

AI-Based Forecasting of Czech Inflation: Quantile Regression Forests with Dynamic Weights

Filip Blaha, Jan Botka, Josef Švéda,
Aleš Michl



The Working Paper Series of the Czech National Bank (CNB) is intended to disseminate the results of the CNB's research projects as well as the other research activities of both the staff of the CNB and collaborating outside contributors, including invited speakers. The Series aims to present original research contributions relevant to central banks. It is refereed internationally. The referee process is managed by the CNB Research Division. The working papers are circulated to stimulate discussion. The views expressed are those of the authors and do not necessarily reflect the official views of the CNB.

Distributed by the Czech National Bank, available at www.cnb.cz

Reviewed by: Jakub Moučka (CNB)

Project Coordinator: Martin Hodula

Issued by: © Czech National Bank, April 2026

AI-Based Forecasting of Czech Inflation: Quantile Regression Forests with Dynamic Weights

Filip Blaha, Jan Botka, Josef Švéda, and Aleš Michl *

Abstract

We construct a quantile regression forest for inflation forecasting in the Czech Republic, inspired by growing literature on the use of Machine Learning in macroeconomics and finance. We contribute to the literature by implementing an optimisation scheme with time-varying weights that incorporates information from the entire distribution to form the point forecast. By dynamically reflecting the distribution of future inflation paths, our framework outperforms both standard mean and median point forecasts and delivers gains relative to conventional linear benchmark models. We also forecast individual inflation subcomponents that enable us to disentangle the drivers of future inflation and its risks. Furthermore, we integrate the Shapley-value decomposition to enhance the interpretability of our results and adjust the model's predictors for a small open economy.

JEL Codes: C53, C55, E31, E37, E52.

Keywords: Czech Republic, forecasting, inflation, machine learning, quantile regression forest, small open economy, time varying weights.

* Filip Blaha, Czech National Bank; Masaryk University; filip.blaha@cnb.cz

Jan Botka, Czech National Bank; Institute of Economic Studies, Charles University; jan.botka@cnb.cz

Josef Švéda, Czech National Bank; Institute of Economic Studies, Charles University; josef.sveda@cnb.cz

Aleš Michl, Czech National Bank; ales.michl@cnb.cz

We gratefully acknowledge comments and suggestions from Martin Hodula, Michele Lenza, Štefan Lyócsa, and Martin Časta. We also acknowledge the use of OpenAI LLM models for editing the paper. The views expressed are those of the authors and not necessarily those of the Czech National Bank. All errors and omissions remain the fault of the authors.

1. Introduction

The inflation surge across advanced economies in 2022 and 2023 exposed the limitations of traditional macroeconomic forecasting models. Many classical frameworks were built on linear and historically stable relationships. As a result, they struggled to capture the sudden and nonlinear shifts in price dynamics and the macroeconomic environment. This episode highlights two related needs for monetary policy. First, policymakers require models that can detect abrupt changes in inflation dynamics and adapt quickly to the new signals. Second, especially as inflation moves back towards target, they must carefully assess the balance of risks around the forecast. It matters whether inflation risks remain tilted upward or instead point to a sharper disinflation. Misjudging this balance threatens either a reigniting inflation surge or easing policy too late.

To address these needs, this paper develops a forecasting framework that captures both the expected inflation path and the associated forecast distribution. We employ an AI-based¹ machine learning approach — a Quantile Regression Forest (QRF) — building on the growing literature on nonlinear macroeconomic forecasting (Chakraborty and Joseph, 2017; Medeiros et al., 2021; Chen et al., 2019), and in particular Lenza et al. (2025). The QRF structure allows us to model nonlinearities and state dependence without imposing restrictive functional forms.

We extend this framework by introducing a Time-Varying Weighted Quantile Combination (TVW-QRF) model that generates tail-sensitive, dynamically updated point forecasts for Czech inflation. While standard QRF applications typically rely on the median or mean of the conditional distribution, our approach assigns time-varying weights to different quantiles based on recent forecast performance. Our contribution thus lies in dynamically exploiting the entire forecast distribution rather than focusing on the mean or median. The weighting mechanism is inspired by Meligkotsidou et al. (2014), who apply a similar approach in quantile linear regression models for equity premium forecasting. We adapt their approach to the distributional structure of the QRF framework.

Beyond the weighting scheme, we tailor the predictor set to the structural features of a small open economy, placing greater emphasis on external price pressures, global demand indicators, and exchange-rate channels. We further disaggregate the Consumer Price Index (CPI) and generate forecasts for the major CPI subcomponents, enabling us to track heterogeneous price dynamics across economic sectors. Finally, to address the common criticism that machine learning tools operate as “black boxes”, we include a Shapley value decomposition that allows us to quantify the contribution of each predictor to the forecast.

Our findings point to several key results. First, dynamically weighted quantile combinations deliver a significant improvement in point forecast accuracy relative to both the QRF mean and median, as well as to other econometric benchmarks. These results are particularly important during periods of heightened inflation and regime switches, when the ability to shift weight toward tail quantiles of the inflation distribution enables the model to capture the evolving risk balance more effectively than static central-tendency forecasts.

Second, the disaggregated CPI forecasts reveal sizeable heterogeneity across price categories. Food and fuel prices contribute disproportionately to short-term fluctuations, reflecting their sensitivity to global commodity markets and supply shocks, while core inflation evolves more

¹ In line with Russell and Norvig (2022), machine learning is considered a subset of AI. In this paper, “AI-based” refers exclusively to machine learning methods.

gradually and displays a smoother, more persistent evolution. This decomposition enables a more precise understanding of the sources of inflation pressures and their expected persistence—information directly relevant for policy considerations.

Moreover, the Shapley-value decomposition enhances the interpretability of our AI-based framework. Among the predictors employed, domestic real activity, together with external predictors such as foreign producer prices and exchange-rate movements, accounts for a significant share of the forecasted risk, which is consistent with the structural characteristics of a small open economy. Domestic expectations also drive our forecast, especially during periods of elevated uncertainty.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 outlines the methodology, and Section 4 describes the data. Section 5 presents the benchmark models used to evaluate our framework's performance. Section 6 reports the main results and discusses their implications. Section 7 concludes.

2. Related Literature

2.1 Nonlinear Inflation Dynamics and Macroeconomic Modelling

An extensive empirical literature shows that the relationship between inflation and economic activity is neither stable nor linear. Instead, inflation dynamics display state dependence, asymmetric responses, and regime shifts, features that complicate both forecasting and policy analysis. This stands in contrast to the traditional representation of this relationship in the Phillips curve (Phillips, 1958), which forms the backbone of many linear models, including the New Keynesian Phillips Curve (NKPC) and variants used in structural and Dynamic Stochastic General Equilibrium (DSGE) frameworks (Galí and Gertler, 1999).² The traditional Phillips curve is built on a stable link between economic activity and price pressures. However, earlier evidence suggests a flattening of the Phillips curve during recent decades (summarised, e.g., by Smith et al., 2025), with several studies attributing this decline in sensitivity to rising trade openness and global integration (Gilchrist and Zakrajsek, 2019; Firat, 2024; Auer et al., 2017). These lead to reduced responsiveness of domestic prices to domestic economic conditions, providing one explanation for the observed weakening of the relationship (Del Negro et al., 2020).

Yet, recent evidence indicates that such flattening is not permanent. Instead, the underlying dynamics resemble state dependence. During the COVID-19 period and the subsequent inflation surge, researchers documented a re-intensification of the Phillips curve (Benigno and Eggertsson, 2023; Inoue et al., 2024). These findings support earlier evidence that the curve's slope depends on the inflation environment. For example, Benigno and Ricci (2011), Babb and Detmeister (2017), and Forbes et al. (2021) show that the Phillips curve steepens at high inflation levels and flattens when inflation is low, while Harding et al. (2022) and Costain et al. (2022) emphasise the role of nominal rigidities. Moreover, the transmission of shocks itself may be nonlinear, with large shocks propagating more quickly and forcefully than small ones (Cavallo et al., 2024).

Nonlinearities in macroeconomic modelling also arise from financial conditions. A key contribution in this area comes from Adrian et al. (2019), who introduce the concept of “vulnerable growth” and

² This does not diminish the value of structural macroeconomic models in the policy toolkit. Their discipline, interpretability and internally coherent narrative remain substantial for monetary policy, economic analysis, and communication. The point is rather that their linear and stable specifications may face limitations when relationships evolve across regimes.

show, using quantile regression, that adverse financial conditions disproportionately affect the lower tail of GDP growth. Similar results are obtained for Europe by Figueres and Jarociński (2020). This line of work demonstrates that macro-financial linkages are highly state-dependent, with stronger effects in periods of stress. Building on this idea, Adams et al. (2021) and López-Salido and Loria (2024) show that tightening financial conditions predict upward unemployment risks and downside inflation risks. Related contributions (Amburgey and McCracken, 2024; Kiley, 2022; Boyarchenko et al., 2023) further link financial instability to macroeconomic tail outcomes.

Other frameworks directly account for nonlinearity and its effect on economic growth. Chavleishvili and Manganelli (2024) employ quantile vector autoregressions to show that financial shocks translate into more substantial real effects when accompanied by negative macroeconomic shocks, while Chavleishvili and Kremer (2023) highlight the role of systemic stress indicators in explaining the severity of the Great Recession. time-varying quantile models such as Korobilis et al. (2021) point to financial predictors, especially household and private-sector credit, as key drivers of inflation tail risks.

2.2 Inflation Forecasting and Machine Learning

If the price dynamics is shaped by nonlinearities, thresholds and state-dependent forces, the question becomes how to forecast inflation in practice. The limitations of traditional parametric frameworks are well known. Simple benchmarks such as random walks, univariate autoregressions or unobserved-components models remain surprisingly difficult to outperform (Atkeson and Ohanian, 2001; Stock and Watson, 2007, 2010; Faust and Wright, 2013). These outcomes, however, do not necessarily show that they are sufficient (Clark et al., 2024). Rather, it reflects the compromise between parsimony and flexibility of linear models.

The difficulty of forecasting the nonlinear inflation behaviour has prompted growing interest in machine learning (ML). ML methods alleviate several of these issues by flexibly estimating nonlinear functions and interactions without requiring explicit parametric assumptions (Buckmann and Joseph, 2022). Although early applications appeared already in the 1990s (summarised, e.g., by Masini et al., 2023), the widespread use of ML in macroeconomic forecasting became feasible only with broader access to large datasets. For example, Garcia et al. (2017) provide one of the first successful applications using Brazilian inflation data, showing that high-dimensional approaches, such as shrinkage and complete subset regression (CSR) methods, improve forecasting performance.

