

Forecasting inflation with the hedged random forest



Elliot Beck | Swiss National Bank (SNB) / University of Zurich

Michael Wolf | ADIA Lab / University of Zurich

Keywords: Exponentially weighted moving average, linear shrinkage, machine learning

JEL codes: C21, C53, C31, C37, E47

Abstract

Accurately forecasting inflation is critical for economic policy, financial markets, and broader societal stability. In recent years, machine learning methods have shown great potential for improving the accuracy of inflation forecasts; specifically, the random forest stands out as a particularly effective approach that consistently outperforms traditional benchmark models in empirical studies. Building on this foundation, this paper adapts the hedged random forest (HRF) framework of Beck et al. (2024) for the task of forecasting inflation. Unlike the standard random forest, the HRF employs non-equal (and even negative) weights of the individual trees, which are designed to improve forecasting accuracy. We develop estimators of the HRF's two inputs, the mean and the covariance matrix of the errors corresponding to the individual trees, that are customized for the task at hand. An extensive empirical analysis demonstrates that the proposed approach consistently outperforms the standard random forest.

Disclaimer: This policy brief is based on [SNB Working Paper 07](#). The views expressed in this paper are those of the author(s) and do not necessarily represent those of the Swiss National Bank. Working Papers describe research in progress. Their aim is to elicit comments and to further debate.

Forecasting Inflation with the Hedged Random Forest

Despite decades of research, improving inflation forecasts beyond simple benchmarks remains a challenge, as documented by Faust and Wright (2013). Recent work by Medeiros et al. (2021), however, shows that machine learning methods — particularly the random forest—can consistently improve forecast accuracy when applied to large macroeconomic datasets.

Building on this insight, we adapt the hedged random forest proposed by Beck et al. (2024) — a weighted extension of the standard random forest — to inflation forecasting. The method assigns non-equal, possibly negative weights to individual trees based on estimated forecast-error moments, with the aim of minimizing the mean-squared forecast error. We propose estimators for these moments tailored to time-series data. In an extensive empirical analysis using US and Swiss inflation data, our forecasts consistently outperform those of the standard random forest.

Improving Inflation Forecasts with the Hedged Random Forest

We assess the forecasting performance of the hedged random forest using FRED-MD data from 1960 to 2025. The dataset comprises over 100 macroeconomic indicators, augmented with autoregressive terms and principal components, resulting in more than 400 predictors. Inflation is measured as the year-over-year percentage change in line with central bank practice. We consider six inflation measures: headline and core versions of U.S. CPI and PCE, and Swiss CPI. Forecast accuracy is evaluated recursively in a rolling-window out-of-sample framework (also known as a backtest analysis).

Table 1 presents **RMSE** and **MAE ratios** of the hedged random forest relative to the standard random forest, with ratios below one indicating improved performance. Both metrics — RMSE and MAE — show consistent gains of the hedged random forest over the standard random forest, with particularly large gains for core inflation measures. To evaluate whether the observed gains in forecast accuracy are statistically significant, we apply a robust version of the Diebold and Mariano (1995) test and compute a p-value for any scenario (loss function, inflation measure, forecast horizon), resulting in a total of $2 \times 6 \times 12 = 144$ p-values. For the squared-error loss, 61 out of 72 ($\approx 85\%$) p-values are below 0.1, whereas for the absolute-error loss, 70 out of 72 ($\approx 97\%$) p-values are below 0.

Table 1. Relative Forecast Accuracy of the Hedged vs. Standard Random Forest

Panel (a): RMSE ratios						
Target	Forecast horizon					Mean
	1	3	6	9	12	
CPIAUCSL	0.993	0.971	0.955	0.963	0.963	0.966
CPILFESL	0.983	0.955	0.934	0.926	0.924	0.940
PCEPI	0.995	0.964	0.946	0.952	0.959	0.958
PCEPILFE	0.972	0.958	0.953	0.938	0.937	0.950
CHECPIALL	0.990	0.981	0.966	0.961	0.950	0.968
CHECPICOR	0.958	0.932	0.913	0.906	0.901	0.918

