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

Marcus Buckmann † Galina Potjagailo ‡ Philip Schnattinger §

Preliminary and incomplete

February 28, 2025

Abstract

We propose the Blockwise Boosted Inflation Model: an interpretable, non-linear, boosted tree method that decomposes inflation dynamics into components akin to an open-economy hybrid Phillips curve. Demand and supply contributions are identified by imposing sign restrictions on the association between inflation and indicators. We model monthly CPI inflation in the United Kingdom. The recent rise in UK inflation is explained by a combination of supply determinants, demand, and changes in the role of lagged inflation and expectations. Supply played a larger role than in the past, and monetary policy and financial determinants a smaller role. Non-linearities that the model learns can be traced over time and help explain the recent rise in inflation. Strong non-linear effects from global supply pressures and cost-related variables contributed to inflation during 2021–22, but these effects quickly reverted to the flat region. On the demand side, we detect a convex Phillips curve relationship for labour market tightness and unemployment. Our model also performs competitively in out-of-sample forecasting against an AR benchmark and other machine-learning models.

Keywords: Inflation, Phillips curve, Boosted decision trees, Machine learning

JEL Codes: E31, E37, C14, C53

*Any views expressed are solely those of the authors and so cannot be taken to represent those of the Bank of England or members of its committees.

†Bank of England, e-mail: marcus.buckmann@bankofengland.co.uk

‡Bank of England, e-mail: galina.potjagailo@bankofengland.co.uk

§Bank of England, e-mail: philip.schnattinger@bankofengland.co.uk

1 Introduction

Disentangling the sources of inflationary pressure is crucial for a monetary authority that aims to stabilise inflation. But it can be particularly challenging when non-linearities are at play, so that standard linear analytical toolkits for assessing inflation dynamics encounter limitations. The recent inflationary period painfully demonstrated this, as central banks initially underestimated the extent and persistence of the inflationary effects in the aftermath of the Covid-19 pandemic and Russia's invasion of Ukraine. Can machine learning help inform policy makers about potentially non-linear inflation dynamics? In economic forecasting, machine learning methods are becoming increasingly appealing for their ability to capture complex non-linearities across a broad range of variables (de Araujo et al., 2024). The key limitation, however, is that these non-parametric methods do not directly reflect economic relations and lack clear intuition. So their use for inflation analysis so far has largely focused on maximising forecasting accuracy rather than explaining the signals learned by the models (Nakamura, 2005; Medeiros et al., 2021; Joseph et al., 2021; Lenza et al., 2023).

In this paper, we propose the Blockwise Boosted Inflation Model (BBIM) that enhances a boosted tree method fed with a range of indicators with a model structure reminiscent of an open-economy hybrid Phillips curve, and with sign restrictions on the direction of the association between inflation and the indicators to separate demand and supply side determinants of inflation. Boosted trees have a strong record in economic forecasting (Ng, 2014; Döpke et al., 2017; Yoon, 2021). Our model extends the appeal of this method beyond forecasting, towards a flexible tool to disentangle inflation determinants and to track non-linearities in the Phillips curve.

We apply the model to understand monthly headline CPI inflation dynamics in the United Kingdom from January 1989 to March 2024. Being a small open economy that heavily relies on goods imports, the UK faced a constellation of particularly large external shocks following the Covid-19 pandemic—which makes the identification of a Phillips curve slope in low-dimension linear regressions challenging and non-linearities likely. We include around 60 indicators that enter into the model with several lags, and we estimate 1-month ahead CPI inflation using 10-fold cross-validation. We show that the recent rise of UK inflation was characterised by supply-side determinants to a larger extent compared to past dynamics, but also reflected non-linear demand effects and an upward shift in signals from short-term inflation expectations. We detect a kinked Phillips curve for measures of labour market slack that came into play, as well as non-linear effects from supply related to global supply pressures.

The blocks in our model cover groups of global and domestic economic activity indicators, cost-push and supply-side variables and lagged inflation dynamics and expectations. In additional specifications, we also control for financial conditions and monetary policy surprises. This block structure is akin to the recently proposed neural network model with a Phillips curve structure by Goulet Coulombe (2022). We capture non-linear associations between the indicators and inflation by training many small decision trees sequentially minimising prediction error. The trees are linearly conditioned on each other within the boosting algorithm and each decision tree is trained using the variables of a single block, which rules out interactions of variabels across blocks. It follows that we can linearly decompose the contributions of the different blocks to the predicted inflation signal.

The boosting model is initiated by fitting the sample mean. We then sequentially calibrate trees from different blocks in random order and update the prediction until the model starts to overfit, as indicated by an increase in the cross-validation error.

The association between economic activity variables and inflation can reflect both demand and supply side determinants. We achieve a degree of economic identification of model blocks by imposing sign restrictions on the tree splits, thus assuring that the signals for inflation that the model learns from an indicator in a given block are economically coherent. Specifically, the demand component is restricted to those splits of the tree that imply a positive association between inflation (the output variable) and real activity indicators (negative association with labour market slack indicators), whereas in the supply component, the reverse associations are imposed. To further inform the components, we also feed identified shock series from the literature directly as indicators into the respective components with a positive sign restriction, such as global supply shocks into supply and global demand shocks into demand (Baumeister and Hamilton, 2019; Känzig, 2021). This approach allows us to separate a meaningful demand-side contribution to inflation that reflects a Phillips-curve type relationship with labour market slack from supply determinants that reflect cost-related price measures and global supply constraints. We also feed identified monetary policy shocks from Braun et al. (2023) into a separate component to control for monetary surprises that might otherwise confound the identification of demand effects.

Two key findings emerge. First, the BBIM provides a meaningful decomposition of UK inflation into contributions from the model components around the sample mean. This is estimated using 10-fold repeated cross-validation. During the recent inflation episode between 2021 and early 2024, supply-side determinants mainly reflecting global supply chain pressures played a relatively stronger role than in the past. Demand effects are meaningful over the sample period, for instance pulling inflation below its mean during the 2008 global financial crisis (GFC); in the recent episode they explain a substantial share of the inflation rise, mainly via the role of labour market tightness and slack. An "inertia" component, that captures the propagation of past inflation changes and expectations and whose contribution was previously rather small and stable, moved up quickly representing shifts in the role of short-term expectations.

Second, non-linearities in inflation determinants mattered in the recent episode. We use the Shapley value framework (Lundberg and Lee, 2017) to reveal the strength and the shape of the (non-linear) associations between inflation and indicators within each component that the model learns. Within the demand block, we detect a strongly non-linear negative Phillips curve slope type relationship between inflation and the unemployment rate, as well as a positive relationship with the ratio of number of vacancies to unemployed workers, i.e. labour market tightness—in line with recent findings from Benigno and Eggertsson (2023) of an L-shaped Phillips curve explaining recent inflation dynamics in the United States. Within the supply block, non-linear effects from a global supply chain pressures index, as well as from food and goods price inflation components, pushed up UK inflation recently, which can reflect non-linear effects from cost-push shocks due to quasi-kinked demand and asymmetric cost-price pass-through (Harding et al., 2023).

These effects reverted back toward the flat region over the course of 2023. Finally, we observe non-linear effects from short-term inflation expectations which remain within the non-linear region up until the end of the sample period, and therefore the inertia component contributes to above-mean inflation even after the global supply shocks have unwound. This can reflect adaptive expectations that take longer to unwind after a strong rise of inflation and that can lead to persistent inflation via their effects on wage and price setting. By contrast, the effects of long-term expectations remain firmly in the flat region over the recent episode, pointing to anchored expectations—quite differently from non-linear effects that the model detects over the inflation episode during the early 1990s before the introduction of inflation targeting.

An unexplained component during the recent rise and extended peak in UK inflation remains, also since the BBIM has only few high inflation observations over the sample period to learn from.

Toward the end of the sample period, the model captures well that most but not all of the non-linear effects unwind relatively quickly. Further, the imposed restrictions on the direction of association between indicators and inflation help achieve a meaningful decomposition: A model with a component fed with unrestricted activity variables struggles to identify a Phillips curve type slope, under-estimates the negative demand effects during the GFC and over-estimates effects from demand relative to supply during the recent inflation episode.

The Phillips curve type components that we define, while enhancing intuition, naturally reflect the assumptions of the modeller. To account for model uncertainty, we estimate a range of specifications with varying Phillips curve component structures, such as omitting the inertia component or replacing the inertia component with a slow-moving time tend; varying the the order in which components enter the boosting model and varying the number of lags of the predictors. Results remain very similar regarding how much the model can explain and the non-linearities it learns from individual indicators. While there is variation regarding the magnitude of the contributions of the different components, we observe a very high correlation across specifications in how each components vary over time.

We also run an out-of-sample forecast exercise that shows a competitive forecast performance of the BBIM against an autoregressive model benchmark and against other standard machine learning tools. The BBIM performs significantly better than the autoregressive model and performs slightly better than a random forecast and Lasso regression. Removing the block structure and sign restrictions from our model and thus increasing its flexibility does not lead to a better forecasting performance.

Our analysis provides an important contribution towards broadening the practical use and relevance of machine learning models for economists. The BBIM can inform policymakers about inflation determinants and the timings when non-linear effects kick in and unwind, in a timely and intuitive manner. We argue that, among the machine learning methods typically employed in economics and finance, boosted trees are particularly well suited for that purpose since they are computationally cheap to train, relatively robust to the choice of hyperparameters, and identification can be well achieved within them. Specifically, imposing sign restrictions on associations between the output variable and input variables is relatively straightforward for tree-based prediction models and implemented in some of the most popular decision tree and boosting implementations (Cano et al., 2019). Also, the sequential training of decision trees that additively form a prediction suits itself well to defining separate model blocks in line with economic intuition that are sequentially trained conditionally on each other, reminiscent of dynamic factor models with a block structure (Potjagailo and Wolters, 2023). Linear boosting models have been employed on several components to retrieve time-varying parameters (Yousuf and Ng, 2021) or variable selection (Obster and Heumann, 2022). But to our knowledge, we are the first to employ component-wise non-linear boosting and to combine it with sign restrictions to achieve economic interpretability in a macroeconomic application.

Sign restrictions to separate supply and demand determinants are typically used within standard linear and time series models. They have recently been employed to disentangle US inflation dynamics, either from disaggregated price items and quantities (Shapiro et al., 2022; Firat and Hao, 2023), or on aggregate macroeconomic series in a dynamic factor model (Eickmeier and Hofmann, 2022), or an SVAR (Kabaca and Tuzcuoglu, 2023; Giannone and Primiceri, 2024). Our use of sign restrictions differs from these approaches as we only restrict associations between inflation and indicators separately for each indicator.