Subsequent work confirms the strong performance of nonlinear ML models. Masini et al. (2023) show that algorithms such as Random Forests (RF), neural networks, and boosting methods perform exceptionally well in environments with structural breaks, nonlinearities, or many predictors. Likewise, Goulet Coulombe et al. (2022) find that nonlinear ML methods provide substantial gains during periods of macroeconomic uncertainty and financial stress.

RF, introduced by Breiman (2001), has attracted particular attention due to its robustness, interpretability, and strong empirical performance. Using US inflation data, Chakraborty and Joseph (2017) show that RF outperforms both linear benchmarks and competing ML methods, especially at short forecast horizons, a finding later confirmed by Chen et al. (2019). Predictor-importance measures in RF repeatedly emphasise the role of commodity prices, exchange rates, and labour-market indicators. Medeiros et al. (2021) further demonstrate that RF improves forecasts for US core inflation even relative to high-dimensional linear methods such as LASSO or neural-network-based architectures.

For distribution forecasting, Meinshausen (2006) extends RF to QRF, allowing the full conditional distribution of inflation to be estimated. A recent paper by Lenza et al. (2025) finds that QRF outperforms linear quantile regression models for core inflation in the euro area, with gains less significant for headline inflation. These results highlight the importance of modelling nonlinear quantile dynamics for at least some components of CPI.

An alternative distributional approach builds directly on quantile regression rather than on tree-based methods used by Lenza et al. (2025). Meligkotsidou et al. (2014) propose a quantile regression framework for equity premium prediction that combines information from multiple conditional quantiles. These include the trimean of Tukey (1977), the three-quantile estimator of Gastwirth (1966), and the five-quantile estimator suggested by Judge et al. (1988), which naturally exploit the entire distributional information. This Composite Quantile Regression approach is robust against outliers and distribution misspecification, effectively capturing tail risk effects and various forms of predictability across the return distribution. Furthermore, the methodology's ability to incorporate a Time-Varying Weighted (TVW) scheme tends to outperform a purely static approach, as it allows the model to adapt to shifts in market sentiment, regulatory conditions, and fundamental macroeconomic relationships over time, thereby providing superior out-of-sample performance.

2.3 Machine Learning in Small Open Economies

The potential of machine learning frameworks is particularly evident in Small Open Economies (SOEs). Their inflation dynamics are shaped not only by domestic activity but also by import price pass-through, exchange rate movements, volatile global commodity prices, and external demand conditions. These channels introduce additional nonlinearities and state dependencies, while short sample sizes and frequent structural breaks further complicate the forecasting exercise. In this environment, flexible methods that can accommodate a richer set of predictors without imposing many functional restrictions may offer advantages relative to conventional econometric models.

Several studies suggest that tree-based models perform well in SOEs. Using a broad sample of advanced economies, Kohlscheen (2023) finds that RF generally outperforms econometric benchmarks and highlights the role of inflation expectations. Positive results for RF are reported in Pakistan, India, and Taiwan (Mirza et al., 2024; Das and Das, 2024; Li et al., 2023), though Li et al. (2023) note that XGBoost performs even better in their setting. In the Netherlands, van de Winkel and Vedder (2024) show that RF improves short-term inflation forecasting accuracy relative to AR benchmarks, whereas Schnorrenberger and Schmidt (2024) find that shrinkage methods dominate tree-based models when only a small predictor set is used. Evidence from Chile similarly suggests that RF performs well when richer datasets are available (Leal et al., 2020).

On the other hand, some studies report weaker RF performance in countries such as Japan and the United Kingdom. Liu (2024) show that LASSO outperforms RF for Japanese core inflation during a period of unprecedented inflation acceleration, and Joseph et al. (2024) find that ridge regression yields more stable forecasts for UK CPI. They argue that RF's nonlinear flexibility can introduce excessive noise when inflation is relatively stable.

In the context of the Czech economy, applications of machine learning to macroeconomic forecasting remain relatively limited. Vochozka et al. (2021) use neural networks to model Czech inflation and show that the Phillips curve provides little explanatory power. Pavlova (2022) demonstrate that ML methods effectively capture nonlinear relationships between labour costs and inflation. RF is also used to forecast Czech GDP by Gawthorpe (2021), while Benecká (2025)

applies RF and QRF to producer-price inflation and finds them superior to traditional benchmarks.

Although prior studies apply RF and QRF to various macroeconomic settings, to the best of our knowledge, no existing work develops a quantile-combination framework with time-varying weights for CPI inflation, and applications of QRF to consumer price inflation in SOEs remain scarce. This paper addresses this gap, offering a flexible point- and distribution-forecasting framework tailored to economies characterised by short time series, external shocks and structural changes.

3. Methodology

We build directly on the work of Lenza et al. (2025), who applied QRF (Meinshausen, 2006), an extension of the RF framework (Breiman, 2001). Rather than producing a single point forecast, QRF estimates the full conditional distribution of inflation, thereby providing an asymmetric measure of forecast uncertainty.

In the original implementation, point forecasts are typically derived from standard summary statistics of this distribution, such as the mean or median. While these measures are informative, they may not fully reflect the balance of risks during periods of elevated volatility, when tail outcomes become more relevant for policy. We therefore extend the framework by incorporating a broader set of distributional information into the construction of the point forecast. Specifically, we develop a TVW Quantile Combination approach within the QRF framework. This method constructs the final forecast as a weighted average of selected quantiles, including those associated with elevated inflation risks. By allowing the weights to adjust over time, the approach remains sensitive to shifts in the distribution of risks. For a central bank concerned not only with the expected inflation path but also with its surrounding uncertainty, such a framework provides a relevant toolkit to support monetary decisions.

3.1 Quantile Regression Forest

Quantile Regression Forest (QRF), introduced by Meinshausen (2006), extends the classical Random Forest framework of Breiman (2001) by providing a nonparametric estimate of the full conditional distribution of the target variable. The method builds an ensemble of regression trees on bootstrap samples of the training data, where each tree is grown using random subsets of predictors at each split, ensuring low correlation across trees and reducing variance.

For a given input vector X_t , the prediction is constructed by aggregating observations that fall into the same terminal leaves across the forest. This induces a data-driven set of adaptive weights $w_i(X_t)$ over the training sample, reflecting the proximity between X_t and past observations in the feature space. Formally, the conditional distribution of inflation π_{t+h} given X_t is estimated as

$$\widehat{F}_t(\pi_{t+h} | X_t)(q) = \sum_{i=1}^N w_i(X_t) \mathbb{I}(\pi_i \leq q) \quad (1)$$

where the weights $w_i(X_t)$ are determined by the frequency with which observation i appears in the same terminal nodes as X_t across the ensemble.

This representation allows QRF to directly estimate conditional quantiles without imposing distributional assumptions. In contrast to the standard Random Forest, which approximates the

conditional mean, QRF recovers the entire predictive distribution, enabling inference on asymmetric risks and tail behaviour. This property is particularly relevant in macroeconomic forecasting, where the distribution of outcomes may shift across regimes and where policy decisions depend not only on central tendencies but also on the balance of risks.

3.2 TVW-QRF: Quantile Combination Framework

Let $\hat{q}_{\alpha_k}^{(h)}(X_t)$ denote the QRF-based forecast of the conditional α_k -quantile of h -step-ahead month-on-month inflation π_{t+h} , given the predictor set X_t . We consider a finite set of quantiles $\mathcal{A}^{(s)} = \{\alpha_1, \dots, \alpha_K\}$ for a given specification $s \in \{1, 2, 3\}$. The final TVW-QRF point forecast is a convex combination of these quantiles:

$$\hat{y}_{t+h|t}^{(s)} = \sum_{k=1}^K w_{t,k}^{(s,h)} \hat{q}_{\alpha_k}^{(h)}(X_t), \quad (2)$$

where $\mathbf{w}_t^{(s,h)}$ is the vector of time varying weights.

Moreover, we test our TVW-QRF against the typical methods aggregating the conditional distribution of the QRF: the QRF Median ($\hat{q}_{0.50}$) and the QRF Mean (expected value).

3.3 Time-Varying Weight Estimation: A Dynamic Optimisation

Following the approach used for quantile linear regression by Meligkotsidou et al. (2014), we estimate the optimal weights $\mathbf{w}_t^{(s,h)}$ by solving a constrained minimisation problem based on a kernel-weighted least squares criterion over a rolling validation window of 12 months.

3.3.1 Optimisation Problem

The weights $\mathbf{w}_t^{(s,h)}$ are determined by minimising the weighted sum of squared errors between the past realised inflation rates and their corresponding quantile combinations within the validation window:

$$\begin{aligned} \mathbf{w}_t^{(s,h)} = \underset{\mathbf{w} \in \mathbb{R}^K}{\operatorname{argmin}} \underbrace{\left(\mathbf{y}_t^{(h)} - \mathbf{Q}_t^{(s,h)} \mathbf{w} \right)^\top \Omega_t \left(\mathbf{y}_t^{(h)} - \mathbf{Q}_t^{(s,h)} \mathbf{w} \right)}_{\text{Weighted Sum of Squared Errors}}, \quad (3) \\ \text{s.t. } \mathbf{1}^\top \mathbf{w} = 1, \quad \mathbf{w} \geq 0, \quad \ell^{(s)} \leq \mathbf{w} \leq \mathbf{u}^{(s)}. \end{aligned}$$

This quadratic optimisation is subject to constraints: (1) Simplex Constraint ($\mathbf{1}^\top \mathbf{w} = 1$ and $\mathbf{w} \geq 0$); and (2) Stability Bounds ($\ell^{(s)}$ and $\mathbf{u}^{(s)}$), enforcing stability and interpretability.

3.3.2 Exponential Kernel Weights (Ω_t)

The diagonal matrix $\Omega_t = \operatorname{diag}(\omega_1, \dots, \omega_{12})$ contains the exponential kernel weights, which apply greater emphasis to the most recent observations in the 12-month validation window (dynamic weighting). The weight ω_i assigned to an observation i periods old is calculated as:

$$\omega_i = \frac{2^{-\frac{i-1}{\text{HL}}}}{\sum_{m=1}^{12} 2^{-\frac{m-1}{\text{HL}}}}, \quad i = 1, \dots, 12, \quad \text{HL} = 6. \quad (4)$$

Here, $\text{HL} = 6$ denotes the half-life in months. The exponential kernel dynamically assigns greater weight to recent data. The half-life of 6 months ensures that an observation is considered only

half as relevant after six months as the most recent data point. This approach allows our quantile combination to react swiftly and precisely to the current, relevant economic information.