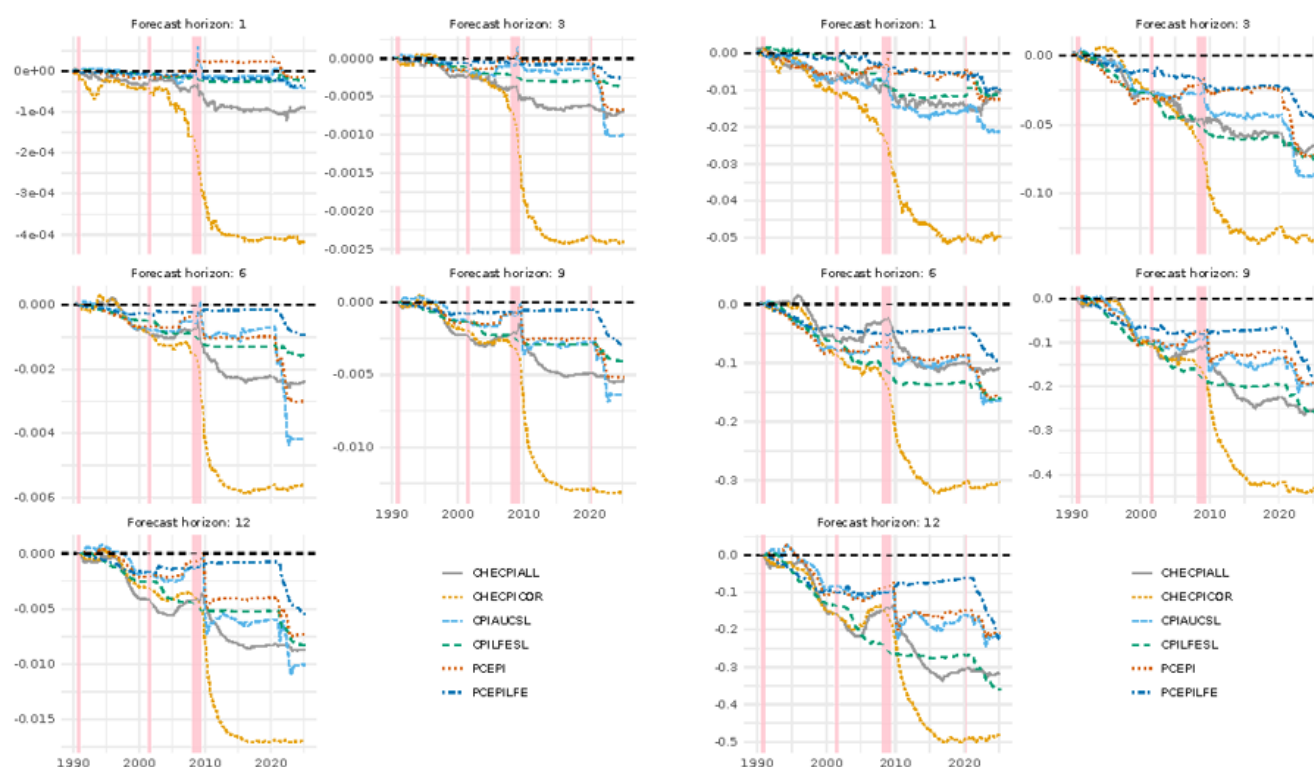
Panel (b): MAE ratios						
Target	Forecast horizon					Mean
	1	3	6	9	12	
CPIAUCSL	0.972	0.953	0.946	0.954	0.960	0.955
CPILFESL	0.969	0.920	0.906	0.903	0.904	0.914
PCEPI	0.979	0.949	0.933	0.941	0.952	0.946
PCEPILFE	0.974	0.947	0.935	0.922	0.927	0.937
CHECPIALL	0.990	0.971	0.964	0.941	0.933	0.958
CHECPICOR	0.953	0.939	0.902	0.901	0.900	0.916

Notes: RMSE and MAE ratios of the hedged random forest relative to the standard random forest for headline and core inflation measures in the US and Switzerland evaluated across forecast horizons of 1, 3, 6, 9, and 12 months, as well as the mean over all twelve horizons (1–12). Ratios below one indicate that the hedged random forest outperforms the standard random forest.

