Bank of England

Blockwise Boosted Inflation: Non-linear determinants of inflation using machine learning

Staff Working Paper No. 1,143

September 2025

Marcus Buckmann, Galina Potjagailo and Philip Schnattinger

Bank of England

Staff Working Paper No. 1,143

Blockwise Boosted Inflation: Non-linear determinants of inflation using machine learning

Marcus Buckmann,[1] Galina Potjagailo[2] and Philip Schnattinger[3]

Abstract

We propose the Blockwise Boosted Inflation Model (BBIM), a boosted tree framework that decomposes inflation dynamics into predictive components aligned with an open-economy hybrid Phillips curve. Demand and supply contributions are identified by imposing monotonicity constraints, ensuring theory-consistent links between inflation and key indicators. Applied to monthly UK CPI inflation, the model shows that the recent surge has been driven mainly by global supply shocks transmitted through supply chains. We also uncover an L-shaped Phillips curve relationship between inflation and labour market tightness, with tight labour markets amplifying recent inflationary pressures. By contrast, earlier episodes saw non-linearities more strongly tied to broader slack, particularly during recessions. The model further accounts for trend shifts informed by inflation expectations. Short-term household expectations have recently displayed persistent non-linear effects, temporarily raising trend inflation and prolonging inflationary pressures, while longer-term expectations remain anchored. Out-of-sample, the BBIM delivers competitive forecasting performance relative to linear benchmarks and unstructured machine learning methods. Our approach provides a flexible yet interpretable framework that combines economic structure with machine learning for policy-relevant analysis of inflation dynamics.

Key words: Inflation, Phillips curve, boosted decision trees, machine learning.

JEL classification: E31, E37, C14, C53.

(1) Bank of England. Email: marcus.buckmann@bankofengland.co.uk
(2) Bank of England. Email: galina.potjagailo@bankofengland.co.uk
(3) Bank of England. Email: philip.schnattinger@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. We are grateful for highly valuable comments by Paul Beaudry, Ambrogio Cesa-Bianchi, Ana B Galvao, Daniel Lewis, Catherine L Mann, Francesca Monti, Adrian Paul, Huw Pill, Ricardo Reis, Anthony Savagar, Tatevik Sekhposyan, Martin Seneca, Bradley Speigner, Silvana Tenreyro, and the participants of the RES 2024 Annual Conference, the ECONDAT 2024 Fall Meeting, the University of Strathclyde Workshop on Macroeconomic Analysis and Forecasting, the BSE Summer Forum: Machine Learning in Economics, the 2025 EEA congress, and seminars at the European Central Bank and Bank of England.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email: enquiries@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1 Introduction

Disentangling the sources of inflationary pressure is a central task for monetary policymakers aiming to stabilise inflation. Yet this task becomes particularly difficult when inflation dynamics are shaped by non-linearities—conditions under which standard linear models often fail to identify underlying drivers or produce reliable forecasts. This challenge became evident in the aftermath of the Covid-19 pandemic and Russia's invasion of Ukraine, when many central banks initially underestimated the magnitude and persistence of the inflation surge.

Could the flexibility of machine learning help policymakers detect and interpret complex, non-linear inflation dynamics? While machine learning has gained traction in economic forecasting for its ability to capture high-dimensional, non-linear relationships (de Araujo et al., 2024), its use has largely centred on maximising predictive accuracy. Meanwhile, its use for uncovering fundamental economic relationships has been limited by the lacking economic interpretability of current machine learning models. These models are often considered to be "black boxes," offering little insight into the underlying economic mechanisms (Nakamura, 2005; Medeiros et al., 2021; Joseph et al., 2021; Lenza et al., 2023). The value of such methods for monetary policy hinges on understanding of the drivers of inflation signals, and whether those are likely to be temporary or persistent.

In this paper, we bridge the gap between predictive accuracy and economic intuition by introducing the Blockwise Boosted Inflation Model (BBIM)—a novel machine learning framework that combines the predictive power of gradient boosting with an economically interpretable structure inspired by the open-economy hybrid Phillips curve. At its core, the model's decision trees capture non-linear associations between inflation and a large set of indicators. Boosting models have shown accurate performance on tabular prediction problems (Grinsztajn et al., 2022; McElfresh et al., 2024) and have delivered strong empirical results in macroeconomic forecasting (Ng, 2014; Döpke et al., 2017; Yoon, 2021).

To generate economically interpretable insights, we impose a blockwise structure on the model, with each block representing a distinct group of indicators linked to specific inflation determinants. The blocks are designed to mirror a hybrid open-economy Phillips curve. A time-varying trend component captures non-linear associations between inflation, a time indicator, inflation expectations, and indicators of domestically generated inflation. Global and domestic supply blocks reflect cyclical effects from cost pressures, commodity price fluctuations, and global supply-chain disruptions. Global and domestic demand blocks capture broader economic activity, with domestic demand further divided to isolate the role of labour market slack.

Within each block, decision trees capture potentially non-linear relationships with inflation. Across blocks, the model is additive, estimated sequentially through boosting iterations that alternate between blocks. This design enhances interpretability by decomposing the overall forecast into economically meaningful components: rather than producing a single aggregate prediction, the model generates blockwise predictive contributions, with each block's marginal effect given by the sum of its trees' predictions. It also allows for the imposition of additional structural restrictions within each block.

Specifically, grouping alone might not suffice for identification, as the association between inflation and activity indicators often reflects a mix of demand and supply forces. To address this, we apply *monotonicity constraints* that ensure relationships between predictors and inflation with the supply and demand blocks are directionally consistent with a theoretically consistent separation of demand- and supply-driven inflation dynamics. Specifically, activity indicators within the demand blocks are constrained to contribute positively to inflation, while in the supply blocks they are constrained to contribute negatively (the reverse sign is imposed if an indicator measures slack, such as unemployment). Indicators of cost shocks and global supply chain pressures are constrained to have a positive effect on inflation within supply. To further support identification, we incorporate externally identified global shocks—such as oil supply or global demand shocks (Baumeister and Hamilton, 2019; Känzig, 2021)—placing them in the relevant blocks and applying similar constraints.

We train the model in up to 200 boosting rounds, where each round adds one tree per block, fitted to the residuals from the previous prediction and subject to any monotonicity constraints within that block. The order of blocks is randomly permuted each round, and a slow learning rate ensures that all blocks condition on one another, giving them similar opportunities to contribute to the overall prediction. As a robustness check, we also explore structured sequences that assume certain blocks contain more exogenous information—for example, training global blocks or the trend block first.

We apply the BBIM to monthly UK CPI inflation data spanning 1989 to 2024—a setting particularly suited to our approach given the UK's openness and exposure to external shocks. The model incorporates more than 50 monthly indicators, each included with multiple lags. We provide a near-term predictive decomposition of inflation into blockwise contributions over time, based on estimation using repeated cross-validation. The model is initialised at the 2% inflation target so that predictive contributions capture deviations from this baseline, and a small learning rate ensures incremental updates that reduce the risk of overfitting. Additionally, we conduct a pseudo out-of-sample forecasting exercise at horizons of up to 12 months to demonstrate that the model remains competitive in predictive performance relative to unstructured linear and machine learning models.

The BBIM provides intuitive, policy-relevant insights regarding the determinants of UK CPI inflation over time. The time-varying trend is generally slow-moving but shifts upward during high-inflation periods. Crucially, these dynamics are not imposed a priori; they emerge from the model learning non-linear associations between inflation, a time indicator, and expectations measures. The model also captures cyclical demand contributions from broader activity, as during the 2008 Global Financial Crisis (GFC) and the post-pandemic recovery. The importance of monotonicity constraints is evident by comparison: without them, the model produces a slow-moving rather than cyclical demand component, missing the disinflationary impact of the GFC. In the recent episode, demand-like inflationary pressures have instead been concentrated in labour market tightness rather than broad activity.

Non-linearities have been central to recent UK inflation dynamics. Using Shapley values (Lundberg and Lee, 2017), we trace the shape and strength of non-linear associations between indicators and their predictive contributions within each model block. Three key findings emerge. First, supply-side non-linearities have been more pronounced in recent years than in earlier episodes, with increases in global supply chain pressures and commodity prices generating stronger inflationary signals, also reflected in disproportionately large contributions from domestic food prices. Second, UK inflation exhibits a non-linear relationship with labour market tightness: predictive contributions rise steeply for high vacancy-to-unemployment gaps and low unemployment gaps, as observed in the recent episode, echoing U.S. evidence of an L-shaped Phillips curve with labour market slack (Benigno and Eggertsson, 2023, 2024; Bernanke and Blanchard, 2025). Third, short-term household inflation expectations show asymmetric effects: weak at low levels but increasingly inflationary once they exceed 4%. These non-linear effects temporarily elevated trend inflation, sustaining price pressures even after supply shocks subsided. This likely reflects adaptive expectations at high inflation, with potential threshold effects (Pfäuti, 2023), as well as heightened sensitivity to salient price changes (Anesti et al., 2024). In contrast, long-term expectations have remained well anchored, unlike the pronounced non-linearities observed during the pre-inflation-targeting era of the early 1990s.

The block structure imposed to inform the model, and the order in which each block is trained, reflect a priori modelling choices that, while enhancing economic intuition, may influence results. To assess robustness, we run a range of specifications varying the training sequence, for example fitting the trend block first or imposing block exogeneity by training external blocks before domestic ones. Both the cyclical characteristics of each block's contribution and the non-linear associations with indicators within blocks remain robust across specifications. While the relative sizes of block contributions vary—blocks trained earlier tend to capture a larger share—the overall picture remains intact. In the recent episode, supply continues to dominate, rising demand-like effects linked to labour market tightness are evident but less pronounced compared to global supply effects, and the expectations-informed trend rises later in the inflation episode—particularly when the trend is trained first. We also test alternative trend specifications. Results are broadly similar but reveal the importance of an expectations-informed trend. Omitting the trend increases the weight of other blocks but underestimates inflation during high-inflation periods. Replacing it with a purely time-based indicator weakens the model's ability to capture a generally flat, slow-moving trend that occasionally rises.

Finally, the out-of-sample forecast exercise shows that the BBIM performs competitively against a range of benchmark models, particularly during the recent high-inflation episode. It significantly outperforms a simple autoregressive model, reducing root mean squared errors by 10–25%, and performs broadly on par with leading alternatives such as Lasso regression and an unobserved components model. During the recent episode, forecast errors are larger overall, and while machine learning models perform relatively better, the BBIM delivers modest improvements over unstructured alternatives. Notably, removing the block structure and monotonicity constraints—thereby increasing model flexibility—does not improve forecast accuracy, suggesting that economically motivated structure can enhance interpretability without sacrificing predictive performance.

Our analysis contributes to expanding the practical relevance of machine learning models in economic applications. While the BBIM remains a predictive rather than fully structural approach, its combination of blockwise structure and monotonicity constraints offers a principled, theory-informed framework for disentangling supply- and demand-driven contributions to inflation. In doing so, it provides policymakers with timely and interpretable insights into the non-linear drivers of inflation dynamics.

We contribute to a growing literature that integrates machine learning methods—such as neural networks and tree-based approaches—with economically motivated constraints to enhance interpretability, rather than relying solely on post-hoc explanations. Buckmann and Potjagailo (2025) survey this emerging field and highlight how economic reasoning can be embedded directly within machine learning models. Closely related to our approach is Goulet Coulombe (2024), who develops a neural network with a block structure designed to reflect a Phillips curve framework for the United States. His model separately identifies a slow-moving slope and a cyclical output gap contribution by imposing assumptions akin to those used in time-varying unobserved components models, but does not distinguish demand and supply drivers based on economically motivated directional restrictions. Block structures have also been used in more traditional macroeconomic models, such as dynamic factor models (Kose et al., 2003; Potjagailo and Wolters, 2023).

We adopt boosted trees, which are well-suited to incorporating economic priors. They are computationally efficient, robust to hyperparameter choices and can accommodate blockwise structures through sequential learning. Prior work used boosting models (Lou et al., 2012) to learn generalised additive models (Hastie and Tibshirani, 1990), which uncover non-linear relationships by decomposing predictions into sums of functions of individual predictors (blocks of size 1), without allowing for interactions between them. Blockwise boosting, with several variables per block, is not prominently discussed in the machine learning literature (but see Hothorn et al., 2010; Mayer et al., 2021) and, to our knowledge, has not been used to introduce economically motivated constraints.

The use of monotonicity constraints—where the relationship between predictors and the target variable is restricted to follow a prespecified direction—is a well-established strategy for enhancing model interpretability. While such constraints have been extensively studied in the general machine learning literature (Nanfack et al., 2022; Cano et al., 2019) and have also been used in economic and financial prediction models to increase the predictive performance by imposing expert knowledge (Li and Tsiakas, 2017; Fisher et al., 2020; Wen et al., 2022; Richman and Wüthrich, 2024), their application in macroeconomic contexts remains limited (but see Chalaux and Turner, 2023). To the best of our knowledge, this paper is the first to integrate component-wise non-linear boosting with monotonicity constraints to recover economically interpretable decompositions of inflation dynamics.

Our analysis also connects naturally to the macroeconometric literature that seeks to identify the drivers of inflation—particularly in the U.S. context—using approaches such as disaggregated price data (Shapiro et al., 2022; Firat and Hao, 2023), dynamic factor models (Eickmeier and Hofmann, 2022), and structural vector autoregressions (SVARs) (Kabaca and Tuzcuoglu, 2023; Giannone and Primiceri, 2024). These methods typically rely on identifying structural demand and supply shocks via sign or zero restrictions (Arias et al., 2018). In contrast, our approach imposes monotonicity constraints directly during model estimation, ensuring that each predictor maintains a theory-consistent directional effect across its domain. This enables a decomposition of predicted inflation into contributions from distinct economic drivers, while also allowing for the tracing of potentially non-linear and time-varying signals associated with each block of determinants. Additionally, the inclusion of a block capturing a stochastic trend component informed by inflation expectations is consistent with unobserved components models of inflation (Chan et al., 2018), and reflects recent findings emphasising the central role of expectations—alongside demand and supply shocks—in explaining the recent inflation surge (Coibion and Gorodnichenko, 2025).

The remainder of the paper is structured as follows. Section 2 outlines the methodological framework: a blockwise boosted tree model structured in line with an open-economy

Phillips curve, with monotonicity constraints to distinguish demand and supply determinants. It also describes the empirical setup for the UK inflation application. Section 3 presents the results, including the inflation decomposition, the learned non-linear associations, and robustness checks across alternative specification. Section 4 concludes.

2 Blockwise Boosted Trees for Modelling Inflation

We first present the blockwise boosted tree method, extending standard boosting, and propose its application to inflation forecasting within an economic structure resembling an open-economy Phillips curve. Section 2.3 introduces monotonicity constraints to identify demand and supply contributions. Section 2.4 outlines model training and parameter choices, and Section 2.5 details the UK inflation application.

2.1 Blockwise boosting

Our method builds on the standard boosting paradigm in machine learning (Friedman, 2001), an ensemble approach that combines predictions from a large number of *base learners* trained sequentially. We use decision trees as base learners, which partition the data into subsets of observations with similar outcome values, assigning the mean outcome in each subset as its prediction. Decision trees are well-suited for boosting due to their ability to capture non-linear relationships and due to their computational efficiency.

Let $f(X_{t-p})$ denote a decision tree, where X_{t-p} represents predictor variables lagged by p periods relative to time t. The variable of interest is inflation at forecast horizon h, denoted by π_{t+h}. A standard boosting model predicts inflation by summing the predictions of M trees:

$$\pi_{t+h} = \sum_{m=1}^{M} f_m(X_{t-p}) + \epsilon_t, \tag{1}$$

where each decision tree m is fitted on the residuals of the previous trees, such that each additional tree improves the prediction of the ensemble slightly.

The blockwise boosting model distinguishes between separate blocks of predictors. Within each block k, the model contains M block-specific decision trees f^k that are trained on the sub-group of predictors X^k_{t-p}. The model allows for interactions among variables *within* blocks but imposes conditional linearity *between* blocks.[1] The model prediction is obtained by summing the predictions of the trees from all K blocks:

$$\pi_{t+h} = \sum_{k=1}^{K} \sum_{m=1}^{M} f^k_m(X^k_{t-p}) + \epsilon_t. \tag{2}$$

The blockwise structure enhances interpretability by expressing the contribution of each block as the sum of its base learners' outputs. These contributions have a meaningful economic interpretation when blocks are constructed based on economic intuition or statistical properties—for example, when predictors within a block jointly capture an

[1]The model does not assume or impose statistically independent block contributions.

unobserved economic driver. This design reflects a macroeconomic principle: economic fluctuations are determined by a small number of structural shocks or latent forces. The identification of interpretable common components parallels dynamic factor models where factors can be extracted from predefined variable groups for interpretation (Kose et al., 2003; Potjagailo and Wolters, 2023), and has been applied in blockwise neural networks (Goulet Coulombe, 2024). For a discussion of the blockwise structure and its capacity to embed economic structure into machine learning, see Buckmann and Potjagailo (2025).

2.2 Blockwise boosting to represent a Phillips curve

Our aim is to represent a structure that reflects a stylised open-economy hybrid new-Keynesian Phillips curve, in line with Gali and Monacelli (2005) and Forbes (2019):

$$\pi_t = \rho\pi_{t-1} + \beta E_t(\pi_{t+1}) + \kappa_1 \hat{y}_t + \kappa_2 \hat{y}_t^F + \gamma c_t^F + u_t, \tag{3}$$

where inflation π_t is modelled as a combination of a term reflecting the role of forward-looking expectations and backward-looking inertia, domestic demand pressures, such as a domestic output gap $\hat{y}_t = y_t - y_t^*$, foreign demand pressures $\hat{y}_t^F$, and foreign cost-push shocks c_t^F. The latter can reflect commodity prices, global supply chain pressures, exchange rates, or global financial conditions. Since we are interested in inflation in the United Kingdom, a small open economy, accounting for foreign demand and supply effects is important to cleanly separate out a Phillips curve association with domestic effects. Finally, the residual $u_t = \phi c_t$ represents domestic cost-push shocks.

How can we incorporate economically intuitive structure in accordance with a Phillips curve as in (3) within a machine learning set-up? In a first step, we can rely on a blockwise boosting structure to sequentially fit decision trees that are calibrated on groups of indicators, where each group or model block intends to summarise one component of the Phillips curve. Specifically, we can group lagged inflation indicators and expectations series (X_t^τ) to reflect a stochastic trend component τ_t that captures backward- and forward-looking dynamics. Next, we can group global and domestic activity indicators $(X_t^y, X_t^{y^F})$ to approximate the association with demand conditions y_t or y_t^F, and finally, group measures of global and domestic costs or supply pressure indicators $(X_t^c,\ X_t^{c^F})$. This results in the following blockwise boosted tree specification:

$$\pi_{t+h} = \sum_{m=1}^{M} \left(f_m^\tau(X_{t-p}^\tau) + f_m^y(X_{t-p}^y) + f_m^{y^F}(X_{t-p}^{yF}) + f_m^c(X_{t-p}^c) + f_m^{c^F}(X_{t-p}^{c^F}) \right) + \epsilon_t. \tag{4}$$

This model imposes conditional linearity between blocks, but allows flexible functional forms between inflation and the indicators within each block. A comparable grouping was introduced by Goulet Coulombe (2024) to represent a Phillips curve structure via a blockwise neural network model for inflation in the United States.

The assumption of linear separability across blocks may be overly restrictive, especially in high-dimensional settings where many indicators lack clear structural interpretation and may not uniquely represent a single component. For example, the relationship between inflation and domestic activity or slack variables X_t^y may reflect both demand-driven pressures—key to identifying the Phillips curve slope via f_m^y —and supply-side forces ideally captured by f_m^c. As a result, trees f_m^y may conflate signals from both sources, weakening the structural interpretability of the demand contribution. Such conflation can also reduce predictive accuracy by diluting otherwise distinct informational signals. To address this, we incorporate additional identifying assumptions via monotonicity constraints.