Our analysis also relates to a nascent strand of the literature that estimates a non-linear slope of the Phillips curve with regard to unemployment and labour market tightness (Benigno and Eggertsson, 2023, 2024; Gitti, 2024), as well as non-linearities induced via large cost-push shocks and global supply chain pressures (Harding et al., 2023; Di Giovanni et al., 2022; Ascari et al., 2024).

Closest to our paper is Goulet Coulombe (2022) which proposes a neural network model with Phillips curve structure for the United States. A slow-moving slope is separately identified from a cyclical output gap contribution— although this component remains unidentified with regard to the underlying shocks. We instead opt for the boosting trees method.

In the machine learning literature boosting models are known to perform well on prediction problems based on tabular data and have not been consistently surpassed by models based on neural networks (Grinsztajn et al., 2022; McElfresh et al., 2024). This contrasts starkly with predictions on text or image data, where neural networks are generally superior (Krizhevsky et al., 2012; Devlin et al., 2018). Especially on small tabular data sets, tree-based methods are known to outperform neural networks. Forecasting competition across thousands of time series have also shown a strong performance of boosting models (Makridakis et al., 2020, 2022).

Machine learning studies focusing on interpretability have also exploited the fact that boosting models can decompose the predictions of a model into the sum of simpler *base models*. By learning each base model on a single variable, the user can directly observe the nonlinear functional forms learned by the model (Lou et al., 2012; Nori et al., 2019; Chang et al., 2021).[1]. This approach however does not allow for any interactions of variables, whereas our component-wise boosting approach accounts for interactions of variables within the same component.

There exists a comprehensive machine learning literature on imposing sign restrictions, which are usually referred to as *monotonicity constraints* and are relevant for the interpretability and fairness of models (Nguyen and Martínez, 2019; Marques-Silva et al.,

[1]See also the earlier literature on generalised additive models (Hastie and Tibshirani, 1990) and the use of neural networks in additive models (Agarwal et al., 2021).

2021; Chen and Zhang, 2023; Sharma and Wehrheim, 2020). An array of different methods has been proposed to impose sign restrictions on neural networks (Cano et al., 2019) but these methods either require specific model structures (You et al., 2017), use heuristics that regularise the weights towards monotonicity without guaranteeing a monotonic solution gupta2019monotonic, or are computationally costly (Liu et al., 2020). The implementation of monotonicity constraints is more straightforward in decision trees – even though there exist a plethora of different approaches (Cano et al., 2019) which differ in performance and computational efficiency.

A disadvantage of the use of tree-based models is that they cannot extrapolate ,i.e. cannot predict values lower or larger than any of the observed values. While this limitation can be significant when forecasting at long horizons it is less crucial for our empirical approaches: cross-validation and short-term forecasting horizons with regular model updates.

The remainder of the paper is structured as follows. Section 2 describes our methodological approach: the boosted tree method organised according to a Phillips curve structure, the set of indicators used and the use of sign restrictions within decision trees, the learning algorithm and the empirical set-up. Section 3 presents the empirical results: the decomposition of UK inflation over time, and the signals and non-linear associations that the model learns. Section 4concludes.

2 Methodology

Applying machine learning to economics usually faces two main challenges. First, machine learning models are often black boxes allowing for limited interpretation of the result. Second, due to their purely data-driven nature, machine learning models may learn signals from data that ignore or contradict economically intuitive relations.

We address these challenges by studying a well-known economic relation, the Phillips curve, and developing a machine learning algorithm that can ascribe relations between a large dataset of individual variables and inflation to the different components of the Phillips curve. We employ sign restrictions on the relation of input variables and inflation to ensure theoretically coherent economic relations are adhered to.

We next develop the model in steps, explaining the advantage of each choice as we proceed.

2.1 Phillips curve blocks

We first start with a classic open economy Phillips curve relation as developed in Gali and Monacelli (2005),

$$\pi_t = \rho\pi_{t-1} + \beta E_t(\pi_{t+1}) + \lambda\hat{y}_t + \phi u_t. \tag{1}$$

Here π_t stands for present-day inflation at time t, $\rho\pi_{t-1} + \beta E_t(\pi_{t+1})$ are past and future expected inflation $\lambda\hat{y}_t = y_t - y_t^*$ is the output gap capturing movements in demand y_t and supply y_t^*, and ϕu_t captures shifts of the Phillips curve often referred to as "cost-push" shocks. This specification of the Phillips curve naturally arises from a linearisation of the New Keynesian model with a recursive inflation state.

While tractable, this stylised form of the Phillips may miss properties of the true empirical Phillips curve. The functional form of many components of the Phillips curve is unlikely to be linear and neither are the coefficients on the Phillips curve likely to be constant. In our empirical specification of the Phillips Curve we therefore split equation (1) into C components of interest,

$$\pi_{t+h} = \sum_{c}^{C} f^c(X^c_{t-p}) + \epsilon_t. \tag{2}$$

These components are equivalent to the variables in equation (1). $f^c(X^c_{t-p})$ describes the empirical function of component c, ϵ_t is a residual.

In our baseline specification, we include a component for inertia, capturing the effect of past and future expected price changes on inflation. This can be viewed as similar to the first two terms in equation (1), though our econometric set-up allows us to include a wider variety of series. Additionally, we include components for global and domestic supply and demand:

$$\begin{aligned}\pi_{t+h} = {} & f^{Inertia}(X^{Inertia}_{t-p})+ \\ & + f^{gDemand}(X^{gDemand}_{t-p}) + f^{dDemand}(X^{dDemand}_{t-p}) \\ & + f^{gSupply}(X^{gSupply}_{t-p}) + f^{dDSupply}(X^{dSupply}_{t-p}) \\ & \ldots + \epsilon_t \end{aligned} \tag{3}$$

2.2 Data and model blocks

We apply our method to data from the United Kingdom. Our sample period spans from February 1988 to August 2024. We argue that our method, which allows for integrating economic intuition is of particular advantage for small open economies like the United Kingdom where the identification of the Phillips curve is often more difficult due to the larger importance of external shocks and their openness to trade Bowdler (2009). Our data can be categorised into four broad categories. Data about UK prices and price expectations, data about economic activity in the UK, data capturing international economic activity, and identified economic shocks.

In total, we include 56 indicators. The inertia component includes inflation measures that reflect underlying inflation dynamics or second-round effects. It also includes time as variable to allow the model to fit a time-varying trend. A range of activity and labour market slack variables inform the demand and supply component. We also feed identified shock series from the literature directly into the model, and we use additional supply-related measures as such global and UK-specific supply chain pressure indices and certain inflation sub-components.

While the inertia component is separated from the other blocks, several series overlap between the supply and demand components. We can only include an indicator in both the demand and supply component if we constrain the association between inflation and the indicator having the opposite sign in supply and demand.

For instance, we restrict trees of the supply component to only allow splits that lead to a negative association between an activity variable (e.g. labour market tightness) and the predicted values of inflation. Similarly, we restrict trees in the supply component to learn a positive association between supply pressure indicators (e.g. GSCPI) and predicted inflation. Tree splits that imply a negative, or non-monotonic association of GSCPI and predicted inflation are discarded. Appendix A.1 provides more intuition on the imposition of sign restrictions in decision trees.

The different groups of variables and restrictions for separating supply and demand side determinants of inflation in the baseline specification of the model is shown in Table 1. The full details on all indicators are provided in Appendix B.

GROUP	INDICATORS	SIGN DEMAND	SIGN SUPPLY
Expectations, services inflation, wage growth	time indicator, 1-y ahead household infl. expectations, 5-y ahead financial market expectations, regular wage growth, services inflation, sub-components by sector		
Global activity	global PMI; US, EA: industrial production: US, EA: imports	+	−
	global activity shock, oil consumption demand shock (Baumeister and Hamilton, 2019)	+	
UK activity	industrial production, index of services; exports, imports, PMIs: services, manufacturing, construction; retail sales; consumer sentiment	+	−
	quarterly (interpolated): consumption, investment	+	−
	Labour market: v/u ratio, employment, inactivity rate	+	−
	Labour market: unemployment rate	−	+
Global supply and costs	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 (Baumeister and Hamilton, 2019)		+
UK supply and costs	CPI components: goods, food, electricity, gas;		+
	PPIs: input, output, gas, electricity; UK spot gas price		+

Table 1: Indicators and their restrictions

2.3 Boosting

Our method builds on the standard boosting paradigm in machine learning (Friedman, 2001). Boosting is an ensemble method that combines the predictions of a large number of *base* learners. The base learners are usually small decision trees that do not perform well by themselves. But in boosting the trees are trained sequentially: The trees are fitted on the residuals of the previous trees such that each additional tree improves the prediction of the ensemble slightly. More formally, let $F(X)$ be the prediction of the boosting model and y be the outcome. The boosting model is trained with the following steps.

1. Initialise model with mean value of target variable: $F_0 = \bar{y}$
2. For $m = 1$ to M:

(a) Compute residuals: $r_{im} = y_i - F_{m-1}(X_i), \forall i$

(b) Fit tree $f_m(X)$ to training set $\{(X_i, r_{im})\}_{i=1}^n$

(c) Update model: $F_m(X) = F_{m-1}(X) + \nu f_m(x)$s

To avoid overfitting, a small learning rate $0 < \nu < 1$ is used which ensures that each base learner changes the prediction of the boosting model only gradually.

2.4 Component-wise boosting

Unlike standard boosting methods, we impose a component structure. We sequentially fit trees that are calibrated on different groups of variables, reflecting the different components of the Phillips curve. Specifically, we train our model in $M = 200$ iterations. In each of the iterations, we fit a tree on the indicators of each of the C components, using a permutation π to randomise the order of the components in each iteration. We initialise the model with the inflation target of 2% such that we measure the contribution of the different components with respect to this baseline. The algorithm can be described as follows:

The algorithm is illustrated in Figure 12 and more formally shown here with π denoting the permutation function:

Initialise model with with: $F_0 = 2\%$ and $j = 1$.
For $m = 1$ to M:

1. For c in π (1:C):

   (a) Compute residuals: $r_i = y_i - F_{j-1}(X_i), \forall i$

   (b) Fit tree $f_m(X^c)$ to training set $\{(X_i^c, r_i)\}_{i=1}^n$

   (c) Update model: $F_j(X_i) = F_{j-1}(X_i) + \nu f_m(X_i^c), \forall i$

   (d) Increment $j = j + 1$

The prediction of our model is the *sum* of the prediction of all $M * C$ decision trees. Thus, we can measure the contribution of a component to the prediction by summing the predictions of those trees training on the component:

$$\pi_{t+h} = \sum_{i=1}^{M} f_i^{Inertia}(X_{t-p}^{Inertia}) + \sum_{i=1}^{M} f_i^{gDemand}(X_{t-p}^{gDemand}) + \sum_{i=1}^{M} f_i^{dDemand}(X_{t-p}^{dDemand}) + ... + \epsilon_t \tag{4}$$

2.5 Decision trees as base learners

We use decision trees as base learners for our boosting model. Their ability to fit arbitrary non-linear functions and their low computational costs make them the most popular base learner in boosting applications. Decision trees partition data points into homogeneous groups of observations that have similar values on the dependent variable. For each group, the tree predicts the mean value on the dependent variable of all observations falling into the group. The larger the decision tree, i.e. the larger the number of groups the observations are split into, the better the fit but the higher the degree of overfitting. Appendix A.1 illustrates in more detail how decision trees are trained.

