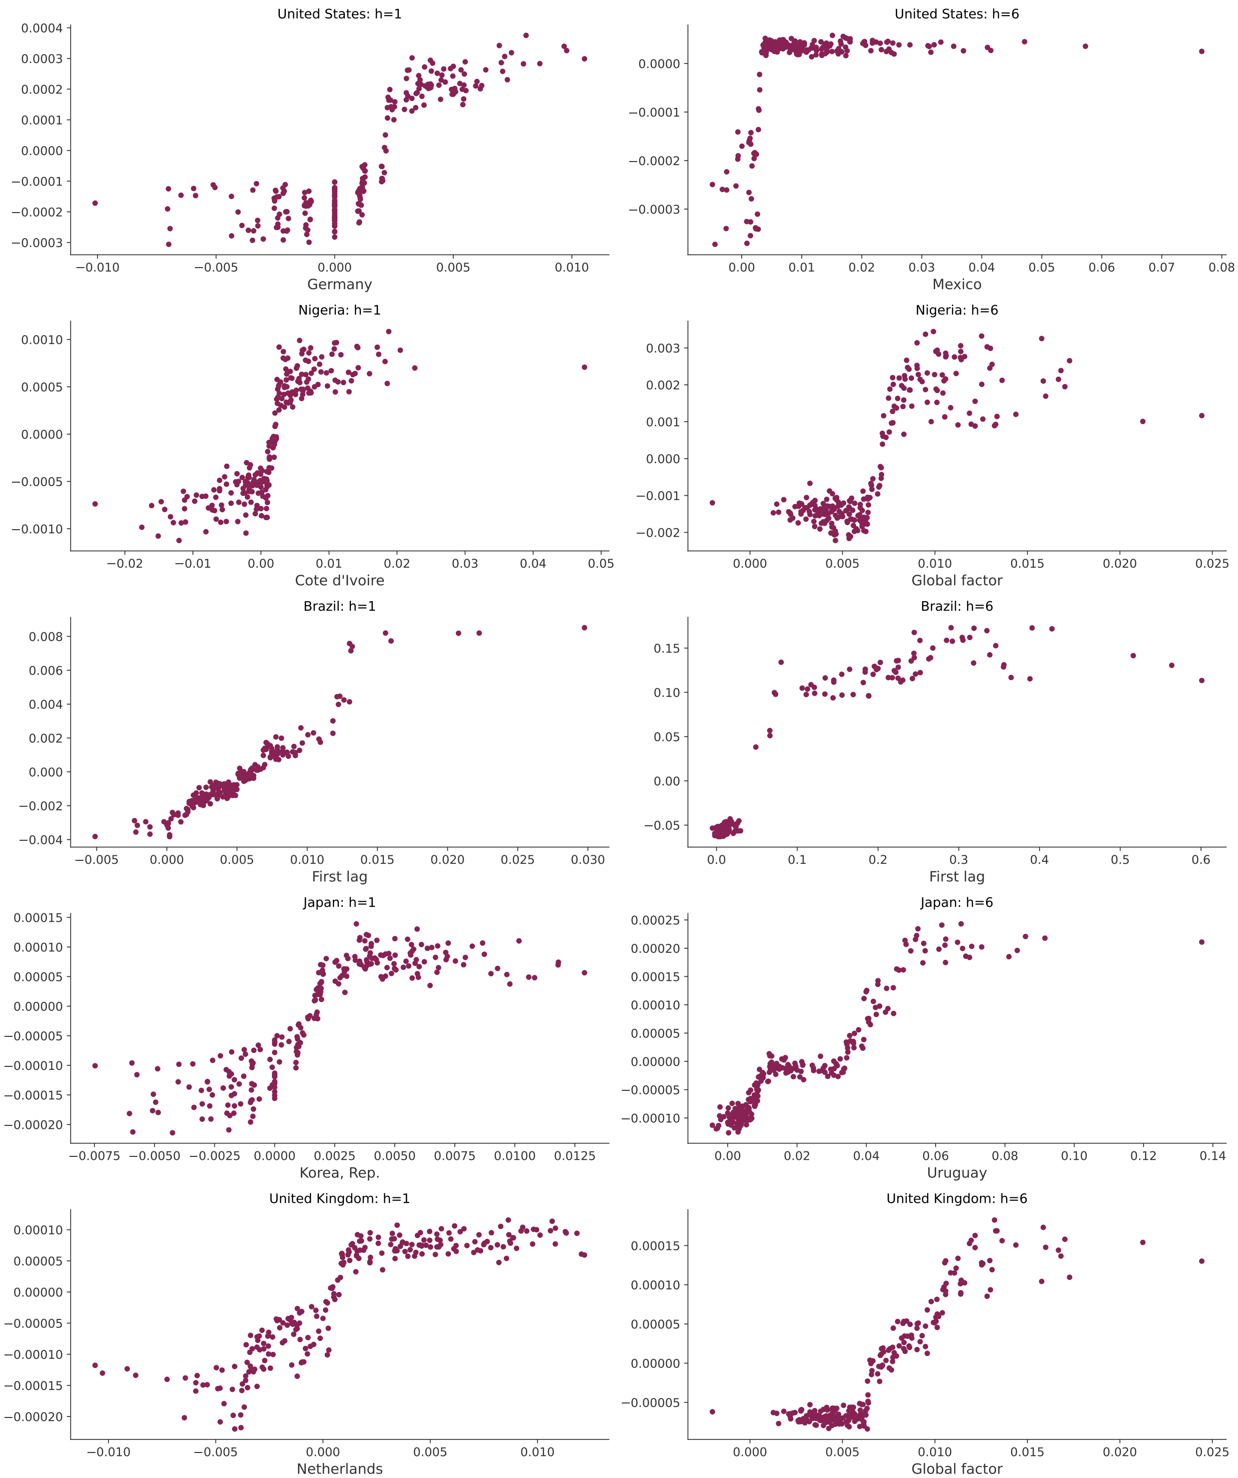

The figure shows the SHAP partial dependence plot for one of the five most important predictors at horizon $h = 1$ (left panel) and $h = 6$ (right panel) for the United States, Nigeria, Brazil, Japan and the United Kingdom. Since we have a rolling window, we select the time period for which the RF model has the largest difference in squared loss differential relative to the RW model and choose the predictor exhibiting the most interesting nonlinear dynamics. For $h = 1$, the RF model has the largest difference in squared loss differential relative to the RW model in February 2018, while it is February 2009 for $h = 6$.

