UNIVERSITÀ DEGLI STUDI DI TORINO
ALMA UNIVERSITAS TAURINENSIS
Founded in 1404

DEPARTMENT OF ECONOMICS AND STATISTICS WORKING PAPER SERIES

Quaderni del Dipartimento di Scienze Economico-Sociali e Matematico-Statistiche

ISSN 2279-7114

# INFLATION AND PRICE FLEXIBILITY

Ivan Petrella

Emiliano Santoro

Yannik Winkelmann

Working paper No. 99 - May 2025

# Inflation and Price Flexibility*

Ivan Petrella†
Collegio Carlo Alberto, University of Turin (ESOMAS)
University of Warwick & CEPR

Emiliano Santoro‡
Catholic University of Milan

Yannik Winkelmann§
University of Tübingen

March 2025

**Abstract**

Using UK consumer price microdata, we report that aggregate *price flexibility* varies substantially over time and induces significant non-linearity in inflation. In a regime of high flexibility, the half-life of inflation drops by 50% and its volatility rises considerably. Such asymmetry arises naturally from state-dependent pricing, for which we find ample evidence in the data, particularly following the Great Recession. Neglecting this property may lead to a systematic underprediction of inflation, as seen in the post-Pandemic inflation surge. Tracking real-time movements in price flexibility is crucial for assessing inflation dynamics and to inform monetary policy decisions.

*This paper is a deeply revised version of "Time-varying Price Flexibility and Inflation Dynamics" of which Lasse de la Porte Simonsen was a co-author. We wish to thank Lasse for his input during the initial stage of this project. We also thank for comments Yoosoon Chang, Andrea Colciago, Luca Dedola, Federico Di Pace, Etienne Gagnon, Ana Galvao, Francesco Lippi, Ricardo Nunes, Chiara Osbat, Roberto Pancrazi, Giorgio Primiceri, Omar Rachedi, Federico Ravenna, Søren Hove Ravn, Joseph Vavra and Fabrizio Venditti.

†Collegio Carlo Aberto, University of Turin (ESOMAS), University of Warwick and CEPR. *Address*: Piazza Vincenzo Arbarello, 8, 10122 Turin, Italy. *E-mail*: ivan.petrella@carloalberto.org.

‡Department of Economics and Finance, Catholic University of Milan. *Address*: Via Necchi 5, 20123 Milan, Italy. *E-mail*: emiliano.santoro@unicatt.it.

§Department of Economics, University of Tübingen. *Address*: Mohlstraße 36, Room 410. *Email:* yannik.winkelmann@uni-tuebingen.de.

# 1 Introduction

Over the past decade, the growing availability of disaggregated consumer price data has allowed economists to closely analyze price-setting behavior, assess the empirical validity of various price-adjustment theories, and derive various measures of aggregate *price flexibility*. The latter, broadly understood as the response of the aggregate price level to macroeconomic shocks, is central to the transmission of monetary policy, and ultimately shapes the trade-off Central Banks face between stabilizing output and inflation. While numerous studies have explored the degree of price sluggishness and its underlying drivers, less emphasis has been placed on how aggregate price flexibility evolves over time. In particular, the literature has largely overlooked how, and to what extent, time-varying price flexibility affects inflation dynamics and the ability of inflation-targeting Central Banks to meet their targets. We address this gap, and show that tracking changes in price flexibility is key in improving inflation projections, deepening our understanding of inflation dynamics, and enhancing the effectiveness of monetary policy.

Using monthly price microdata underlying the UK Consumer Price Index (CPI) from 1996 to 2024, we estimate the generalized $Ss$ model developed by Caballero and Engel (2007). Along with encompassing different price-setting protocols, this model is well suited to examine time variation and comovement among various price-setting statistics. Estimation involves fitting both the *distribution of price gaps* (i.e., the wedge between actual and optimal reset prices) and the *adjustment hazard* (i.e., the probability of a good's price changing as a function of its price gap). We exploit our time-varying estimates of the $Ss$ model to establish a connection between inflation and the underlying process of price setting. To this end, we back out predetermined price adjustments—the so-called *intensive margin*—and adjustments triggered or canceled by the shock—the *extensive margin*.[1] While the intensive margin was typically the primary driver of price flexibility up until the Great Recession, state dependence

[1] Adjustments occurring over the intensive margin characterize both time- and state-dependent models. The extensive margin, instead, is a defining feature of state-dependent models.

in price setting—reflected in the extensive margin—becomes largely dominant thereafter, so that larger price adjustments become more likely to be enacted. Such development induces considerable volatility in inflation dynamics and appears quite important, particularly at the onset of the Great Recession and of the post-Pandemic period, both episodes being characterized by spikes in inflation volatility.

Aggregate price flexibility varies significantly over time, peaking during 2008-2011—more than 50% higher than the pre-recession level—before halving by 2016 and then climbing steadily to its latest peak after 2020. Concurrently, inflation has fluctuated sharply since the onset of the Great Recession, being almost twice as volatile, even excluding the post-Pandemic sample. We highlight that changes in price flexibility shape inflation dynamics, so that similar inflationary shocks—such as exchange rate fluctuations and commodity price changes—may lead to very different inflation outcomes depending on the price-flexibility regime in place. Failure to recognize such state dependence may help explain why the Bank of England has frequently struggled to meet its 2% inflation target in the last two decades.