2.3 Monotonicity constraints to separate supply and demand

For interested readers, we first illustrate the general mechanics of monotonicity constraints before presenting our main BBIM specification, which incorporates these constraints, in Section 2.3.2.

2.3.1 Illustration of monotonicity constraints

A monotonicity constraint on predictive model $f(\cdot)$ for variable y with predictors $x_1, ..., x_n$ requires that:

$$f(x_1, x_2, \cdots, x_i, \cdots) \geq f(x_1, x_2, \cdots, x_i', \cdots)$$

when $x_i > x_i'$ is a positive (increasing) constraint; or

$$f(x_1, x_2, \cdots, x_i, \cdots) \leq f(x_1, x_2, \cdots, x_i', \cdots)$$

when $x_i < x_i'$ is a negative (decreasing) constraint.

For instance, with a positive constraint imposed on variable i, tree f will only allow splits that predict $f(x) \geq f(x')$, if $x_i > x_i'$, and $x_j = x_j', \forall j \neq i$. Figure 1 illustrates this for a simple tree with a single indicator.

The blue-framed box highlights a disallowed split that violates the monotonicity constraint and is therefore discarded. Such splits would allow the tree to fit non-monotonic patterns or associations inconsistent with the modeller's priors (bottom left panel). By excluding them, the resulting tree captures only the monotonically increasing portion of the relationship (bottom right panel).

This setup allows an indicator to appear in two identified blocks, while restricting its association with inflation to opposite directions across the blocks. Monotonicity constraints, especially when imposed on various indicators within the identified blocks, enhance identification by guiding the model toward theory-consistent relationships. By training many trees across a broad set of indicators, the model leverages both temporal and cross-variable variation to disentangle supply- and demand-driven contributions. Although these constraints do not identify structural shocks directly, they support a theory-informed decomposition of predicted inflation, for instance into supply and demand components. For further discussion on aligning models with economic theory using monotonicity constraints, see Buckmann and Potjagailo (2025).[2]

[2]Unlike sign (or zero) restrictions in VARs (Arias et al., 2018), which identify structural shocks by imposing sign constraints on impulse responses over specific horizons using rotations of contemporaneous correlations, monotonicity constraints are embedded directly in the estimation of the predictive model.

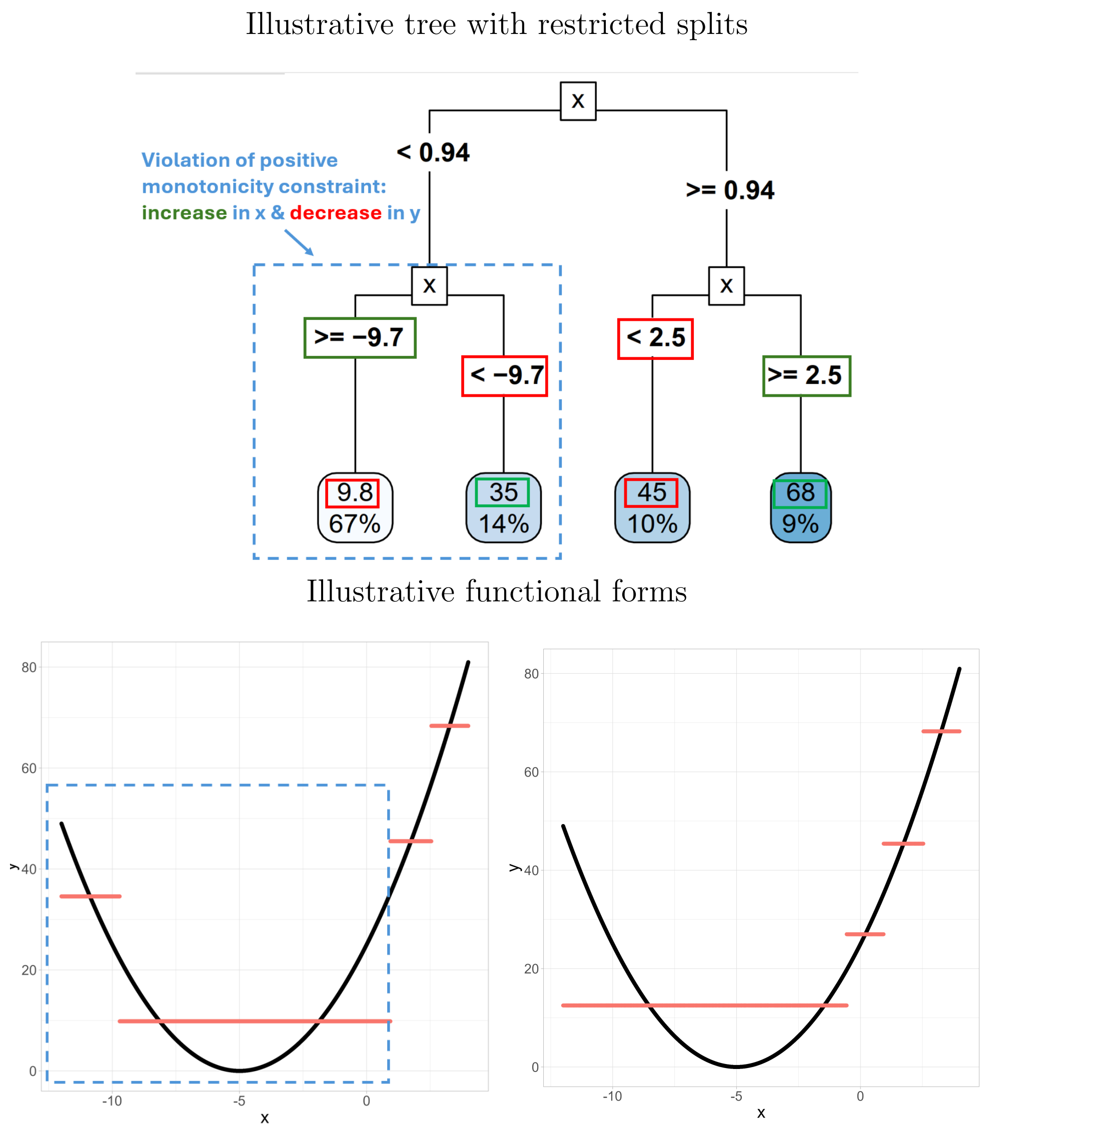

Figure 1: Constraining a decision tree towards fitting a monotonically positive association.
Notes: Top panel: Decision tree with split violating positive monotonicity constraint in blue-framed box. Bottom panel: Black line shows ground-truth function, red lines shows prediction by a decision trees without (left panel) and with (right panel) positive monotonicity constraint imposed.

2.3.2 Blockwise Boosted Inflation Model

Enhancing the model in equation (5) with monotonicity constraints yields our proposed Blockwise Boosted Inflation Model (BBIM), which avoids attributing all activity-related variation solely to demand and instead distinguishes between (global and domestic) demand (D^g, D^d) and supply (S^g, S^d) blocks.

$$\pi_{t+h} = \sum_{m=1}^{M} \left(f_m^{\tau}(X_{t-p}^{\tau}) + f_m^{D^g}(X_{t-p}^{D^g}) + f_m^{D^d}(X_{t-p}^{D^d}) + f_m^{S^g}(X_{t-p}^{S^g}) + f_m^{S^d}(X_{t-p}^{D^d}) \right) + \epsilon_t \tag{5}$$

The demand block is restricted to have positive relationships with activity indicators (and negative ones with labour market slack), while the supply block is restricted to negative relationships with activity and positive ones with supply pressures. To further distinguish global from domestic drivers, we include externally identified global oil supply and global demand shocks, constrained to load positively in the global supply and demand blocks, respectively. These shocks support interpretation, though the monotonicity constraints do most of the work in separating the demand component (see Section 3.3). Indicators in the trend block remain unrestricted, though imposing positive constraints produced similar results. Section 2.5 details the indicators and restrictions used in our application.

2.4 Model training and parameterisation

During training, the blockwise boosting model, as described in equations (2) and (5), is initialised with a baseline prediction F_0 and then fitted sequentially. The total number of trees in the ensemble is $K \times M$, where K denotes the number of blocks and M the number of boosting rounds. In each boosting round, the ensemble F is updated by adding one tree $f_m^k(X^k)$ from each block k, respectively. F_j denotes the boosting ensemble containing j trees, where $0 < j \leq M \times K$. The order of blocks within each round is randomly permuted via permutation P. Each new tree is trained—subject to any monotonicity constraints imposed within block k—on the residuals from the previous prediction, $r = \pi_i - F_{j-1}$. A small learning rate ν ensures that each tree contributes only incrementally to the overall prediction, thereby reducing the risk of overfitting: $F_j = F_{j-1} + \nu f_m^k(X^k)$. The training procedure is summarised in Algorithm 1.

Algorithm 1 Training of blockwise boosting model

$F_0 \leftarrow$ baseline prediction
$j \leftarrow 1$
for $m = 1$ **to** M **do**
  $P \leftarrow \textsc{RandPerm}(\{1, \ldots, K\})$ ▷ Shuffle block order
  for all $k \in P$ **do**
    $r \leftarrow \pi - F_{j-1}$ ▷ Compute pseudo-residuals
    Fit f_m^k on (X^k, r)
    $F_j \leftarrow F_{j-1} + \nu\, f_m^k(X^k)$ ▷ Update ensemble
    $j \leftarrow j + 1$
  end for
end for

Specifically, we train the model for $M = 200$ rounds during cross-validation, but stop early if the cross-validation error increases with additional boosting rounds. For forecasting, we use $M = 100$ rounds without early stopping because we lack a sufficiently large validation set to identify the stopping point without leaking information from the test sample.[3] The model is initialised with the inflation target, $F_0 = 2\%$, so that predictive contributions reflect deviations from this baseline. The learning rate is set to $\nu = 0.02$.

To mitigate overfitting during training, we apply several strategies in addition to the small learning rate and stopping rule discussed above. First, each decision tree is fitted on a random subsample containing 50% of the training observations. Second, within

[3]Figure C10 in the appendix shows that the cross-validation and forecasting error increase only slowly when using a large number of boosting rounds. Using early stopping in cross-validation, our baseline model stops training after 56 boosting rounds, on average.

each block, trees are trained on a random subset of 25% of the predictors. Third, tree complexity is constrained by limiting the maximum depth to 3 and requiring at least 5 observations per node. The results are robust to these hyperparameter choices, as shown in Appendix C.5.

Note that the random permutation of blocks in each boosting round operationalises the assumption that every block conditions on the others during training. Combined with the slow learning rate, this ensures that blocks have the same opportunity to contribute to the overall prediction. In alternative specifications, we also explore more structured training sequences that impose explicit assumptions about block ordering—for example, training the global blocks or the trend block first, before proceeding with random permutations of the remaining blocks.

We implement BBIM in Python and use the `xgboost` library (Chen and Guestrin, 2016) to train individual decision trees. Although the library is primarily designed for gradient boosting, we restrict it to training single trees by setting `n_estimators` $= 1$ and `learning_rate` $= 1$. This library is computationally efficient and effectively imposes monotonicity constraints with the `monotone_constraints` parameter.

2.5 Empirical implementation for UK inflation

2.5.1 Data and block specification

Our variable of interest is monthly CPI inflation in the United Kingdom. The sample period spans from February 1988 to December 2024. The baseline model includes 54 monthly indicators, grouped into distinct blocks and incorporated contemporaneously and with two lags each.[4] A few indicators have missing values at the start or end of the sample, which we impute using median values.[5] Series that are not available in seasonally adjusted form (mainly CPI sub-components) are seasonally adjusted using X-12-ARIMA. To ensure stationarity, most variables are expressed as month-on-month log changes. Full details on indicators and grouping by blocks, data coverage, transformations, and data sources are provided in Appendix B.

A summary of the groups of indicators in each model block and the associated monotonicity constraints is provided in Table 1. The trend component includes household one-year-ahead inflation expectations, five-year market-based expectations derived from inflation swap rates, and indicators of domestic inflation persistence, such as services inflation and aggregate as well as industry-level wage growth. A linear time trend is also included to allow for a smoothly evolving inflation trend.

The demand and supply components are informed by a broad set of activity indicators, including indices of production and services, PMIs, trade volumes, retail sales, and consumer confidence—as well as labour market slack measures, notably deviations of the unemployment rate and the vacancy-to-unemployment ratio from their trends as a proxy for labour market tightness.[6] We also incorporate identified global activity (demand) and

[4] We experimented with alternative lag numbers, with results similar, see Figure C5 in appendix.

[5] We also tested Expectation Maximisation, which gave overall similar results, but tended to attribute stronger Shapley contributions to imputed values. This is likely because the latent-variable method captures co-movement that the model interprets as meaningful signals.

[6] As trend estimates around which the gap measures are computed, we use estimates for the long-run equilibrium unemployment rate (u^*) and a trend measure for the v-u-ratio produced by the Bank of

BLOCK	INDICATORS	SIGN	∈D&S
Trend	time indicator, 1-y ahead household infl. expectations, 5-y ahead financial market expectations, regular wage growth, services inflation, sub-components by sector		
Global demand	global PMI; UK exports, US, EA: industrial production: US, EA: imports	+	yes
	global activity shock, oil consumption demand shock (BH2019)	+	
Domestic demand	industrial production, index of services; UK imports, PMIs: services, manufacturing, construction; retail sales; consumer sentiment, quarterly (interpolated): consumption, investment	+	yes
	Labour market: v/u ratio gap,	+	yes
	Labour market: unemployment rate gap	−	yes
Global supply	global PMI; US, EA: industrial production, UK: imports	−	yes
	commodity prices: energy, non-energy, metals, food, agriculture	+	
	global supply chain pressures: GSCPI (Fed), SCI (BoE)	+	
	US PPI, EA PPI	+	
	oil supply news shock (Känzig, 2021), global oil supply shock (BH2019)	+	
Domestic supply	industrial production, index of services; UK exports, PMIs: services, manufacturing, construction; retail sales; consumer sentiment, quarterly (interpolated): consumption, investment, TFP	−	yes
	Labour market: v/u ratio gap	−	yes
	Labour market: unemployment rate gap	+	yes
	CPI components: goods, food, electricity, gas;	+	
	PPIs: input, output, gas, electricity; UK spot gas price	+	

Notes: Last column indicates groups of activity indicators that are included in both demand and supply blocks, under monotonicity constraints with opposite signs.

Table 1: Groups of indicators and restrictions imposed to separate supply and demand.

oil supply series (Baumeister and Hamilton, 2019, referred to as BH2019 in what follows) and oil supply news shock series (Känzig, 2021) directly into the corresponding blocks. Supply blocks are further enriched with supply chain pressure indices and selected CPI sub-components. Both demand and supply are disaggregated into global and domestic sub-blocks. For instance, the global supply block includes the global supply chain pressure index, global oil price shocks, and commodity prices, while the domestic supply block covers CPI food and energy components, domestic supply pressure indicators, and negatively restricted economic activity variables. Monotonicity constraints on UK trade variables help distinguish global from domestic blocks: UK exports load positively on global demand and negatively on domestic supply, while UK imports load positively on domestic demand and negatively on global supply.

2.5.2 Estimation and forecast evaluation

To derive the inflation decompositions and functional forms in our main results, we use 10-fold cross-validation. This approach enables us to train the model across the full sample period while estimating consistent functional relationships without having to adjust for model instability or shifts over time (see also Bluwstein et al., 2023; Buckmann et al.,

England. In alternative specifications, we also used the unemployment rate and v-u-ratio directly in levels and log levels, or alternatively in differences, to check that our results were not dependent on specific trend estimates.

2023). We fit annualised monthly CPI inflation 1-month ahead in our main specification, but also run specifications for longer horizons, with $h = 1, 3, 6, 9, 12$.[7] To enhance the model's robustness to randomness in the estimation process—arising from data sampling in cross-validation, tree building, and the random order of components in each boosting iteration—we repeat the cross-validation procedure 10 times and average the predictions.

We also evaluate forecasting performance against various linear and standard machine learning benchmarks, in a pseudo out-of-sample exercise, with more details and results discussed in Section 3.6.

2.5.3 Shapley values for ex-post signal communication

While our model linearly decomposes predicted inflation into components, this does not directly reveal the learned relationships between individual variables and the output. To understand these functional forms, we use Shapley values, a standard interpretability tool in machine learning (Štrumbelj and Kononenko, 2014; Lundberg and Lee, 2017). Following Buckmann et al. (2023), each predicted value $\hat{y}_i$ is decomposed as $\hat{y}_i = \sum_j \phi_i^j + \phi^0$, where ϕ^j is the Shapley value of predictor j and ϕ^0 is the baseline (usually the mean prediction from the training sample). Indicators not entering in any tree split have zero Shapley values, while larger values indicate stronger contributions. We compute Shapley values for each tree and aggregate them across trees belonging to the same component. To assess the average importance of each variable over the sample, we calculate mean absolute Shapley values. Estimation uses the Python `shap` library (Lundberg and Lee, 2017) with the efficient `TreeExplainer` method (Lundberg et al., 2018).

3 Results

3.1 Inflation determinants

Figure 2 presents the decomposition of the Blockwise Boosted Inflation model predicting CPI one month ahead. This specification focuses on month-on-month annualised CPI inflation, and includes blocks for global and domestic demand and supply, identified using monotonicity constraints, as well as the expectations-informed trend block.

The model provides a meaningful decomposition of UK CPI inflation over the sample period. During the early 1990s high-inflation episode, the trend component (purple) was driven by rises in both long-term and short-term expectations, gradually declining as expectations re-anchored. While generally slow-moving, the trend adjusts relatively quickly when expectations shift. Crucially, these dynamics are not imposed a priori but emerge from the model learning non-linear associations, with higher expectation realizations exerting stronger inflationary effects. Supply (orange and yellow) contributed modestly in the early 1990s, exerted downward pressure from the late 1990s to mid-2000s—likely reflecting global supply chain integration—and again in the mid-2010s after the drop in global oil prices. Demand (green) typically contributes positively, though the

[7] We also estimated a specification with year-on-year CPI inflation as the target and used indicators in year-on-year growth rates. Figure C6 shows that decompositions are similar to those for annualised monthly CPI growth. Non-linear were also broadly consistent—though month-on-month transformations yield slightly clearer signals.

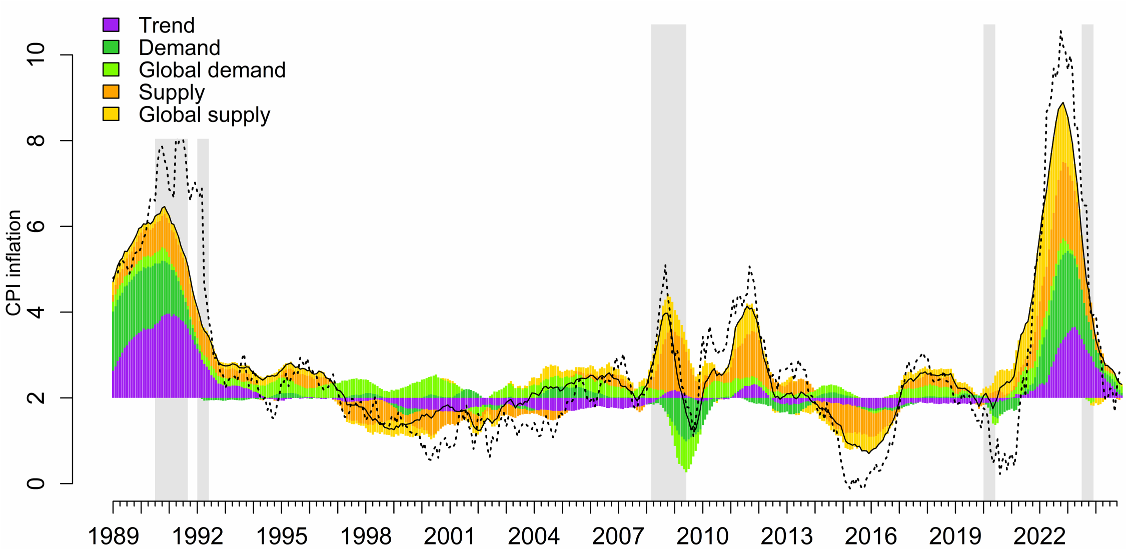

Figure 2: Decomposition of 1-month-ahead CPI inflation from main BBIM specification.
Notes: Blockwise predictive contributions (coloured bars) to 1-month-ahead annualised month-on-month CPI inflation (black line), around inflation target (2%). Model estimated via cross-validation over 1988M2–2024M12. Dashed line: actual CPI inflation, annualised and smoothed over 12 months. Grey bars indicate recessions.

model captures sharp negative demand effects during the Global Financial Crisis and near-zero contributions during the Covid-19 pandemic.