Figure 6: Variable importance by predictor groups

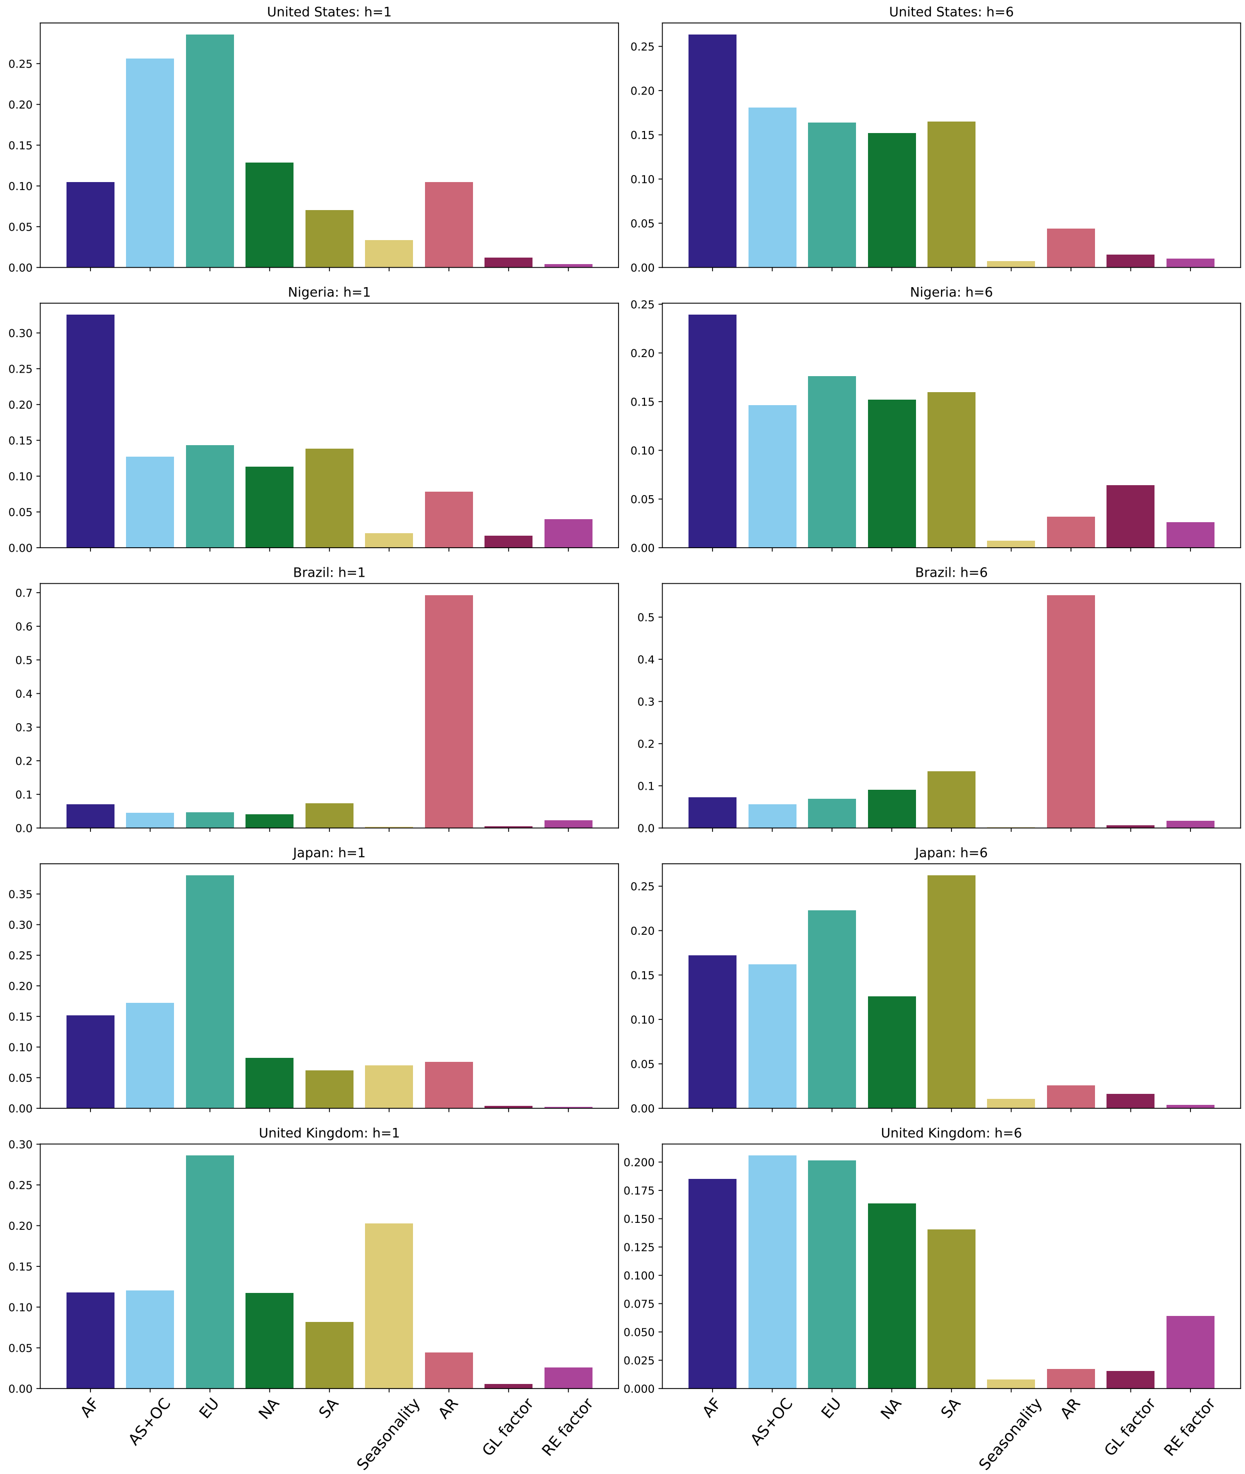

The figure shows variable importance (VI) measures of each group of predictors for the RF model at horizon $h = 1$ (left panel) and $h = 6$ (right panel) for the United States, Nigeria, Brazil, Japan and the United Kingdom. The groups are Africa (AF), Asia & Oceania (AS & OC), Europe (EU), North America (NA), South America (SA), seasonal dummies (Seasonality), autoregressive terms (AR), the global factor (GL factor), and regional factor (RE factor). The VI measure is computed using Equation 3.9. The values are scaled to sum to one.

7. Concluding remarks

We show that inflation series across the globe (global inflation) contain a significant amount of predictive information for individual inflation series. This predictive information is best exploited using nonlinear machine learning methods to successfully select the most important predictors for each country while accounting for nonlinear links in international inflation across time.

Our results support the global inflation hypothesis, which posits that global developments are the primary drivers of inflation worldwide. Predictability is stronger in countries with higher trade integration, implying the importance of considering global factors in forecasting inflation. We also show that most of the information embedded in the high-dimensional panel of inflation series can be subsumed into a global and regional factor. In this sense, our results complement and expand on the findings of Ciccarelli and Mojon (2010) by showing that these common predictive components of inflation extend beyond their original sample of 22 rich OECD countries. This highlights the need for monetary policy coordination to be a global effort.

Our study also extends the recent literature on the role of machine learning in inflation forecasting by showing that its strong performance with domestic predictors extends to the use of international predictors, such as inflation series. Although we do not investigate the potential improvement from combining domestic and international predictors, this is an avenue for future research. The nonparametric nonlinear relationships that machine learning models estimate in the data are difficult to incorporate into structural models. Nonetheless, the significant advantages that this flexibility of functional form demonstrates in the data imply that integrating nonlinear relationships in structural models of inflation transmission across countries is an interesting avenue for future research.