We establish that regime shifts in price flexibility are key to understand inflation dynamics. The half-life of inflation is 50% higher during periods of relatively low flexibility. Otherwise, inflation tends to be more volatile, less persistent, and is typically higher when price flexibility is relatively high. We show that both Bank of England and broader market participants fail to account for such state dependence when projecting future inflation. Inflation forecasts are generally unbiased when aggregate price flexibility is low or average, yet a significant negative bias emerges during periods of high price flexibility, even excluding the post-Pandemic sample. This bias is not only statistically significant, but also economically important. We conduct a counterfactual experiment in which the Bank of England's forecasts from the last quarter of 2020 are adjusted for the high-flexibility bias. Our findings indicate that the return of inflation to target should have been reassessed, offering a cautionary perspective on the transitory nature inflation deviations from target, even in the early stages of the post-Pandemic inflation surge.

Our work has important implications for monetary authorities aiming to stabilize inflation. We show that in periods of relatively low price flexibility, inflationary shocks are likely to dissipate slowly, while the same shock would result in a larger inflation response—but also revert more quickly—under high price flexibility. A failure to recognize that can help explain why the Bank of England, along with other Central Banks, was caught off guard by the rapid inflation surge following COVID-19, and equally surprised by the swift decline beginning in the second half of 2023. Such state dependence is likely to influence the trade-off Central Banks face between stabilizing output and inflation. While this insight naturally emerges in state-dependent models of price setting, it is only minimally incorporated into Central Bank practices and communications.

We find significant non-linearities in inflation dynamics, even during a period of relative stability—apart from the post-Pandemic inflation surge. These non-linearities are likely to become even more pronounced in a more uncertain macroeconomic environment, particularly if the era of the Great Moderation gives way to a period of heightened inflation volatility, driven by adverse supply-side developments—deglobalization, reshoring, lengthening of the supply chains etc.—that place greater strain on Central Banks' policy trade-offs. In such conditions, inflation risk becomes more difficult to manage, as both the persistence and magnitude of inflation fluctuations grow increasingly unpredictable. Accurately characterizing these elements is therefore essential, as it provides a key input for assessing inflation risk and guiding the central bank's ability to navigate an increasingly complex policy landscape.

**Related literature** Our work relates to numerous studies examining the link between micro price changes and aggregate inflation.[2] Closest to our analysis is Berger and Vavra (2018), who document time variation in price flexibility. Karadi and Reiff (2019) and Cavallo et al. (2024) emphasize how large shocks significantly alter price flexibility, affecting aggregate inflation (see also Ascari and Haber, 2022). Relatedly, Karadi et al. (2024) and Ghassibe and Nakov (2025)

[2] See, among others, Bils and Klenow (2004), Alvarez et al. (2006), Klenow and Kryvtsov (2008), Nakamura and Steinsson (2008), Gagnon (2009), Costain and Nakov (2011), Midrigan (2011), Alvarez and Lippi (2014), Berardi et al. (2015), Nakamura et al. (2018), Carvalho and Kryvtsov (2021) and Gautier et al. (2024).

show that price flexibility depends on the type of shocks. Thus, time-varying price flexibility may reflect shifts in the size and composition of shocks hitting the economy. Regardless of its drivers, our key contribution is to demonstrate that incorporating time variation in price flexibility improves our understanding of inflation dynamics. We show that inflation is less persistent and more volatile when price flexibility is relatively high, and failing to account for this leads to systematic prediction biases. In line with Chu et al. (2018), who investigate the predictive power contained in the distribution of price changes, we demonstrate that price flexibility—capturing key information in micro price data—provides valuable insights for inflation projections.

Our study also builds on research that models the relationship between the distribution of price changes and price flexibility (see, e.g., Alvarez et al., 2016, Midrigan, 2011, Vavra, 2014). These models typically assume specific shocks to price-setting units, whereas our approach remains agnostic, avoiding strong assumptions about price adjustment mechanisms *ex ante*. Our estimates of the time-varying hazard function reveal both (i) a nonzero probability of price adjustment—even for minimal deviations from optimal prices—consistent with time-dependent pricing, and (ii) U-shaped hazard functions— consistent with state-dependent pricing.[3]

We further contribute to empirical work using UK consumer price microdata. Bunn and Ellis (2012) were among the first to examine price-setting frequency and hazard functions using data from the Office for National Statistics (ONS), while Dixon et al. (2020) focus on the Great Recession's impact on pricing (see also Dixon and Tian, 2017). These studies attribute little importance to endogenous macroeconomic effects on pricing, while our evidence points to a certain prominence of state dependence in price setting, and more so in the sample that is not accounted for in their analysis, during which the extensive margin of price adjustment overcomes the intensive one in the contribution to price flexibility. In fact, the novelty of the approach rests on tracking time changes in both margins of adjustment, rather than focusing on their average relative

[3]Evidence supporting the coexistence of time and state dependence is extensive; see, e.g., Nakamura and Steinsson (2010b), as well as Lein (2010), Carlsson and Skans (2012), and Dixon and Grimme (2022) using firm level data.

importance.

**Structure** The rest of the paper is organized as follows. Section 2 discusses the data and provides some motivating preliminary analysis. Section 3 reviews the generalized $Ss$ model and takes it to the data. Section 4 examines time variation in price flexibility, as well as the relative importance of adjustments along the intensive and the extensive margin over time. Section 5 discusses the implications of state dependence in price flexibility for inflation dynamics and forecasting. Section 6 concludes.