In the recent inflation episode, supply factors were the main drivers of the initial surge. Global supply effects rose sharply in 2021–2022 and spilled over into domestic supply, though both components unwound over 2023. Domestic demand contributed strongly from 2021 to mid-2023, driven primarily by tight labour markets rather than broader activity measures. Later in the episode, the inflation trend rose, reflecting mainly rising short-term expectations that had non-linear effects. Part of inflation in 2022–early 2023 remains unexplained.

Block contributions reflect meaningful signals from individual indicators, based on mean absolute Shapley values (Figure 3). The trend is largely driven by short-term household expectations, with smaller contributions from five-year market-based expectations and services inflation; wage growth is weak. Effects from short-term expectations dominated the recent trend rise more so than in the past, while long-term expectations remained anchored and thus contributed less than in the early 1990s. Global supply is driven by supply chain pressures, partner production, and commodity prices, while domestic supply captures the transmission of global shocks into food and goods prices. Global demand contributions stem from shocks to oil demand, economic activity, trade, and production in major partners.

Domestic demand is primarily influenced by labour market indicators, services indices, retail sales, PMIs, and imports. Notably, labour market indicators were the main demand-sided drivers in the recent episode. This is highlighted in an extended BBIM specification (Figure 4), where the domestic demand block is split into labour market tightness—comprising the positively constrained vacancy-to-unemployment gap and negatively constrained unemployment gap—and other demand, which includes broader ac-

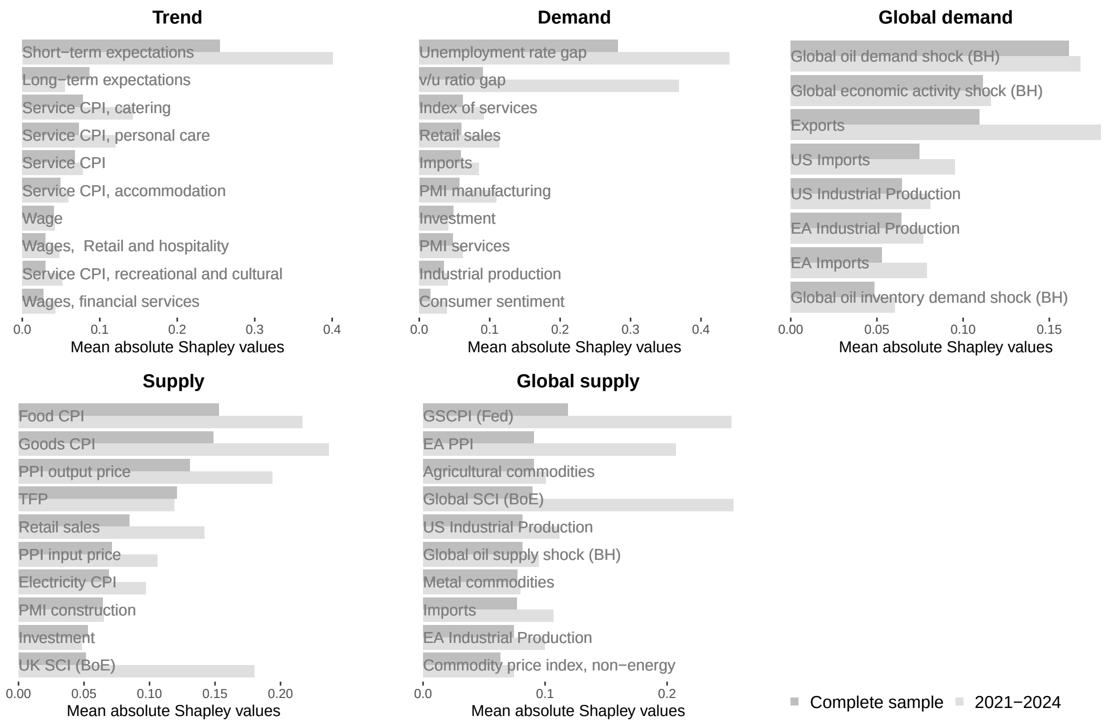

Figure 3: Mean absolute Shapley values within each component.
Notes: Ten indicators with the largest average Shapley contributions in each block, averaged over 1989-2024 (dark grey) and 2021-2024 (light grey).

tivity indicators constrained positively.[8]

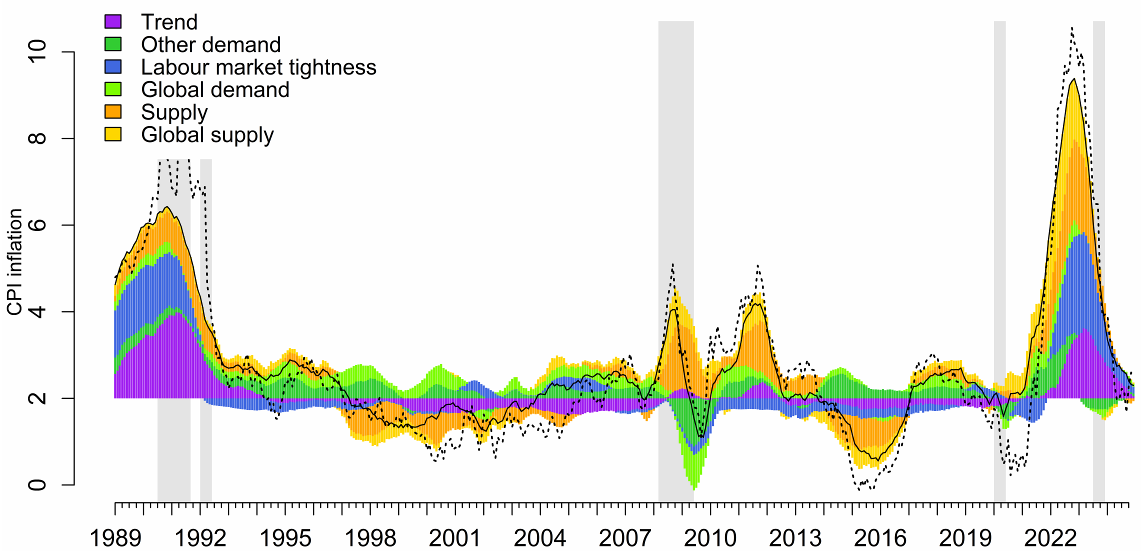

Figure 4: BBIM specification with domestic demand split into "labour market tightness" and "other demand".

Demand effects linked to labour market tightness (blue bars) were the dominant driver in the recent inflation episode and in the early 1990s, whereas other demand (green bars) played a larger role in periods such as the GFC. In the recent episode, the vacancy-

[8]The global and domestic supply blocks and the trend component are unchanged in this exercise. Figure C7 in the appendix shows the average Shapley contributions for each block in this model.

to-unemployment ratio gap emerged as a key signal, alongside a negative unemployment gap. By contrast, in the early 1990s, the v/u ratio gap was less relevant, with the negative unemployment gap around a structurally high equilibrium rate u^* driving the demand-like labour market contribution. During recessionary periods like the GFC, it was broader contractions in domestic activity that acted in a deflationary manner.

3.2 The role of non-linearities in inflation determinants

3.2.1 Non-linearities in the Phillips curve, supply, and inflation expectations

Next, we examine the direction and shape of signals extracted by the model from specific indicators within each block. For indicators with monotonicity constraints, the model learns strictly non-decreasing or non-increasing relationships, though the slope and shape can vary. For unconstrained indicators, the model learns fully flexible functional forms. Comparing Shapley values to actual indicator values allows us to assess the non-linear patterns captured over time.

Figure 5 illustrates these learned functions for selected indicators and lags of each block of our baseline model.[9] We plot indicator values (horizontal axis) against their Shapley value contributions to CPI inflation predictions (vertical axis). Dots represent monthly observations, with early 1990s data in dark grey and 2021–2024 data in colour. Non-linear patterns arise when stronger signals align with very high or low indicator values. Several key indicators exhibit strong non-linear signals during the recent UK inflation episode.

First, we detect a non-linear Phillips curve relationship between labour market slack and inflation within the domestic demand block (panel (a) of Figure 5). During the recent period—particularly in 2022 (red dots)—the unemployment rate remained below 5%, creating a negative unemployment gap relative to trend, while the vacancy-to-unemployment (v/u) ratio was high relative to its slow-moving trend. For these pronounced realisations, the model produces markedly stronger inflationary signals, with Shapley contributions rising from near zero to 0.2–0.5 percentage points. These non-linear effects, though less intense, persisted into 2023 and largely flattened by late 2024. The steep relationship between labour market indicators and inflation reflects L-shaped Phillips curve dynamics that a linear model cannot capture, explaining the increased demand-like contribution from tight labour markets shown in Figure 4. This aligns with recent U.S. and cross-country evidence on non-linear Phillips curves amplifying inflation under tight labour conditions (Benigno and Eggertsson, 2023, 2024).

By contrast, broader activity indicators show no notable non-linear demand effects in the recent episode. PMIs, trade, industrial production, and investment exhibit stronger effects primarily during contractions, such as the GFC—both within domestic and global demand (panel (b) of Figure 5, Figure C2). Some non-linearities reflect asymmetric responses to outliers: global activity shocks, whether expansionary or contractionary, affect UK inflation most strongly within a moderate range, while very large shocks—e.g., the Covid-19 contraction—are down-weighted and have muted deflationary effects. A similar pattern occurs for UK retail sales and PMIs, and for US production and imports.

[9] We focus on selected indicators and on more distant lags for variables typically released with delay; full results for all indicators and all three lags are shown in Figure C1 in the appendix. Non-linear associations remain very similar for the extended model shown in Figure 4.

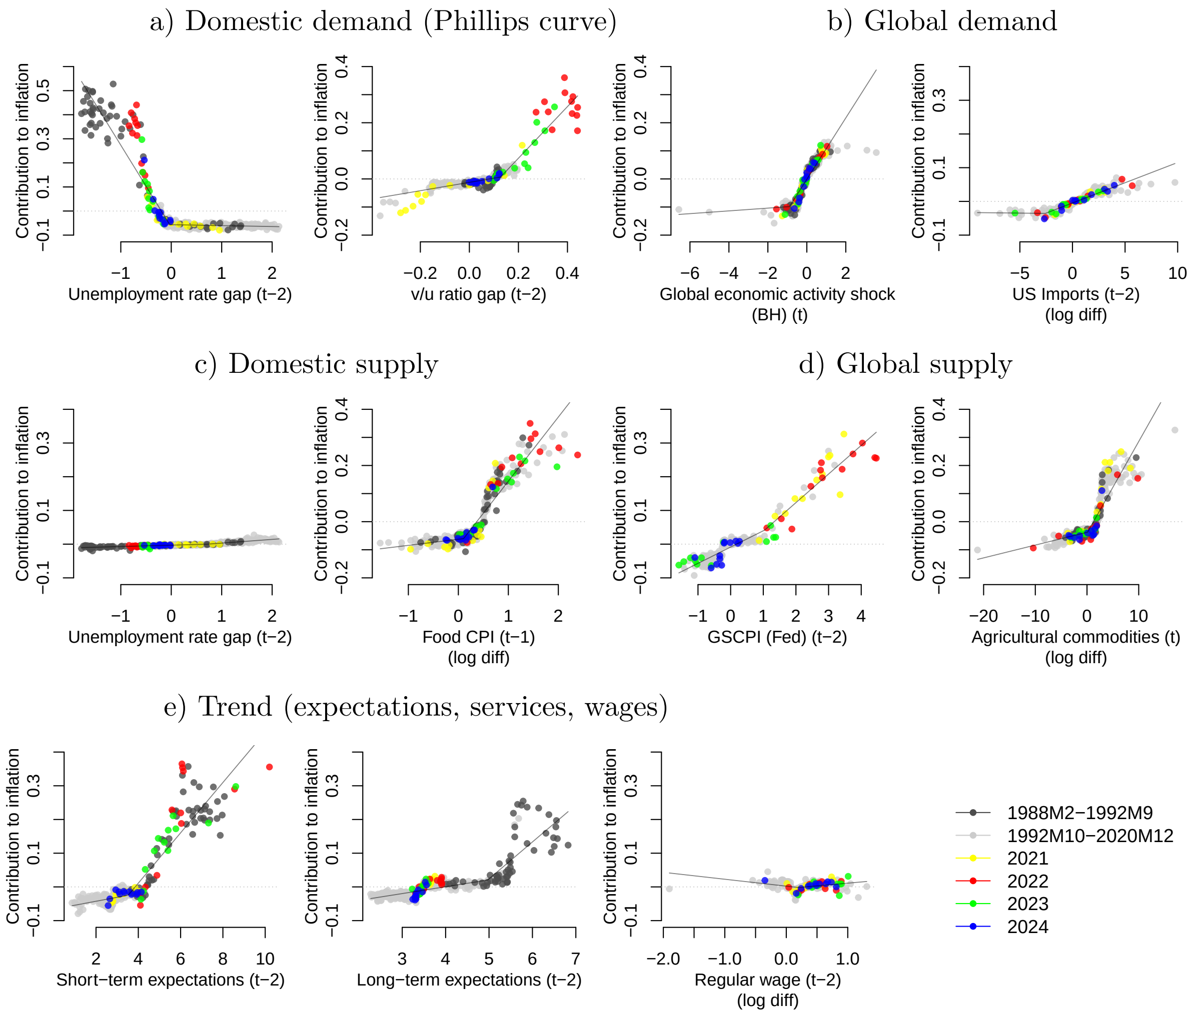

Figure 5: Functional forms learnt within each block.
Note: Baseline BBIM (see Figure 2). Scatter plots show the realised value of each indicator at a selected lag (x-axis) against its Shapley value contribution to the model's CPI inflation prediction (y-axis), within a given block (panels a–c). Monthly predictions are shown without smoothing. For comparison, a fitted linear model with a single breakpoint (i.e. two slopes) is overlaid in grey, following Muggeo (2003).

These results align with Harding et al. (2022), highlighting muted disinflation during contractions as firms resist lowering prices to preserve mark-ups.

Second, on the supply side (panels (c) and (d) of Figure 5), the global supply chain pressure index (GSCPI) had little impact before the pandemic but became highly influential during 2021–2022, consistent with international evidence of a non-linear inflationary effect from global supply disruptions (Di Giovanni et al., 2022; Comin et al., 2023; Ascari et al., 2024). We also observe non-linear effects from global agricultural and food commodity price inflation, followed—after a lag—by similar patterns in UK food CPI inflation, which significantly contributed to aggregate UK inflation through 2023. These patterns likely reflect delayed transmission of global shocks and a stronger pass-through of rising input costs (Cavallo et al., 2024). Energy-related variables—including global energy commodities, oil supply news shocks, and domestic electricity and gas prices (see Figure C1 in the appendix)—exhibit more diffuse functional forms, with less clear evidence of non-linear effects beyond potential indirect channels through supply chains. Labour market slack does not show a strong signal on the supply side.

Third, we find that the recent rise in the inflation trend component was primarily driven by strong effects from short-term household inflation expectations (panel (e)). The model detects little signal when expectations are below 4%, but identifies steadily increasing contributions at higher levels—where UK household expectations remained through late 2023. This non-linearity may reflect greater attentiveness and stronger adjustment of expectations when inflation is elevated (Pfäuti, 2023), particularly in response to salient price increases in categories like food and energy (Anesti et al., 2024). This can, in turn, fuel persistent wage and price pressures (Lorenzoni and Werning, 2023), though our model finds only weak signals from wage growth. By contrast, long-term financial market expectations had strong effects in the early 1990s but have remained flat in recent years, suggesting they remain well anchored.

3.2.2 Shutting down of non-linearities within blocks

How relevant are the detected non-linearities for inflation prediction and decomposition? To assess this, we run a counterfactual exercise where non-linearities are removed block by block. Specifically, we replace decision trees with linear least-squares models in a given block—while keeping other blocks unchanged. Monotonicity constraints, as in the baseline, are still imposed. Figure 6 compares non-linear and linear blockwise contributions for the extended BBIM specification that separates demand into labour market tightness and "other demand".[10]

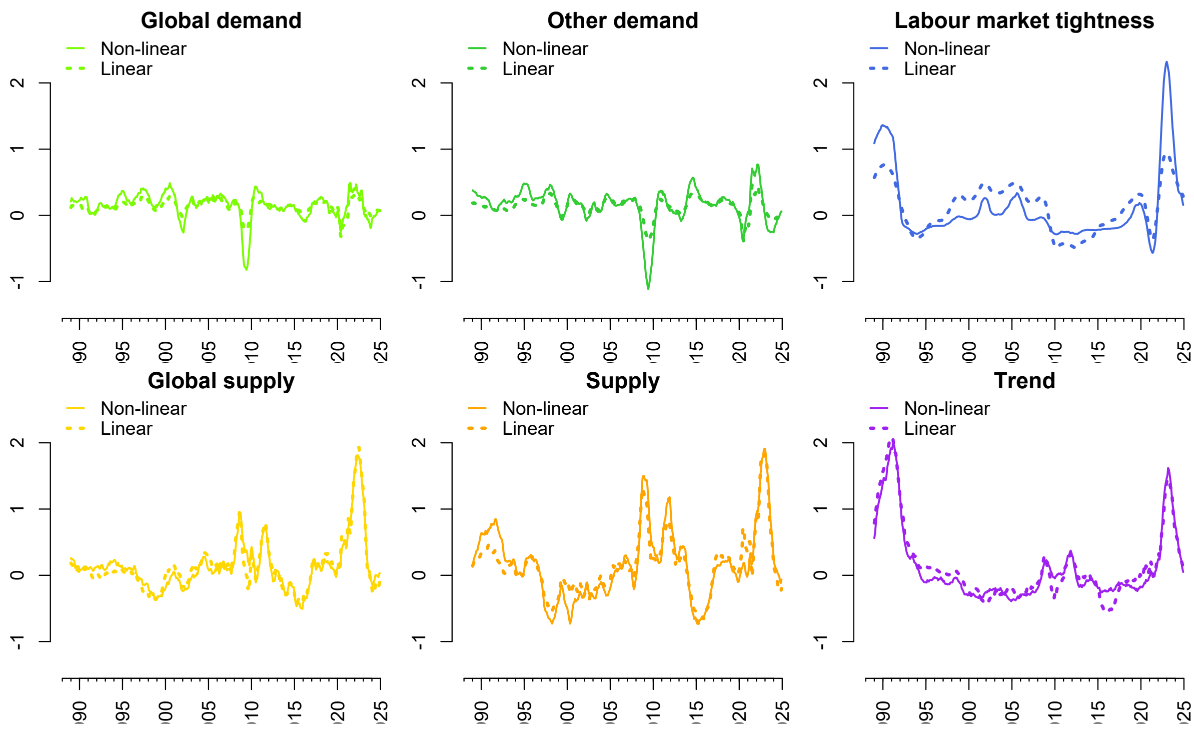

Figure 6: Shutting down non-linearities within a given block.
Notes: Blockwise predictive contributions to 1-month ahead annualised CPI inflation. Dashed lines: contributions from specifications where decision trees are replaced with linear least-squares models within the respective component, compared to the block contribution from the fully non-linear model (solid lines).

[10] Shutting down non-linearities in one block had only small effects on other blocks, so we only show here the respective block from each specification. Full decompositions are available upon request.