2.6 Implementational details

To reduce overfitting we employ several strategies commonly used when training tree ensembles, including boosting models. Firstly, we do not fit each decision tree on the complete training data but on a random sample of 50% of the training observations. Secondly, we also sub-sample the variables the tree can learn from to 25% of the predictors in the respective component. Thirdly, we constrain the complexity of the tree by setting the maximum depth to 3 and by requiring that each node in a tree contains at least 5 observations.

2.7 Estimating Shapley values

While our model linearly decomposes the predicted values for inflation into components, the components do not directly reveal the learned relationship between individual variables and the output. To understand the learned functional forms, we use Shapley values, a standard interpretability framework in machine learning (Štrumbelj and Kononenko, 2014; Lundberg and Lee, 2017). We decompose the predicted value $\hat{y}_i$ into the sum of the Shapley values of predictors ϕ_j. Thus, $\hat{y}_i = \sum_j \phi_i^j + \phi^0$, where ϕ^0 is the baseline value, which is usually the mean predicted value in the training sample. If a variable is not split in a decision tree, its Shapley value is 0, the stronger the contribution to the prediction the higher the Shapley value. WE compute the Shapley values separatly for each tree in the boosting model and then aggregate the Shapley values across all trees of the same component.

In this study, we employ Shapley values in two ways. First, we measure a variable's average contribution to the prediction by averaging its absolute Shapley values of an indicator over the relevant period. Second, we reveal the functional form learned by the model by plotting the Shapley values of a variable as a function of the values on that variable that the model has learned from (see Buckmann et al., 2022).

2.8 Empirical approach

Indicators are included contemporaneously and with two lags and are transformed to be stationary. Our baseline empirical approach is 10-fold cross-validation. This allows us to train a model across the whole sample period and estimate consistent functional forms without the need to account for model shifts (see also Bluwstein et al., 2023; Buckmann et al., 2022). To obtain stable estimates and gauge the sensitivity of the model to stochastic processes in the estimation (sampling data points, and indicators in trees, permuting order of components, ...), we repeat the cross-validation 10 times and report the mean prediction and mean Shapley values.

We also test the forecasting performance of our model using out-of-sample testing. updating the model every quarter.

3 Results

We first present the results from the baseline model estimated from 1989, including global and domestic demand and supply and inertia that reflects expectations and lagged underlying inflation. Results are based on repeated 10-fold cross-validation, sample period 1989M1 to 2024M3. Figure 1 shows the decomposition of the inflation signal, smoothed across 12 months.

The model provides a meaningful inflation decomposition over the sample period. In the high inflation episode in the early 1990s, inertia (purple) – and therein mainly expectations explain the bulk of inflation. As expectations re-anchored following that episode the inertia contribution comes down substantially. Hence, while inertia is an overall rather slow-moving component, it can move quite swiftly at times potentially reflecting regime shifts in expectations and agents' beliefs.

Supply effects (orange) played some role in the early inflation episode but dragged on inflation between the end of the 1990s and mid-2000s, when increased global supply chain integration pushed down UK goods price inflation, and during the mid-2010s when global oil prices dropped. Demand (green) mostly contributes positively to CPI inflation. The sign restriction identification helps to detect a strong negative contribution from demand during the Global Financial Crisis (GFC), and a fall in the demand contribution to zero during the Covid-19 pandemic.

For the recent inflation episode, there remains an unexplained component during 2022 and early 2023 that the model is not able to capture. Demand and supply factors are seen as the main drivers of inflation on the way up, followed by a stronger role of inertia in the second half of the recent episode. Demand has a large contribution to the predictions mainly in 2021 until mid-2022. The supply contribution rapidly rose over 2022, before mostly unwound over 2023. The inertia contribution picked up over 2022, accounted for 2pp to the inflation prediction in 2023 but has nearly reached zero at the end of the sample period.

Components also reflect meaningful signals from individual indicators, as shown via absolute mean Shapley values in Figure 2.

3.1 Tracking non-linearities

The model detects non-linearities in various, but not all indicators in the recent period, and that non-linearities contributed to stronger model predictions. Figures 3 shows the learnt functional forms for a few key indicators. These are represented as scatter plots between the variable's input values (vertical axis) and the contributions to the predictions (Shapley values, vertical axis). Each dot represent a single month in the data.[2] While the the model is constrained to learn monotonically non-decreasing or non-increasing functions when sign restrictions are applied any functional form can be learnt when no sign restrictions are imposed. As there is no straightforward way to integrate the Shapley values across the different lags of the indicators, we here only show the functional form for the lag which shows the strongest predictive signals.

[2]Note that we do not apply smoothing here and that we exclude data points where we imputed the respective variable.

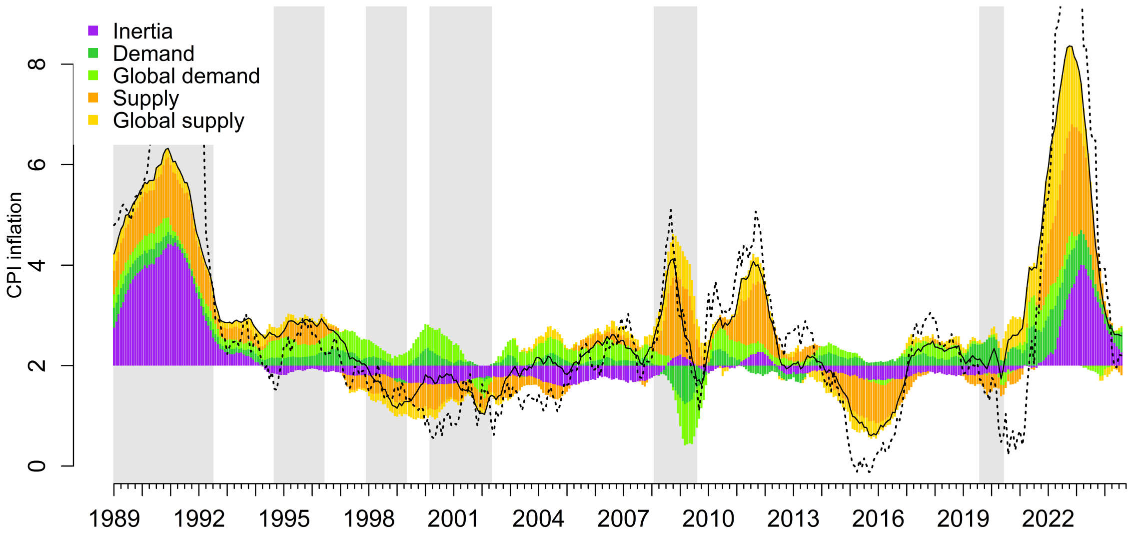

Figure 1: Decomposition of CPI inflation into contributions from model components.

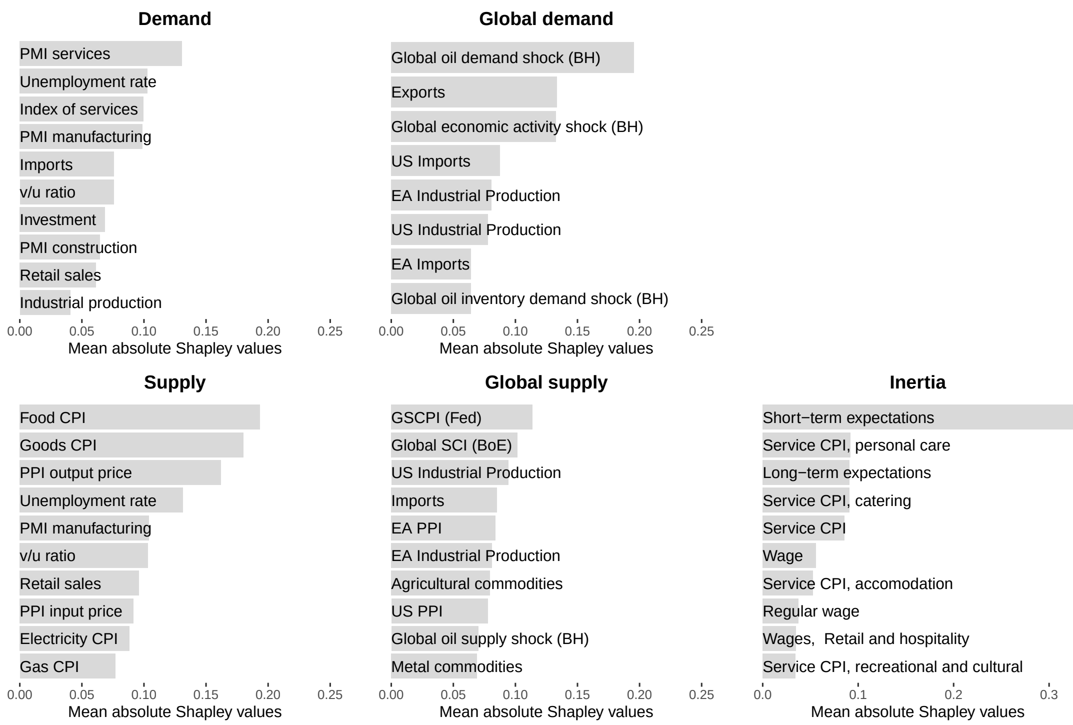

Figure 2: Mean absolute Shapley values within each component. 10 main indicators in each component.

Figure 3 depcits the learnt functional forms for a few key predictors in the baseline model. The colours are used to highlight the early period of high inflation (1989–1992) and the rise and fall of inflation in 2021–2024. To emphasise the non-linearities, we fit a linear model with a single breakpoint (i.e. two different slopes) to the relationship between input and Shapely values of the indicators (Muggeo, 2003).