Moreover, it accounts for the well-established notion that macroeconomic relationships and parameters evolve over time and, therefore, tackles the possible limits of our forecast. In practical forecasting exercises, the model is inevitably exposed to omitted-variable bias, limited sample size, and imperfect identification of predictors that may signal regime transitions. Moreover, the predictor space can occasionally attain values that the tree-based model has not previously encountered, which further contributes to systematic forecast bias. In periods of transition into a high-inflation regime that is absent from the historical training data, this bias tends to be negative, with the median prediction systematically underestimating the true inflation rate. The TVW procedure mitigates this effect by shifting attention toward higher predictive quantiles, thereby partially compensating for the model's structural underestimation in the upper tail of the distribution. In an idealised setting with an infinite sample and complete information, the median or mean forecast would be sufficient and unbiased. Under realistic data limitations, however, the quantile-weighted approach provides a more robust approximation by explicitly correcting for asymmetries and regime-dependent distortions that standard central-tendency forecasts might struggle to capture.

3.4 Specification of Weight Schemes

We evaluate the TVW-QRF across three different specifications, s , determined by the chosen set of quantiles ($\mathcal{A}^{(s)}$) and their associated stability bounds ($\ell^{(s)}$ and $u^{(s)}$):

$$\text{TVW1: } \mathcal{A}^{(1)} = \{0.25, 0.50, 0.75\}, \quad \ell^{(1)} = (0.20, 0.40, 0.20)^\top, \quad u^{(1)} = (0.40, 0.60, 0.40)^\top, \quad (5)$$

$$\text{TVW2: } \mathcal{A}^{(2)} = \{1/3, 0.50, 2/3\}, \quad \ell^{(2)} = (0.20, 0.30, 0.20)^\top, \quad u^{(2)} = (0.40, 0.50, 0.40)^\top, \quad (6)$$

$$\begin{aligned} \text{TVW3: } \mathcal{A}^{(3)} &= \{0.10, 0.25, 0.50, 0.75, 0.90\}, \\ \ell^{(3)} &= (0, 0.15, 0.30, 0.15, 0)^\top, \quad u^{(3)} = (0.15, 0.35, 0.50, 0.35, 0.15)^\top. \end{aligned} \quad (7)$$

More specifically, following the approach of Meligkotsidou et al. (2014), we use TVW1 based on Tukey's Trimean (Tukey, 1977) and TVW2 based on the Gastwirth three-quantile estimator (Gastwirth, 1966). The TVW3 specification is based on the five-quantile estimator proposed by Judge et al. (1988).

We evaluate the model using an expanding window forecast. Starting from an initial date T_0 , the QRF model is initially trained on all available historical data up to time $t-h$. At this forecast origin t , the QRF model generates the h -step-ahead forecast $\hat{y}_{t+h|t}^{(s)}$, where the most recent 12 months (up to $t-h$) are reserved for the validation of weights $\mathbf{w}_t^{(s,h)}$. The entire sample window then expands by one period, the QRF model is retrained on the newly available data, and the next forecast is generated. Such a recursive process ensures that, at every step, the prediction is generated using the information available to a forecaster at that time (up to $t-h$).

3.5 Predictor Importance via Shapley Values

Interpreting nonlinear ML models often requires tools that decompose individual predictions into contributions attributable to each predictor. We employ Shapley values, a concept from cooperative game theory developed by Shapley (1953). Given a predictor space $N = \{1, \dots, p\}$, quantile α_k and

horizon h , the Shapley value $\phi_{j,k,h}(x)$ of predictor j represents its average marginal contribution to the model output $\widehat{q}_{\alpha_k}^{(h)}(x)$ across all possible predictor coalitions.

Since the exact computation of $\phi_{j,k,h}$ faces computational constraints, we use approximation methods. Specifically, we implement the Tree SHAP Monte Carlo approach (Štrumbelj and Kononenko, 2014; Lundberg et al., 2020) for every prediction. This method efficiently approximates Shapley values through sampling of predictor permutations and averaging their marginal contributions.

Cross-temporal and cross-horizon comparability requires establishing a model-independent baseline b and reference dataset X_{bg} . We set the baseline at $b = 0$. Therefore, when $\phi_{j,k}(x) > 0$, predictor j at its observed value (with non-coalition predictors sampled from X_{bg}) drives the conditional quantile Q_{α_k} towards a pro-inflationary direction. Conversely, $\phi_{j,k}(x) < 0$ signals deflationary pressure. For the reference dataset X_{bg} , we select observations from 2005-2019, capturing the stable post-EU membership and pre-pandemic era. The Monte Carlo procedure ($n_{sim} = 400$) replaces missing predictors in x with row-wise samples from X_{bg} , maintaining predictor dependencies. This fixed reference set, rather than period-specific training data, ensures consistent counterfactual scenarios across different predictions.

Final adjustment ensures exact decomposition at each prediction point:

$$\widehat{q}_{\alpha_k}^{(h)}(x) - b = \sum_{j=1}^p \phi_{j,k,h}(x) \quad (8)$$

Predictor j 's relative contribution is measured through its proportional share, which we aggregate over time and horizons:

$$s_{j,k,h,t}(x) = \frac{|\phi_{j,k,h,t}(x)|}{\sum_{n=1}^p |\phi_{n,k,h,t}(x)|} \quad (9)$$

We organise the predictors into six economic categories based on economic theory and variable characteristics (Table A6). Reporting individual contributions is not informative given the large number of predictors and their strong multicollinearity. We therefore aggregate Shapley contributions within each group and report their combined share in the forecast. This yields a decomposition in terms of broad inflation drivers. The results are descriptive rather than causal.

4. Data

4.1 Predictors

The empirical foundation of our forecasting framework relies on a comprehensive dataset comprising 72 macroeconomic and financial indicators spanning the period from May 2002 to September 2025.

A detailed overview of all predictors, including their definitions and transformations, is provided in Table A6 in the Appendix. The predictor set is structured into seven primary categories (G1–G7), similar to the dataset constructed by Lenza et al. (2025), designed to capture key economic

channels relevant to inflation dynamics: inflation expectations, real activity, confidence and sentiment, producer price indices, financial sector and commodity prices. Given that the Czech Republic is considered an SOE with deep trade linkages, our dataset also emphasises external pressures, including indicators of foreign demand and price signals from key trading partners, with particular focus on Germany.

First, we seasonally adjusted all time series using the X-13 ARIMA-SEATS method (Findley et al., 1998; U.S. Census Bureau, 2023). Then we tested the stationarity of all series using the Augmented Dickey-Fuller (ADF) test (Dickey and Fuller, 1979) and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test (Kwiatkowski et al., 1992). Based on the testing results, and in line with usual practice, each series is transformed into one of the following forms: (i) no transformation in case that predictor is already stationary, (ii) log-difference ($\Delta \ln(X_t)$), or (iii) simple difference (ΔX_t).

We also made several additional adjustments to our data. First, the fuel price predictor (see Predictor 64 in Table A6) was constructed from four highly correlated time series for various fuel types, obtained from the Czech Statistical Office (CZSO, 2025). It includes Petrol 95 Natural (CZK/l), Petrol 98 Super plus (CZK/l), Diesel (CZK/l), and Liquefied Petroleum Gas (LPG) (CZK/l). We standardised the time series, applied principal component analysis (Jolliffe, 2002), and retained the first component.³

Second, missing initial observations for selected predictors—most notably the Rushin Index (Predictor 14)—are imputed using the factor-based procedure of Chen and Labonne (2023).

Third, incorporating Refinitiv’s commodity futures, namely natural gas and Brent crude oil (see Predictor 65 and Predictor 66 in Table A6), requires a specific time shift to maintain predictive coherence. A futures contract ($F_{t,t+12}$) observed at time t for delivery in twelve months serves as the market’s expectation for the price P_{t+12} . Consequently, to use this predictor in a predictive model, the series is shifted forward by 12 periods in the regression sample. This ensures that the future market expectation observed at time t serves as the appropriate explanatory predictor for inflation at time $t + 12$.

Fourth, predictors capturing labour cost dynamics—specifically the LUCI Total Index (Predictor 15), its subcomponent LUCI Wages and Labour Costs (Predictor 16), and Nominal Unit Labour Costs (Predictor 17)—are not available at a monthly frequency for the Czech Republic. We therefore apply the temporal disaggregation method of Chow and Lin (1971) to interpolate quarterly observations into monthly series.

4.2 Predicted Variable

Unlike Lenza et al. (2025), who focused on annualised h -period inflation, our main predicted variable is the classical month-to-month inflation rate (π_{t+h}) for each individual month at horizons $h \in \{3, 6, 9, 12\}$ months.

The index used for headline inflation is the CPI, constructed by the Czech Statistical Office (CZSO). Moreover, we forecast the month-on-month change in its four subcomponents, which allows us to track the dynamics and nonlinearities within individual price categories. These subcomponents (core inflation, food inflation, administrative prices, and fuel inflation) are computed and published

³ The first component explained over 93% of the total variance.

by the Czech National Bank (CNB). Unlike the Harmonised Index of Consumer Prices (HICP) targeted by the European Central Bank (ECB), the CPI (and correspondingly core inflation) also includes imputed rent, reflecting the cost of owner-occupied housing. The month-on-month inflation rate (π_t) is defined as:

$$\pi_t = \left(\frac{CPI_t}{CPI_{t-1}} - 1 \right) \times 100\% \quad (10)$$

5. Benchmark Models

As a first class of benchmarks, we employ a standard Autoregressive process (AR) of third order, a Random Walk, and an Autoregressive Integrated Moving Average (ARIMA) process.

The AR process of order p , which serves as our main linear reference, is formally defined as:

$$\pi_t = c + \sum_{i=1}^p \phi_i \pi_{t-i} + \varepsilon_t \quad (11)$$

The Random Walk (RW) model, often serving as the fundamental persistence benchmark (where the best forecast is the last observed value), is also included. This model assumes that the inflation rate follows:

$$\pi_t = \pi_{t-1} + \varepsilon_t \quad (12)$$

The ARIMA(p, d, q) model generalizes these processes by combining the autoregressive (ϕ), integrated (differencing, d), and moving average (θ) components. Using the backward shift operator B , the ARIMA model is formally defined as:

$$\phi(B)(1 - B)^d \pi_t = c + \theta(B)\varepsilon_t \quad (13)$$

where π_t is the inflation rate, $\phi(B)$ and $\theta(B)$ are the autoregressive and moving average operators, and ε_t is the white noise error term. The ARIMA extension allows for modelling non-stationarity through differencing, while the AR and MA components capture autocorrelation patterns and error term dependencies. In our case, the (3, 1, 3) specification is applied.