# 2 Microdata on consumer prices

We use ONS microdata underpinning the UK CPI for the period from 1996:M2 to 2024:M8. Prices are collected on a monthly basis for more than $700$ categories of goods and services, and are published with a one-month lag. The data exclude centrally collected price quotes. On our end, we discard price quotes that have not been validated by the system or accepted by the ONS, as a preliminary step. The overall number of price quotes accounts for approximately 60% of those included in the CPI.

Each price quote is classified by region, location, outlet, and item. Due to a confidentiality agreement between the ONS and individual shops, only the region, outlet, and item classifications are published. As a result, some price quotes may not be uniquely identified. This typically occurs when the ONS samples the same item in multiple locations within the same region for outlets that are part of a chain. To ensure that price trajectories can be uniquely identified, we use 'base prices', defined as the January price for each item under consideration.[4] Even after conditioning on base prices, a small portion of price trajectories (about $0.6\%$ on average) remain non-uniquely identified, and we opt to discard them.[5] Individual price quotes are weighted according to the

[4]Base prices are typically used to normalize price quotes, calculate price indices, and adjust for changes in the quality and/or quantity of a given good.

[5]Due to particularly low coverage, Housing and Housing Services (COICOP 4) and Education (COICOP 10) are also excluded from the sample. Additionally, price changes exceed-

ONS *stratification* weights (see Chapter 7 of the ONS CPI Manual, ONS, 2019). A provides further details on the dataset's construction.

After processing price quotes as described above, it is crucial to distinguish between regular and temporary price changes, such as *sales*, which behave in a significantly different manner, relative to regular prices (see Eichenbaum et al., 2011, Kehoe and Midrigan, 2015). Sales are a common tool for temporary price cuts. Kryvtsov and Vincent (2021) document their countercyclicality, showing that firms use sales to respond to temporary negative shocks. Consequently, excluding sales removes large price drops at the onset of a downturn and the subsequent sharp increases when the economy recovers. However, sales often reflect marketing strategies rather than fundamental price-setting behavior. Guimaraes and Sheedy (2011) highlight that even when firms can freely adjust prices, sales act as strategic substitutes—firms discount less when competitors already do. As a result, sales do not necessarily translate into meaningful aggregate price flexibility, and monetary policy continues to exert strong real effects. Such disconnect suggests that including sales may overstate the economy's actual responsiveness to shocks. By focusing on regular prices, we aim to capture the persistent component of price setting that drives inflation dynamics.

To this end, we first exclude all price quotes marked with a sales indicator by the ONS.[6] As a second step, we apply a symmetric V-shaped filter, as defined by Nakamura and Steinsson (2010b), to detect implicit sales. According to this filter, the sale price of item $i$ at time $t$, $P^s_{i,t}$, is identified as follows: (i) it is lower than the previous period's price (i.e., $P^s_{i,t} < P_{i,t-1}$), and (ii) the price in the following period reverts to the previous price (i.e., $P_{i,t+1} = P_{i,t-1}$). A recovery price $P^r_{i,t}$ is instead identified by the following criteria: (i) it is higher than the previous period's price (i.e., $P^r_{i,t} > P_{i,t-1}$) and (ii) it equals the price two periods earlier (i.e., $P^r_{i,t} = P_{i,t-2}$). Once a price quote is identified as a sale or recovery price, we discard it from the sample.[7]

ing $300\%$, which are likely due to measurement errors, are removed. These occur very rarely ($< 0.02\%$).

[6] For a price to be classified as part of a sale, the ONS requires that the discount be available to all potential customers—excluding quantity discounts and membership deals—and that it represents only a temporary or end-of-season price reduction. This definition excludes clearance sales of products that have reached the end of their life cycle.

[7] See also Nakamura and Steinsson (2008) and Vavra (2014). An alternative approach, em-

Item substitutions present another challenge when identifying price trajectories, as they introduce ambiguity in determining whether a price change reflects a quality adjustment or a pure price change. Product substitutions occur when an item in the sample is discontinued at an outlet, and the ONS selects a replacement item to continue tracking prices. Consequently, price changes following product turnovers may either reflect uncaptured quality differences (Bils, 2009) or simply represent an opportunity to reset prices, independent of the underlying sources of price rigidity, as argued by Nakamura and Steinsson (2008). In line with previous studies, we terminate a price trajectory whenever it encounters a substitution flag (see, e.g., Berardi et al., 2015, Berger and Vavra, 2018, Kryvtsov and Vincent, 2021).

After these steps, we define the price change of item $i$ at time $t$ as $\Delta p_{i,t} = \log(P_{i,t}/P_{i,t-1})$.[8]

## 2.1 Stylized facts

This section presents some key facts about the behavior of the ONS microdata, and their implications for inflation dynamics. Aggregate inflation may rise as a result of larger average price increases or because of more frequent price changes at the micro level. Specifically, inflation can be written as the product of the frequency of adjustment ($fr_t$)—defined as the share of prices being adjusted in every month—and the average price change in every month ($\Delta p_t$)—for those goods and services changing prices in any given month:

$$\pi_t = fr_t \times \Delta p_t. \tag{1}$$

The frequency is computed as $\sum_i \omega_{i,t} \mathbb{1}_{\{\Delta p_{i,t} \neq 0\}}$, with $\omega_{i,t}$ denoting the CPI weight associated with good $i$ at time $t$, and $\mathbb{1}_{\{\Delta p_{i,t} \neq 0\}} = 1$ if $\Delta p_{i,t} \neq 0$, and zero otherwise. The average price, instead, is computed as $fr_t^{-1} \sum_i \omega_{i,t} \mathbb{1}_{\{\Delta p_{i,t} \neq 0\}} \Delta p_{i,t}$.

ployed by Klenow and Kryvtsov (2008), replaces sale prices with the last regular price until a new regular price is observed. Our results are robust to this approach.