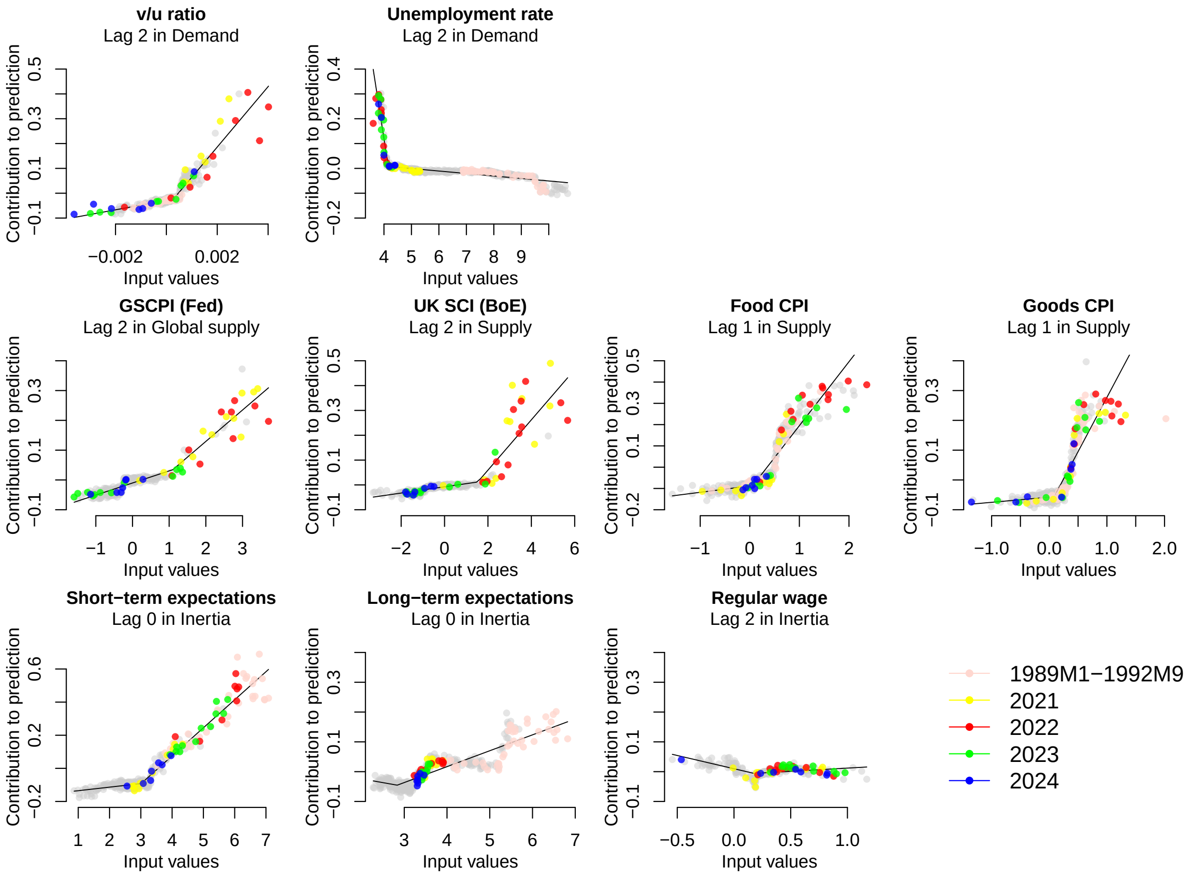

Figure 3: Functional forms learnt by the baseline model.

A non-linear Phillips curve with labour market slack

In the recent episode, the UK economy moved into the non-linear region of the Phillips curve both with respect to labour market tightness (v/u ratio) and with respect to unemployment (first row of Figure 3). This can explain the relatively large role of demand determinants identified by the model during the recent inflation rise. This result is in line with an L-shaped Phillips curve having played a role (Benigno and Eggertsson, 2023). Non-linear effects from labour market tightness mattered in 2021 and 2022 but have largely returned to the flat region. For unemployment, the signals remain in the steep region but started to return towards the flatter region.

Non-linearities in supply

On the supply side, we see clear non-linear associations with a range of indicators, most strikingly with the measure of global supply pressures, which have by now unwound (second row of Figure 3). The global supply chain pressure index (GSCPI) contributed very little to model predictions over the pre-pandemic sample, but its role strongly increased over 2021 and 2022. This is in line with previous evidence finding a non-linear role of global supply chains for inflation (Comin et al., 2023). These non-linear effects have fully returned to the flat region over the course of 2023. Non-linearities for cost-related price measures such as food CPI inflation (goods CPI, PPI), kicked in somewhat later and still contributed substantially to inflation in 2023, but recently have come down toward the flat region too. These can reflect indirect effects from global supply pressures but can also indicate that rising costs affecting goods producers might be passed on into price changes relatively strongly.

For energy prices, the evidence on the role of non-linearities in the recent episode is less clear-cut (electricity, gas, global oil supply news) and functional forms are more dispersed. While energy might have transmitted to inflation non-linearly via supply chains, the model does not attribute this directly to energy.

The recent rise in inertia relates non-linear effects from short-term expectations.

The effects from long term expectations were strong in the early 1990s, but remained flat recently (third row of Figure 3).

For inertia, we observe non-linear effects from short-term inflation expectations since 2023 that are currently still at play although slowly unwinding. Especially short-term expectations contributed strongly to the inflation prediction in 2023, and still into early 2024, although most recently their contribution is reduced. Short-term expectations might have non-linear effects when inflation is high because for instance household expectations can be particularly sensitive to the strong food and energy shocks.Also, firms' short-term expectations can become more relevant and more responsive to past outturn after an inflation surge (Cornea et al., 2013; Werning, 2022). This can add to domestic wage and price for some time (Lorenzoni and Werning, 2023)—albeit we find less evidence for non-linear wage effects having been at play recently. The effects from 5-year ahead market inflation expectations have remained flat, in stark contrast to the earlier inflation episode in the early 1990s, suggesting that expectations have remained anchored.

3.2 Model uncertainty around a suite of specifications

3.2.1 The role of the inertia component

In the baseline specification of our model, the inertia component includes 14 indicators that reflect past and expected inflation and second-round effects. Here we investigate the robustness of the inflation components to alternative specifications of the inertia component. First, we test a specification where we completely remove the inertia component from the model. Second, we only retain the two expectations in the inertia component, the 1-year ahead household expectations and 5-year ahead financial market expectations. Third, we place only the time variable in the inertia component, removing all other indicators, thus making the model learn a smooth time trend. Fourth, we combine the latter two specifications informing the time trend with the two expectation series.

Figure 4 (top-left panel) compares the contribution of the inertia component to the prediction for these different specifications. It shows that the predictions of all the models are very similar, with only the model without inertia differing significantly in the early period of high inflation.

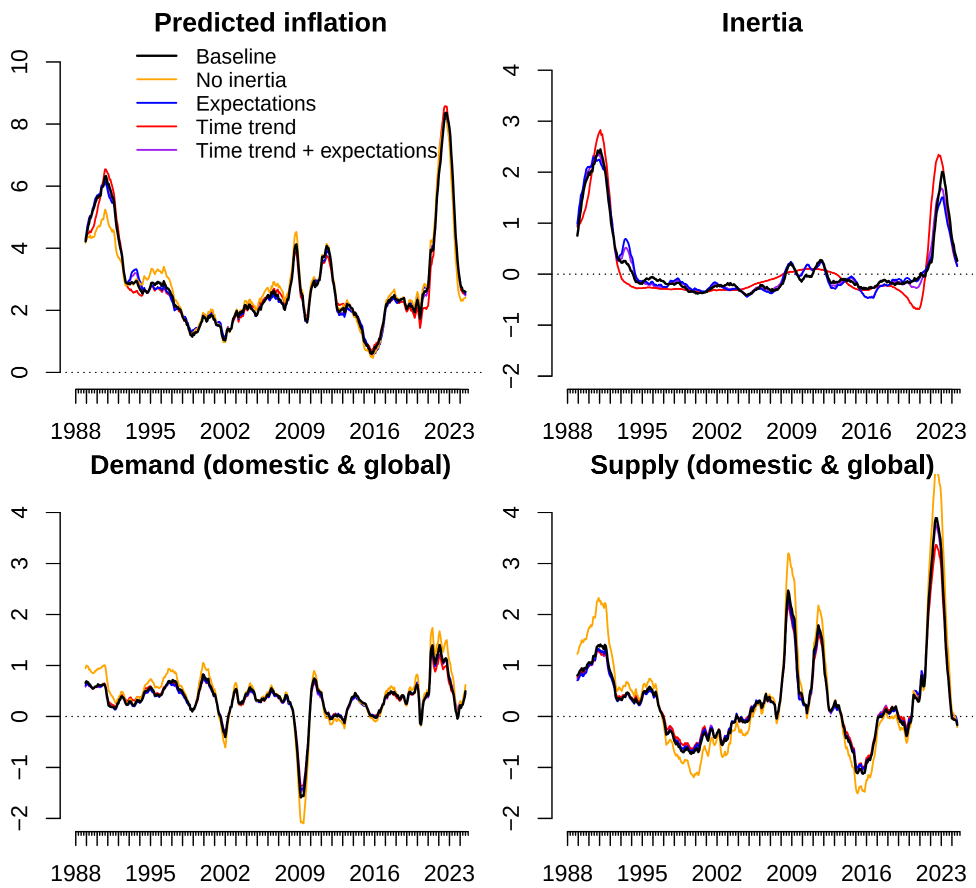

Figure 4: Alternative assumptions for the inertia component.

For all models with an inertia component, the we observe similar signals for that component (Figure 4, top-right panel). Only the time trend specification is smoother and does not explain smaller short-period changes in inflation as observed in the global financial crisis.[3].

[3]Note that the time variable is also included in the baseline specification, we observe that removing

The demand and supply components vary only very slightly for models that have an inertia component (bottom panel). Only when removing the inertia component we see stronger component values, particular for supply. This is expected as the supply component, additionally to activity variables, contains inputs prices which generally show a higher correlation with headline inflation.

3.2.2 Order of components

With our baseline approach, in each boosting iteration, we fit one tree of each component before we reshuffle the order of components (see Section 2.4. Here, we investigate the robustness of our baseline result to the order in which we learn from the components.