References

Atkeson, A. and L. Ohanian (2001). Are Phillips Curves Useful for Forecasting Inflation? *Federal Reserve Bank of Minneapolis Quarterly Review* 25: 2–11.

Attinasi, M. G., M. Balatti, et al. (2021). Globalisation and its implications for inflation in advanced economies. *ECB Economic Bulletin* 4:

Beck, N. and J. N. Katz (1995). What to do (and not to do) with time-series cross-section data. *American political science review* 89:3, 634–647.

Bernanke, B. S. and J. Boivin (2003). Monetary policy in a data-rich environment. *Journal of Monetary Economics* 50:3, 525–546.

Breiman, L. (2001). Random forests. *Machine learning* 45:1, 5–32.

Cagan, P. (1956). The monetary dynamics of hyperinflation. *Studies in the Quantity Theory if Money.*

Ciccarelli, M. and B. Mojon (2010). Global Inflation. *The Review of Economics and Statistics* 92:3, 524–535.

Coulombe, P., M. Leroux, D. Stevanovic, and S. Surprenant (Aug. 2020). How is Machine Learning Useful for Macroeconomic Forecasting?

Coulombe, P. G. (2022). A Neural Phillips Curve and a Deep Output Gap. *Working paper.*

Garcia, M. G., M. Medeiros, and G. F. Vasconcelos (2017). Real-time inflation forecasting with high-dimensional models: The case of Brazil. *International Journal of Forecasting* 33:3, 679–693.

Giacomini, R. and H. White (2006). Tests of conditional predictive ability. *Econometrica* 74:6, 1545–1578.

Ha, J., M. A. Kose, and F. Ohnsorge (July 2021). *One-Stop Source: A Global Database of Inflation.* Koç University-TUSIAD Economic Research Forum Working Papers 2107. Koc University-TUSIAD Economic Research Forum.

Harvey, D. I., S. J. Leybourne, and P. Newbold (1998). Tests for forecast encompassing. *Journal of Business & Economic Statistics* 16:2, 254–259.

Inoue, A. and L. Kilian (2008). How Useful Is Bagging in Forecasting Economic Time Series? A Case Study of U.S. Consumer Price Inflation. *Journal of the American Statistical Association* 103:482, 511–522.