Accounting for non-linearities in demand has the strongest impact on the predictive decomposition. When global demand and broad ("other") domestic demand blocks are modeled linearly, the model underestimates the negative demand contribution during the Global Financial Crisis. More recently, non-linearities in these components have played a smaller role. By contrast, removing non-linearities from the labour market tightness block has a substantial effect in both high-inflation episodes. In the early 1990s, these effects mainly reflected a negative unemployment gap, with unemployment falling below a high trend rate. In the recent episode, a combination of very low unemployment and a high vacancy-to-unemployment ratio generated strong non-linear effects. During the intervening low-inflation period, the linear component instead captured slow-moving rather than cyclical fluctuations in labour market tightness. This illustrates that non-linearities in inflation's sensitivity to labour market slack and broader demand are both important, but their relevance can shift over time: broader demand non-linearities primarily explain disinflationary pressures during recessions, while tight labour markets non-linearly amplified UK inflation during high-inflation periods.

In contrast, removing non-linearities in the other blocks has more limited effects. Linearising the supply blocks (yellow lines) causes the model to miss some supply-side contributions early in the sample and overstate others, but the differences are generally small, especially in the recent period. This is consistent with the functional forms discussed earlier, where supply indicators exhibit some non-linearity—mainly through down-weighting of outliers—but their overall relationship to inflation could also be approximated by a linear component. When the trend block is linearised (blue lines), the estimated trend becomes more volatile, reducing the model's ability to differentiate between a stable, flat trend in a low-inflation environment and periods of genuinely shifting inflation trends.

3.3 Role of monotonicity constraints for decomposing supply and demand determinants

We have seen that accounting for non-linearity helps isolate a distinct role for demand. However, the identification of demand versus supply ultimately also relies on the monotonicity constraints imposed within each block. To assess their importance, we compare our baseline BBIM to an unidentified blockwise Phillips curve-style boosting model, as defined in Equation 4. This alternative model simply groups all (global and domestic) activity indicators together—implicitly treating them as demand-driven—while assigning all (global and domestic) supply and cost indicators to a separate block.

Figure 7 compares this unidentified specification to the demand and supply components from our baseline model. The activity block in the unidentified model (red line, left panel) fails to reflect key demand dynamics evident in the baseline—such as the negative demand contributions during the GFC and Covid-19 recessions—and instead captures broader, slower-moving cycles. It does, however, pick up the recent rise in demand, indicating that the strong non-linear effect of labour market tightness emerges even without additional identification. Conversely, the supply and input cost block (blue line, right panel) tracks supply fluctuations reasonably well but tends to understate the magnitude of supply contributions relative to the baseline.[11]

[11] Here, we exclude the identified shock series to avoid providing explicit identifying information. Results are similar when included (Figure C3), suggesting that monotonicity constraints drive most of the identification, with shocks mainly refining the magnitude of large global fluctuations.

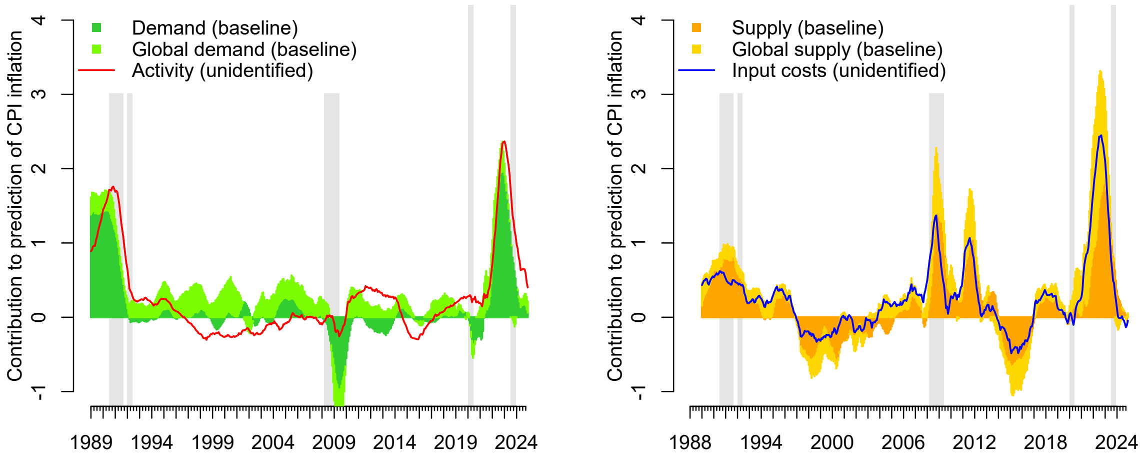

Figure 7: Comparison with unidentified model without monotonicity constraints.
Notes: Predictive contributions to 1-month-ahead annualised CPI inflation from baseline BBIM with monotonicity constraints—separating demand (green bars, left) and supply (yellow bars, right)—compared to unidentified blockwise boosting that groups activity indicators into a single block (red line, left) and all supply and cost indicators into another (blue line, right). Both models include a trend component (not shown).

Further, Figure C4 in the appendix shows that the functional forms estimated by the unidentified model are less interpretable. The non-linear Phillips curve relationship with the vacancy-to-unemployment and unemployment gaps are notably weaker. Moreover, relationships with broader activity indicators, such as retail sales, U.S. imports, and U.S. industrial production, display mixed signs across lags, failing to yield a coherent or economically interpretable pattern.

3.4 Alternative block structures

We assess the robustness of the inflation decomposition and learned functional forms to alternative assumptions regarding block ordering during boosting and the modelling of the trend component. In the main BBIM specification, each boosting iteration fits one tree per block, with the block order randomly reshuffled thereafter (see Algorithm 1). We consider two alternative ordering strategies. First, a fixed ordering where 100 trees from the expectations-informed trend block are estimated before updating the remaining blocks in random order—prioritising persistent, forward-looking components. Second, a reversed structure that fits 100 trees each from the global demand and supply blocks first, conditioning on external pressures before domestic ones. Additionally, we explore three alternative specifications for the trend: an uninformed trend using only a time variable; a trend informed solely by the two expectations series and the time variable, excluding other persistence indicators; and a specification without a trend component, attributing inflation entirely to global and domestic demand and supply. Figure 8 shows the resulting prediction of inflation and blockwise contributions.

While block ordering influences the relative size of contributions—with earlier-trained blocks receiving more weight—the overall inflation prediction remains broadly stable. In the block-exogenous ordering (orange lines), global demand and supply become more prominent, while domestic components and the trend contribute less. Instead, prioritising

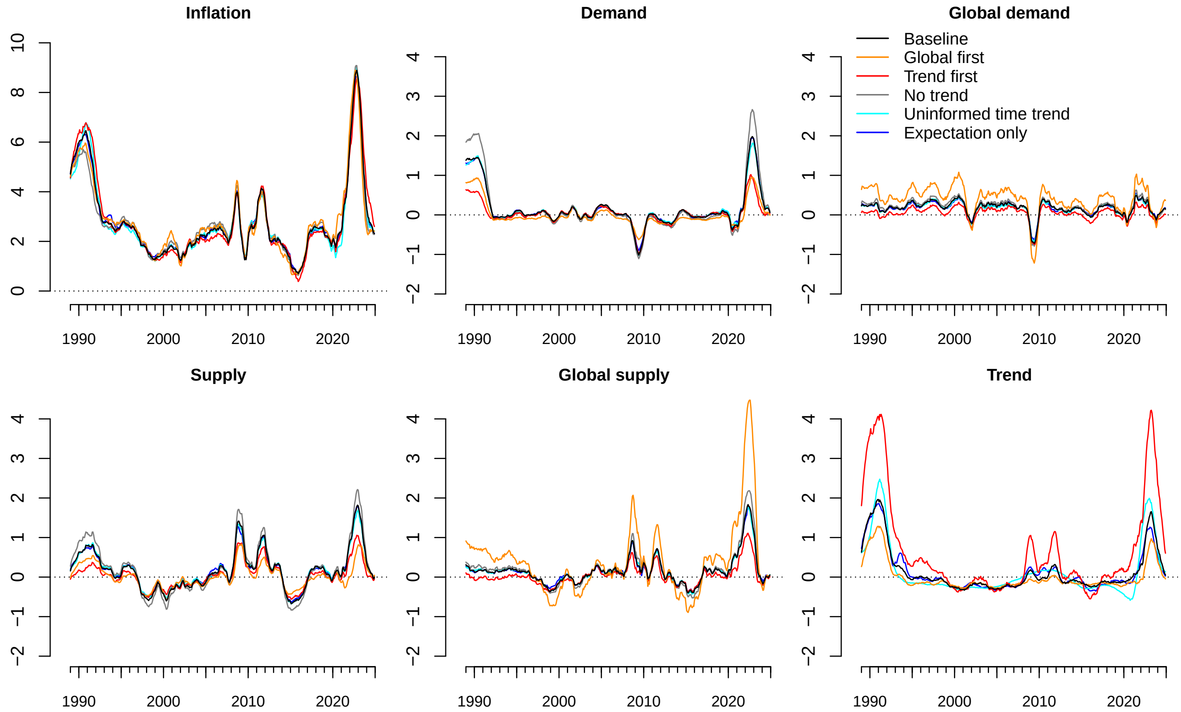

Figure 8: Model specifications with alternative block orderings and trend specifications.
Notes: Predictive contributions to 1-month-ahead annualised CPI inflation (first panel) and blockwise components, comparing baseline BBIM (black) with alternative specifications for block orderings—global blocks fitted first (orange), trend block fitted first (red)—and trend: no trend (grey), trend based on a time variable only (teal), trend based on both a time variable and expectations series (blue).

the trend block (red lines) substantially increases the trend contribution while reducing the role of demand and supply. When removing the trend altogether (grey lines), the model assigns stronger contributions to both supply and demand—but this leads to a more pronounced underestimation of inflation in the early 1990s. The model with an uninformed time-trend (teal lines) captures broad dynamics but overstates trend fluctuations during the Covid-19 period and the recent inflation surge. A model with a trend based solely on the two expectations series (blue lines) closely matches the baseline, though with a slightly weaker recent trend contribution—consistent with the baseline additionally drawing on services inflation as a driver of recent trend dynamics.

Overall, the cyclical patterns of block contributions are preserved, and the learned functional forms remain stable across specifications (see Figure C2 in the appendix), suggesting that the model captures consistent underlying dynamics regardless of block ordering.

3.5 Extended model with financial conditions and monetary policy surprises

In an extended model, we add two blocks that capture financial conditions and monetary policy surprises. The financial block includes UK house prices, corporate bond spreads, the FTSE stock price index, the real effective exchange rate and the US-GBP spot exchange rate. The monetary policy block uses identified UK policy surprises—target rate, path, and QE factors—from Braun et al. (2023); it thus does not capture systematic monetary policy but instead instances where monetary policy surprises on the upside or downside. To pin down the direction of this block's contribution to reflect expansionary surprises, we impose negative monotonicity constraints on the association between inflation and the target and path factors, respectively. As data for these additional indicators have a shorter coverage, this specification covers the shorter sample 1997–2024. Figure 9 presents the resulting decomposition, and Figure C8 displays the mean absolute Shapley values.

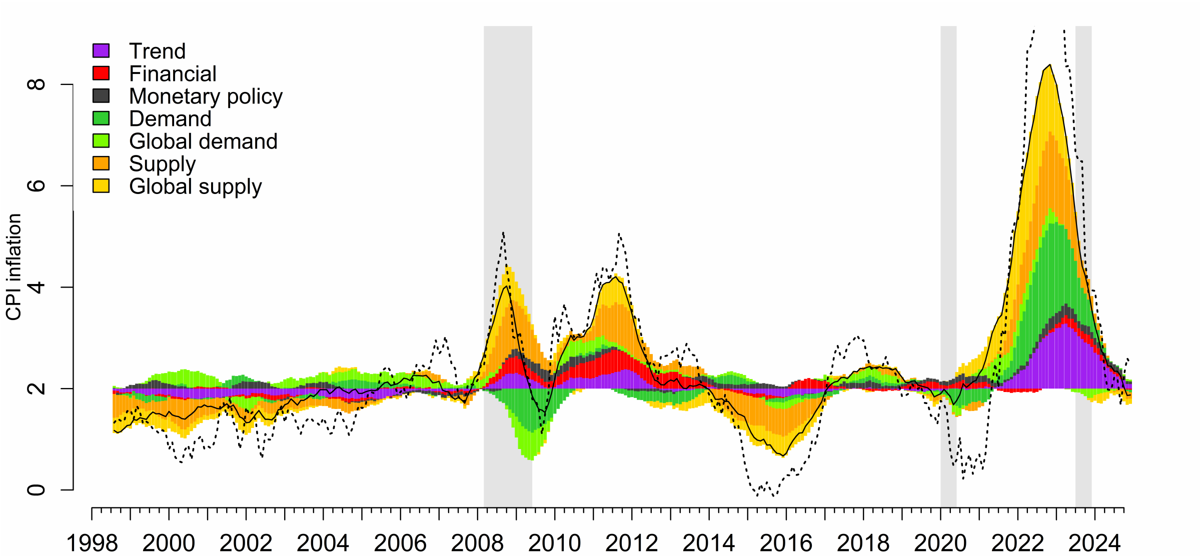

Figure 9: Extended model with financial and monetary policy surprises blocks.
Notes: Predictive contributions to 1-month-ahead annualised CPI inflation from extended BBIM with additional blocks containing monetary policy surprises (black) and financial conditions (red). Estimation with cross-validation over sample period 1997M9-2024M12.

Within the financial block, house prices and corporate bond spreads are the most influential indicators on average. In the monetary policy block, the QE and path factors dominate, capturing the impact of unconventional monetary policy surprises.[12] Looking at the predictive decomposition over time, tight financial conditions contributed to inflation following the Global Financial Crisis (GFC)—mainly reflected in rising corporate bond spreads, which may have prompted firms to raise prices, consistent with Gilchrist et al. (2017). Exchange rate movements following the 2016 Brexit referendum also provide inflationary signals.

Expansionary monetary policy surprises contributed modestly to inflation predictions post-GFC, to a lesser extent in 2017–2019 and again during the Covid-19 stimulus in 2020. In late 2022 and early 2023, monetary surprises are as mildly inflationary by the model—likely reflecting slower than expected tightening given the economic context at the time. Contributions from the demand, supply, and trend components remain broadly in line with the baseline model, though post-GFC supply effects are smaller—now partly captured by the financial block.

3.6 Out-of-sample forecasts

The key strength of the BBIM lies in its intuitive cross-validation decomposition of inflation drivers and its ability to trace non-linear signals. Nonetheless, strong out-of-

[12] We also estimated a specification with only the path and target factors; results were similar, indicating that including the QE factor does not drive the findings.

sample forecasting performance remains an important evaluation criterion. Ideally, the introduction of economic structure should not come at the expense of predictive accuracy—typically a strength of machine learning models. We therefore assess the BBIM's forecasting performance in a pseudo out-of-sample exercise, comparing it against linear and standard machine learning benchmarks. Starting with an initial training sample ending in 1999M12, we generate monthly forecasts from 2000M1 to 2024M12 using an expanding window, retraining the model every three months.[13] The forecast averages predictions from 10 boosting models, each trained on a random 80% subsample of the respective training set. Forecasts are produced for horizons $h = 1, 3, 6, 9, 12$ months ahead. Figure 10 presents annualised h-month-ahead out-of-sample forecasts, generated at the start of each year in the test sample, alongside the actual CPI inflation realisation (panel a), as well as corresponding out-of-sample forecasts for each of the block components, alongside the annualised 1-month-ahead cross-validation decomposition (panel b).

At the 1-month horizon, the out-of-sample forecasts for CPI inflation and the block-wise component predictions closely replicate the cyclical patterns observed in the cross-validation results. In the early 2000s inflation is mostly over-predicted, which reflects that at this point the model has a relatively short training sample that in large parts covers a high-inflation period whereas realised inflation turned out lower given a low inflation trend following the introduction of inflation targeting as well as global supply factors weighing on inflation. Subsequently the model predicts quite well at shorter horizons, albeit with a lag at higher horizons. The recent episode of high inflation is under-predicted out-of-sample, driven by lower contributions of domestic demand and global supply—precisely the components where we observe strong non-linearities. Out of sample, the model captures these non-linearities only to a limited and noisier extent, even at short horizons. At higher horizons, predicted fluctuations are more muted and at times lagged, but broadly aligned with the cross-validation results.

Next, we evaluate the BBIM's forecasting performance relative to a range of alternative models. Forecast accuracy is measured by root mean squared forecast errors (RMSFE) relative to a simple AR(2) benchmark, with statistical significance assessed using the Diebold and Mariano (1995) test. We compare the baseline BBIM to variants without block structure or monotonicity constraints, as well as to standard linear and machine learning benchmarks, including an unobserved components model (UC, Stock and Watson, 2007), Lasso regression, random forest (Breiman, 2001), and standard boosting. Implementation details are provided in Appendix A.1. Most alternative models use the same indicator set as the BBIM but omit its structural constraints. For standard boosting without block structure, we also test specifications using only a subset of indicators—specifically those informing the trend in our baseline (time indicators, expectations, wage growth, and services inflation), excluding activity and supply variables.

Table 2 reports forecast evaluation results for annualised monthly CPI inflation at horizons $h = 1, 6, 12$ months. At short horizons (one to six months), the baseline BBIM delivers the strongest results, reducing RMSFE by around 10% pre-pandemic and up to 25% during the recent high-inflation period—though model differences are less significant in the latter due to smaller sample size. Alternative boosting specifications without structural constraints and using fewer indicators perform notably worse, especially post-

[13] As out-of-sample forecasting is not the primary focus of our analysis, we abstract from data revisions and publication delays in this stylised setup.