We start with an approach where we first learn 100 trees of the inertia component before fitting trees of the other components. This reflects a modelling approach where we control for an inflation trend first before pinning down supply and demand factors. Second, we test a specification where we fit the global supply and demand components (100 trees each) first before learning from the remaining components. This reflects a modelling approach where we first control for exogenous factors before considering domestic determinants of inflation. Third, we also test a specification where we remove the global supply and demand components.

Figure 5 compares the component contribution to the prediction for the different specifications. As expected, when trees of a specific component are fitted first, this component gets gets more weight. However, while the signals are stronger they correlate highly with the signals of the other specifications. This also holds when removing the global variables, which gives more weight to the remaining components but does not qualitatively change the activation patterns on these components.

3.2.3 Identification of demand and supply

We also test a model without sign restriction identification, simply with a component combining all activity variables, and another component including the cost and supply indicators. Figure 6 compares this specification to our baseline. The activity component does not show the intuitive demand signals we observe with our baseline specification, such as the negative demand effects in the GFC or Covid (left panel). Furthermore, the input shocks show less pronounced signals than the supply component in our baseline specification (right panel).

3.3 Extended model with financial conditions and monetary policy surprises blocks

We extend our baseline model with two additional components, financial conditions and monetary policy. The monetary policy component contains the UK monetary policy surprises (target, path & QE factor) estimated by (Braun et al., 2023). As these are only available from June 1997, we start the estimate in that year. The decomposition is shown in Figure 7. Lose financial conditions slightly pushed down inflation in the run-up to the GFC and thereafter contributed to inflation as tighter financial conditions likely incentivised firms to raise prices. Exchange rate effects, within the financial component, also

it does not have an observable effect on the inertia component

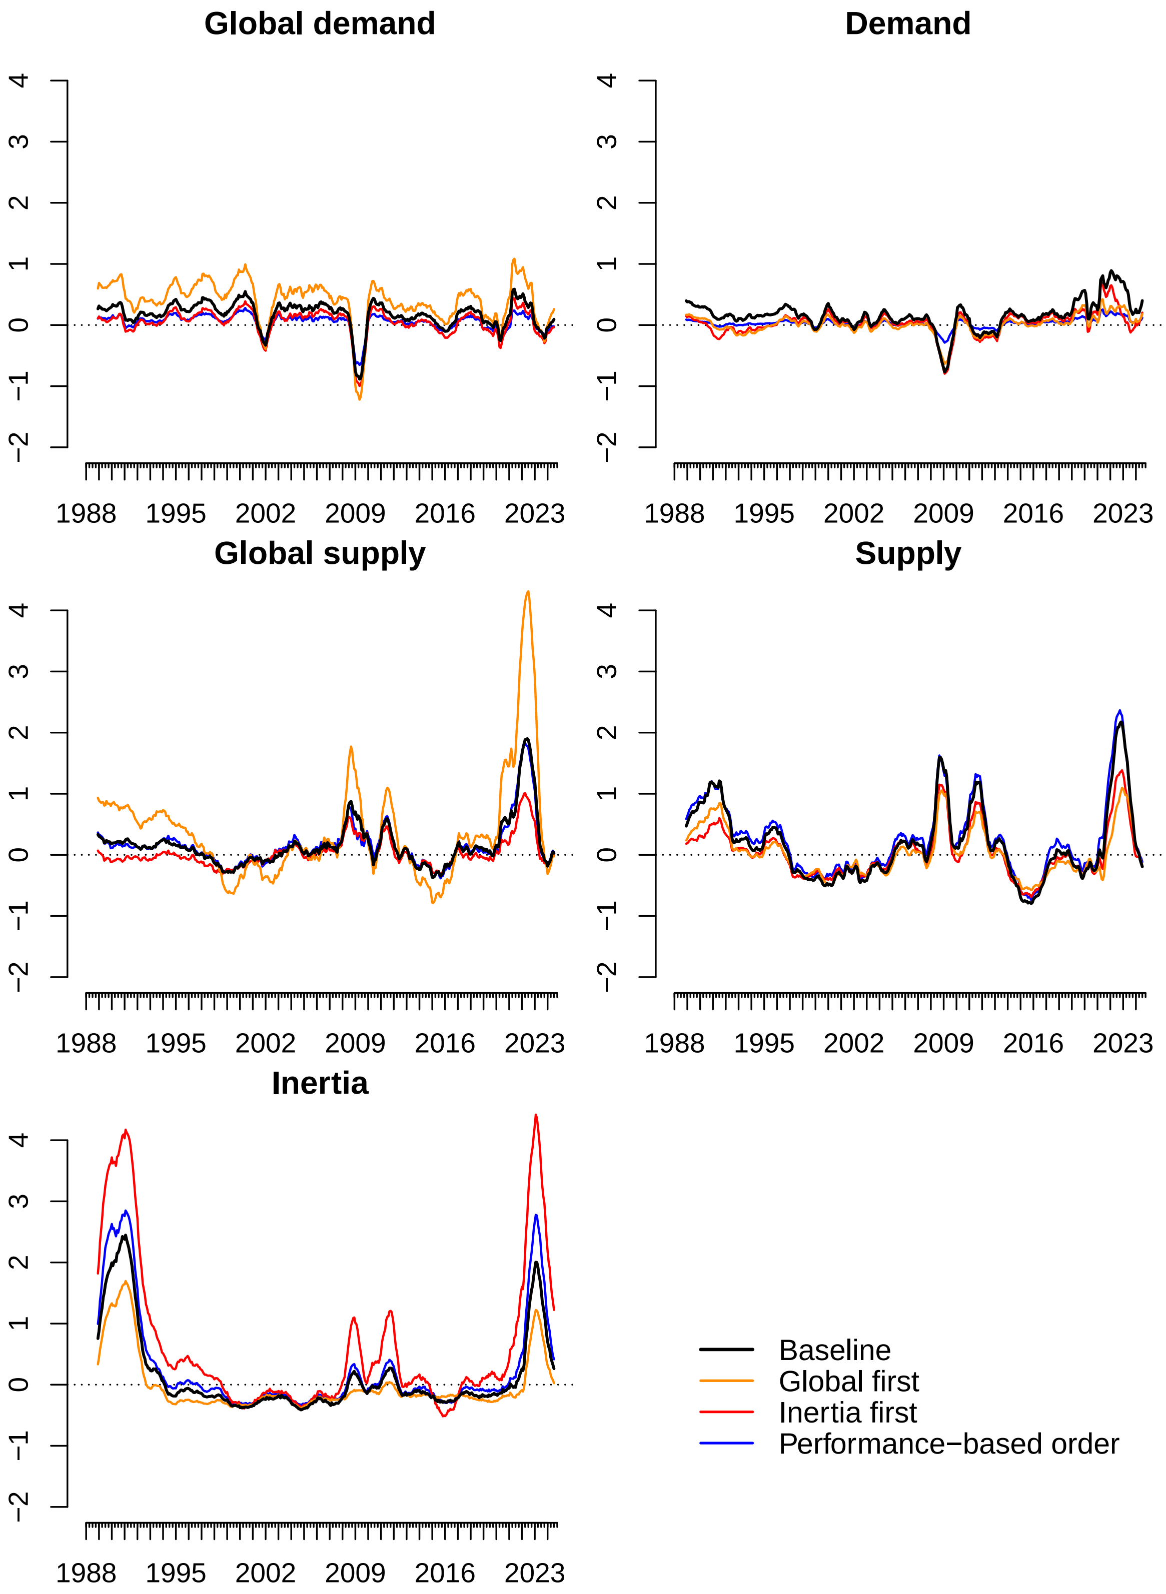

Figure 5: Contribution to the prediction of all components in different specifications.

pushed up inflation in the early following the Brexit vote. Monetary policy surprises were slightly inflationary following the GFC, during 2017-2019 before QE was fully unwound, and in 2020 during the Covid-19 stimulus. At the end of 2022 and early 2023 monetary policy surprises might have systematically surprised to the expansionary side, pushing

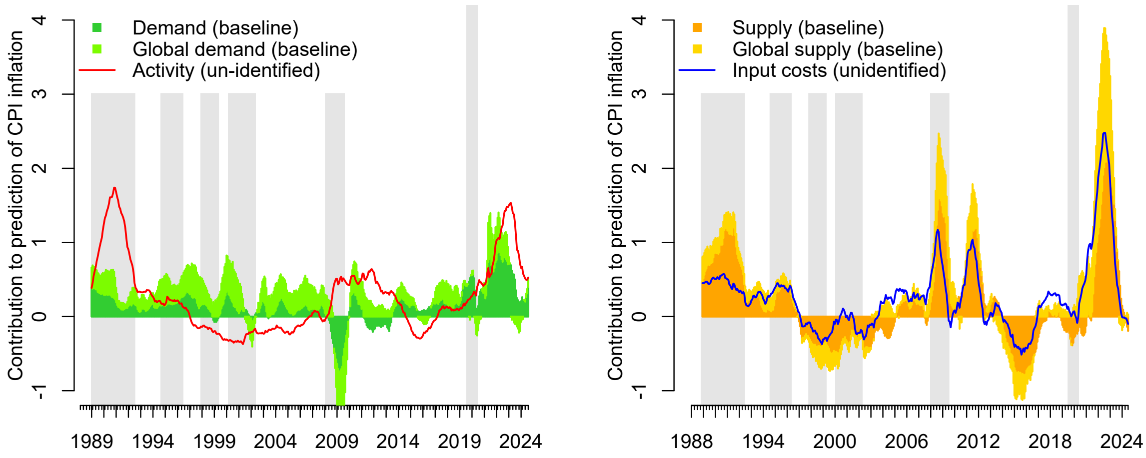

Figure 6: Comparison with model without sign restrictions.

up inflation according to the model.