[8]We also compute price changes as $\Delta p_{i,t} = 2\frac{P_{i,t} - P_{i,t-1}}{P_{i,t} + P_{i,t-1}}$. This alternative measure is bounded and less sensitive to outliers. The results using this alternative measure of price change remain virtually unchanged from those we report.

As month-on-month (MoM) inflation variations exhibit pronounced seasonality and significant high-frequency noise, it is customary—including in Central Bank communication—to focus on year-on-year (YoY) inflation, i.e. a 12-month moving average of annualized MoM inflation. While this approach helps mitigate seasonality and noise in the data, the backward-looking moving average (MA) filter introduces a phase shift in inflation dynamics, causing a slight delay in the observed peaks and troughs (see, e.g., Harvey, 1993, p.189-198). To enhance comparability with YoY inflation, we apply the same filter to all statistics derived from microdata, as they also exhibit substantial high-frequency variation and seasonality.

The top panels of Figure 1 report $\widetilde{fr}_t$ and $\widetilde{\Delta p}_t$, respectively (for a generic variable $\kappa_t$, $\widetilde{\kappa}_t = \frac{1}{12}\sum_{j=0}^{11} \kappa_{t-j}$ represents its 12-month moving-average transformation). As expected, the average price change tracks CPI inflation closely. The frequency of adjustment also exhibits notable fluctuations, including a temporary increase during and immediately after the Great Recession (consistent with the findings of Dixon et al., 2020), followed by a downward trend starting with the inflation decline in 2012. During this period, it dropped well below its previous sample average before undergoing a significant reversal after the COVID-19 Pandemic.[9] The frequency of price adjustments increased significantly during the most recent inflationary episode. However, even at its peak, only about 14% of prices were adjusted each month, despite inflation exceeding 10%. This rate is considerably lower than what was documented during high-inflation periods in the US (see Nakamura et al., 2018). In fact, despite exceptional inflationary pressures, the recent surge in the frequency of price adjustment in the UK remains lower than during the Great Recession, when inflation was well below the level observed in the most recent years.

The way inflation rises—whether through larger price increases or more frequent price adjustments—has important implications for inflation dynamics.

[9] The average frequency of price adjustment prior to its drop is slightly lower than the estimates reported by previous studies on UK price microdata conducted over roughly the same time span. This reflects the fact that we exclude from our sample both sales and utility prices (COICOP 4), the latter being a particularly volatile component of the CPI index. Specifically, Bunn and Ellis (2012) include utility prices and sales, while Dixon and Tian (2017) and Dixon et al. (2020) include sales.