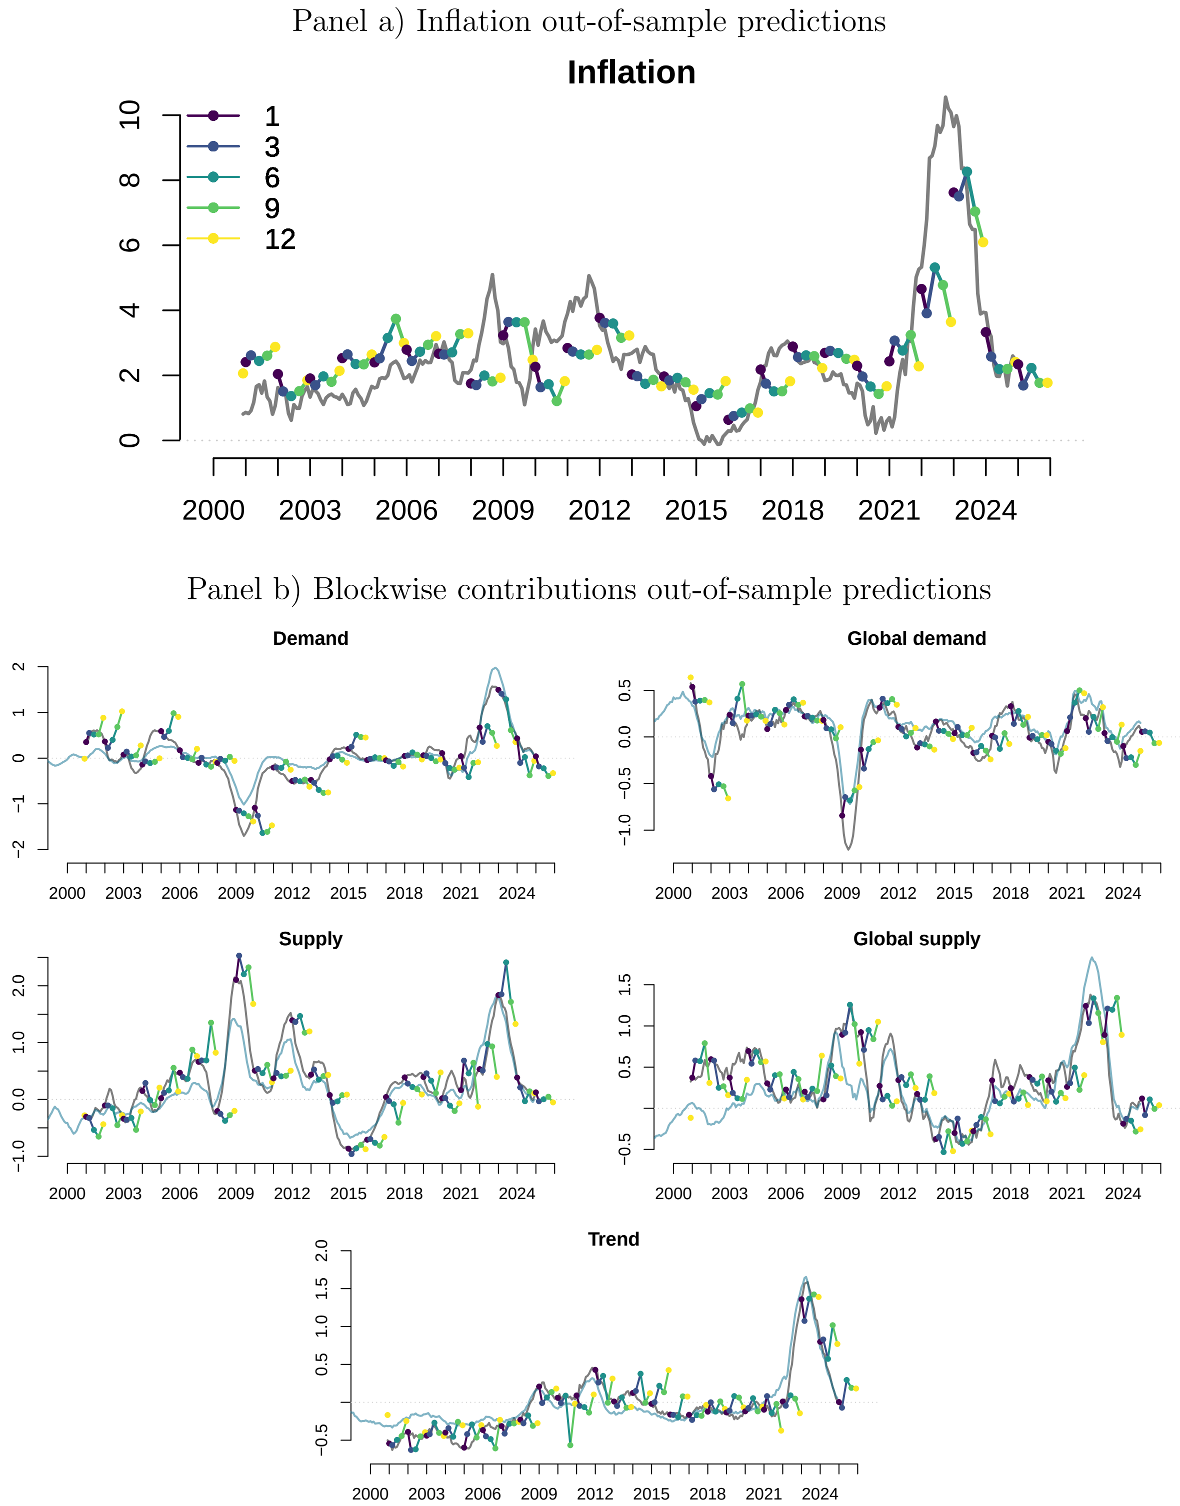

Figure 10: Pseudo out-of-sample forecasts over horizons.

Notes: Predictions of annualised monthly CPI inflation (panel a) and of blockwise components (panel b). Grey lines show actual monthly CPI inflation (annualised). Teal lines in panel b) indicate 1-month-ahead cross-validated predictive contributions by block ($h = 1$), as in Figure 2. Coloured dots represent out-of-sample forecasts at 1-, 3-, 6-, 9-, and 12-month horizons, generated at the start of each year.

pandemic, indicating that the BBIM's added structure improves both interpretability and out-of-sample forecasting performance. Compared to other models, the BBIM performs

	2000-2024	2000-2019	2020-2024
		monthly CPI inflation	
		h=1	
AR(2), absolute RMSFE	3.77	3.15	5.61
BBIM (baseline)	**0.87** (0.03)**	0.91** (0.01)	**0.82* (0.05)**
BBIM (without trend)	0.90** (0.04)	0.96 (0.27)	**0.82** (0.05)**
Blockwise boosting (unidentified)	0.90* (0.05)	0.93 (0.18)	0.85* (0.07)
Boosting	0.89** (0.05)	0.90* (0.07)	0.87* (0.08)
Boosting (only trend indicators)	0.92 (0.15)	0.92* (0.04)	0.93 (0.21)
Random forest	0.89* (0.06)	0.90 (0.13)	0.86* (0.08)
Lasso regression	0.88* (0.06)	**0.89* (0.06)**	0.86* (0.09)
UC	0.92 (0.15)	0.90 (0.15)	0.95 (0.18)
		h=6	
AR(2), absolute RMSFE	4.22	3.40	6.52
BBIM (baseline)	**0.83* (0.07)**	0.92 (0.35)	0.72* (0.05)
BBIM (without trend)	0.89 (0.11)	1.04 (0.81)	**0.69* (0.05)**
Blockwise boosting (unidentified)	0.84* (0.07)	0.88 (0.11)	0.80* (0.06)
Boosting	0.84* (0.06)	0.87** (0.02)	0.81* (0.06)
Boosting (only trend indicators)	0.86* (0.05)	0.86** (0.01)	0.87* (0.07)
Random forest	0.84* (0.06)	0.87** (0.03)	0.80* (0.06)
Lasso regression	0.87* (0.07)	0.89* (0.06)	0.84* (0.08)
UC	0.86* (0.07)	**0.85** (0.04)**	0.87* (0.08)
		h=12	
AR(2), absolute RMSFE	4.46	3.51	7.08
BBIM (baseline)	**0.85 (0.13)**	0.89* (0.06)	0.80 (0.13)
BBIM (without trend)	0.88 (0.14)	0.99 (0.49)	**0.77 (0.12)**
Blockwise boosting (unidentified)	0.86* (0.10)	0.84** (0.02)	0.88 (0.13)
Boosting	**0.85* (0.10)**	**0.83** (0.02)**	0.86 (0.12)
Boosting (only trend indicators)	0.87 (0.12)	0.84** (0.03)	0.90 (0.15)
Random forest	0.88* (0.09)	0.84** (0.02)	0.92 (0.14)
Lasso regression	0.89 (0.13)	0.89* (0.05)	0.89 (0.16)
UC	0.90 (0.17)	**0.83** (0.02)**	0.96 (0.23)
		year-on-year CPI inflation	
		h=12	
AR(2), absolute RMSFE	2.24	1.14	4.46
BBIM (baseline)	0.82* (0.08)	1.12 (0.53)	**0.72**** (0.01)
BBIM (without trend)	**0.80* (0.08)**	1.27 (0.97)	**0.62** (0.01)**
Boosting	0.85 (0.12)	0.94 (0.29)	0.83* (0.10)
Boosting (only trend indicators)	0.98 (0.49)	1.01 (0.41)	0.98 (0.49)
Blockwise boosting (unidentified)	0.87 (0.12)	0.97 (0.31)	0.84* (0.10)
Random forest	0.83* (0.09)	0.95 (0.25)	0.80** (0.05)
Lasso regression	1.00 (0.53)	1.20 (0.89)	0.94 (0.51)
UC	0.99* (0.10)	1.01 (0.54)	0.98** (0.03)

Table 2: Out-of-sample forecast evaluation comparison.

Notes: Upper panel: Forecasting month-on-month CPI inflation at horizons $h = 1, 6, 12$. Bottom panel: Forecasting year-on-year CPI inflation at $h = 12$. Top row: Absolute forecast error of AR(2) across horizons and evaluation samples; other rows: relative RMSFE of alternative models. BBIM (baseline): model specification as in Figure 7. Blockwise boosting (unidentified): as in Figure 7, using activity and input cost blocks without monotonicity constraints or identified shocks. Boosting: standard gradient boosting with decision trees; no monotonicity constraints, no blocks, all indicators. Boosting (only trend indicators): standard gradient boosting, no blocks, only including time indicator, inflation expectations, wage growth, services inflation. p-values from Diebold-Mariano tests shown in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

a bit worse than Lasso regression and UC before the pandemic, while outperforming all alternatives, including unstructured machine learning models, more clearly in the recent period. An unidentified variant of the BBIM without monotonicity constraints or separate demand-supply blocks performs slightly worse, suggesting that identifying a clean demand signal supports forecast accuracy. At the 12-month horizon, results are similar pre-pandemic, with the BBIM slightly falling behind UC and the unstructured machine learning models. During the recent period, BBIM performs best but absolute errors grow and model differences become statistically insignificant.

For the 12-month horizon, we also report results for a specification that directly predicts year-on-year CPI inflation, with indicators expressed in year-on-year growth rates. In this setup, the BBIM struggles to outperform the AR(2) benchmark at short horizons but shows marked improvements at longer horizons (9–12 months), achieving roughly a 20% RMSFE reduction relative to the AR(2). In the post-pandemic period, the BBIM also outperforms the Random Forest, the strongest unstructured benchmark.

4 Conclusion

We have introduced the Blockwise Boosted Inflation model, a framework for modelling inflation dynamics that combines the flexibility of machine learning with a structure in line with economic theory. By organising indicators into interpretable blocks—demand, supply, and trend—and by imposing monotonicity constraints consistent with economic priors, the approach enables a theory-informed decomposition of inflation drivers. These constraints, applied during model estimation, strengthen identification specifically of the demand contribution, and allow for transparent attribution of inflation movements, even in high-dimensional settings.

Applied to UK inflation, our results show that monotonicity constraints are crucial for achieving a meaningful separation between supply- and demand-driven pressures. Within the demand component, we identify a non-linear Phillips curve-type relationship between inflation and labour market slack, which contributed more strongly during the recent inflation episode. We also observe non-linear effects from global supply chain pressures, which amplified global supply influences and their pass-through to domestic costs, as well as from short-term inflation expectations, which drove a rise in the inflation trend. In terms of predictive performance, the structured model matches or exceeds standard machine learning benchmarks while providing significantly greater interpretability of inflation dynamics.

Overall, our approach demonstrates how structured machine learning can serve as a powerful tool for macroeconomic modelling—providing accurate forecasts alongside economically meaningful insights, with potential for broader future applications, beyond inflation to other macroeconomic variables or country contexts.

References

Anesti, N., V. Esady, and M. Naylor (2024). Food prices matter most: Sensitive household inflation expectations. London School of Economics and Political Science CFM Working Paper.

Arias, J. E., J. F. Rubio-Ramírez, and D. F. Waggoner (2018). Inference based on structural vector autoregressions identified with sign and zero restrictions: Theory and applications. Econometrica 86(2), 685–720.

Ascari, G., D. Bonam, and A. Smadu (2024). Global supply chain pressures, inflation, and implications for monetary policy. Journal of International Money and Finance 142, 103029.

Baumeister, C. and J. D. Hamilton (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. American Economic Review 109(5), 1873–1910.

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 WP No 31197.

Benigno, P. and G. B. Eggertsson (2024). Slanted-l phillips curve. In AEA Papers and Proceedings, Volume 114, pp. 84–89. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Bernanke, B. and O. Blanchard (2025). What caused the us pandemic-era inflation? American Economic Journal: Macroeconomics 17(3), 1–35.

Bluwstein, K., M. Buckmann, A. Joseph, S. Kapadia, and Ö. Şimşek (2023). Credit growth, the yield curve and financial crisis prediction: Evidence from a machine learning approach. Journal of International Economics 145, 103773.

Braun, R., S. Miranda-Agrippino, and S. Saha (2023). Measuring monetary policy in the UK: the uk monetary policy event-study database.

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

Buckmann, M., A. Joseph, and H. Robertson (2023). An interpretable machine learning workflow with an application to economic forecasting. International Journal of Central Banking 19(4), 449–552.

Buckmann, M. and G. Potjagailo (2025). Infusing economically motivated structure into machine learning methods. Bank of England Staff Working Paper Series (1144).

Cano, J.-R., P. A. Gutiérrez, B. Krawczyk, M. Woźniak, and S. García (2019). Monotonic classification: An overview on algorithms, performance measures and data sets. Neurocomputing 341, 168–182.

Cavallo, A., F. Lippi, and K. Miyahara (2024). Large shocks travel fast. American Economic Review: Insights 6(4), 558–574.

Chalaux, T. and D. Turner (2023). Doombot: a machine learning algorithm for predicting downturns in OECD countries. Documents de travail du Département des Affaires économiques de l'OCDE.

Chan, J. C., T. E. Clark, and G. Koop (2018). A new model of inflation, trend inflation, and long-run inflation expectations. Journal of Money, Credit and Banking 50(1), 5–53.

Chen, T. and C. Guestrin (2016). Xgboost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 785–794.

Coibion, O. and Y. Gorodnichenko (2025). Inflation, expectations and monetary policy: What have we learned and to what end? National Bureau of Economic Research Working Paper 33858.

Comin, D. A., R. C. Johnson, and C. J. Jones (2023). Supply chain constraints and inflation. National Bureau of Economic Research WP 31179.

de Araujo, D. K. G., S. Doerr, L. Gambacorta, and B. Tissot (2024). Artificial intelligence in central banking. BIS Bulletin 2024.

Di Giovanni, J., Kalemli-Özcan, A. Silva, and M. A. Yildirim (2022). Global supply chain pressures, international trade, and inflation. NBER Working Paper No w30240.

Diebold, F. X. and R. S. Mariano (1995). Comparing predictive accuracy. Journal of Business and Economic Statistics 13(3), 253–263.

Döpke, J., U. Fritsche, and C. Pierdzioch (2017). Predicting recessions with boosted regression trees. International Journal of Forecasting 33(4), 745–759.

Eickmeier, S. and B. Hofmann (2022). What drives inflation? disentangling demand and supply factors. Deutsche Bundesbank Discussion Paper No 46/2022.

Firat, M. and O. Hao (2023). Demand vs. supply decomposition of inflation: Cross-country evidence with applications. IMF Working Paper No 23/205.

Fisher, J. D., D. W. Puelz, and C. M. Carvalho (2020). Monotonic effects of characteristics on returns. Ann. Appl. Stat. 14(4): 1622-1650 (December 2020).

Forbes, K. J. (2019). Has globalization changed the inflation process? BIS working paper No 791.

Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. Annals of Statistics, 1189–1232.

Gali, J. and T. Monacelli (2005). Monetary policy and exchange rate volatility in a small open economy. The Review of Economic Studies 72(3), 707–734.

Giannone, D. and G. Primiceri (2024). The drivers of post-pandemic inflation. National Bureau of Economic Research WP 32859.

Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajšek (2017). Inflation dynamics during the financial crisis. American Economic Review 107(3), 785–823.

Goulet Coulombe, P. (2024). A neural Phillips curve and a deep output gap. Journal of Business and Economic Statistics.

Grinsztajn, L., E. Oyallon, and G. Varoquaux (2022). Why do tree-based models still outperform deep learning on typical tabular data? Advances in Neural Information Processing Systems 35, 507–520.

Harding, M., J. Lindé, and M. Trabandt (2022). Resolving the missing deflation puzzle. Journal of Monetary Economics 126, 15–34.

Hastie, T. J. and R. J. Tibshirani (1990). Generalized Additive Models. Chapman & Hall/CRC.

Hothorn, T., P. Bühlmann, T. Kneib, M. Schmid, and B. Hofner (2010). Model-based Boosting 2.0. Journal of Machine Learning Research 11(71), 2109–2113.

Joseph, A., G. Potjagailo, E. Kalamara, C. Chakraborty, and G. Kapetanios (2021). Forecasting UK inflation bottom up. Bank of England Staff Working Paper No 915.

Kabaca, S. and K. Tuzcuoglu (2023). Supply drivers of us inflation since the pandemic. Bank of Canada Staff Working Paper 2023-19.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from OPEC announcements. American Economic Review 111(4), 1092–1125.

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

Lenza, M., I. Moutachaker, and J. Paredes (2023). Density forecasts of inflation: A quantile regression forest approach. ECB Working Paper No 2830.

Li, J. and I. Tsiakas (2017). Equity premium prediction: The role of economic and

statistical constraints. Journal of Financial Markets 36, 56–75.
Lorenzoni, G. and I. Werning (2023). Wage price spirals. Forthcoming, Brookings Papers in Economic Activity.
Lou, Y., R. Caruana, and J. Gehrke (2012). Intelligible models for classification and regression. In Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 150–158.
Lundberg, S. M., G. G. Erion, and S.-I. Lee (2018). Consistent individualized feature attribution for tree ensembles. arXiv preprint arXiv:1802.03888.
Lundberg, S. M. and S.-I. Lee (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems 30.
Mayer, M., S. C. Bourassa, M. Hoesli, and D. Scognamiglio (2021). Structured Additive Regression and Tree Boosting. Swiss Finance Institute Research Paper (21-83).
McCracken, M. W. and S. Ng (2016). Fred-md: A monthly database for macroeconomic research. Journal of Business & Economic Statistics 34(4), 574–589.
McElfresh, D., S. Khandagale, J. Valverde, V. Prasad C, G. Ramakrishnan, M. Goldblum, and C. White (2024). When do neural nets outperform boosted trees on tabular data? Advances in Neural Information Processing Systems 36.
Medeiros, M. C., G. F. Vasconcelos, Á. Veiga, and E. Zilberman (2021). Forecasting inflation in a data-rich environment: the benefits of machine learning methods. Journal of Business & Economic Statistics 39(1), 98–119.
Muggeo, V. M. (2003). Estimating regression models with unknown break-points. Statistics in Medicine 22(19), 3055–3071.
Nakamura, E. (2005). Inflation forecasting using a neural network. Economics Letters 86(3), 373–378.
Nanfack, G., P. Temple, and B. Frénay (2022). Constraint enforcement on decision trees: A survey. ACM Computing Surveys (CSUR) 54(10s), 1–36.
Ng, S. (2014). Boosting recessions. Canadian Journal of Economics/Revue canadienne d'économique 47(1), 1–34.
Pfäuti, O. (2023). The inflation attention threshold and inflation surges. arXiv preprint arXiv:2308.09480.
Potjagailo, G. and M. H. Wolters (2023). Global financial cycles since 1880. Journal of International Money and Finance 131, 102801.
Richman, R. and M. V. Wüthrich (2024). Smoothness and monotonicity constraints for neural networks using ICEnet. Annals of Actuarial Science, 1–28.
Seabold, S. and J. Perktold (2010). Statsmodels: Econometric and statistical modeling with Python. In 9th Python in Science Conference.
Shapiro, A. H. et al. (2022). Decomposing supply and demand driven inflation. FRB of San Francisco Working Paper 2022-18.
Stock, J. H. and M. W. Watson (2007). Why has us inflation become harder to forecast? Journal of Money, Credit and Banking 39, 3–33.
Štrumbelj, E. and I. Kononenko (2014). Explaining prediction models and individual predictions with feature contributions. Knowledge and Information Systems 41, 647–665.
Wen, D., M. He, L. Liu, and Y. Zhang (2022). Forecasting crude oil prices: do technical indicators need economic constraints? Quantitative Finance 22(8), 1545–1559.
Yoon, J. (2021). Forecasting of real gdp growth using machine learning models: Gradient boosting and random forest approach. Computational Economics 57(1), 247–265.

A Methodology

A.1 Implementation details for benchmarks

We use the *UnobservedComponents* class from the Python library `statsmodels` (Seabold and Perktold, 2010) to calibrate the unobserved component (UC) model. We set the parameters *stochastic_level* and *stochastic_trend* to True as this improves the out-of-sample performance.

We also tested the standard Unobserved Components Stochastic Volatility (UCSV) model (Stock and Watson, 2007) using the implementation by Medeiros et al. (2021). However, this model performed worse than the UC model described above and is computationally substantially more expensive.

Hyperparameters of the random forest and Lasso are selected using 10-fold block cross-validation within the training set and updated every five years to reduce computation time. The AR(2) benchmark and the UC model are re-estimated monthly, all other models every three months, thus leveraging the lower computational burden of the simpler benchmarks to their advantage.