Last but not least, as a key linear competitor, we use the Ensemble Linear Quantile Regression (LQR), following the methodology described by Lenza et al. (2025). LQR is a linear model that explicitly estimates the α -quantile (\hat{q}_α) by minimizing the asymmetrically weighted absolute errors (Koenker and Bassett, 1978):

$$\min_{\beta} \left[\sum_{i:\pi_i \geq X_i' \beta} \alpha |\pi_i - X_i' \beta| + \sum_{i:\pi_i < X_i' \beta} (1 - \alpha) |\pi_i - X_i' \beta| \right] \quad (14)$$

To mimic a linear QRF benchmark, the LQR ensemble aggregates forecasts from approximately 500 individual quantile regression models. Each of these 500 LQR models is randomly fit using four predictors plus three lagged terms.

6. Results

6.1 Weighted Approaches Evaluation

We train our model from May 2002 to December 2010, then forecast for a given forecasting horizon, move by one month, and repeat this process until our data sample is exhausted. By using this expanding-window out-of-sample strategy, we aim to mimic real-time forecasting.

We report the results in Table 1. The TVW point forecasts consistently outperform the baseline mean and median⁴ across all forecast horizons, as measured by the Root Mean Squared Error (RMSE). The same pattern holds for each subcomponent of headline inflation, with the exception of administrative prices, which are, in general, very difficult to forecast from past (see Appendix Table A1).

The findings indicate that relevant predictive information resides beyond the central tendency of the QRF distribution and is not fully captured by the mean or median. The time-varying weighting scheme appears to exploit this information by dynamically reallocating weights across quantiles in response to evolving economic conditions. Among the alternative specifications, TVW3 performs best and is statistically superior to both TVW1 and TVW2 across horizons and components. We, therefore, retain TVW3 as the benchmark specification for the remainder of the analysis.

Table 1: RMSE Comparison: TVW Models vs. QRF Derivatives (Headline Inflation)

Horizon	Median	Mean	TVW1	TVW2	TVW3	
3 Months	0.731	0.713	0.694	0.701	0.669	**
6 Months	0.708	0.723	0.688	0.693	0.661	*
9 Months	0.748	0.737	0.732	0.738	0.712	
12 Months	0.707	0.700	0.687	0.690	0.662	*

Note: Boldface indicates the best model according to the lowest RMSE. †, *, **, and *** denote statistical significance of the improvement in RMSE at the 10%, 5%, 1%, and 0.1% levels, respectively, based on the Diebold–Mariano test.

6.2 Comparative Benchmark Performance

Our benchmark comparison shows a strong empirical performance of our model relative to conventional linear benchmarks. The TVW3 specification mostly outperforms all econometric benchmarks across all forecast horizons ($h = 3, 6, 9$ and 12 months) for headline (Table 2), core and food inflation (see Appendix Table A2), demonstrating that our dynamic, quantile-based approach effectively exploits nonlinear information and inflation asymmetry that traditional models fail to capture.⁵ For fuel inflation and administrative prices (see Appendix Table A2), our model also

⁴ Lenza et al. (2025) employs the median as the point forecast within the QRF framework.

⁵ LQR Ensemble shows a slightly smaller RMSE for a twelve-month ahead prediction of headline inflation. Random Walk performs exactly the same for the 3-month-ahead prediction of core inflation.

outperforms linear benchmarks across most horizons, with some exceptions where the benchmark models provide forecasts of similar quality.

For the QRF derivatives, the results show that the TVW3 specification consistently outperforms alternative variants, with statistically significant improvements at most horizons, indicating robust gains from the proposed weighting scheme.

Table 2: RMSE Comparison: TVW3 vs. Econometric Benchmarks (Headline Inflation)

Horizon	LQR Ensemble	RW	AR	ARIMA	TVW3
3 Months	0.739	0.896	0.710	0.716	0.669
6 Months	0.748	0.715	0.740	0.736	0.661 *
9 Months	0.765	1.005	0.753	0.787	0.712 .
12 Months	0.659	0.874	0.765	0.772	0.662


Note: Boldface indicates the best model according to the lowest RMSE. ., *, **, and *** denote statistical significance of the improvement in RMSE at the 10%, 5%, 1%, and 0.1% levels, respectively, based on the Diebold–Mariano test.

Overall, the results indicate that our TVW-QRF specification outperforms classical linear benchmarks across most horizons for all inflation subcomponents. This provides evidence of a nonlinear structure and shifting tail risk in Czech inflation dynamics, which our time-varying QRF specification can better capture.

6.3 Distribution Calibration and Evaluation

Let $F_t(\cdot)$ denote the forecast Cumulative Distribution Function (CDF). We assess our distribution forecast calibration by testing whether the Probability Integral Transform (PIT)⁶ of the realised values of π_t follows a uniform distribution $U(0, 1)$ (Diebold et al., 1998). For this task, we adopt the procedure developed by Rossi and Sekhposyan (2019), which is based on the Kolmogorov–Smirnov test and also provides a graphical representation of calibration results. The diagonal line in Figure 1 represents perfect calibration, while the dashed edge lines indicate the 95% confidence interval; the red line corresponds to our forecast.

⁶ This represents the CDF corresponding to $F_t(\cdot)$.

Figure 1: Calibration Results for Headline Inflation (LQR vs QRF)

Note: The solid black line represents the empirical CDF of the PIT values. The dashed red lines denote the 95% confidence bands for the Kolmogorov–Smirnov test, following Rossi and Sekhposyan (2019). The left column represents the results for the benchmark model LQR, and the right column for the QRF.

For headline inflation, we do not reject the null hypothesis for either the QRF or LQR models, indicating that the forecasts from both models are well calibrated across all horizons ($h = 3, 6, 9, 12$). This is evidenced by the empirical CDF of the PIT values (solid black line) lying primarily within the 95% confidence bands (dashed red lines). For individual inflation subcomponents, minor deviations emerge at longer horizons (see Appendix Figures A1, A2, A3, A4). However, these departures are limited in magnitude and do not materially alter the overall calibration performance. Overall, the results suggest that QRF distribution forecasts remain well-behaved, particularly at shorter horizons.

We assess the probabilistic performance across models using the Continuous Ranked Probability Score (CRPS). As a proper scoring rule (Gneiting and Raftery, 2007), it measures the distance between the forecast CDF and the realised inflation. Unlike point forecast metrics such as the RMSE, it evaluates the entire distribution, jointly capturing sharpness and calibration. Lower values, therefore, indicate superior forecast performance.

The CRPS results for headline inflation (Table 3) indicate that the QRF distribution forecast performs particularly well at the three- and nine-month horizons. At the six- and twelve-month horizons, the performance of QRF and the LQR ensemble is broadly comparable.

We find a similar pattern across inflation subcomponents, although the relative performance varies (see Appendix Table A3). For core inflation, QRF delivers consistently lower CRPS values across three-, six- and nine-month horizons. For fuel and food inflation, differences between the models are modest and economically small, and for administrative prices, LQR achieves slightly lower CRPS

values across the horizons.

Table 3: Continuous Ranked Probability Score (CRPS) Approximation (Headline Inflation)

Horizon	LQR	QRF
3 Months	0.298	0.284
6 Months	0.296	0.295
9 Months	0.316	0.308
12 Months	0.294	0.294

Note: Boldface indicates the best model according to the lowest CRPS approximation.

6.4 Forecast Distribution and Realised Inflation

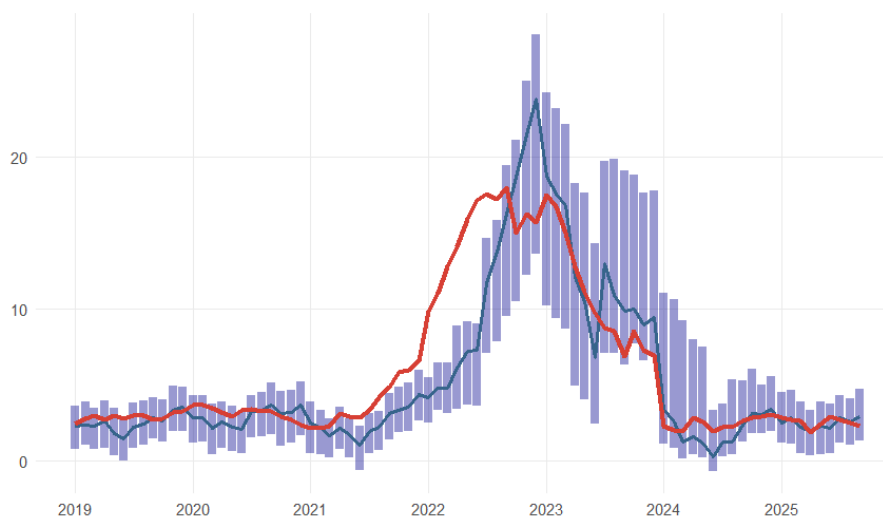
The forecast distribution allows us to assess how inflation risks, as captured by the model, evolve around the point forecast. It indicates whether inflation is more likely to exceed or fall short of the projection. Figure 2 presents the out-of-sample distribution forecast, representing the risk, for the headline year-on-year (YoY) inflation at a six-month horizon ($h = 6$).⁷ The YoY forecast is constructed using realised inflation at $t - 6$ as the base period and the predicted cumulative inflation over the subsequent six months. The blue area represents the 90% prediction interval constructed from the 5th and 95th conditional quantiles. The 5th quantile captures downside inflation risk, while the 95th quantile captures upside inflation risk. The red line denotes the realised inflation. For most of the sample, the realised inflation rate lies within the interval, which indicates that the model captures the prevailing uncertainty reasonably well.

A clear exception emerges during the inflation surge of 2021–2023. In this period, realised inflation exceeds the upper tail. This episode reflects a sharp and unprecedented shift in inflation dynamics that most contemporary forecasting models failed to anticipate. Since the QRF model draws conclusions only from historical relationships, which were marked by low and stable inflation, it was not able to fully accommodate such extreme values in real time.⁸ Importantly, once observations from the high-inflation period enter the training sample, the model adapts. The updated distribution becomes wider and more responsive to large inflationary movements, which suggests improved robustness to future shocks.⁹

⁷ We show in Table A4 that the TVW3 specification gives the best performance and also beats the benchmark models (Table A5).

⁸ This highlights a known limitation of tree-based methods: they do not extrapolate beyond the range of observed historical paths. While linear models are capable of some extrapolation, they did not perform better because of the ineffective reflection of state-dependent relationships between the explained variable and predictors.

⁹ We construct an in-sample analysis to support our claim that the model is now well prepared to forecast periods of high inflation (see Figure A5).

Figure 2: Model Performance: Headline YoY Inflation ($h=6$)


Note: The figure illustrates the realised YoY inflation (red line), our point forecast of the TVW3 specification (dark blue line), and the corresponding 90% prediction interval (QRF 5th and 95th quantiles, light blue shaded area).

Although inflation has declined from its peak after 2022 and moved closer to target, the balance of risks has not fully normalised. To assess how much renewed upward pressure may still be present, we examine the relationship between realised inflation and the upside inflation risk.