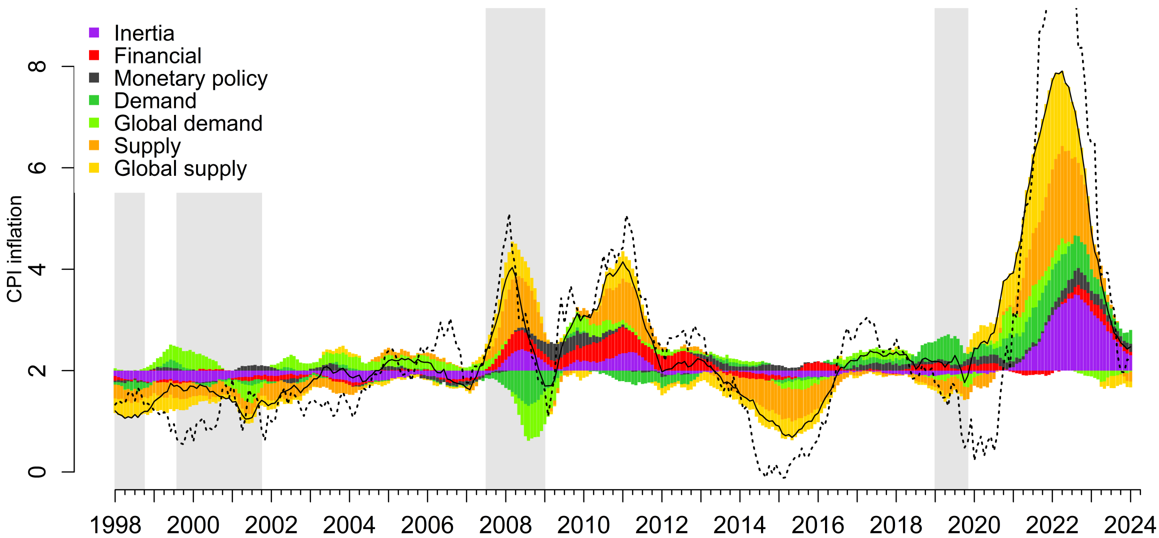

Figure 7: Decomposition of CPI inflation into contributions from extended model components.

3.4 Out-of-sample forecasts

We compare the performance of our boosting model against an AR(2), a random forest and a Lasso regression. Apart from AR(2) all models use the same set of indicators as our boosting model but without any block structure or sign restrictions.

We forecast inflation one month ahead, starting in January 2020, using an expanding training sample. We re-train the machine learning models every quarter. The computationally cheap AR(2) model is re-estimate every month. We conduct a hyperparameter search for the regularisation parameter of the Lasso regression and the random forest at the beginning of the forecasting sample (January 2020) but do not re-estimate the hyperparameters to save computation time. Table 2 presents the results. Our baseline model performns significantly better than the AR and performs slightly better than the random forest and Lasso regression. In particular in the recent period of high inflation. Removing the inertia component (BBIM (no inertia) or using the specification without sign restrictions and the activity series (BBIM (activities & input costs), see Figure 6) does not have a strong effect on performance. However we see that these specifications perform slightly worse after 2020.

Figure 8 compares the forecasting performance of our baseline to the benchmarks at different forecast horizons. While all ML models show a similar performance in the early sample period, we observe that our model performs better than the other methods between 2020–2024.

Figure 9 compares the contribution of the component to the prediction when doing forecasting vs. cross-validation. In the recent inflation surge, the forecasted values are generally lower than the predictions obtained using cross-validation but the series are generally well aligned.

	Complete sample	2000-2019	2020-2024
AR(2)	1.00 (-)	1.00 (-)	1.00 (-)
Random forest	0.89*** (0.00)	0.90*** (0.00)	0.87** (0.02)
Lasso regression	0.87*** (0.00)	**0.88***** (0.00)	0.82*** (0.01)
BBIM (baseline)	**0.86***** (0.00)	**0.88***** (0.00)	**0.80***** (0.00)
BBIM (activity & input costs)	0.89*** (0.00)	0.89*** (0.00)	0.87** (0.03)

Notes: Mean absolute error relative to mean absolute error of AR(2). In parentheses: p-value of Diebold-Mariano test. ***, **, * indicate significance at 1%, 5%, or 10%. Sample period up to 2024M8.

Table 2: Absolute forecast error (1-month ahead) relative to AR2.

4 Conclusion

Standard linear models are not well equipped to model inflation in the face of large shocks. Machine learning methods enhanced with economic intuition such as the Boosted Inflation Model provide an appealing tool. We show that non-linearities in the Phillips curve and in the effects from supply shocks and global supply constraints strongly amplified UK CPI inflation dynamics in the recent episode. In part, they have also contributed to swift disinflation along the “steep” slope. The lingering contribution from inertia suggests that inflation might take longer to unwind in the medium term because of short-term

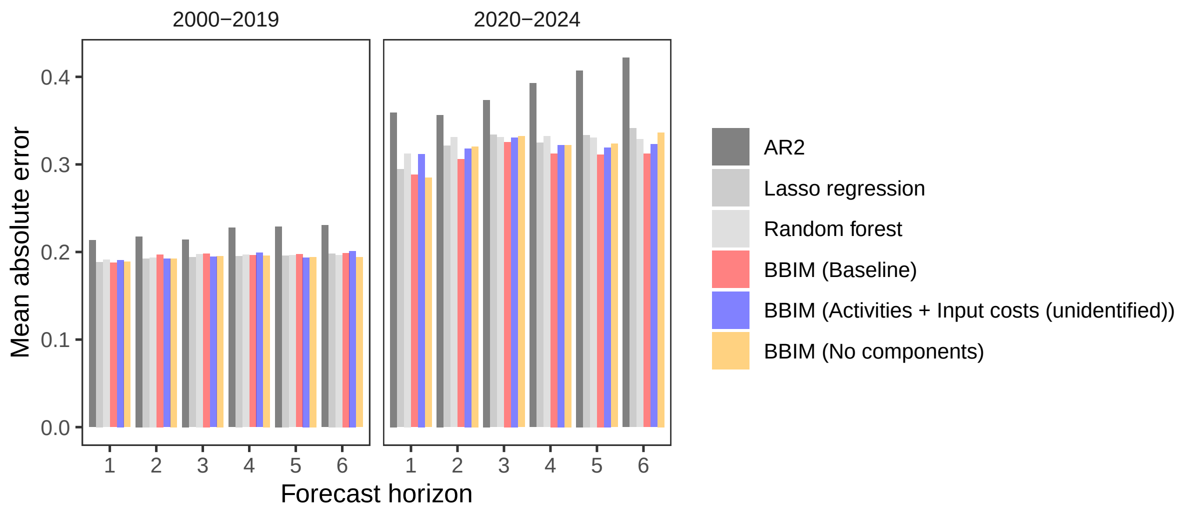

Figure 8: Forecast performance at different horizons for two sample periods: 2000–2019 and 2020–2024.

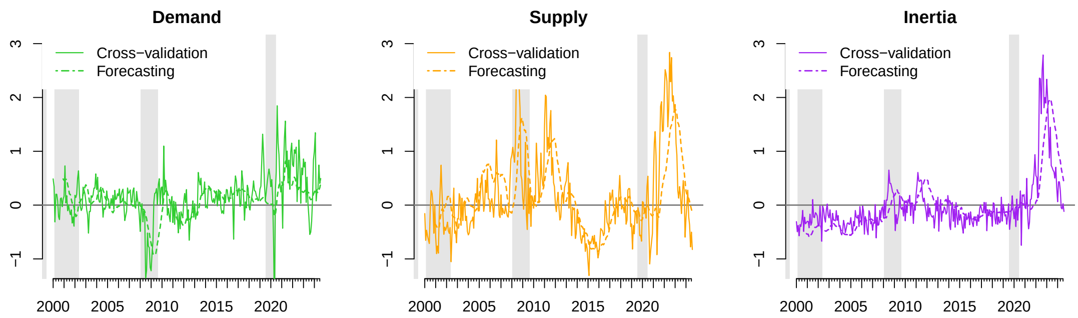

Figure 9: OOS forecast components (h=1) qualitatively similar to cross-validation results.

expectations effects. Long-term expectations effects remain flat, so there is little evidence for a strongly persistent regime shift.

References

Agarwal, R., L. Melnick, N. Frosst, X. Zhang, B. Lengerich, R. Caruana, and G. E. Hinton (2021). Neural additive models: Interpretable machine learning with neural nets. Advances in neural information processing systems 34, 4699–4711.

Ascari, G., D. Bonam, and A. Smadu (2024). Global supply chain pressures, inflation, and implications for monetary policy. Journal of International Money and Finance, 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). The Slanted-L Phillips Curve. NBER Working Paper No w32172.

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.

Bowdler, C. (2009). Openness, exchange rate regimes and the phillips curve. Journal of International Money and Finance 28(1), 148–160.

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

Buckmann, M., A. Joseph, and H. Robertson (2022). An interpretable machine learning workflow with an application to economic forecasting. Technical report, Bank of England.

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.

Chang, C.-H., S. Tan, B. Lengerich, A. Goldenberg, and R. Caruana (2021). How interpretable and trustworthy are GAMs? In Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining, pp. 95–105.

Chen, D. and L. Zhang (2023). Monotonicity for ai ethics and society: An empirical study of the monotonic neural additive model in criminology, education, health care, and finance. arXiv preprint arXiv:2301.07060.

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.

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

Cornea, A., C. Hommes, and D. Massaro (2013). Behavioral heterogeneity in us inflation dynamics. Technical report, Tinbergen Institute Discussion Paper.

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

Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.

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.

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.

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.

Gitti, G. (2024). Nonlinearities in the regional phillips curve with labor market tightness. Unpublished manuscript.

Goulet Coulombe, P. (2022). A neural phillips curve and a deep output gap. Available at SSRN 4018079.

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 (2023). Understanding post-covid inflation dynamics. Journal of Monetary Economics 140, S101–S118.

Hastie, T. and R. Tibshirani (1990). Generalized additive models. Taylor and Francis.

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.

Krizhevsky, A., I. Sutskever, and G. E. Hinton (2012). Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems 25.

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

Liu, X., X. Han, N. Zhang, and Q. Liu (2020). Certified monotonic neural networks. Advances in Neural Information Processing Systems 33, 15427–15438.

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.

Makridakis, S., E. Spiliotis, and V. Assimakopoulos (2020). The m4 competition: 100,000 time series and 61 forecasting methods. International Journal of Forecasting 36(1), 54–74.