We consider the the following hyperparameter values for `scikit-learn`'s random forest implementation:
$max\_depth \in [2, 3, 4, 7, 10]$,
$max\_features \in [1, 3, 5, 7, 10, 15, 20, 30]$,
$min\_samples\_leaf \in [1, 2, 4, 10, 15]$.

For Lasso regression, we use the default regularisation path of the `scikit-learn`'s *LassoCV* implementation with 100 values for α.

B Data Description

Variable	Mean	SD	First Obs.	Latest Obs.	Transf.	Blocks$^{(\text{constraint})}$	Source
Time index			1988M02	2024M12	1	Trend	
Wage (Average Weekly Earnings, AWE)	4.1	10.2	1988M02	2024M12	5	Trend	ONS
Wage (AWE), regular pay (excl. bonuses)	3.3	3.7	2000M02	2024M12	5	Trend	ONS
Wage (AWE), services	3.4	3.9	2000M02	2024M12	5	Trend	ONS
Wages (AWE), financial services	3.7	7.1	2000M02	2024M12	5	Trend	ONS
Wage (AWE), manufacturing	3.0	5.5	2000M02	2024M12	5	Trend	ONS
Wage (AWE), retail and hospitality	3.3	8.0	2000M02	2024M12	5	Trend	ONS
Service CPI	4.0	3.2	1988M02	2024M12	5	Trend	ONS
Service CPI, catering	4.0	5.5	1988M02	2024M12	5	Trend	ONS
Service CPI, accommodation	3.7	13.1	1996M02	2024M12	5	Trend	ONS
Service CPI, recreational and cultural	3.9	5.8	1988M02	2024M12	5	Trend	ONS
Service CPI, personal care	2.4	5.4	1988M02	2024M12	5	Trend	ONS
5y5y inflation expectations (swaps)	3.6	0.9	1988M02	2024M12	1	Trend	Bloomberg, Tradeweb
1-year-ahead household infl. expectations	3.2	1.4	1988M02	2024M12	1	Trend	Citi Group
Unemployment rate gap	0.3	0.9	1988M02	2024M12	1	Demand^{-}, Supply^{+}	ONS
vacancy/unemployment (v/u) ratio gap	0.0	0.1	1988M03	2024M12	4	Demand^{+}, Supply^{-}	ONS
Retail sales	1.8	19.3	1988M02	2024M12	5	Demand^{+}, Supply^{-}	ONS
Index of production	0.5	20.2	1997M02	2024M12	5	Demand^{+}, Supply^{-}	ONS
Index of services	2.1	17.6	1997M02	2024M12	5	Demand^{+}, Supply^{-}	ONS
UK imports, real	4.7	53.9	1988M02	2024M12	5	Demand^{+}, Global Supply^{-}	ONS
PMI manufacturing	51.7	4.3	1991M07	2024M12	1	Demand^{+}, Supply^{-}	S&P Global UK
PMI services	54.3	4.9	1996M07	2024M12	1	Demand^{+}, Supply^{-}	S&P Global UK
PMI construction	53.7	6.7	1997M04	2024M12	1	Demand^{+}, Supply^{-}	S&P Global UK
Investment*	1.5	13.5	1988Q1	2024Q4	5	Demand^{+}, Supply^{-}	ONS
Consumer sentiment	-11.9	12.1	1988M02	2024M12	1	Demand^{+}	GfK
Global economic activity shock	0.0	0.7	1988M02	2024M12	1	Global Demand^{+}	BH2019
Global oil demand shock	0.1	3.7	1988M02	2024M12	1	Global Demand^{+}	BH2019
Global oil inventory demand shock	-0.1	1.1	1988M02	2024M12	1	Global Demand^{+}	BH2019
US imports, real	3.0	27.6	2000M02	2024M12	5	Global Demand^{+}	Refinitiv, DS
Euro area imports, real	0.5	24.9	2000M02	2024M12	5	Global Demand^{+}	Refinitiv, DS
US industrial production	1.4	12.1	1988M02	2024M12	5	Global Demand^{+}, Global Supply^{-}	Refinitiv, DS
Euro area industrial production	0.6	21.1	1991M02	2024M12	5	Global Demand^{+}, Global Supply^{-}	Refinitiv, DS
UK exports, real	3.4	63.2	1988M02	2024M12	5	Global Demand^{+}, Supply^{-}	ONS
Goods CPI	1.9	4.7	1988M02	2024M12	5	Supply^{+}	ONS
Food CPI	2.8	6.7	1988M02	2024M12	5	Supply^{+}	ONS
Gas (CPI)	4.7	43.2	1988M02	2024M12	5	Supply^{+}	ONS
Gas spot price	6.7	209.8	1996M05	2024M12	5	Supply^{+}	Bloomberg
Electricity (CPI)	4.7	26.4	1988M02	2024M12	5	Supply^{+}	ONS
Electricity (PPI)	5.2	33.7	1996M02	2024M12	5	Supply^{+}	ONS
PPI input price	2.8	8.7	1988M02	2024M12	5	Supply^{+}	ONS
PPI output price	2.3	6.6	1988M02	2024M12	5	Supply^{+}	ONS
UK-specific supply chain index (SCI)	-0.3	1.4	1998M01	2024M12	1	Supply^{+}	BoE
Total factor productivity (TFP)	1.8	11.9	1988M02	2023M04	5	Supply^{-}	
Commodity price index, energy	3.8	94.1	1988M02	2024M12	5	Global Supply^{+}	World Bank
Commodity price index, non-energy	2.2	31.0	1988M02	2024M12	5	Global Supply^{+}	World Bank
Commodity price index, food	2.1	35.2	1988M02	2024M12	5	Global Supply^{+}	World Bank
Commodity price index, metal	2.6	59.8	1988M02	2024M12	5	Global Supply^{+}	S&P Global UK
Commodity price index, agriculture	1.9	40.7	1988M02	2024M12	5	Global Supply^{+}	S&P Global UK
US PPI	2.4	8.0	1988M02	2024M12	5	Global Supply^{+}	Refinitiv, DS
Euro area PPI	1.8	6.3	1995M02	2024M12	5	Global Supply^{+}	Refinitiv, DS
Global supply chain pressure index (GSCPI)	0.0	1.0	1998M01	2024M12	1	Global Supply^{+}	Fed. Reserve Bank, NY
Global supply chain index (GSCI)	0.0	1.3	2007M05	2024M12	1	Global Supply^{+}	BoE
Global oil supply news shock	0.0	0.6	1988M02	2024M12	1	Global Supply^{+}	Känzig (2021)
Global oil supply shock	-0.1	1.2	1988M02	2024M12	1	Global Supply^{+}	BH2019
Extended model specification							
Real exchange rate index	-0.2	18.2	1990M02	2024M12	5	Financial	FAME
GBP-USD spot exrate	0.0	0.0	1988M02	2024M12	2	Financial	FAME
FTSE UK focused	2.5	47.2	1995M02	2024M12	5	Financial	FAME
House price index	4.8	10.8	1991M02	2024M12	5	Financial	Nationwide
Corporate bond spread	156.0	75.5	1998M01	2023M10	1	Financial	FAME
MP shock, target	0.0	0.1	1997M06	2024M07	1	Monetary Policy^{-}	Braun et al. (2023)
MP shock, path	0.0	0.0	1997M06	2024M07	1	Monetary Policy^{-}	Braun et al. (2023)
MP QE shock	0.0	0.0	1997M06	2024M07	1	Monetary Policy	Braun et al. (2023)

Table B1: Data series - descriptive stats, sample coverage, transformations, model blocks and constraints, data sources.

Notes: Series transformations to achieve stationarity follow McCracken and Ng (2016) codes: 1 – no transformation, 2 – first difference, 4 – logs, 5 – first log difference. In the alternative specification using year-on-year CPI inflation, we take year-on-year growth rates for series with code 5. * - Investment series is quarterly, interpolated to monthly. CPI components such as service, food, energy CPIs are seasonally adjusted using X12-ARIMA, other series come directly in seasonally adjusted form. Mean and SD are calculated after transformation. When series appear in multiple blocks, these are listed, separated by commas. Monotonicity constraints indicated as $^{+}$ (positive) or $^{-}$ (negative). Source acronyms: ONS - UK Office for National Statistics, BoE - Bank of England, BH2019 - Baumeister and Hamilton (2019), Bloomberg - Bloomberg Finance L.P., Refinitiv, DS - LSEG Workspace, Refinitiv, DataStream.

C Additional results

C.1 Results for baseline specification

C.1.1 Functional forms for all lags and more indicators

Demand

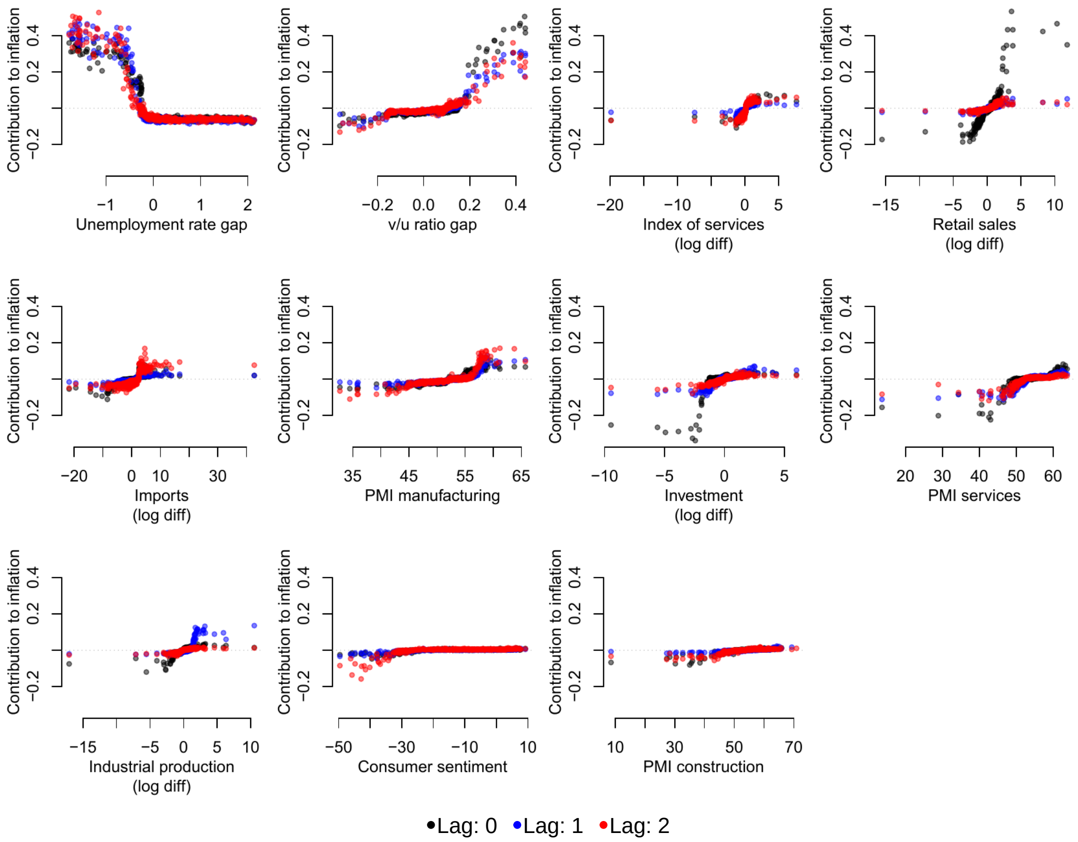

Global demand

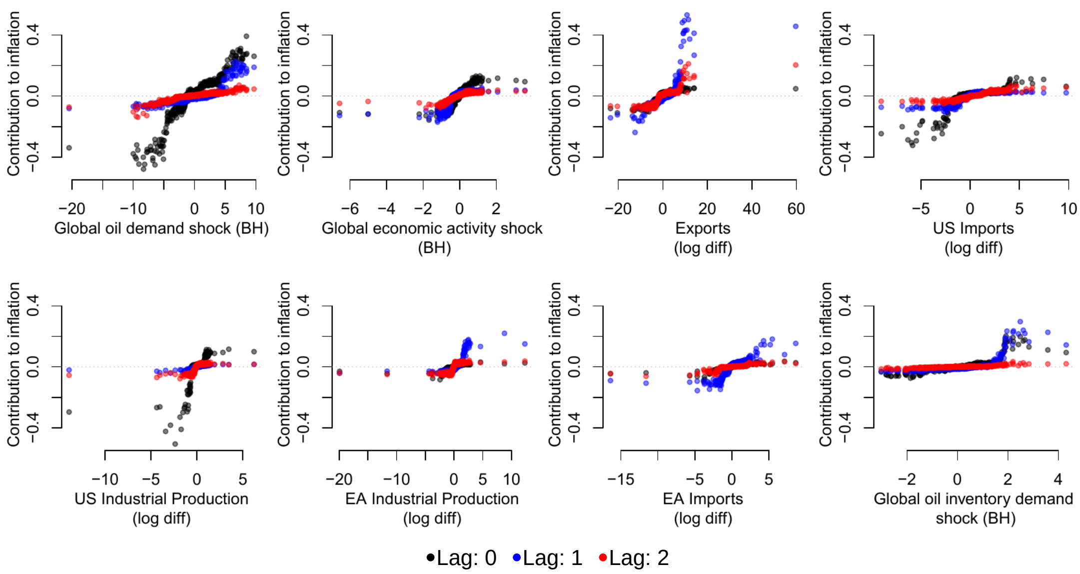

Figure C1 (page 1 of 4): Functional forms of all indicators in the demand blocks of the baseline model. Colours show the different lags of the indicators.

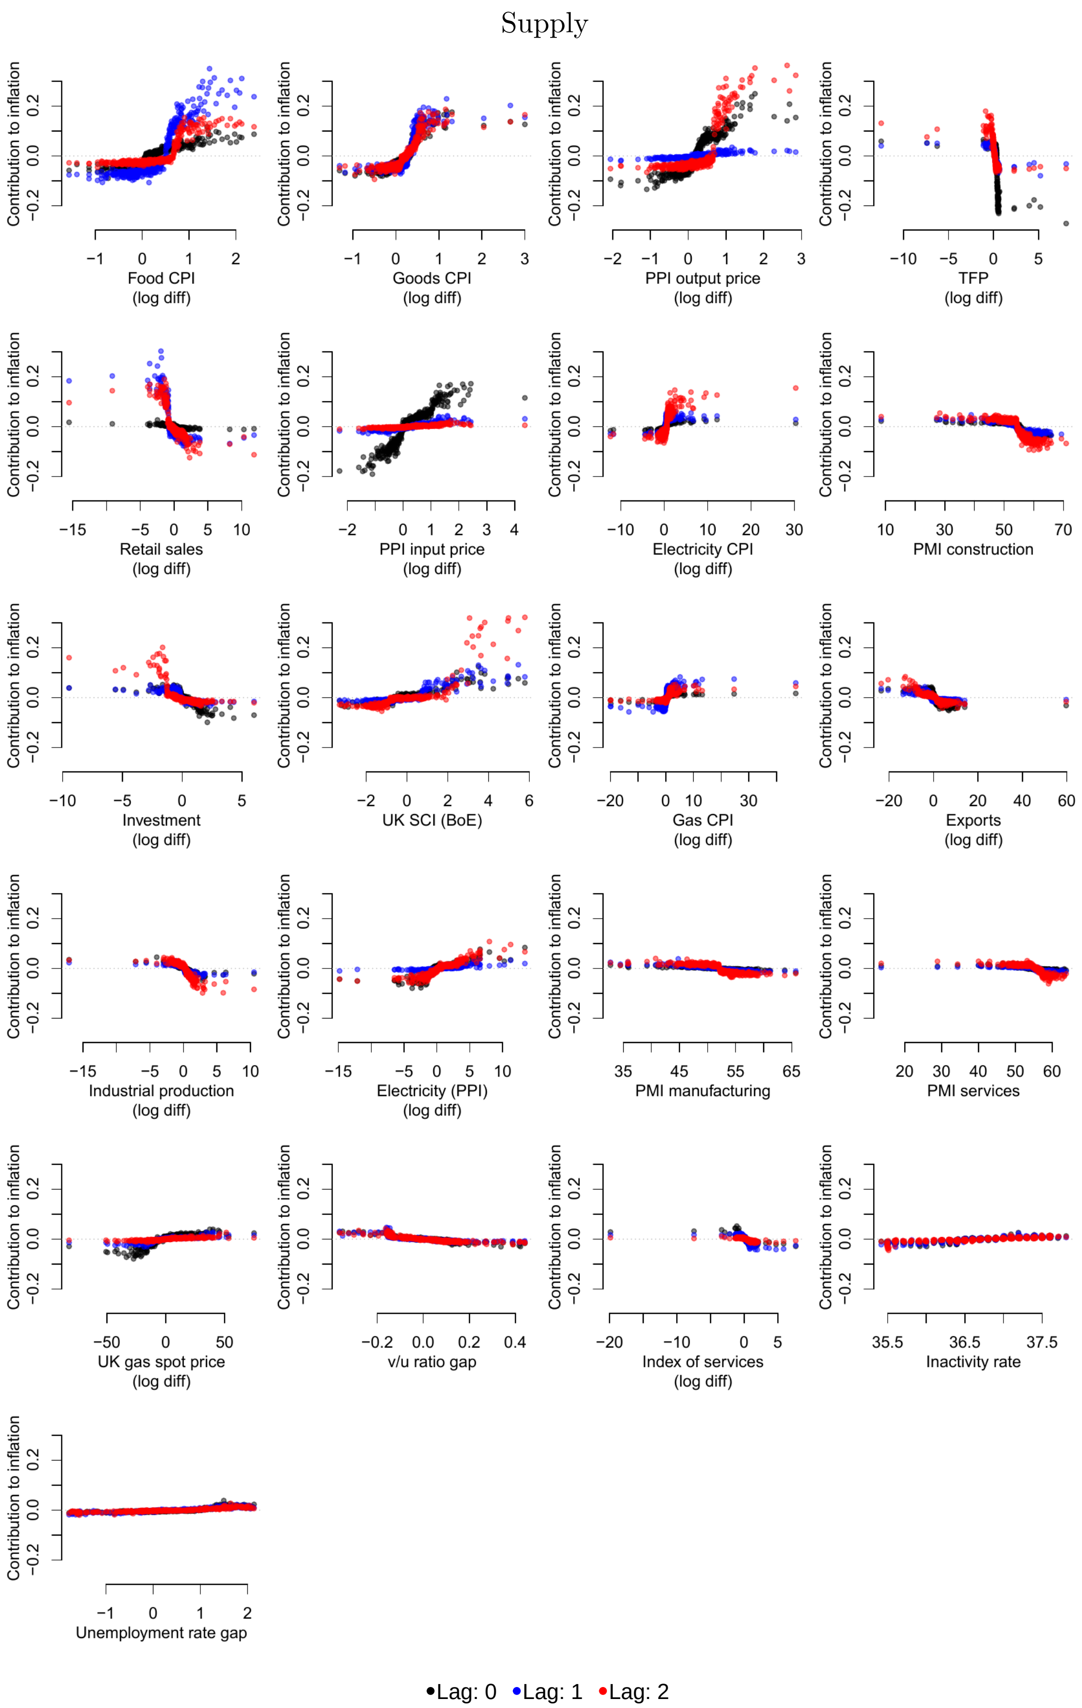

Figure C1 (page 2 of 4): Functional forms of all indicators in the domestic supply block of the baseline model. Colours show the different lags of the indicators.