Figure 3 illustrates this evolution. Specifically, we compute the difference between realised year-on-year inflation and the corresponding 95th conditional quantile. We interpret this measure as an indicator of upside inflation pressure. Positive values indicate that realised inflation remains below the 95th quantile, while negative values imply that realised inflation exceeded it.

Before 2022, the upside inflation pressure was small and relatively stable, reflecting that the risk of an inflation surge was limited. This changed at the end of 2021, when realised inflation briefly exceeded the 95th conditional quantile. During the period of elevated inflation between mid-2022 and 2024, the model signalled strong upside inflation pressure.

While inflation has since declined from its peak, the behaviour of upside inflation pressure suggests that the risk environment has not fully normalised. As shown in Figure 3, the average level (black dashed lines) of upside pressure since mid-2024 remains somewhat higher than in the pre-pandemic period. This suggests that risks remain tilted to the upside relative to earlier years. Even if point forecasts have moved close to target, policymakers should remain cautious in reducing policy rates too quickly, as residual upside pressures could re-ignite inflation.

Figure 3: Evolution of Upside Inflation Pressure ($h=6$)

Note: The figure illustrates the difference between realised YoY inflation and the corresponding 95th quantile, interpreted as upside inflation pressure (red line). The dashed black line represents the average inflation pressure, split into two periods: before the high-inflation episode (2019–2021) and the period of return to the inflation target (from mid-2024 onwards).

6.5 Disaggregated Upside Inflation Risk and Historical Decomposition


We next examine the sources of upside inflation risk (the upper 95th quantile) to better interpret our results for policy making analysis. We proceed in two complementary steps. First, we disaggregate headline inflation into its main CPI components and examine their contribution to the upside inflation risk. This identifies which price categories account for higher inflation. Second, we apply a Shapley value decomposition to descriptively attribute upside inflation risk to the underlying predictors and to trace how their contributions shift between periods of relative price stability and episodes of elevated inflation.

6.5.1 Disaggregated Inflation Risk

We begin with the disaggregated structure of the CPI. Specifically, we compute the upside inflation risk for three key components: *i*) core inflation, *ii*) food inflation, and *iii*) fuel inflation—alongside headline inflation. To give the subcomponent forecasts an economic interpretation, we weight the forecasts by their respective shares in the CPI basket.

Figure 4 shows that during the inflation surge of 2022–2023, fuel prices contributed in some periods most strongly to the 95th quantile forecasts. The dominant role reflects their exposure to global commodity markets, exchange rate movements, and supply disruptions.

Core inflation, together with food inflation, behaves differently. Its 95th quantile evolves more gradually and exhibits smaller short-term fluctuations. However, once elevated, upside risks tend to persist. By late 2023, fuel prices still account for a substantial portion of the remaining upside inflation risk, while core inflation contributes more steadily but with less magnitude. The results suggest that a large portion of the upside risk lies outside the influence of monetary policy. As a consequence, there is a lower need for a tighter monetary policy because of the upside risks.

Figure 4: 95th Quantile Forecast Value Comparison Across Inflation Components

Note: The figure displays the predicted 95th quantile for YoY headline inflation and its three subcomponents over the out-of-sample period.

6.5.2 Predictor Importance using Shapley Values

The forecasts presented above tell us how the risks evolve. The next question is what drives these risks within the models. To answer this, we decompose the predicted quantiles using Shapley values. This decomposition is descriptive, not causal. Shapley values attribute each forecast to the average marginal contribution of predictors within the trained QRF model, not to how inflation transmits through the economy.¹⁰ Shapley values decomposition should, therefore, serve as a complement to structural models, not its substitute.

We focus primarily on the 95th conditional quantile, which represents the upside inflation risk, and complement it with results for the 5th quantile, representing the downside risk. Because the decomposition is applied directly to the conditional quantile, it identifies drivers of the risks and not the drivers of the point forecast.


The 95th quantile decomposition (Figure 5) illustrates which groups of predictors contribute most to forecasting upside inflation risk. Each colour corresponds to a predictor group as defined in Table A6, and its share reflects the group's relative predictive importance. We see evolving dynamics throughout the years. From 2012 to 2017, the upside risk was largely associated with real activity (light blue), pointing to domestic demand as the main source of the risk. Later, before the COVID pandemic, foreign influence (turquoise) gained prominence, suggesting an increased role for external shocks.

¹⁰ Due to the large number of variables and the interconnectedness of economic indicators, some predictors may exhibit similar trends and be significantly correlated. While this does not pose any problem for the prediction itself, which is not affected by multicollinearity, similar variables might steal each other's contribution to the forecast and thus affect the Shapley values' share. Grouping variables resolves within-group collinearity but does not address similarity between variables across groups.

During the COVID pandemic, the risk was spread out, and no group seemed to be its driver. This changed dramatically during the peak inflation period, when inflation expectations (orange) became the dominant contributor to the upper tail. In this phase, the model increasingly linked upside risk to potential shifts in expectations, which might be interpreted as the possibility of unanchored inflation expectations. However, commodity and energy variables (light red) play a comparatively modest role in the decomposition, despite the result from above showing the prominent role of fuel prices and their widely recognised macroeconomic importance during this period. This likely reflects the fact that inflation expectations indicators embed information from multiple channels, including commodity prices. Any causality remarks should therefore be taken with a grain of salt.

Although the importance of expectations has decreased since 2023, it remains above its pre-2022 levels. The latest forecasts suggest that real activity (light blue) and producer prices (PPI, light green) now account for a larger share of the upside risk. This shift suggests that demand-side conditions and cost pressures currently play a larger role in explaining the upside risks.


Figure 5: Predictor Contributions to the 95th Quantile Forecast for Headline Inflation ($h=6$)



Note: The figure shows the importance of different economic groups in the prediction of the 95th quantile based on the shares of absolute Shapley values.

Turning to the downside inflation risk, Figure 6 reports the Shapley decomposition of the 5th conditional quantile. The importance dynamics across predictor groups is much more stable over time. Downside risks are predominantly shaped by PPI (light green), and real activity (light blue) predictors.

During the 2022–2023 inflation surge, inflation persistence became relevant for a short period. In an environment characterised by a dynamics not previously observed in the training sample, the model placed greater weight on persistence. This reflects an adaptive learning mechanism rather than a structural claim about transmission. However, after the inflation returned to the target, the largest importance is again associated with PPI.

Figure 6: Predictor Contributions to the 5th Quantile Forecast for Headline Inflation ($h=6$)

Note: The figure shows the importance of different economic groups in the prediction of the 5th quantile based on the shares of absolute Shapley values.

7. Conclusion

This paper addresses a practical monetary policy question: how to forecast inflation when economic relationships are nonlinear and state-dependent? In such environments, linear models built on stable historical relationships struggle to capture shifts in dynamics. At the same time, policymakers do not act on point forecasts alone. They assess the risks surrounding the projection. Sound policy requires understanding not only the most likely path of inflation, but also reflects whether risks are tilted upward or downward.

We respond to this need by developing a quantile regression forest framework tailored to the Czech economy and improved it with a time-varying weighted quantile combination. Our methodology exploits the full conditional distribution of inflation. By reallocating weights across quantiles over time, the model adapts to changes in the risk environment and incorporates tail information directly into the construction of the point forecast. Empirically, the dynamically weighted specification outperforms both the standard QRF mean and median, as well as conventional linear benchmarks, across forecast horizons.

Beyond point forecast accuracy, we quantify the risks surrounding the projection. We construct an indicator of upside inflation pressure based on the distance between realised inflation and the upper conditional quantile. We find that, even after inflation declined from its peak, the average level of upside inflation pressure remains slightly higher than in the pre-surge period.

We further break headline inflation into its main subcomponents to identify the source of these pressures. The results reveal that they have recently been largely driven by fuel prices. This component reflects global commodity developments, where the direct influence of monetary policy is limited. To enhance interpretability, we applied a Shapley-value decomposition to the forecasts. The contribution of predictors varies across states of the economy. In the most recent period,

domestic real activity indicators, producer prices, and inflation expectations have played a prominent role in shaping the prediction of upside inflation risk.

The framework has clear limitations. It did not fully anticipate the magnitude of the 2021 inflation surge. This reflects a structural property of tree-based models: they do not extrapolate beyond patterns present in the training sample. Given a historical record dominated by low and stable inflation, extreme values were largely absent from the training data. However, once high-inflation observations entered the training set, the model adapted and is now equipped to handle potential inflation surges in future.

The contribution of this paper lies in organising information about inflation and the surrounding risks. As inflation returns toward target, understanding the sources of the remaining risks becomes increasingly important. A tool that shows whether risks stem from globally driven signals or from more persistent domestic pressures can help inform this assessment. The TVW-QRF model does not prescribe policy choices nor replace policy judgement, but offers a structured way to read developments in inflation and provides an interpretation of how the balance of risks evolves over time.

References

- ADAMS, P. A., T. ADRIAN, N. BOYARCHENKO, AND D. GIANNONE (2021): “Forecasting Macroeconomic Risks.” *International Journal of Forecasting*, 37(3):1173–1191.
- ADRIAN, T., N. BOYARCHENKO, AND D. GIANNONE (2019): “Vulnerable Growth.” *American Economic Review*, 109(4):1263–1289.
- AMBURGEY, A. AND M. W. MCCracken (2024): “Growth-at-Risk Is Investment-at-Risk.” FRB St. Louis Working Paper No. 2023-020, Federal Reserve Bank of St. Louis.
- ATKESON, A. AND L. E. OHANIAN (2001): “Are Phillips Curves Useful for Forecasting Inflation?” *Federal Reserve Bank of Minneapolis Quarterly Review*, 25(1):2–11.
- AUER, R., C. E. V. BORIO, AND A. J. FILARDO (2017): “The Globalisation of Inflation: The Growing Importance of Global Value Chains.” BIS Working Paper No. 602, Bank for International Settlements.
- BABB, N. R. AND A. K. DETMEISTER (2017): “Nonlinearities in the Phillips Curve for the United States: Evidence Using Metropolitan Data.” *Finance and Economics Discussion Series*, 2017-070. Board of Governors of the Federal Reserve System.
- BENECKÁ, S. (2025): “Forecasting Disaggregated Producer Prices: A Fusion of Machine Learning and Econometric Techniques.” CNB Working Paper No. 2/2025, Czech National Bank.
- BENIGNO, P. AND G. B. EGGERTSSON (2023): “It’s Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve.” NBER Working Paper No. 31197, National Bureau of Economic Research.
- BENIGNO, P. AND L. A. RICCI (2011): “The Inflation-Output Trade-Off with Downward Wage Rigidities.” *American Economic Review*, 101(4):1436–1466.
- BOYARCHENKO, N., R. K. CRUMP, L. ELIAS, AND I. L. GAFFNEY (2023): “What Is “Outlook-at-Risk?”” *Liberty Street Economics 20230215*, Federal Reserve Bank of New York.
- BREIMAN, L. (2001): “Random Forests.” *Machine Learning*, 45(1):5–32.
- BUCKMANN, M. AND A. JOSEPH (2022): “An Interpretable Machine Learning Workflow with an Application to Economic Forecasting.” *BoE Working Paper No. 984*, Bank of England.
- CAVALLO, A., F. LIPPI, AND K. MIYAHARA (2024): “Large Shocks Travel Fast.” *American Economic Review: Insights*, 6(4):558–574.
- CHAKRABORTY, C. AND A. JOSEPH (2017): “Machine Learning at Central Banks.” BoE Working Paper No. 674, Bank of England.
- CHAVLEISHVILI, S. AND M. KREMER (2023): “Measuring Systemic Financial Stress and Its Risks for Growth.” ECB Working Paper Series No. 2842, European Central Bank
- CHAVLEISHVILI, S. AND S. MANGANELLI (2024): “Forecasting and Stress Testing with Quantile Vector Autoregression.” *Journal of Applied Econometrics*, 39(1):66–85.
- CHEN, B. AND P. LABONNE (2023): “FBI: Factor-Based Imputation for Missing Data.” Software/Code. R implementation. Available at: (<https://github.com/cykbennie/fbi>).