Makridakis, S., E. Spiliotis, and V. Assimakopoulos (2022). M5 accuracy competition: Results, findings, and conclusions. International Journal of Forecasting 38(4), 1346–

1364.
Marques-Silva, J., T. Gerspacher, M. C. Cooper, A. Ignatiev, and N. Narodytska (2021). Explanations for monotonic classifiers. In International Conference on Machine Learning, pp. 7469–7479. PMLR.
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.
Ng, S. (2014). Boosting recessions. Canadian Journal of Economics/Revue canadienne d'économique 47(1), 1–34.
Nguyen, A.-p. and M. R. Martínez (2019). Mononet: towards interpretable models by learning monotonic features. arXiv preprint arXiv:1909.13611.
Nori, H., S. Jenkins, P. Koch, and R. Caruana (2019). Interpretml: A unified framework for machine learning interpretability. arXiv preprint arXiv:1909.09223.
Obster, F. and C. Heumann (2022). Sparse-group boosting–unbiased group and variable selection. arXiv preprint arXiv:2206.06344.
Potjagailo, G. and M. H. Wolters (2023). Global financial cycles since 1880. Journal of International Money and Finance 131, 102801.
Shapiro, A. H. et al. (2022). Decomposing supply and demand driven inflation. FRB of San Francisco Working Paper 2022-18.
Sharma, A. and H. Wehrheim (2020). Higher income, larger loan? monotonicity testing of machine learning models. In Proceedings of the 29th ACM SIGSOFT International Symposium on Software Testing and Analysis, pp. 200–210.
Štrumbelj, E. and I. Kononenko (2014). Explaining prediction models and individual predictions with feature contributions. Knowledge and information systems 41, 647–665.
Werning, I. (2022). Expectations and the rate of inflation. Technical report, National Bureau of Economic Research.
Yoon, J. (2021). Forecasting of real gdp growth using machine learning models: Gradient boosting and random forest approach. Computational Economics 57(1), 247–265.
You, S., D. Ding, K. Canini, J. Pfeifer, and M. Gupta (2017). Deep lattice networks and partial monotonic functions. Advances in neural information processing systems 30.
Yousuf, K. and S. Ng (2021). Boosting high dimensional predictive regressions with time varying parameters. Journal of Econometrics 224(1), 60–87.

A Methods

A.1 Intuition for regression trees

Regression trees are universal function approximators. They approximate functions by splitting the variable space of the independent variables into a set of intervals with predicted outcomes for a y variable. The deeper a tree, the more splits, the better the fit. This is illustrated for the sigmoid function with one input X, $Y = \frac{1}{1+\exp(-X)}$ below in Figure 10.

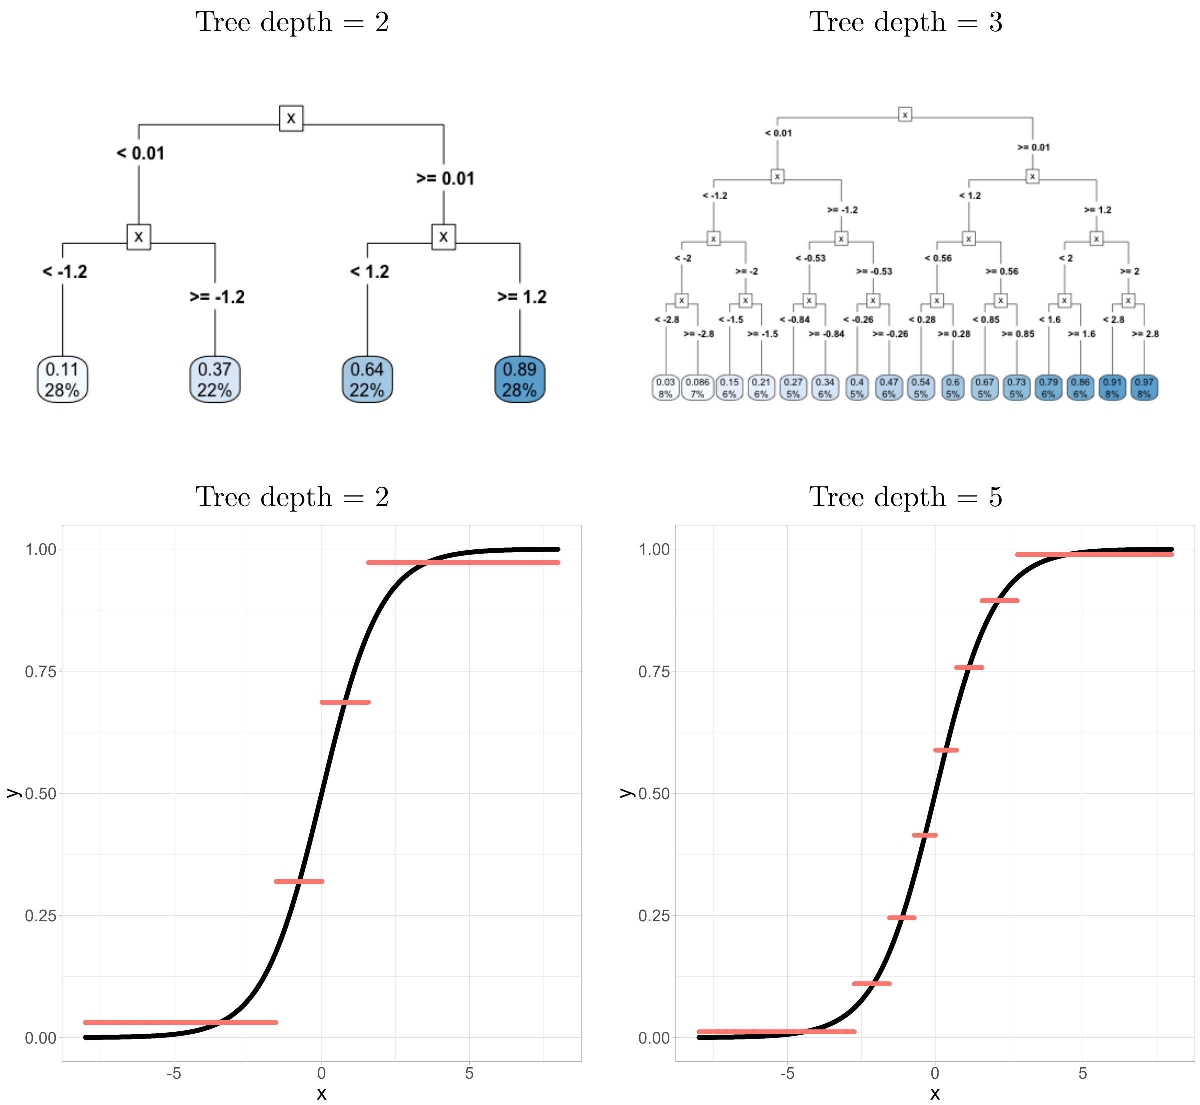

Figure 10: Illustration of fitting a function with one input with a regression tree.

To prevent over-fitting, trees will usually be restricted in their depth. However, that can't prevent the model from fitting a tree against trialled and tested theoretical economic predictions. To prevent this we can implement sign restrictions. This is illustrated with Figure 11. With a sign restriction, in this case, a positive sign restriction we disallow splits that would predict for any $X_1 < X_2$ that $f(X_1) > f(X_2)$. The top panel illustrates with the framed box such a disallowed split. After removing these splits the function is approximated with a purely increasing regression tree. A sign restricted tree is exemplified in the lower panels of figure 11 where we fit a parabola with positive sign restrictions.

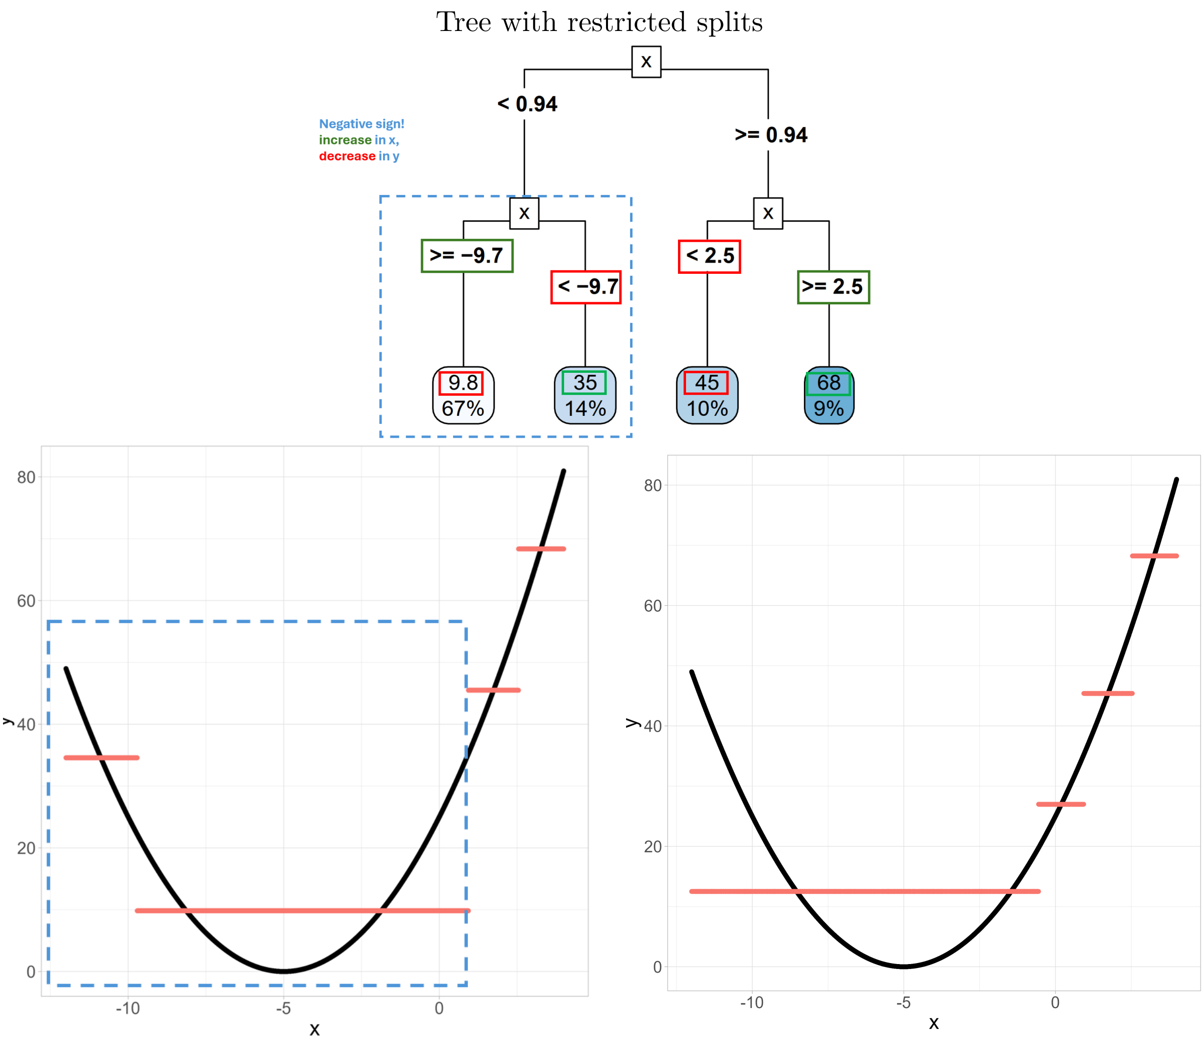

Figure 11: Illustration of fitting a function with one input with a regression tree with sign restrictions.