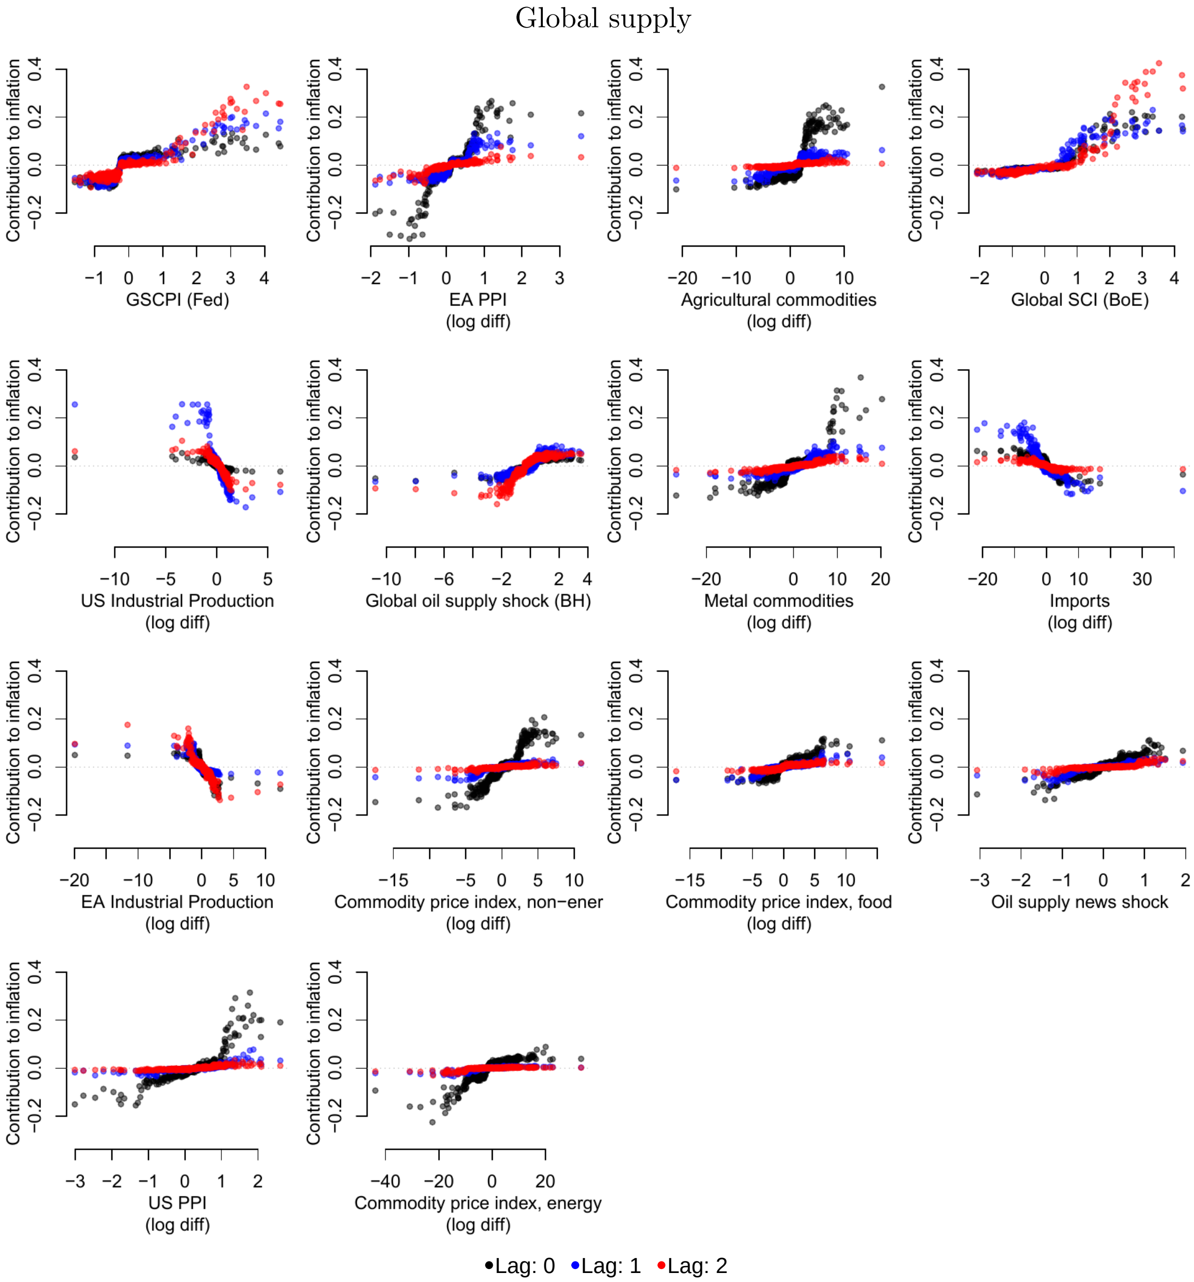

Figure C1 (page 3 of 4): Functional forms of all indicators in the global supply block of the baseline model. Colours show the different lags of the indicators.

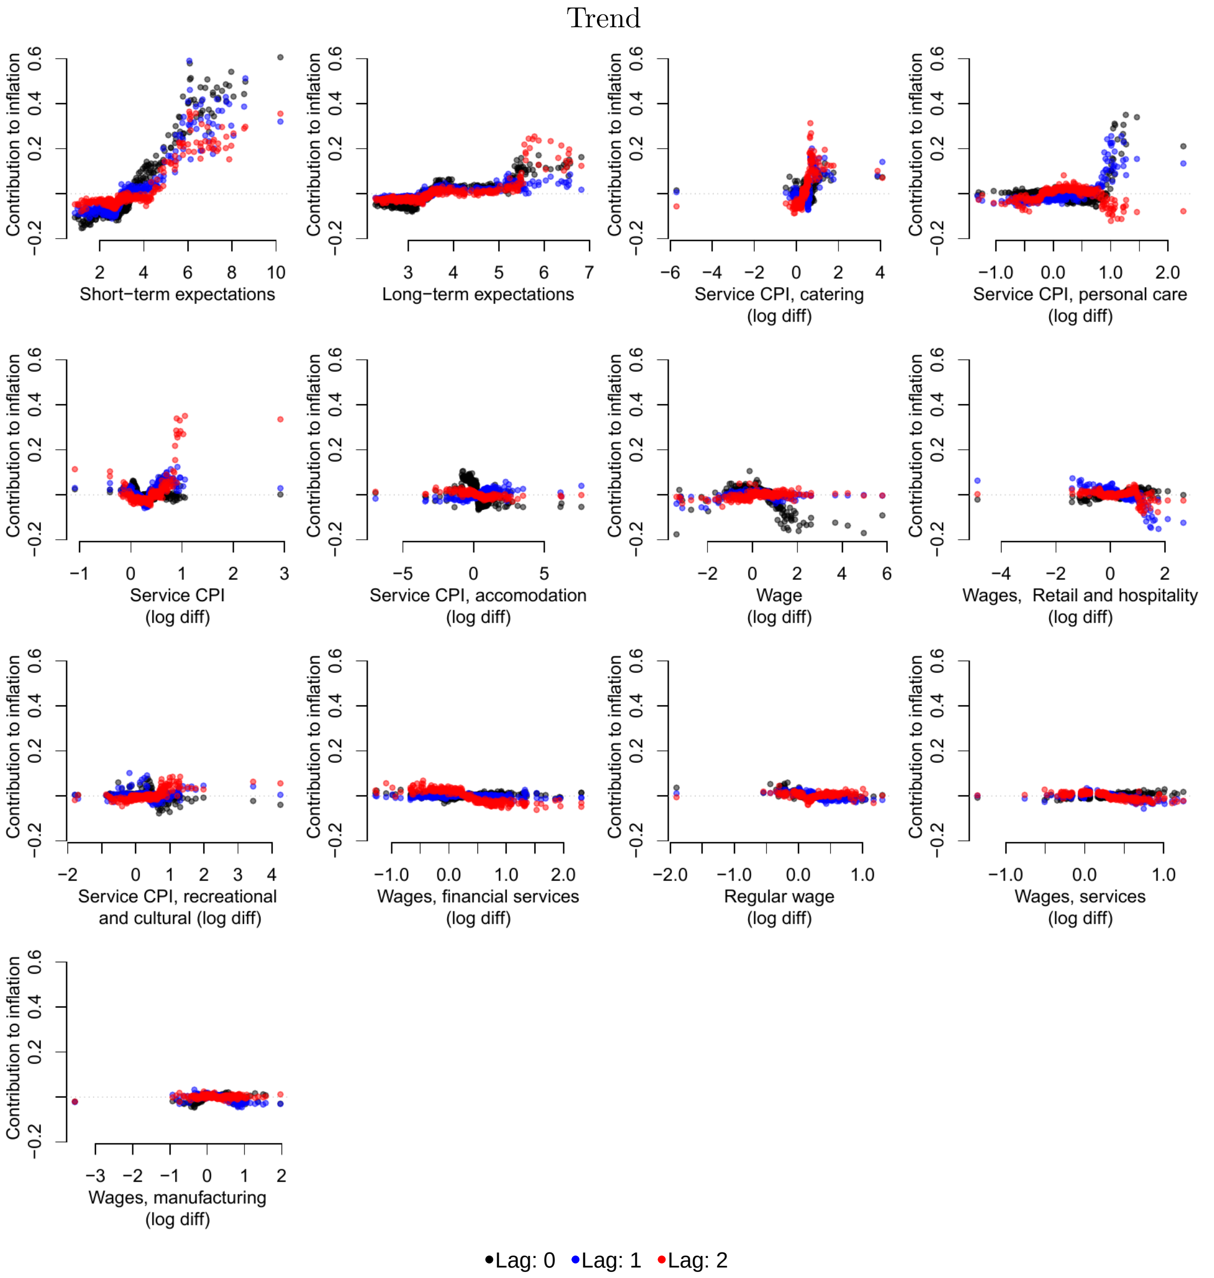

Figure C1 (page 4 of 4): Functional forms of all indicators in the trend block of the baseline model. Colours show the different lags of the indicators.

C.2 Results for alternative specifications to baseline

C.2.1 Functional forms across alternative specifications

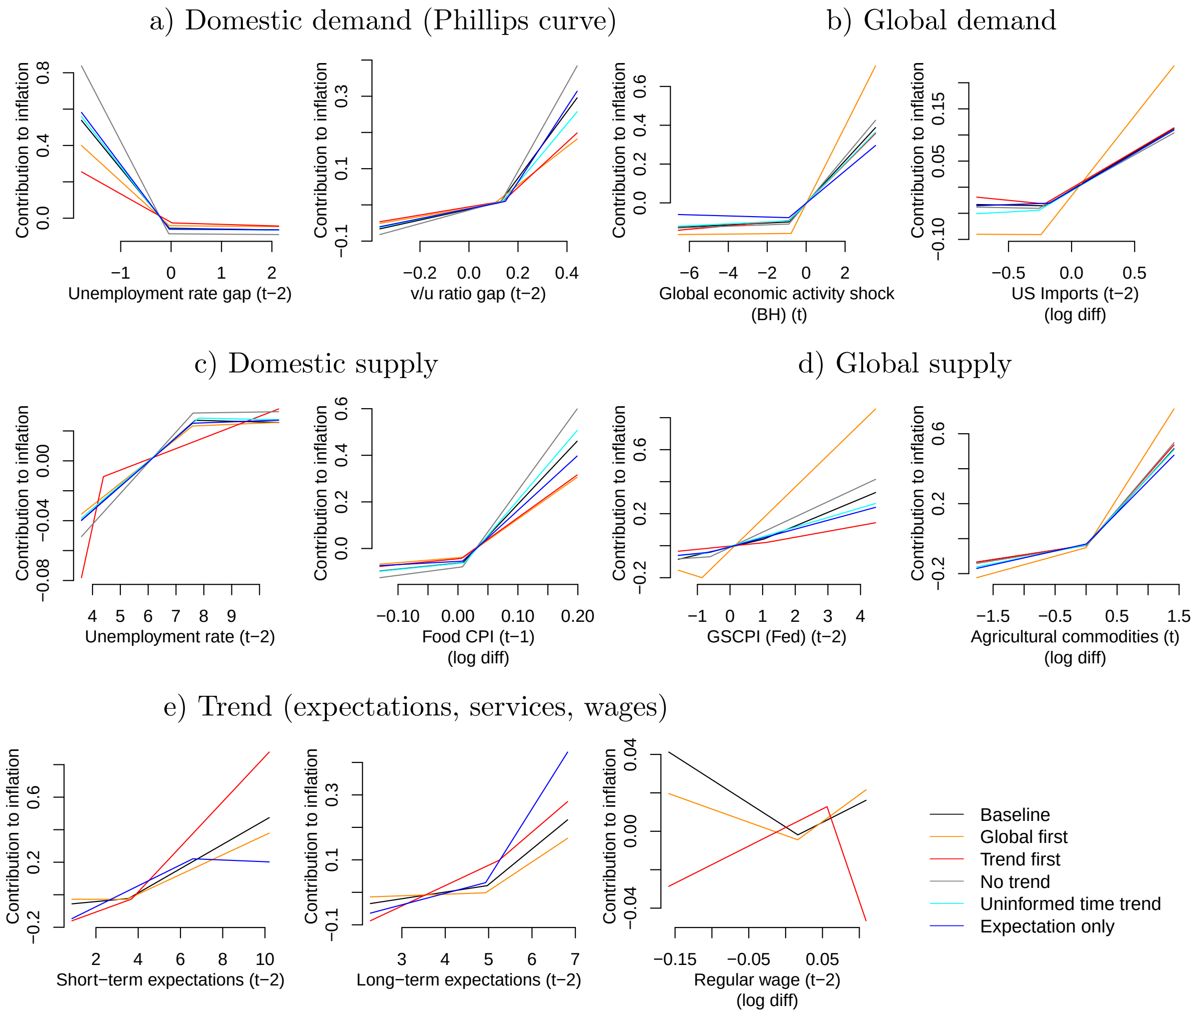

Figure C2: Selected functional forms learnt within each block, across model specifications.
Note: Baseline model, see Figure 2. Scatter plots of the realised values of a given indicator at a selected lag (horizontal axis), against its contribution to the model's CPI inflation prediction measured by Shapley values (vertical axis), within a given block (panel a) to c)). We show a fitted linear model with a single breakpoint (i.e. two different slopes) to the relationship between input and Shapely values of the indicators (Muggeo, 2003). Observations with imputed values for the respective variable are excluded from estimation.

C.2.2 Role of identification via constraints and identified shock series

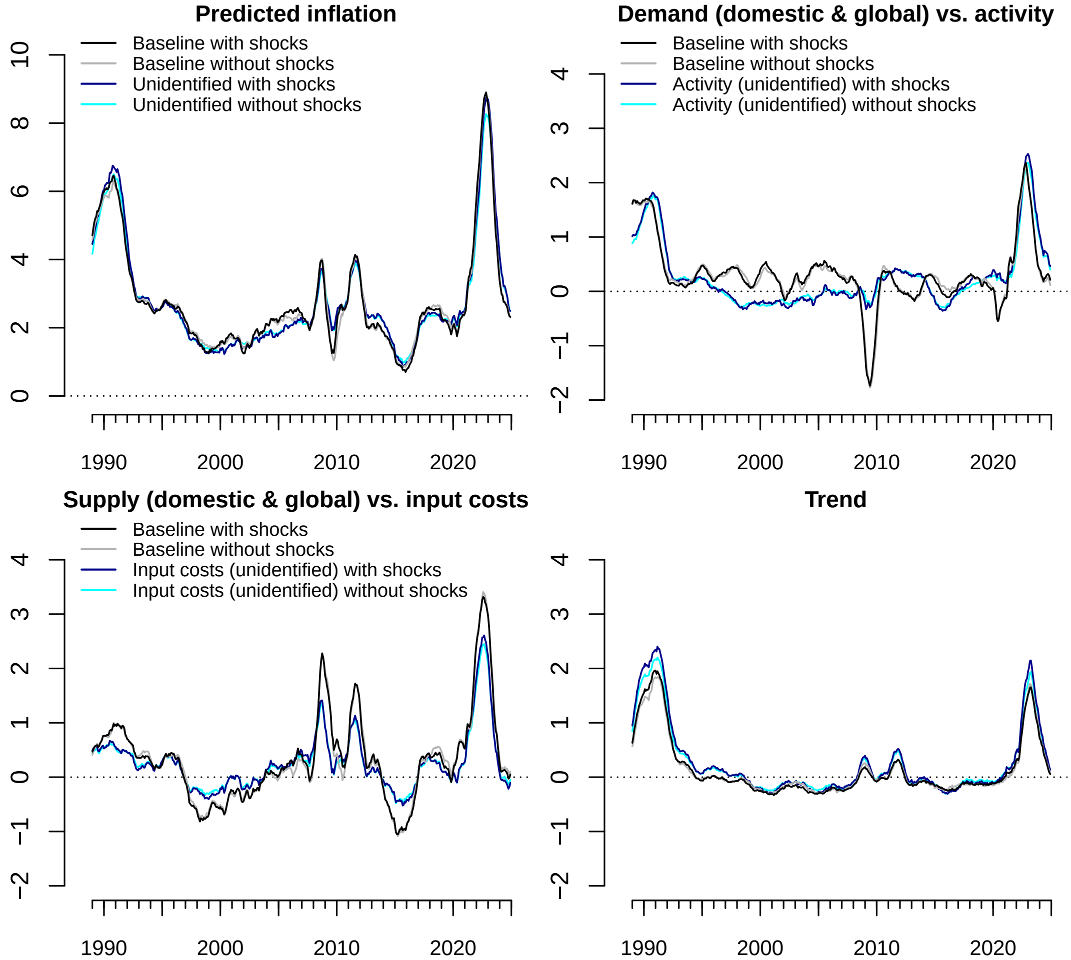

Figure C3: Predicted inflation and blockwise predictive contributions for alternative models with and without monotonicity constraints and identified shock series.