- CHEN, J. C., A. DUNN, K. HOOD, A. DRIESSEN, AND A. BATCH (2019): *Off to the Races: A Comparison of Machine Learning and Alternative Data for Predicting Economic Indicators*. In *Big Data for Twenty-First-Century Economic Statistics*, pages 373–402. University of Chicago Press.
- CHOW, G. C. AND A.-L. LIN (1971): “Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series.” *The Review of Economics and Statistics*, 53(4):372–375.
- CLARK, T. E., F. HUBER, G. KOOP, AND M. MARCELLINO (2024): “Forecasting U.S. Inflation Using Bayesian Nonparametric Models.” *The Annals of Applied Statistics, Institute of Mathematical Statistics*, 18(2):1421–1444.
- COSTAIN, J., A. NAKOV, AND B. PETIT (2022): “Flattening of the Phillips Curve with State-Dependent Prices and Wages.” *The Economic Journal*, 132(642):546–581.
- CZSO (2025): “Average Consumer Fuel Prices in the Czech Republic.” Inflation, Consumer Prices; Prices, Inflation; Statistics; Public database, Czech Statistical Office.
- DAS, P. K. AND P. K. DAS (2024): “Forecasting and Analyzing Predictors of Inflation Rate: Using Machine Learning Approach.” *Journal of Quantitative Economics*, 22(2):493–517.
- DEL NEGRO, M., M. LENZA, G. E. PRIMICERI, AND A. TAMBALOTTI (2020): “What’s Up with the Phillips Curve?” NBER Working Paper 27003, National Bureau of Economic Research.
- DICKEY, D. A. AND W. A. FULLER (1979): “Distribution of the Estimators for Autoregressive Time Series with a Unit Root.” *Journal of the American Statistical Association*, 74(366a): 427–431.
- DIEBOLD, F. X., T. A. GUNTHER, AND A. S. TAY (1998): “Evaluating Density Forecasts with Applications to Financial Risk Management.” *International Economic Review*, 39(4):863–883.
- FAUST, J. AND J. H. WRIGHT (2013): *Forecasting Inflation*. In *Handbook of Economic Forecasting*, volume 2 of, pages 2–56. Elsevier, eds. by Graham Elliott, and Allan Timmermann.
- FIGUERES, J. M. AND M. JAROCIŃSKI (2020): “Vulnerable Growth in the Euro Area: Measuring the Financial Conditions.” *Economics Letters*, 191.
- FINDLEY, D. F., B. C. MONSELL, W. R. BELL, M. C. OTTO, AND B. C. CHEN (1998): “New Capabilities and Methods of the X-12-ARIMA Seasonal Adjustment Program.” *Journal of Business & Economic Statistics*, 16(2):127–152.
- FIRAT, M. (2024): “The China Shock, Market Concentration and the U.S. Phillips Curve.” IMF Working Paper No. 4765603, International Monetary Fund.
- FORBES, K. J., J. GAGNON, AND C. G. COLLINS (2021): “Low Inflation Bends the Phillips Curve around the World: Extended Results.” PIIE Working Paper No. 21-15, Peterson Institute for International Economics.
- GALÍ, J. AND M. GERTLER (1999): “Inflation Dynamics: A Structural Econometric Analysis.” *Journal of Monetary Economics*, 44(2):195–222.
- GARCIA, M. G., M. C. 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.

- GASTWIRTH, J. L. (1966): “On Robust Procedures.” *Journal of the American Statistical Association*, 61:929–948.
- GAWTHORPE, K. (2021): “Random Forest as a Model for Czech Forecasting.” *Prague Economic Papers*, 30(3):336–357.
- GILCHRIST, S. AND E. ZAKRAJSEK (2019): “Trade Exposure and the Evolution of Inflation Dynamics.” Finance and Economics Discussion Series 2019-007, Board of Governors of the Federal Reserve System.
- GNEITING, T. AND A. E. RAFTERY (2007): “Strictly Proper Scoring Rules, Prediction, and Estimation.” *Journal of the American Statistical Association*, 102(477):359–378.
- GOULET COULOMBE, P., M. LEROUX, D. STEVANOVIC, AND S. SURPRENANT (2022): “How is Machine Learning Useful for Macroeconomic Forecasting?” *Journal of Applied Econometrics*, 37(5):920–964.
- HARDING, M., J. LINDÉ, AND M. TRABANDT (2022): “Resolving the Missing Deflation Puzzle.” *Journal of Monetary Economics*, 126:15–34.
- INOUE, A., B. ROSSI, AND Y. WANG (2024): “Has the Phillips Curve Flattened?” French Stata Users’ Group Meetings 2024 22, Stata Users Group.
- JOLLIFFE, I. T. (2002): *Principal Component Analysis*. Springer, 2nd edition.
- JOSEPH, A., G. POTJAGAILO, C. CHAKRABORTY, AND G. KAPETANIOS (2024): “Forecasting UK inflation bottom up.” *International Journal of Forecasting*, 40(4):1521–1538.
- JUDGE, G. G., R. C. HILL, W. E. GRIFFITHS, H. LÜTKEPOHL, AND T. C. LEE (1988): “Introduction to the Theory and Practice of Econometrics.” *John Wiley & Sons*.
- KILEY, M. T. (2022): “Unemployment Risk.” *Journal of Money, Credit and Banking*, 54(5): 1407–1424.
- KOENKER, R. AND G. BASSETT (1978): “Regression Quantiles.” *Econometrica: Journal of the Econometric Society*, 46(1):33–50.
- KOHLSCHEEN, E. (2023): “What Does Machine Learning Say about the Drivers of Inflation?” BIS Working Paper No. 980, Bank for International Settlements.
- KOROBILIS, D., B. LANDAU, A. MUSSO, AND A. PHELLA (2021): “The Time-Varying Evolution of Inflation Risks.” ECB Working Paper No. 2600, European Central Bank.
- KWIATKOWSKI, D., P. C. PHILLIPS, P. SCHMIDT, AND Y. SHIN (1992): “Testing the Null Hypothesis of Stationarity against the Alternative of a Unit Root: How Sure Are We that Economic Time Series Have a Unit Root?” *Journal of Econometrics*, 54(1-3):159–178.
- LEAL, F., C. MOLINA, AND E. ZILBERMAN (2020): “Inflation Forecast in Chile with Machine Learning Methods.” BCCh Working Paper No. 860, Banco Central de Chile.
- LENZA, M., I. MOUTACHAKER, AND J. PAREDES (2025): “Density Forecasts of Inflation: A Quantile Regression Forest Approach.” *European Economic Review*, 178(C).
- LI, Y.-S., P.-F. PAI, AND Y.-L. LIN (2023): “Forecasting Inflation Rates by Extreme Gradient Boosting with the Genetic Algorithm.” *Journal of Ambient Intelligence and Humanized Computing*, 14(3):2211–2220.

- LIU, Y. (2024): “Mending the Crystal Ball: Enhanced Inflation Forecasts with Machine Learning.” IMF Working Paper No. 24/206, International Monetary Fund.
- LUNDBERG, S. M., G. ERION, H. CHEN, A. DEGRAVE, J. M. PRUTKIN, B. NAIR, R. KATZ, J. HIMMELFARB, N. BANSAL, AND S.-I. LEE (2020): “From Local Explanations to Global Understanding with Explainable AI for Trees.” *Nature Machine Intelligence*, 2(1): 56–67.
- LÓPEZ-SALIDO, D. AND F. LORIA (2024): “Inflation at Risk.” *Journal of Monetary Economics*, 145(Supplement).
- MASINI, R. P., M. C. MEDEIROS, AND E. F. MENDES (2023): “Machine Learning Advances for Time Series Forecasting.” *Journal of Economic Surveys*, 37(1):76–111.
- MEDEIROS, M. C., G. F. R. VASCONCELOS, 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.
- MEINSHAUSEN, N. (2006): “Quantile Regression Forests.” *Journal of Machine Learning Research*, 7(6):983–999.
- MELIGKOTSIDOU, L., E. PANOPOULOU, I. D. VRONTOS, AND S. D. VRONTOS (2014): “A Quantile Regression Approach to Equity Premium Prediction.” *Journal of Forecasting*, 33: 558–576.
- MIRZA, N., S. K. A. RIZVI, B. NAQVI, AND M. UMAR (2024): “Inflation Prediction in Emerging Economies: Machine Learning and FX Reserves Integration for Enhanced Forecasting.” *International Review of Financial Analysis*, 94(2024).
- PAVLOVA, A. (2022): “Non-linearity between Price Inflation and Labor Costs: The Case of Central European Countries.” IES Working Paper No. 25/2022, Institute of Economic Studies, Faculty of Social Sciences, Charles University.
- PHILLIPS, A. W. (1958): “The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957.” *Economica*, 25(100):283–299.
- ROSSI, B. AND T. SEKHPOSYAN (2019): “Alternative Tests for Correct Specification of Conditional Predictive Densities.” *Journal of Econometrics*, 208(2):638–657.
- RUSSELL, S. J. AND P. NORVIG (2022): *Artificial Intelligence: A Modern Approach* Prentice Hall Series in Artificial Intelligence. Pearson, Boston, Fourth edition.
- SCHNORRENBERGER, R. AND A. SCHMIDT (2024): “Harnessing Machine Learning for Real-Time Inflation Nowcasting.” DNB Working Paper 806, De Nederlandsche Bank.
- SHAPLEY, L. (1953): *A Value for n-Person Games*. In *Contributions to the Theory of Games*, volume 2 of, pages 307–317. Princeton University Press.
- SMITH, S. C., A. TIMMERMANN, AND J. H. WRIGHT (2025): “Breaks in the Phillips Curve: Evidence From Panel Data.” *Journal of Applied Econometrics*, 40(2):131–148.
- STOCK, J. H. AND M. W. WATSON (2010): “Modeling Inflation After the Crisis.” NBER Working Paper No. 16488, National Bureau of Economic Research.
- STOCK, J. H. AND M. W. WATSON (2007): “Why Has U.S. Inflation Become Harder to Forecast?” *Journal of Money, Credit and Banking*, 39:3–33.