Persistent Outperformance in Changing Economic Environments

To assess the stability of the outperformance of the hedged random forest over time, we analyze the cumulative sum of squared/absolute forecast-error differences (CSSED/CSAED) of the hedged random forest relative to the standard random forest. Figure 1 displays the corresponding trajectories, which show a persistent advantage of the hedged random forest throughout the out-of-sample evaluation period. Performance gains are particularly pronounced (i) following periods of recessions and (ii) during periods of heightened volatility, that is, precisely when accurate forecasts are most valuable for monetary policy. This pattern reflects the method's ability to adapt to changing economic conditions and highlights its robustness across a wide range of macroeconomic environments.

Figure 1.



Notes: CSSED (left panel) and CSAED (right panel) trajectories of relative cumulative performance of the HRF and the RF over time, across different inflation forecast horizons and inflation measures. Negative values indicate that the HRF has outperformed the RF in a cumulative sense until this point in time, whereas downward slopes identify periods where the HRF outperforms the RF in a local sense.

References

- Beck, E., Kozbur, D., and Wolf, M. (2024). The hedged random forest. Working paper. Available at SSRN: <https://ssrn.com/abstract=5032102>.
- Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13:253–263.
- Faust, J. and Wright, J. (2013). Forecasting inflation. In *Handbook of Economic Forecasting*, volume 2, chapter 1, pages 2–56. Elsevier.
- Medeiros, M. C., Vasconcelos, G. F., Veiga, A., and Zilberman, E. (2021). Forecasting inflation in a data-rich environment: The benefits of machine learning methods. *Journal of Business & Economic Statistics*, 39(1):98–119.

About the author(s)

Elliot Beck is a Data Scientist in the Department of Economics at the Swiss National Bank, where he works in the Economic Data Science team, and a Ph.D. candidate in Financial Econometrics at the University of Zurich. He holds a Master's degree in Statistics from ETH Zurich and a Bachelor's degree in Finance from the University of Zurich. His research focuses on machine learning methods for economic forecasting, with an emphasis on non-parametric approaches and their application to inflation forecasting. He has co-authored work on hedged random forests, adapting the method to time series data, and his broader interests include financial econometrics, data science, and natural language processing.

Michael Wolf is a Professor of Econometrics and Applied Statistics at the University of Zurich, and holds a Ph.D. in Statistics from Stanford University. Before joining the University of Zurich's Department of Economics, he held positions at the University of California, Los Angeles (UCLA), Universidad Carlos III de Madrid, and Universitat Pompeu Fabra in Barcelona. His research interests include resampling-based inference, multiple testing methods, the estimation of large-dimensional covariance matrices, and financial econometrics. His work has been published in leading journals such as *The Annals of Statistics*, *Biometrika*, *Econometrica*, *Journal of the American Statistical Association*, and *The Review of Financial Studies*.

SUERF Policy Notes and Briefs disseminate SUERF Members' economic research, policy-oriented analyses, and views. They analyze relevant developments, address challenges and propose solutions to current monetary, financial and macroeconomic themes. The style is analytical yet non-technical, facilitating interaction and the exchange of ideas between researchers, policy makers and financial practitioners.

SUERF Policy Notes and Briefs are accessible to the public free of charge at <https://www.suerf.org/publications/suerf-policy-notes-and-briefs/>.

The views expressed are those of the authors and not necessarily those of the institutions the authors are affiliated with.

© SUERF – The European Money and Finance Forum. Reproduction or translation for educational and non-commercial purposes is permitted provided that the source is acknowledged.

Editorial Board: Ernest Gnan, David T. Llewellyn, Donato Masciandaro, Natacha Valla

Designed by the Information Management and Services Division of the Oesterreichische Nationalbank (OeNB)

SUERF Secretariat

c/o OeNB, Otto-Wagner-Platz 3A-1090 Vienna, Austria

Phone: +43 1 40 420 7206

E-Mail: suerf@oenb.at

Website: <https://www.suerf.org/>