Figure 1: FREQUENCY, AVERAGE PRICE CHANGES, AND DISPERSION

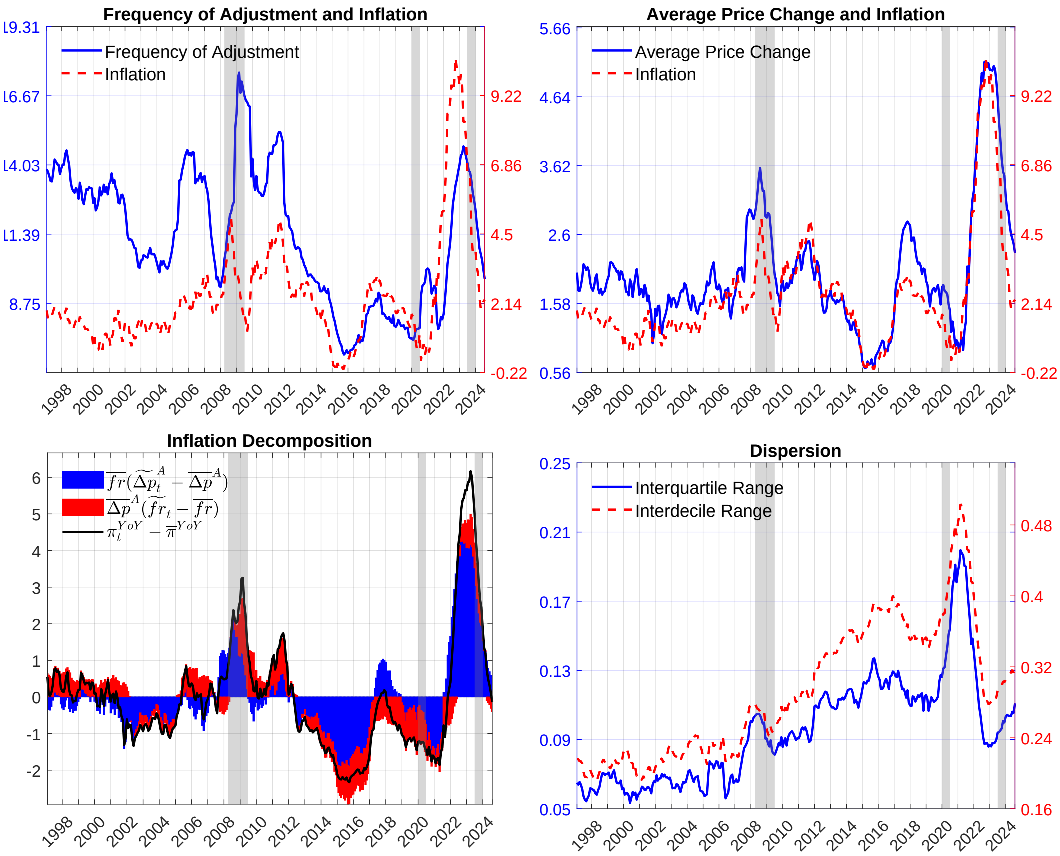


Notes: The frequency of price adjustment, $fr_t$ , measures the share of prices being adjusted in every month, and is computed as $\sum_i \omega_{i,t} \mathbb{1}_{\{\Delta p_{i,t} \neq 0\}}$, where $\omega_{i,t}$ denotes the CPI weight associated to good $i$ at time $t$, and $\mathbb{1}_{\{\Delta p_{i,t} \neq 0\}} = 1$ if $\Delta p_{i,t} \neq 0$ and zero otherwise. The average price, instead, is denoted by $\Delta p_t$ and is computed as $fr_t^{-1} \sum_i \omega_{i,t} \mathbb{1}_{\{\Delta p_{i,t} \neq 0\}} \Delta p_{i,t}$. All series are reported in percentage terms. In the bottom-left panel of the figure we decompose the deviation of inflation from its sample average between the contribution of the variation in the average price change (holding the frequency fixed) and that of the variation in the frequency of adjustment (holding the average price change fixed). Specifically, since $\pi_t = fr_t \Delta p_t$, one can take the following decomposition: $\pi_t^{YoY} - \overline{\pi}^{YoY} = \overline{fr}(\widetilde{\Delta p}_t - \overline{\Delta p}) + \overline{\Delta p}(\widetilde{fr}_t - \overline{fr}) + (\widetilde{\Delta p}_t - \overline{\Delta p})(\widetilde{fr}_t - \overline{fr})$. The inflation rate graphed in the upper panels of the figure is the official CPI inflation rate published by the ONS. The shaded vertical bands denote the duration of recessionary episodes.

For instance, when firms respond to large nominal cost shocks by adjusting prices more frequently, inflation initially rises more sharply but subsides more quickly than if the frequency had remained unchanged. Taking a first-order approximation of Equation (1), and expressing it in terms of YoY inflation, we obtain

$$\pi_t^{YoY} - \overline{\pi}^{YoY} \approx \overline{fr}(\widetilde{\Delta p}_t^A - \overline{\Delta p}^A) + \overline{\Delta p}^A(\widetilde{fr}_t - \overline{fr}), \tag{2}$$

where $\Delta p_t^A \equiv 12\Delta p_t$ represents the annualized average price change, and the bar denotes sample averages. Therefore, the first term on the right side of the equation reflects the contribution to inflation of variation in the average price change, while the second term captures the contribution of changes in the frequency of adjustment.[10] Notably, only about half of inflation variability is explained by the average price change, the remaining part being accounted for by changes in the frequency (either directly or indirectly, through its positive comovement with the average price change). A relatively large contribution of the frequency is particularly evident in the post-Great Recession sample. The post-Pandemic inflation surge is, instead, attributed mostly to the average price-change component.[11] However, changes in the frequency of price adjustment, coupled with the higher order term, still account for approximately one-third of inflation at the peak.

The bottom-right panel of the figure plots different measures of dispersion of the distribution of (non-zero) price changes. Both the interquantile and the interdecile range display a large increase in the aftermath of the Great Recession, to then skyrocket and abruptly decline in coincidence with the onset and the attenuation of the COVID-19 emergency, respectively.[12] A key observation from the graphical analysis is that the dispersion of price changes and the frequency of adjustment tend to move in opposite directions. For example, in the first decade of the sample, the average frequency of price adjustment is roughly 50% higher, whereas the average interquartile range of price changes is twice as large in the last decade, as compared with the first one. Similarly,

[10]Higher-order terms account for variation in inflation due to the covariation between the frequency of price adjustment and the average price change. A detailed derivation of this decomposition is available in B.

[11]This is in line with evidence of Montag and Villar (2022) for the US and Dedola et al. (2024) for the Euro Area.

[12]Also the standard deviation displays a similar pattern. However, this measure is often influenced by outliers. This type of problem does not plague the interquantile and the interdecile range.

in the post-COVID-19 period, the peak in the dispersion of price changes occurred alongside a very low frequency of adjustment, just before the inflation surge in the second half of 2021. Conversely, the peak in the frequency of price adjustment in late 2022 coincided with a very low level of dispersion in price changes. These opposite movements suggest major shifts between a regime of relatively small—yet, frequent—price changes, and one of much larger—yet, more infrequent—adjustments in prices.[13]

As stressed by Vavra (2014) and Berger and Vavra (2018), the specific properties of the frequency and the dispersion of price changes, as well as their joint dynamics, are key to unveiling the endogenous and exogenous determinants of price adjustment, and to tracking time variation in the pass-through of nominal shocks to inflation. The remainder of the analysis will be devoted to examine these aspects.