TUKEY, J. W. (1977): “Exploratory Data Analysis.” *Addison-Wesley*.

U.S. CENSUS BUREAU (2023): “X-13 ARIMA-SEATS Seasonal Adjustment Program.” Website. Available at: (<https://www.census.gov/srd/x13as/>).

VAN DE WINKEL, E. AND C. VEDDER (2024): “Is Machine learning Beneficial for Macroeconomic Forecasting with Limited Observations? Evidence from the Netherlands.” CPB Discussion Paper 456, CPB Netherlands Bureau for Economic Policy Analysis.

VOCHOZKA, M., T. KRULICKY, AND L. NOVOTNA (2021): “Assessing Inflation in CR Using Artificial Intelligence.” In *SHS Web of Conferences*, volume 129, 03034. EDP Sciences.

ŠTRUMBELJ, E. AND I. KONONENKO (2014): “Explaining Prediction Models and Individual Predictions with Feature Contributions.” *Knowledge and Information Systems*, 41(3):647–665.

Appendix A: Supplementary Materials

Table A1: RMSE Comparison Across Inflation Components and Forecast Horizons (h)

Inflation Component	Horizon	Median	Mean	TVW1	TVW2	TVW3
Core inflation	3 Months	0.413	0.413	0.397	0.399	0.378**
	6 Months	0.355	0.349	0.340	0.346	0.319*
	9 Months	0.479	0.484	0.458	0.465	0.431*
	12 Months	0.378	0.376	0.357	0.365	0.333*
Fuel inflation	3 Months	3.228	3.305	3.129	3.177	3.046***
	6 Months	3.501	3.692	3.453	3.464	3.355***
	9 Months	3.388	3.503	3.346	3.349	3.217***
	12 Months	3.460	3.480	3.354	3.385	3.219***
Food inflation	3 Months	0.936	0.933	0.900	0.915	0.872*
	6 Months	0.977	0.979	0.950	0.959	0.909***
	9 Months	0.984	0.988	0.959	0.968	0.923**
	12 Months	0.922	0.915	0.878	0.896	0.838**
Administrative prices	3 Months	3.151	3.539	3.321	3.326	3.218
	6 Months	3.206	3.394	3.217	3.215	3.142
	9 Months	3.260	3.314	3.236	3.256	3.190
	12 Months	3.201	3.207	3.204	3.195	3.152

Note: Boldface indicates the best model according to the lowest RMSE. †, *, **, and *** denote statistical significance of the improvement in RMSE at the 10%, 5%, 1%, and 0.1% levels, respectively, based on the Diebold–Mariano test.

Table A2: RMSE Comparison: TVW3 vs. Econometric Benchmarks

Inflation Component	Horizon	LQR Ensemble	RW	AR	ARIMA	TVW3
Core inflation	3 Months	0.454	0.378	0.436	0.423	0.378
	6 Months	0.380	0.370	0.482	0.470	0.319
	9 Months	0.495	0.506	0.495	0.488	0.431***
	12 Months	0.368	0.469	0.503	0.499	0.333
Fuel inflation	3 Months	3.292	4.714	3.290	3.220	3.046
	6 Months	3.363	4.635	3.308	3.348	3.355
	9 Months	3.439	4.893	3.357	3.395	3.217
	12 Months	3.413	4.852	3.368	3.406	3.219
Food inflation	3 Months	0.972	1.453	0.974	1.043	0.872
	6 Months	0.972	1.033	0.974	1.125	0.909*
	9 Months	0.989	1.466	0.972	1.060	0.923*
	12 Months	0.899	1.023	0.976	0.976	0.838***
Administrative prices	3 Months	3.196	4.494	3.340	3.204	3.218
	6 Months	3.226	3.499	3.202	3.221	3.142
	9 Months	3.256	4.658	3.260	3.355	3.190
	12 Months	3.118	4.472	3.318	3.398	3.152

Note: Boldface indicates the best model according to the lowest RMSE. †, *, **, and *** denote statistical significance of the improvement in RMSE at the 10%, 5%, 1%, and 0.1% levels, respectively, based on the Diebold–Mariano test.

Table A3: Continuous Ranked Probability Score (CRPS) Approximation Across Inflation Components and Forecast Horizons (h)

Inflation Component	Horizon	LQR	QRF
Core inflation	3 Months	0.236	0.224
	6 Months	0.197	0.191
	9 Months	0.264	0.256
	12 Months	0.177	0.196
Fuel inflation	3 Months	1.618	1.613
	6 Months	1.697	1.728
	9 Months	1.726	1.694
	12 Months	1.727	1.727
Food inflation	3 Months	0.518	0.508
	6 Months	0.522	0.527
	9 Months	0.532	0.526
	12 Months	0.494	0.497
Administrative prices	3 Months	0.759	0.787
	6 Months	0.784	0.819
	9 Months	0.782	0.795
	12 Months	0.796	0.807

Note: Boldface indicates the best model according to the lowest CRPS approximation.

Table A4: RMSE for YoY Transformation: QRF and TVW Models


Horizon	QRF (Median)	QRF (Mean)	TVW1	TVW2	TVW3
6 Months	2.304	2.224	2.089	2.173	1.973

Note: The table reports RMSE for YoY headline inflation forecasts based on QRF and TVW specifications. YoY inflation rates are constructed using an out-of-sample forecast of semi-annual inflation at horizon $t + 6$ combined with realised inflation over the preceding six months ($t - 6$). The resulting 12-month inflation path is recalculated into an index form, from which YoY inflation rates are computed. Boldface indicates the best model according to the lowest RMSE.


Table A5: RMSE for YoY Transformation: Econometric Benchmarks

Horizon	LQR	RW	AR(3)	ARIMA	TVW3
6 Months	2.567	2.653	2.449	2.944	1.973


Note: The table reports RMSE for YoY headline inflation forecasts from standard econometric benchmark models and the best-performing QRF-based specification (TVW3). YoY inflation rates are constructed using an out-of-sample forecast of semi-annual inflation at horizon $t + 6$ combined with realised inflation over the preceding six months ($t - 6$). The resulting 12-month inflation path is recalculated into an index form, from which YoY inflation rates are computed. Boldface indicates the best model according to the lowest RMSE.

Figure A1: Calibration Results for Food Inflation (LQR vs QRF)

Note: The solid black line represents the empirical CDF of the PIT values. The dashed red lines denote the 95% confidence bands for the Kolmogorov–Smirnov test, following Rossi and Sekhposyan (2019). The left column represents the results for the benchmark model LQR, and the right column for the QRF.


Figure A2: Calibration Results for Fuel Inflation (LQR vs QRF)

Note: The solid black line represents the empirical CDF of the PIT values. The dashed red lines denote the 95% confidence bands for the Kolmogorov–Smirnov test, following Rossi and Sekhposyan (2019). The left column represents the results for the benchmark model LQR, and the right column for the QRF.


Figure A3: Calibration Results for Administrative prices (LQR vs QRF)

Note: The solid black line represents the empirical CDF of the PIT values. The dashed red lines denote the 95% confidence bands for the Kolmogorov–Smirnov test, following Rossi and Sekhposyan (2019). The left column represents the results for the benchmark model LQR, and the right column for the QRF.

Figure A4: Calibration Results for Core Inflation (LQR vs QRF)



Note: The solid black line represents the empirical CDF of the PIT values. The dashed red lines denote the 95% confidence bands for the Kolmogorov–Smirnov test, following Rossi and Sekhposyan (2019). The left column represents the results for the benchmark model LQR, and the right column for the QRF.

Figure A5: In-Sample Analysis of Semi-Annual Forecast ($h=6$)

Note: The figure illustrates the realised YoY inflation (red line), our point forecast of the TVW3 specification (dark blue line), and the corresponding 90% prediction interval (QRF 5th and 95th quantiles, light blue shaded area).

Table A6: Predictors ($p=72$) Used in the QRF Model

No.	Category	Transf.	Description (Source)
<i>G1: Real Activity Czech Republic</i>			
1.	G1: Real Activity CZ	0	EC business and consumer surveys - Production development observed over the past 3 months
2.	G1: Real Activity CZ	0	EC business and consumer surveys - Production expectations over the next 3 months
3.	G1: Real Activity CZ	3	EC business and consumer surveys - Business situation development over the past 3 months
4.	G1: Real Activity CZ	3	EC business and consumer surveys - Evolution of demand over the past 3 months
5.	G1: Real Activity CZ	3	EC business and consumer surveys - Business activity (sales) development over the past 3 months
6.	G1: Real Activity CZ	3	EC business and consumer surveys - Building activity development over the past 3 months
7.	G1: Real Activity CZ	3	EC business and consumer surveys - Factors limiting building activity: shortage of labour
8.	G1: Real Activity CZ	3	EC business and consumer surveys - Factors limiting building activity: shortage of material/equipment
9.	G1: Real Activity CZ	3	EC business and consumer surveys - Factors limiting building activity: financial constraints

Table A6 (Continuation)

No.	Category	Transf.	Description (Source)
10.	G1: Real Activity CZ	3	EC business and consumer surveys - Financial situation over the next 13 months
11.	G1: Real Activity CZ	0	Unemployment
12.	G1: Real Activity CZ	0	Index of industrial production
13.	G1: Real Activity CZ	0	Building permits - CZSO ¹¹
14.	G1: Real Activity CZ	0	Rushin Index ¹²
15.	G1: Real Activity CZ	0	LUCI (Labour Utilisation Composite Index): Total ¹³
16.	G1: Real Activity CZ	0	LUCI (Labour Utilisation Composite Index): Wages and Labour costs
17.	G1: Real Activity CZ	0	Unit labor costs: Nominal unit labor costs
<i>G2: Foreign influence</i>			
18.	G2: Foreign influence	0	EC business and consumer surveys - Industrial confidence indicator (DE)
19.	G2: Foreign influence	3	EC business and consumer surveys - Services confidence indicator (DE)
20.	G2: Foreign influence	3	EC business and consumer surveys - Retail confidence indicator (DE)
21.	G2: Foreign influence	3	EC business and consumer surveys - Construction confidence indicator (DE)
22.	G2: Foreign influence	2	E C business and consumer surveys - Retail confidence indicator (PL)
23.	G2: Foreign influence	2	Harmonised index of consumer prices (DE)
24.	G2: Foreign influence	0	Unemployment (DE)
25.	G2: Foreign influence	2	External economic relations: Balance of trade (FOB/FOB)
26.	G2: Foreign influence	2	Import prices
27.	G2: Foreign influence	0	Inflation expectations: Domestic inflation expectations in Eurozone - balance of answers
28.	G2: Foreign influence	0	Inflation expectations: Perceived inflation in Eurozone - balance of answers
<i>G3: Confidence/Sentiment Czech Republic</i>			
29.	G3: Confidence/Sentiment CZ	3	EC business and consumer surveys - Industrial confidence indicator
30.	G3: Confidence/Sentiment CZ	3	EC business and consumer surveys - Expectation of the demand over the next 3 months

¹¹ The time series was transformed from the cumulative form into the number of building permits each month.