Kamber, G. and B. Wong (2020). Global factors and trend inflation. *Journal of International Economics* 122: 103265.

Kingston, G. H. and S. J. Turnovsky (1978). A small economy in an inflationary world: Monetary and fiscal policies under fixed exchange rates. *The Economic Journal* 88:349, 18–43.

Kose, M. A., C. Otrok, and C. H. Whiteman (2003). International business cycles: World, region, and country-specific factors. *american economic review* 93:4, 1216–1239.

Levin, A., V. Wieland, and J. C. Williams (2003). The performance of forecast-based monetary policy rules under model uncertainty. *American Economic Review* 93:3, 622–645.

Lundberg, S. M. and S.-I. Lee (2017). A unified approach to interpreting model predictions. *Advances in neural information processing systems* 30:

Medeiros, M. C. and E. F. Mendes (2016). ℓ_1-regularization of high-dimensional time-series models with non-Gaussian and heteroskedastic errors. *Journal of Econometrics* 191:1, 255–271. ISSN: 0304-4076.

Medeiros, M. C., G. F. Vasconcelos, Á. Veiga, and E. Zilberman (2021). Forecasting Inflation in a Data-Rich Environment: The Benefits of Machine Learning Methods. *Journal of Business & Economic Statistics* 39:1, 98–119.

Mishkin, F. S. (2007). Inflation dynamics. *International Finance* 10:3, 317–334.

Molnar, C. (2020). *Interpretable Machine Learning: A Guide for Making Black Box Models Explainable*. Creative Common License.

Mumtaz, H. and P. Surico (2012). Evolving international inflation dynamics: world and country-specific factors. *Journal of the European Economic Association* 10:4, 716–734.

Neely, C. J. and D. E. Rapach (2011). International comovements in inflation rates and country characteristics. *Journal of International Money and Finance* 30:7, 1471–1490.

Quaedvlieg, R. (Jan. 2021). Multi-Horizon Forecast Comparison. *Journal of Business & Economic Statistics* 39:1, 40–53.

Shapley, L. S. (1953). *A value for n-person games, Contributions to the Theory of Games, 2, 307–317.*

Shinkai, Y. (1973). A model of imported inflation. *Journal of Political Economy* 81:4, 962–971.

Stock, J. and M. Watson (2002a). Macroeconomic Forecasting Using Diffusion Indexes. *Journal of Business & Economic Statistics* 20:2, 147–162.

Welch, I. and A. Goyal (2008). A Comprehensive Look at The Empirical Performance of Equity Premium Prediction. *Review of Financial Studies* 21:4, 1455–1508.

Zou, H. and T. Hastie (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67:2, 301–320.

A. Summary statistics

Table 7: Descriptive statistics 1980-2019

Country	Mean	Std	Min	Max	AC1
Europe					
Austria	0.203	0.404	-1.109	2.577	0.112
Belgium	0.217	0.313	-0.609	1.628	0.308
Cyprus	0.248	0.925	-2.518	3.373	0.073
Denmark	0.238	0.421	-0.657	2.535	0.279
Finland	0.242	0.404	-0.749	1.928	0.332
France	0.232	0.359	-1.006	1.690	0.441
Germany	0.168	0.337	-1.010	1.712	0.003
Greece	0.641	1.392	-2.410	5.129	0.157
Hungary	0.755	1.153	-1.592	9.434	0.400
Iceland	0.898	1.416	-1.473	9.970	0.675
Ireland	0.288	0.826	-1.610	7.466	-0.066
Italy	0.354	0.433	-0.581	2.387	0.636
Luxembourg	0.218	0.548	-1.682	1.888	-0.199
Malta	0.193	0.673	-2.284	3.036	0.054
Netherlands	0.182	0.429	-1.323	1.219	0.254
Norway	0.296	0.502	-1.284	3.212	0.173
Portugal	0.513	0.809	-1.423	4.371	0.569
Slovenia	2.637	5.674	-2.071	45.856	0.890
Spain	0.358	0.570	-1.925	2.833	0.352
Sweden	0.264	0.547	-1.352	3.165	0.189
Switzerland	0.130	0.359	-1.052	1.480	0.234
United Kingdom	0.281	0.428	-0.703	3.353	0.290
Average	0.434	0.860	-1.383	5.466	0.280
North America					
Bahamas	0.252	0.434	-1.726	2.753	0.104
Barbados	0.336	0.805	-4.145	6.687	0.152
Canada	0.247	0.391	-1.043	2.594	0.302
Costa Rica	1.058	1.235	-1.705	10.178	0.666
Dominica	0.214	0.947	-4.407	5.068	-0.178
Dominican Republic	0.951	1.624	-3.333	10.926	0.683
El Salvador	0.643	0.939	-3.433	4.560	0.503
Grenada	0.255	0.673	-2.676	4.179	0.224