# 3 Framing the analysis

To explore the origins of time variation in the moments of the price-change distribution and how they may reflect different price adjustment protocols, we draw on the generalized $Ss$ setting developed by Caballero and Engel (2007). This model has two clear advantages that make it particularly indicated to discipline our data. First, it is consistent with lumpy and infrequent price adjustments—which are typically seen as distinctive traits of price setting—along with encompassing several pricing protocols,[14] without necessarily being constrained to match any of those specifically. Second, as we allow for time variation in the determinants of price adjustment, we can estimate the model over each cross

[13]In A we show that composition effects have no role in generating the facts presented in this section: here we compare the moments of the distribution of price changes with their counterparts obtained by averaging the corresponding moments of the price quotes, for each of the 25 COICOP group categories.

[14]To focus on two somewhat extreme examples, the generalized $Ss$ model can account for both price setting à la Calvo (1983)—where firms are selected to adjust prices at random and price flexibility is fully determined by the frequency of adjustment—as well as for schemes à la Caplin and Spulber (1987)—where adjusting firms change prices by such large amounts that the aggregate price is fully flexible, regardless of the frequency of adjustment. In fact, Berger and Vavra (2018) show how this empirical setting provides a good fit to the data generated by different structural models (e.g., Golosov and Lucas, 2007, Nakamura and Steinsson, 2010a).

section of price microdata, matching different price-setting statistics.

This framework assumes that, due to price rigidities, the log of firm $i$'s actual price may deviate from the log of its target or reset price, denoted as $p^*_{it}$. As a result, the observed distribution of price changes arises from the interaction between the distribution of price deviations from their optimal (reset) values and a hazard function that determines the probability of a price change given the deviation.

To formalize this, we define the price gap as $x_{it} \equiv p_{it-1} - p^*_{it}$, where a positive (negative) price gap indicates a falling (rising) price upon adjustment. A price is adjusted when the associated price gap becomes sufficiently large. After the adjustment, $p_{it} = p^*_{it}$. Therefore, when a price change is enacted, the change in price reflects the original misalignment in prices, i.e. $\Delta p_{it} = -x_{it}$.[15] If $l_{it}$ represents the number of periods since the last price change, then the adjustment reflects the accumulated shocks: $\Delta p_{it} = \sum_{j=0}^{l_{it}} \Delta p^*_{it-j}$, where $\Delta p^*_{it} = \mu_t + \upsilon_{it}$, with $\mu_t$ representing a shock to nominal demand and $\upsilon_{it}$ an idiosyncratic shock.

Caballero and Engel (2007) assume the presence of $iid$ idiosyncratic shocks to the adjustment cost, which gives rise to price stickiness. By integrating over the possible realizations of these shocks, the adjustment hazard function, $\Lambda_t(x)$, is obtained. This function represents the probability—at time $t$—that a firm will adjust its price before knowing the current adjustment cost draw, given that it would have adjusted by $x$ in the absence of adjustment costs (i.e., if the adjustment cost draw was zero). Denoting with $f_t(x)$ the cross-sectional distribution of price gaps immediately before an adjustment takes place at time $t$, aggregate inflation can be recovered as

$$\pi_t = -\int x \Lambda_t(x) f_t(x)\, dx. \tag{3}$$

Notice that the Calvo pricing protocol implies the same hazard across $x$'s (i.e., $\Lambda'_t(x) = \frac{\partial \Lambda'_t}{\partial x} = 0$ and $\Lambda_t(x) = \Lambda_t > 0,\ \forall x$). Conversely, upward-sloping hazard functions provide direct evidence of state dependence in price setting and imply that prices further away from their reset values are more likely to change.

[15]This assumption is empirically validated in the analysis of Karadi et al. (2023) using scanner price data, where the optimal reset price is proxied by the price of close competitors.

## 3.1 Taking the model to the data

To take the model to the data, we need to specify a functional form for the distribution of price gaps and the hazard function. We postulate that the distribution of price gaps at time $t$, $f_t(x)$, can be accounted for by the Asymmetric Power Distribution (APD henceforth; see Komunjer, 2007). The probability density function of an APD random variable is defined as

$$f_t(x) = \begin{cases} \frac{\delta(\varrho_t,\nu_t)^{1/\nu_t}}{\psi_t \Gamma(1+1/\nu_t)} \exp\left[-\frac{\delta(\varrho_t,\nu_t)}{\varrho_t^{\nu_t}} \left|\frac{x-\theta_t}{\psi_t}\right|^{\nu_t}\right] & \text{if } x \leq \theta_t \\ \frac{\delta(\varrho_t,\nu_t)^{1/\nu_t}}{\psi_t \Gamma(1+1/\nu_t)} \exp\left[-\frac{\delta(\varrho_t,\nu_t)}{(1-\varrho_t)^{\nu_t}} \left|\frac{x-\theta_t}{\psi_t}\right|^{\nu_t}\right] & \text{if } x > \theta_t \end{cases}, \tag{4}$$

with $\delta(\varrho_t, \nu_t) = \frac{2\varrho_t^{\nu_t}(1-\varrho_t)^{\nu_t}}{\varrho_t^{\nu_t}+(1-\varrho_t)^{\nu_t}}$. The parameters $\theta_t$ and $\psi_t > 0$ capture the location and the scale of the distribution, whereas $0 < \varrho_t < 1$ accounts for the degree of asymmetry. Last, the parameter $\nu_t > 0$ measures the degree of tail decay: for $\infty > \nu_t \geq 2$ the distribution is characterized by short tails, whereas it features fat tails when $2 > \nu_t > 0$. This functional form nests a number of standard specifications, such as the Normal ($\nu_t = 2$), Laplace ($\nu_t = 1$) and Uniform ($\nu_t \to \infty$). Moreover, it can capture intermediate cases between the Normal and the Laplace distribution, consistent with the steady-state distribution of price changes according to Alvarez et al. (2016).