A.2 Implementational details

We implemented our model in Python. We use the decision tree implementation in `xgboost` (Chen and Guestrin, 2016). While this package's main purpose is training boosting models, we use it to train individual trees by setting `n_estimators` = 1 and `learning_rate` = 1. We use this library as it is computationally efficient and implements sign restrictions. To impose sign restrictions, we use the `monotone_constraints` parameter. To estimate Shapley values we use the `shap` implementation (Lundberg and Lee, 2017). Specifically, we use computationally efficient `TreeExplainer` approach (Lundberg et al., 2018).

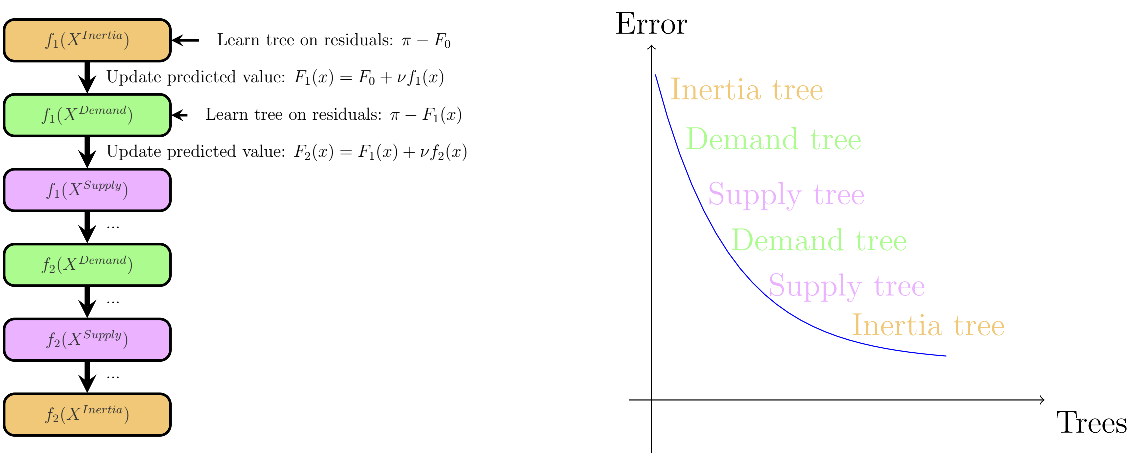

Figure 12: Illustration of the algorithm

B Data Description

We estimate our model specifications based on 69 series for which main moments are described in Tables 3 and 4. We estimate the model on monthly data. Before going into the model we transform a subset of the series. Our codes for transformation of the series follow here McCracken and Ng (2016). Transformation code 1 means no transformation, while 2 means taking the first difference and 5 means taking the log difference of the underlying series. These transformations are done to make the series stationary and meaningful for inflation.

variable	Mean	Stdv	First Obs.	Latest Obs.	trans	Blocks
Consumer sentiment	-11.09	11.61	1980-01-01	2024-12-01	1	DEMAND
Regular wage	437.82	88.47	2000-01-01	2024-10-01	5	DEMAND
Wage	345.79	164.99	1980-01-01	2025-03-01	5	DEMAND
Wages, Retail and hospitality	290.23	58.91	2000-01-01	2024-10-01	5	DEMAND
Wages, financial services	526.78	124.73	2000-01-01	2024-10-01	5	DEMAND
Wages, manufacturing	509.19	95.56	2000-01-01	2024-10-01	5	DEMAND
Wages, services	422.99	88.71	2000-01-01	2024-10-01	5	DEMAND
Imports	28686.15	17002.71	1980-01-01	2024-10-01	5	DEMAND,GLOBAL_SUPPLY
Index of services	81.39	12.72	1997-01-01	2024-10-01	5	DEMAND,SUPPLY
Industrial production	99.64	6.27	1997-01-01	2024-10-01	5	DEMAND,SUPPLY
Investment	78977.95	19749.28	1980-01-01	2024-07-01	5	DEMAND,SUPPLY
PMI construction	53.65	6.72	1997-04-01	2024-12-01	5	DEMAND,SUPPLY
PMI manufacturing	51.65	4.28	1991-07-01	2024-12-01	5	DEMAND,SUPPLY
PMI services	54.33	4.87	1996-07-01	2024-12-01	5	DEMAND,SUPPLY
Retail sales	76.03	17.34	1988-01-01	2024-11-01	5	DEMAND,SUPPLY
Unemployment rate	7.04	2.41	1980-01-01	2025-02-01	1	DEMAND,SUPPLY
v/u ratio	0.30	0.21	1980-01-01	2024-09-01	4	DEMAND,SUPPLY
Corporate bond spread	155.94	75.59	1998-01-01	2023-10-01	1	FINANCIAL
FTSE UK focused	86.93	17.88	1995-01-01	2024-12-01	5	FINANCIAL
GBP-USD spot exrate	1.59	0.24	1980-01-01	2024-12-01	2	FINANCIAL
House price index	295.04	138.59	1991-01-01	2024-12-01	5	FINANCIAL
Real exchange rate index	89.02	9.04	1990-01-01	2024-12-01	5	FINANCIAL
EA Exports	91.35	10.70	2000-01-01	2024-10-01	5	GLOBAL_DEMAND
EA Imports	93.34	8.33	2000-01-01	2024-10-01	5	GLOBAL_DEMAND
Global economic activity shock (BH)	-0.02	0.68	1980-01-01	2024-08-01	1	GLOBAL_DEMAND
Global oil demand shock (BH)	-0.15	3.68	1980-01-01	2024-08-01	1	GLOBAL_DEMAND
Global oil inventory demand shock (BH)	-0.01	1.11	1980-01-01	2024-08-01	1	GLOBAL_DEMAND
US Exports	85.37	17.31	2000-01-01	2024-10-01	5	GLOBAL_DEMAND
US Imports	79.56	15.17	2000-01-01	2024-10-01	5	GLOBAL_DEMAND
EA Industrial Production	90.96	8.70	1991-01-01	2024-11-01	5	GLOBAL_DEMAND,GLOBAL_SUPPLY
US Industrial Production	82.16	18.93	1980-01-01	2024-11-01	5	GLOBAL_DEMAND,GLOBAL_SUPPLY
Exports	22091.92	11154.69	1980-01-01	2024-10-01	5	GLOBAL_DEMAND,SUPPLY
Agricultural commodities	242.14	73.88	1980-01-01	2024-12-01	5	GLOBAL_SUPPLY
Commodity price index, energy	61.18	39.14	1980-01-01	2024-12-01	5	GLOBAL_SUPPLY
Commodity price index, food	76.14	26.96	1980-01-01	2024-12-01	5	GLOBAL_SUPPLY
Commodity price index, non-energy	69.96	25.38	1980-01-01	2024-12-01	5	GLOBAL_SUPPLY
EA PPI	86.90	13.73	1995-01-01	2024-11-01	5	GLOBAL_SUPPLY
GSCPI (Fed)	0.01	1.00	1998-01-01	2024-12-01	1	GLOBAL_SUPPLY
Global SCI (BoE)	0.00	1.26	2007-05-01	2024-12-01	1	GLOBAL_SUPPLY
Global oil supply shock (BH)	-0.08	1.38	1980-01-01	2024-08-01	1	GLOBAL_SUPPLY

Table 3: Series on which we estimate the different model specifications. Note some series enter in multiple categories, in which case they are separated by commas.

variable	Mean	Stdv	First Obs.	Latest Obs.	trans	Blocks
Metal commodities	247.68	124.31	1980-01-01	2024-12-01	5	GLOBAL_SUPPLY
Oil supply news shock	0.00	0.58	1980-01-01	2024-06-01	1	GLOBAL_SUPPLY
US PPI	155.07	45.15	1980-01-01	2024-12-01	5	GLOBAL_SUPPLY
Long-term expectations	3.82	1.13	1985-02-01	2024-11-01	1	INERTIA
Service CPI	78.67	29.29	1988-01-01	2025-09-01	5	INERTIA
Service CPI, accomodation	92.68	27.41	1996-01-01	2025-09-01	5	INERTIA
Service CPI, catering	80.83	29.32	1988-01-01	2025-09-01	5	INERTIA
Service CPI, personal care	88.76	16.30	1988-01-01	2024-12-01	5	INERTIA
Service CPI, recreational and cultural	78.71	29.24	1988-01-01	2025-09-01	5	INERTIA
Short-term expectations	3.81	2.16	1980-01-01	2024-12-01	1	INERTIA
MP QE shock (BMS)	0.00	0.03	1997-06-01	2024-07-01	1	MONETARY_POLICY
MP shock, path (BMS)	-0.01	0.04	1997-06-01	2024-07-01	1	MONETARY_POLICY
MP shock, target (BMS)	0.00	0.05	1997-06-01	2024-07-01	1	MONETARY_POLICY
Electricity (PPI)	96.05	55.05	1996-01-01	2024-12-01	5	SUPPLY
Electricity CPI	89.68	54.54	1988-01-01	2028-12-01	5	SUPPLY
Food CPI	83.32	23.27	1988-01-01	2025-09-01	5	SUPPLY
Gas CPI	73.46	45.82	1988-01-01	2028-12-01	5	SUPPLY
Goods CPI	93.14	15.48	1988-01-01	2025-09-01	5	SUPPLY
Inactivity rate	36.82	0.57	1980-01-01	2024-09-01	1	SUPPLY
PPI input price	83.81	28.66	1984-01-01	2024-12-01	5	SUPPLY
PPI output price	82.07	23.74	1980-01-01	2024-12-01	5	SUPPLY
Real regular wage	4.58	0.19	2000-01-01	2024-10-01	5	SUPPLY
Real wage	4.20	0.84	1980-01-01	2024-12-01	5	SUPPLY
Real wage, retail and hospitality	3.04	0.12	2000-01-01	2024-10-01	5	SUPPLY
Real wage, financial services	5.48	0.35	2000-01-01	2024-10-01	5	SUPPLY
Real wage, manufacturing	5.34	0.19	2000-01-01	2024-10-01	5	SUPPLY
Real wage, services	4.42	0.21	2000-01-01	2024-10-01	5	SUPPLY
UK Gas	48.32	46.94	1996-04-01	2024-12-01	5	SUPPLY
UK SCI (BoE)	-0.28	1.42	1998-01-01	2024-12-01	1	SUPPLY

Table 4: Series on which we estimate the different model specifications. Note some series enter in multiple categories, in which case they are separated by commas.

C Further results