Continued on next page

Table 7 continued from previous page

Country	Mean	Std	Min	Max	AC1
Guatemala	0.720	1.131	-3.083	7.997	0.366
Haiti	0.905	1.463	-4.619	12.371	0.274
Honduras	0.756	0.789	-1.361	5.202	0.508
Jamaica	1.113	2.189	-6.899	18.232	0.053
Mexico	1.594	2.127	-0.740	14.376	0.902
Panama	0.174	0.403	-0.857	5.188	0.293
St. Kitts and Nevis	0.233	0.715	-2.134	4.981	0.024
St. Lucia	0.241	0.925	-2.547	4.749	-0.029
United States	0.249	0.342	-1.934	1.509	0.540
Average	0.585	1.008	-2.744	7.150	0.317
		South America			
Bolivia	2.822	9.041	-2.578	103.950	0.748
Brazil	5.680	9.317	-0.511	60.098	0.943
Chile	0.726	0.903	-1.311	7.872	0.636
Colombia	1.035	0.971	-1.139	4.485	0.781
Ecuador	1.597	2.494	-0.706	13.232	0.287
y Paraguay	0.926	1.282	-5.050	9.699	0.359
Peru	3.877	10.450	-0.578	160.334	0.604
Suriname	1.930	4.031	-5.891	34.164	0.665
Trinidad and Tobago	0.554	0.819	-2.027	6.322	0.116
Uruguay	1.984	2.121	-0.734	14.334	0.717
Average	2.113	4.143	-2.053	41.949	0.586
		Asia and Oceania			
Fiji	0.351	0.783	-2.083	5.376	0.068
India	0.628	0.818	-2.120	4.474	0.368
Indonesia	0.628	1.143	-1.034	11.869	0.516
Israel	1.687	3.530	-1.338	24.285	0.837
Japan	0.074	0.419	-1.004	2.031	0.157
Jordan	0.349	1.874	-13.991	13.601	0.017
Korea, Rep.	0.353	0.597	-1.189	4.237	0.419
Malaysia	0.231	0.483	-1.954	3.894	0.064
Myanmar	1.201	1.992	-7.020	9.682	0.392
Nepal	0.650	1.525	-3.090	7.949	0.391
Pakistan	0.647	0.790	-1.557	4.404	0.174
Philippines	0.622	0.996	-2.353	8.829	0.498
Samoa	0.477	1.847	-4.964	11.952	0.065
Singapore	0.152	0.458	-1.621	2.000	-0.033