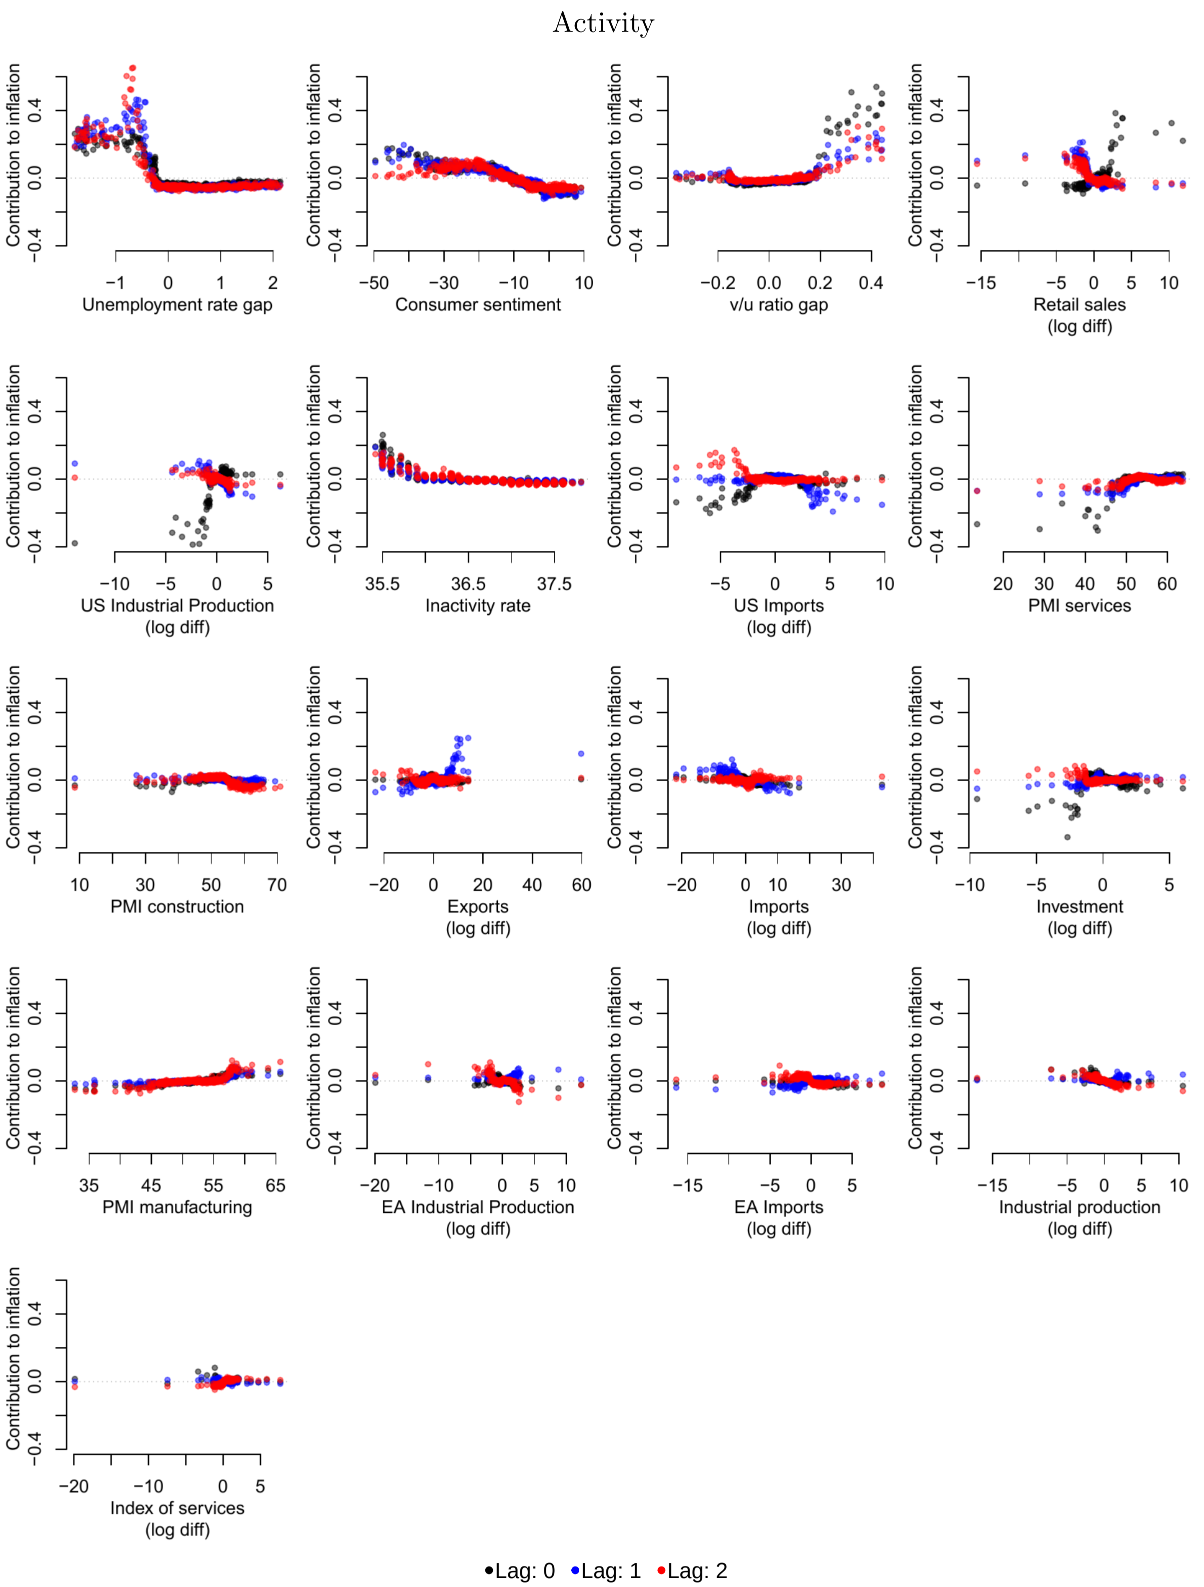

Figure C4: Functional forms of all indicators in the activity block of the model without monotonicity constraints (see Figure 7). Colours show the different lags of the indicators.

C.2.3 Changing the number of lags of indicators included

In the baseline specification, we include indicators with three lags: at times $t, t-1, t-2$. Figure C5 shows the predictive contribution of the different blocks when altering the number of lags.

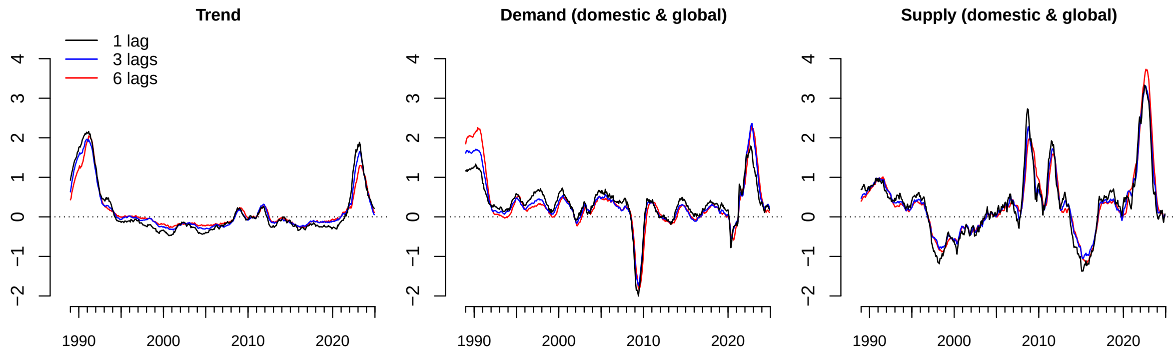

Figure C5: Blockwise contributions to inflation prediction, different numbers of lags of indicators included. Model with three lags corresponds to baseline specification.

C.2.4 Model for year-on-year CPI inflation

In the specification shown in Figure C6, the target variable is 1-month ahead year-on-year CPI inflation. All indicators that are transformed to monthly growth rates in the baseline model (see Table B1, transform code 5) are instead transformed into year-on-year growth rates.

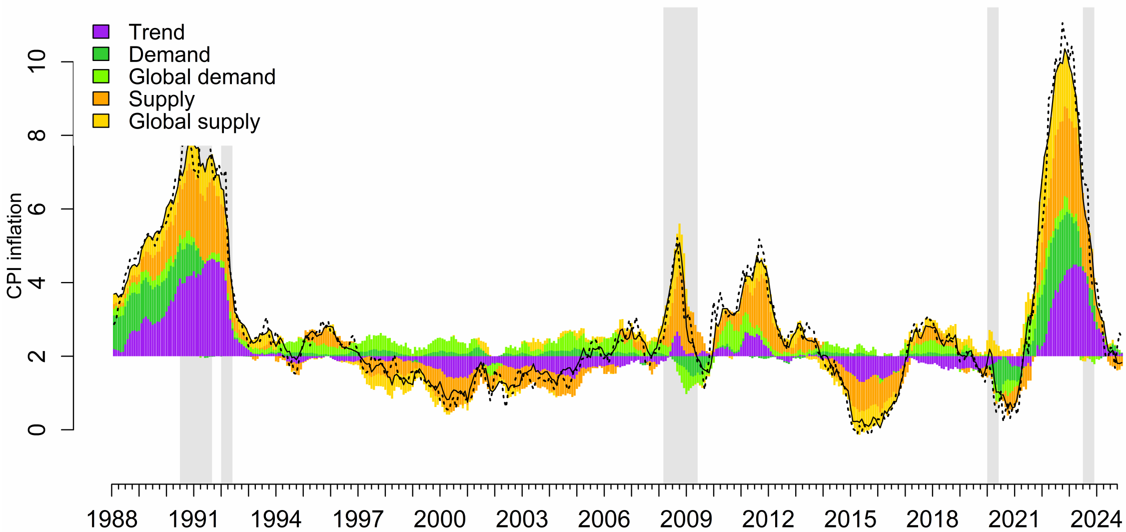

Figure C6: Decomposition of 1-month-ahead CPI inflation from baseline blockwise Boosted Inflation Model.

C.3 Extended model separating role of labour market tightness

In this specification, the "labour market tightness" block comprises the unemployment rate gap (−) and the vacancy-unemployment ratio gap (+) and their respective lags. The "other demand" block comprises all other activity indicators (+) as in the baseline specification.

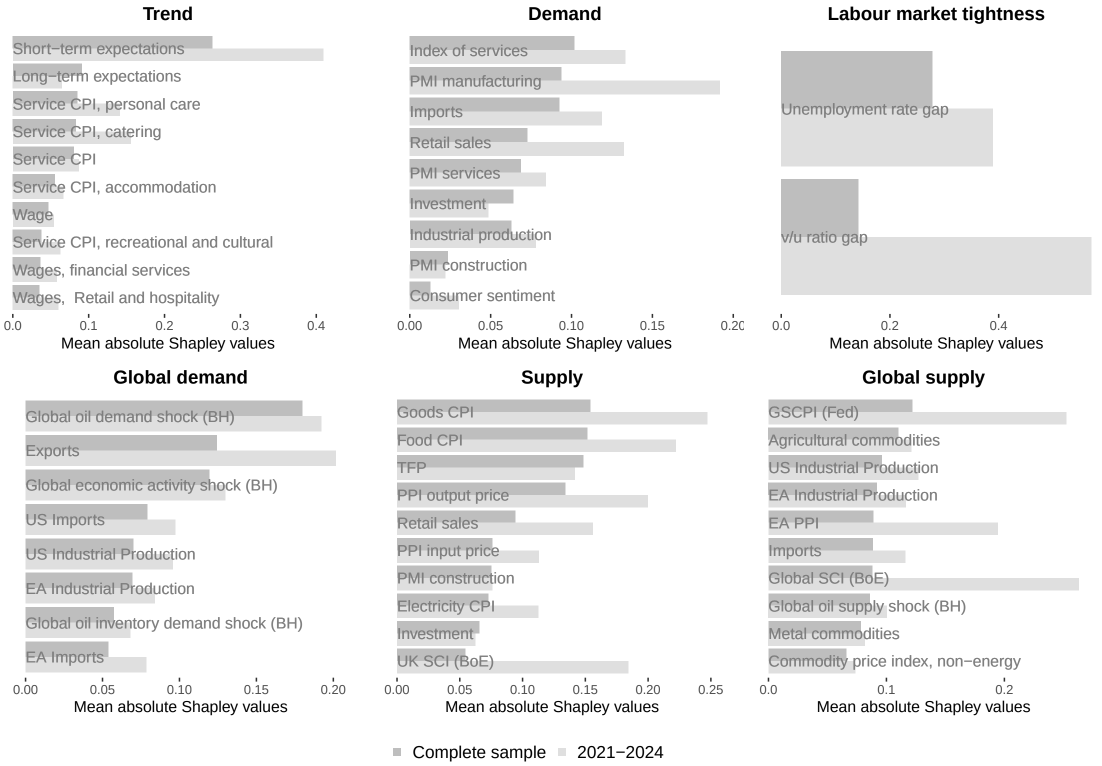

Figure C7: Mean absolute Shapley values, with labour market tightness block.
Notes: Ten largest average contributions in each block, averaged over 1989-2024, ordered by size of average contribution (global demand block has 8 indicators in total).

C.4 Extended model with financial conditions and monetary policy shocks

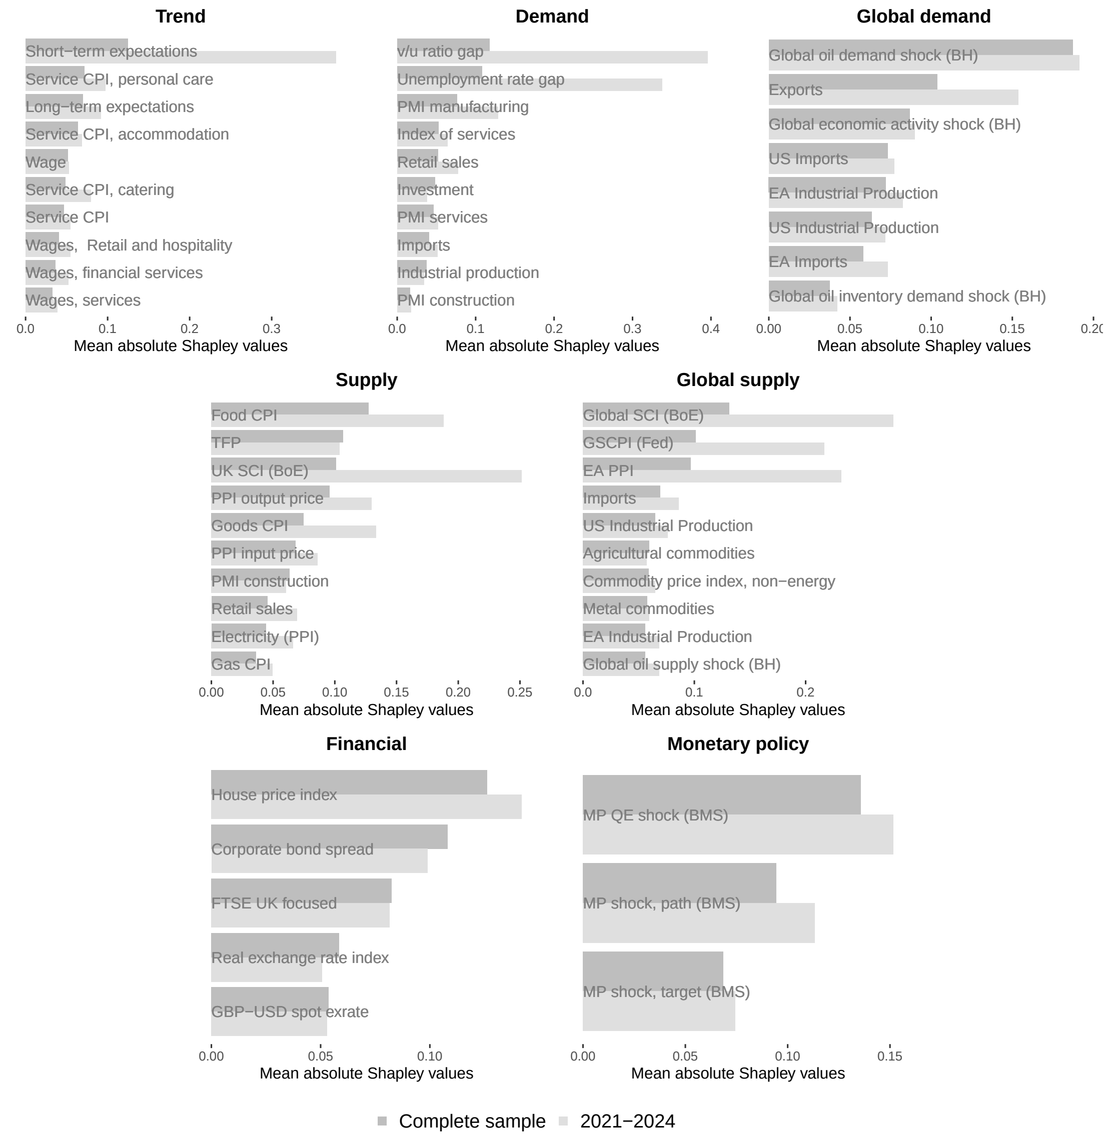

Figure C8: Extended model with financial conditions and monetary policy blocks, sample period 1997–2024. Mean absolute Shapley values within each component. Ten largest average contributions in each component.

C.5 Robustness of the BBIM to the choice of hyperparameters

We test how the choice of key hyperparameters affects our main result. We alter the depth of the individual trees (1, 2, 4, baseline: 3), the minimum number of observations in a node to consider another split (2, 20, baseline: 5), and whether trees are fit on a random subsample of the training observations (baseline: 50%) and predictors (baseline: 25%) or not.

Figure C9 shows the results. While the predicted inflation as well as the component values are highly consistent across the different hyperparameter settings, we observe some notable differences. First, when limiting the trees to a depth of 1, the trend contribution is substantially higher. Note that a boosting model with a tree depth if 1, while still allowing for non-linearities, does not allow for interactions between variables. Second, when limiting splits to nodes with more than 20 observations, we observe some changes in the demand component.

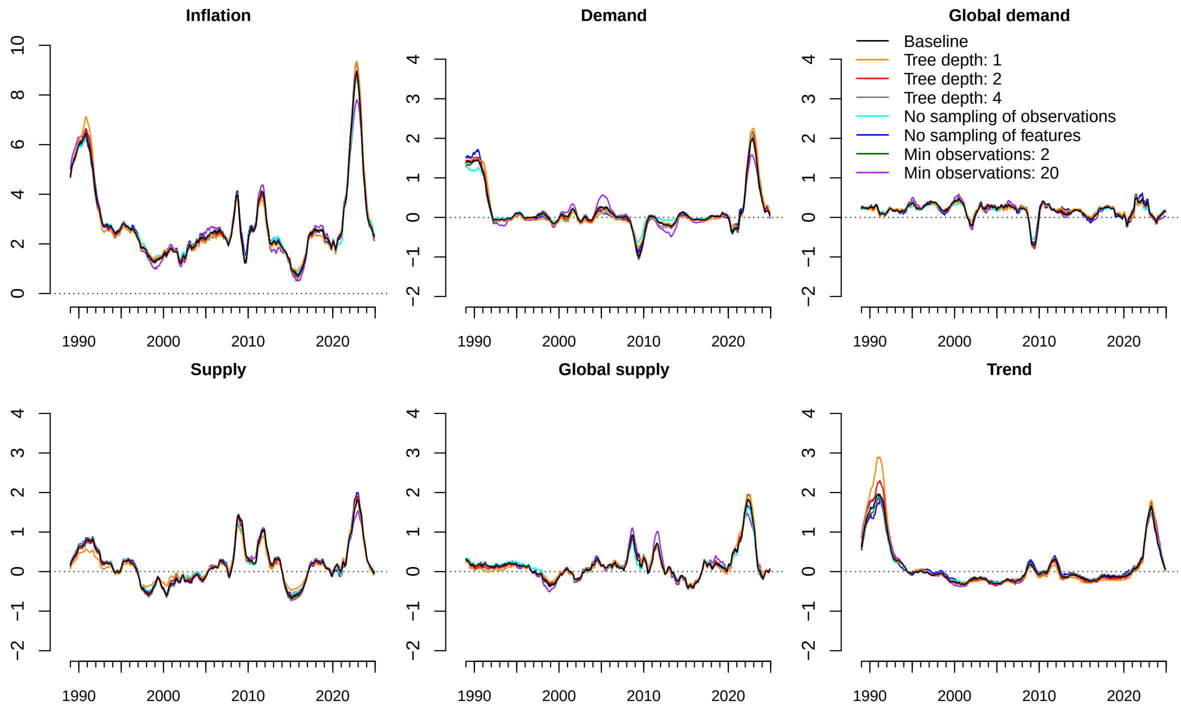

Figure C9: Predicted inflation and blockwise predictive contributions for models with alternative hyperparameter settings.

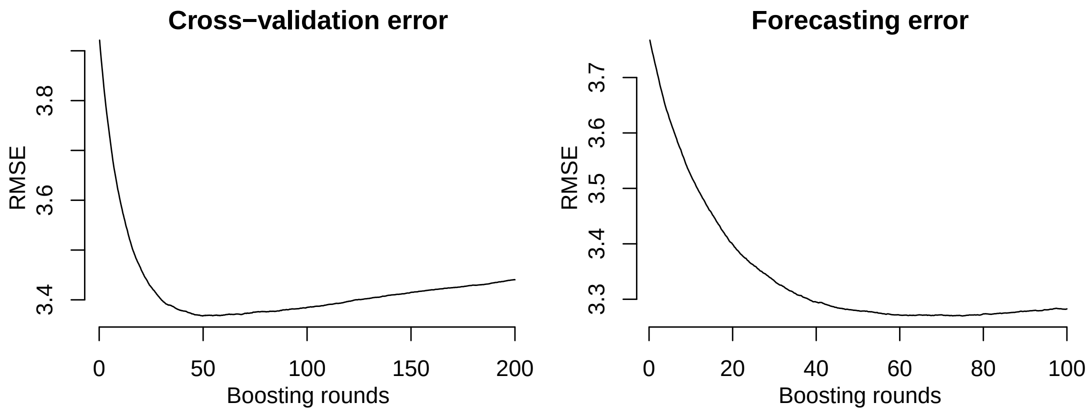

Figure C10: Cross-validation and forecasting error as a function of the number of boosting rounds. We apply early stopping in cross-validation but not in forecasting, where we limit the number of rounds to 100.