We then assume that the hazard function can be characterized by an asymmetric quadratic function:

$$\Lambda_t(x) = \min\left\{a_t + b_t x^2 \mathbb{1}_{\{x>0\}} + c_t x^2 \mathbb{1}_{\{x<0\}}, 1\right\}, \tag{5}$$

where $\mathbb{1}_{\{z\}}$ is an indicator function taking value $1$ when condition $z$ is verified, and zero otherwise. This parsimonious specification nests the Calvo pricing protocol for $b_t = c_t = 0$, while allowing for asymmetric costs of adjustment, which has recently been supported by Luo and Villar (2021) and Karadi et al. (2023).[16]

[16]We have also checked that our results are robust to plausible variations to the specification of these functional forms. Using a mixture of two Normal distributions, or a mixture of a Laplace and a Normal distribution for the price gap, as well as an asymmetric inverted normal function for the hazard function, delivers results that are qualitatively similar to those reported in the next section.

Unlike Berger and Vavra (2018), we allow for asymmetry in both the hazard function and the distribution of price gaps. In A, we show that the distribution of price changes exhibits significant asymmetry, with skewness varying notably over time. Moreover, we document substantial differences in the frequency of price adjustments, average price changes, and dispersion between positive and negative price changes. By allowing asymmetry in both the price gap distribution and the hazard function, we avoid imposing restrictive assumptions about the sources of non-linearity.

## 3.2 Estimation and identification

Alvarez et al. (2023) highlight that the moments of the price gap distribution, together with the frequency of price changes, provide enough information to identify the distribution of price gaps and the hazard function. Therefore, given the parametric specifications of $f_t(x)$ and $\Lambda_t(x)$, we estimate seven parameters for each cross section of price microdata, so as to match the following moments of the distribution of price changes: mean, median, standard deviation, interquartile range, difference between the 90th and 10th quantile of the distribution, as well as (quantile-based) skewness and kurtosis (as in Groeneveld and Meeden, 1984).[17] We also match the frequency and the average size of prices movements, conditioning on positive and negative price changes. Last, we match the observed rate of inflation. The estimates are obtained by simulated minimum distance, using the identity matrix to weight different moments.[18] C reports the estimates of the model.

While Alvarez et al. (2023) prove that $\Lambda_t(x)$ and $f_t(x)$ are fully encoded in the distribution of price changes and $fr_t$ for the specific case of symmetric functional forms, it remains an open question whether these results extend to a more general framework with a generic (asymmetric) distribution of price gaps and

[17]We match quantilic moments, as the 3rd and 4th moments of the cross-sectional distribution are quite sensitive to outliers.

[18]Altonji and Segal (1996) highlight that matching the unweighted distance between moments often performs better in small samples, as compared with using optimal weights. The moments of the simulated distribution are estimated by drawing $100,000$ price quotes. We use the Genetic Algorithm to minimize the quadratic distance between data moments and simulated moments, so as avoid ending up in local minima (see, e.g., Dorsey and Mayer, 1995).

an asymmetric hazard function such as the ones we consider. To address this, D reports a series of exercises that highlight how close we come to identify the shape of the price gap distribution and the hazard function. As a first exercise, we evaluate the systematic impact of each of the estimated parameters on the moments that we are matching. To this end, we vary the parameters of $f_t(x)$ and $\Lambda_t(x)$—one at a time, while keeping all other coefficients at their baseline estimates—and examine their impact on key moments of the price change distribution, as well as on the resulting rate of inflation. All in all, marginal changes in the parameters typically correspond to large variation in the moments we match, indicating the latter carry valuable information to identify the parameter of interest. We then ask whether moment matching allows us to appropriately identify/distinguish the shape of the price gap distribution from that of the hazard function. To see this, we simulate price-change data from the model, under different parameterizations, and then contrast the true price gap distribution and the hazard function to their estimated counterparts. The overall discrepancy is minimal, and the model does a good job at separately identifying the parameters of $f_t(x)$ and $\Lambda_t(x)$.

# 4 Inspecting price setting in a time-varying environment

The generalized $Ss$ model emphasize the importance of tracking changes in the distribution of price gaps and the hazard function. Caballero and Engel (2007) show that, within their accounting framework, one can derive a measure of aggregate price flexibility. The latter measures the extent to which a marginal shift in the price gap distribution (such as one stemming from a common macroeconomic shock that equally affects all price setters) translates into contemporaneous inflation:

$$\mathcal{F}_t = \lim_{\mu_t \to 0} \frac{\partial \pi_t}{\partial \mu_t} = \underbrace{\int \Lambda_t(x) f_t(x)\, dx}_{\text{Intensive Margin}} + \underbrace{\int x \Lambda_t'(x) f_t(x)\, dx}_{\text{Extensive Margin}}. \tag{6}$$