¹² Rushin Index is index of economic activity for Czech Republic constructed by CNB. It captures the common dynamics of four alternative high-frequency indicators and six standard macroeconomic monthly indicators of activity.

¹³ LUCI is an index compiled by the CNB that aggregates information from twenty labor market indicators using the principal component analysis method.

Table A6 (Continuation)

No.	Category	Transf.	Description (Source)
31.	G3: Confidence/Sentiment CZ	3	EC business and consumer surveys - Expectation of the employment over the next 3 months
32.	G3: Confidence/Sentiment CZ	0	EC business and consumer surveys - Services confidence indicator
33.	G3: Confidence/Sentiment CZ	3	EC business and consumer surveys - Expectations of the number of orders placed with suppliers over the next 3 months
34.	G3: Confidence/Sentiment CZ	3	EC business and consumer surveys - Business activity expectations over the next 3 months
35.	G3: Confidence/Sentiment CZ	3	EC business and consumer surveys - Employment expectations over the next 3 months
36.	G3: Confidence/Sentiment CZ	3	EC business and consumer surveys - Retail confidence indicator
37.	G3: Confidence/Sentiment CZ	3	EC business and consumer surveys - Employment expectations over the next 3 months - construction
38.	G3: Confidence/Sentiment CZ	3	EC business and consumer surveys - Construction confidence indicator
39.	G3: Confidence/Sentiment CZ	3	EC business and consumer surveys - Financial situation over the last 13 months
40.	G3: Confidence/Sentiment CZ	3	EC business and consumer surveys - Savings over the next 13 months
41.	G3: Confidence/Sentiment CZ	3	EC business and consumer surveys - Unemployment expectations over the next 13 months
42.	G3: Confidence/Sentiment CZ	3	EC business and consumer surveys - Major purchases over the next 13 months
<i>G4: PPI (Producer Price Index)</i>			
43.	G4: PPI	0	YoY change in Producer prices of agricultural products
44.	G4: PPI	0	YoY change in Producer prices of industrial products
45.	G4: PPI	2	Producer prices of industrial products - Total
46.	G4: PPI	2	Producer prices of industrial products - Mineral resources
47.	G4: PPI	2	Producer prices of industrial products - Manufactured goods
48.	G4: PPI	2	Producer prices of industrial products - Electricity, gas, and steam
49.	G4: PPI	2	Producer prices of industrial products - Water
50.	G4: PPI	2	Producer prices of agricultural products - Animal based products
51.	G4: PPI	2	Producer prices of agricultural products - Plant based products
52.	G4: PPI	2	Producer prices of agricultural products - Agricultural production and fish

Table A6 (Continuation)

No.	Category	Transf.	Description (Source)
<i>G5: Financial</i>			
53.	G5: Financial	0	Yield of 10-year gov. debt, monthly average
54.	G5: Financial	0	PRIBOR: Domestic interbank deposit market rate - 3 months, value at the end of the month
55.	G5: Financial	0	2 Week repo rate, value at the end of the month
56.	G5: Financial	2	Real effective exchange rate: Deflated by PPI
57.	G5: Financial	2	Real effective exchange rate: Deflated by CPI
58.	G5: Financial	2	Client loans: Residents - Non-financial corporations, balance
59.	G5: Financial	2	Client loans: Residents - Households, balance
<i>G6: Commodities and Energy</i>			
60.	G6: Commodities/Energy	0	Commodity prices: Brent Crude Oil, USD/barrel
61.	G6: Commodities/Energy	2	Commodity prices: Average natural gas price in Europe
62.	G6: Commodities/Energy	2	Commodity prices: Industrial metals price index
63.	G6: Commodities/Energy	2	Commodity prices: Food commodity price index
64.	G6: Commodities/Energy	2	Principal Components of Fuel prices - CZSO ¹⁴
65.	G6: Commodities/Energy	0	Refinitiv - TRPC Natural Gas
66.	G6: Commodities/Energy	0	Refinitiv - ICE Europe Brent Crude Electronic Energy Future
<i>G7: Inflation Expectations</i>			
67.	G7: Inflation Expectations	3	Selling price expectations over the next 3 months
68.	G7: Inflation Expectations	3	Price expectations over the next 3 months - construction
69.	G7: Inflation Expectations	0	Inflation expectations: Inflation expectations in the horizon 3 years by financial markets - in percentage points
70.	G7: Inflation Expectations	0	Inflation expectations: Inflation expectations in the horizon 1 year by financial markets in percentage points
71.	G7: Inflation Expectations	0	Inflation expectations: Perceived inflation - balance of answers
72.	G7: Inflation Expectations	0	Inflation expectations: Expected domestic inflation - balance of answers

Note: Transformations: 0 denotes no transformation (level); 2 denotes log-difference ($\ln(X_t) - \ln(X_{t-1})$); and 3 denotes simple difference ($X_t - X_{t-1}$). Sources: Eurostat, ARAD CNB, CZSO, Refinitiv. We construct our dataset based on Lenza et al. (2025) and further enrich the variables by those relevant for the SOE. Concerning the variable transformation, we proceed by testing our time series for stationarity using ADF and KPSS tests.

¹⁴ We use four highly correlated time series - Petrol 95 O Natural (CZK/l), Petrol 98 O Super plus (CZK/l), Diesel (CZK/l) and Liquefied Petroleum Gas (LPG) (CZK/l). The Dataset includes only PCA's first component, which explains more than 93% of the variance confirming similarity of these series.

CNB Working Paper Series (since 2025)

WP 9/2026	Filip Blaha Jan Botka Josef Švéda Aleš Michl	<i>AI-Based Forecasting of Czech Inflation: Quantile Regression Forests with Dynamic Weights</i>
WP 8/2026	Martijn Boermans Laurens Swinkels	<i>Hedging Against Inflation: International Evidence on Investor Clientele Effects</i>
WP 7/2026	Zuzana Gric Jan Janků Simona Malovaná	<i>Euroisation and the Bank Lending Channel of Monetary Policy: Evidence from Czechia</i>
WP 6/2026	František Mašek	<i>Households' Inflation Expectations and Consumption in Macroeconomic Models: A Negative Real Income Channel</i>
WP 5/2026	Volha Audzei Michal Franta	<i>International Spillovers from Euro Area Monetary Policy to Advanced Small Open Economies: Investment Behavior of Czech Firms</i>
WP 4/2026	Josef Simpartl	<i>Social Media as a Monetary Policy Tool? Evidence from a Survey Experiment</i>
WP 3/2026	Tomáš Šestořád Jan Vlček Karel Musil	<i>Monetary Policy and the Output Gap in DSGE Models for Small Open Economies: Insights from the Czech Republic</i>
WP 2/2026	Martin Hodula Simona Malovaná Lukáš Pfeifer	<i>When Monetary and Macroprudential Policies Tighten Together: Evidence from the Czech Mortgage Market</i>
WP 1/2026	Jiří Gregor Jan Janků	<i>Chasing Lower Rates: How Households Balance Refinancing Incentives and Debt Constraints</i>
WP 13/2025	Oxana Babecká Kucharčuková Jan Brůha Petr Štěřba	<i>Web Reviews as a New Leading Indicator for Nowcasting Travel Expenditure in Balance of Payments Statistics</i>
WP 12/2025	Jan Janků Simona Malovaná Josef Bajzík Klára Moravcová Ngoc Anh Ngo	<i>Credit Shocks Fade, Output Shocks Persist: A Meta-Analysis of 2,600 VAR Estimates Across 63 Countries</i>
WP 11/2025	Martin Hodula Simona Malovaná	<i>When Foreign Rates Matter More: Domestic Investor Responses in a Small Open Economy</i>
WP 10/2025	Volha Audzei Jan Brůha Ivan Sutóris	<i>Does Firms' Financing in Foreign Currency Matter for Monetary Policy?</i>
WP 9/2025	Volha Audzei Sergey Slobodyan	<i>Dynamic Sparse Adaptive Learning</i>
WP 8/2025	Michal Franta Jan Vlček	<i>Inflation at Risk: The Czech Case</i>
WP 7/2025	František Brázdík Karel Musil Tomáš Pokorný Tomáš Šestořád Jaromír Tonner Jan Žáček	<i>Upgrading the Czech National Bank's Core Forecasting Model g3+</i>
WP 6/2025	Raphael Auer David Köpfer Josef Švéda	<i>The Rise of Generative AI: Modelling Exposure, Substitution, and Inequality Effects on the US Labour Market</i>
WP 5/2025	Josef Švéda	<i>Navigating Banking Resilience: Bail-Ins & Bailouts in the Czech Banking Sector</i>

CNB Working Paper Series (since 2025)

WP 4/2025	Ivan Trubelík Tomáš Karhánek Simona Malovaná Aleš Michl	<i>Instant Payments in Czechia: Adoption and Future Trends</i>
WP 3/2025	Daniel Štodt	<i>Non-Linearity of Government Spending Multiplier: The Case of a Small Open Economy</i>
WP 2/2025	Soňa Benecká	<i>Forecasting Disaggregated Producer Prices: A Fusion of Machine Learning and Econometric Techniques</i>
WP 1/2025	Martin Hodula Lukáš Pfeifer	<i>Payment Holidays, Credit Risk, and Borrower-Based Limits: Insights from the Czech Mortgage Market</i>

CZECH NATIONAL BANK

Na Příkopě 28

115 03 Prague 1

Czech Republic

RESEARCH DIVISION

www.cnb.cz

research@cnb.cz

ISSN 1803-7070

ČNB ČESKÁ
NÁRODNÍ
BANKA