Continued on next page

Table 7 continued from previous page

Country	Mean	Std	Min	Max	AC1
Solomon Islands	0.685	1.265	-3.714	9.108	0.050
Sri Lanka	0.773	1.378	-3.575	5.682	0.333
Taiwan, China	0.172	0.884	-2.610	4.404	-0.015
Thailand	0.279	0.563	-3.015	3.737	0.297
Turkey	2.573	2.648	-1.675	20.479	0.517
Average	0.664	1.263	-3.153	8.315	0.269
Africa					
Algeria	0.680	1.958	-5.099	10.795	-0.006
Botswana	0.690	0.586	-1.445	3.683	0.277
Burundi	0.757	2.045	-5.039	15.763	-0.014
Burkina Faso	0.240	1.938	-12.437	11.754	-0.145
Cameroon	0.367	1.353	-7.696	9.169	0.065
Cote d'Ivoire	0.331	1.267	-5.300	7.685	0.041
Egypt, Arab Rep.	0.917	1.824	-6.899	9.844	-0.093
Eswatini	0.736	1.964	-14.377	18.137	-0.225
Ethiopia	0.702	2.179	-5.960	11.996	0.243
Gambia	0.691	1.325	-3.532	11.654	0.296
Ghana	1.817	2.559	-9.856	12.063	0.687
Kenya	0.903	1.505	-3.121	10.978	0.264
Madagascar	1.003	1.713	-9.471	10.589	0.359
Malawi	1.447	3.093	-9.249	13.410	0.337
Mauritius	0.492	0.886	-2.503	4.567	0.123
Morocco	0.295	0.799	-2.053	3.555	0.283
Niger	0.215	2.282	-18.424	14.433	0.189
Nigeria	1.387	2.817	-9.531	22.314	0.148
Senegal	0.276	1.509	-3.916	11.882	0.259
Seychelles	0.337	1.848	-18.250	21.617	0.148
South Africa	0.685	0.589	-1.142	3.599	0.386
Sudan	2.548	4.282	-14.007	24.935	0.294
Tunisia	0.433	0.478	-1.690	2.564	0.365
Average	0.780	1.774	-7.435	11.999	0.186
Global average	0.782	1.563	-3.610	12.036	0.294

The table reports the mean, standard deviation (Std), minimum (Min), maximum (Max), and first-order autocorrelation (AC1) for monthly headline inflation for each of the 91 countries in the full sample period (1980-2019). Inflation is computed as $\pi_t = 100 \cdot (log(P_t) - log(P_{t-1}))$, where P_t represents the headline CPI measure for each country in the table.

Table 8: Descriptive statistics 1980-1999

Country	Mean	Std	Min	Max	AC1
Europe					
Austria	0.247	0.452	-0.996	2.577	0.169
Belgium	0.275	0.342	-0.459	1.628	0.345
Cyprus	0.371	0.996	-2.518	3.373	-0.060
Denmark	0.344	0.443	-0.657	2.535	0.309
Finland	0.364	0.446	-0.450	1.928	0.357
France	0.347	0.365	-0.332	1.690	0.721
Germany	0.215	0.315	-0.512	1.712	0.272
Greece	1.131	1.438	-2.410	5.129	0.090
Hungary	1.172	1.441	-1.592	9.434	0.302
Iceland	1.427	1.783	-1.473	9.970	0.653
Ireland	0.433	1.050	-0.714	7.466	-0.165
Italy	0.571	0.487	-0.477	2.387	0.592
Luxembourg	0.276	0.386	-1.682	1.829	0.233
Malta	0.228	0.687	-1.718	2.638	0.139
Netherlands	0.209	0.368	-1.323	1.089	0.295
Norway	0.424	0.515	-0.450	3.212	0.189
Portugal	0.877	0.859	-0.780	4.371	0.579
Slovenia	5.049	7.251	-2.071	45.856	0.869
Spain	0.546	0.509	-0.348	2.833	0.316
Sweden	0.419	0.614	-0.853	3.165	0.195
Switzerland	0.225	0.341	-0.623	1.480	0.200
United Kingdom	0.398	0.503	-0.616	3.353	0.338
Average	0.707	0.981	-1.048	5.439	0.315
North America					
Bahamas	0.348	0.436	-1.726	2.608	0.180
Barbados	0.365	0.909	-4.154	6.687	0.057
Canada	0.338	0.390	-0.816	2.594	0.326
Costa Rica	1.577	1.493	-1.705	10.178	0.600
Dominica	0.314	1.151	-4.407	5.068	-0.188
Dominican Republic	1.278	1.843	-1.514	10.926	0.703
El Salvador	1.094	1.050	-3.433	4.560	0.417
Grenada	0.351	0.817	-2.676	4.179	0.206
Guatemala	1.003	1.483	-3.083	7.997	0.318
Haiti	0.958	1.871	-4.619	12.371	0.204
Honduras	1.008	0.990	-1.361	5.202	0.459
Jamaica	1.546	2.958	-6.899	18.232	-0.011
Mexico	2.829	2.428	0.408	14.376	0.860

Continued on next page