In turn, aggregate price flexibility can naturally be decomposed into an intensive and an extensive margin component.[19] The intensive margin ($Int$) measures the average frequency of adjustment, and accounts for the part of inflation that reflects price adjustments that would have happened even in the absence of the nominal shock. The extensive margin ($Ext$) accounts for the additional inflation contribution of firms whose decision to adjust is either triggered or canceled by the nominal shock. Therefore, it comprises both firms that would have kept their price constant and instead change it, as well as firms that would have adjusted their price but choose not to do it.[20] Therefore, when the hazard function is not flat, inflation's response to shocks depends not only on the frequency of price changes in each period but also on which prices adjust and by how much. If the prices that do change are those with the most significant misalignments, the aggregate price level remains highly flexible, possibly exhibiting high volatility, and responding strongly to aggregate shocks even if relatively few firms adjust their prices.

Figure 2 reports the estimated index of price flexibility (left panel), as well as its decomposition into the intensive and the extensive margin of price adjustment (right panel). Aggregate price flexibility is, on average, nearly twice as high as what would be implied by the frequency of price changes alone. In fact, our estimates of the hazard function are far from flat, providing strong evidence that price setting is state-dependent.[21] This corroborates the findings of Karadi et al. (2023), who—using scanner price data and a different methodology—also estimate a U-shaped hazard function for the Euro Area.

$\mathcal{F}_t$ exhibits substantial variation over time, with the intensive margin accounting for approximately 50% of its variance, the extensive margin contributing by 33%, and the remaining share due to the positive comovement between

[19]It is also important to stress that, since $\mathcal{F}_t$ is simply derived from the accounting identity (3), its validity as a measure of aggregate flexibility does not require that we take a stand on a specific model of price setting.

[20]In this respect, it is useful to recall that, being characterized by a constant hazard function, Calvo price setting implicitly assumes that the extensive margin is null.

[21]At the same time, we cannot rule out the presence of time dependence in price setting, as the probability of adjusting prices for very small misalignments is not negligible, i.e., $\Lambda_t(0) \neq 0$ for all $t$. In fact, shifts in the lower bound of the hazard function are a key driver of the overall variation in price flexibility.

Figure 2: PRICE FLEXIBILITY AND DIFFERENT MARGINS OF PRICE ADJUSTMENT

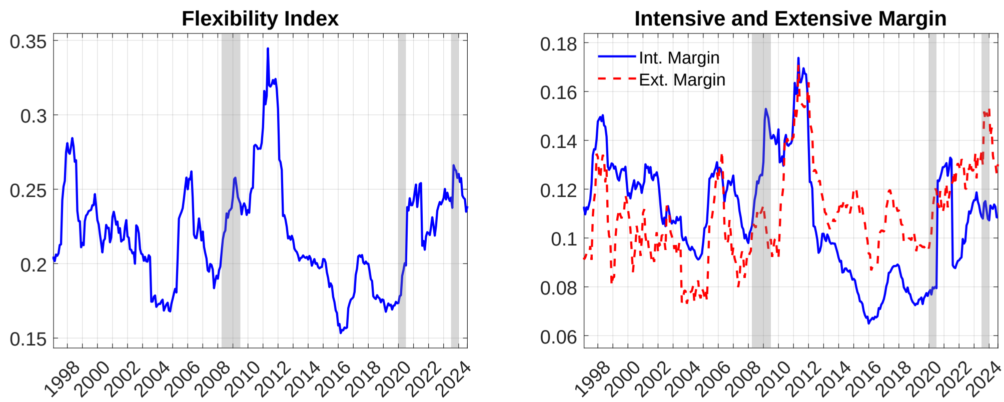


Notes: The left panel reports the estimated index of price flexibility, which is decomposed in the right panel between the intensive and the extensive margin of price adjustment. The shaded vertical bands indicate the duration of recessionary episodes. To aid the interpretability of the results, we report a 12-month backward-looking moving average of the estimated aggregate price flexibility and its components. This approach smooths out high-frequency variability and seasonality in the estimates and, most importantly, aligns them with YoY inflation.

the two components. Aggregate price flexibility rises significantly during the Great Recession, with secondary peaks observed in the subsequent recessions. After the Great Recession, both the intensive and extensive margins contract, though the decline in the former is much more abrupt. During this period of contraction in aggregate price flexibility, the extensive margin becomes the dominant force. Even after both margins revert in 2016, the extensive margin remains the predominant driver.

What explains the significant role of the extensive margin in price adjustment over the past two decades? Figure 3 examines changes in the price gap distribution and the hazard function in two cases of notable shifts in aggregate price flexibility. Panel (a) compares 2011 to 2016, highlighting the post-Great Recession decline in price flexibility. The hazard function shifts downward, indicating that only substantial price misalignments are likely to be corrected. As a result, the probability of price adjustment declines, increasing the relative importance of the extensive margin and explaining the observed rise in price dispersion over that period (see Figure 1), despite the similarity in price gap distributions. A lower hazard function is consistent with rising market power, which reduces the cost of deviating from optimal prices. In fact, studies on UK markups by DeLoecker and Eeckhout (2018) and Bell and Tomlinson (2018) suggest that while markups have been stable in the decade 1996 to 2007, they display a substantial rise thereafter.

Panel (b) examines the post-COVID inflation surge, during which large shocks caused a pronounced shift in the price gap distribution. With prices substantially below their optimal levels, a high probability of price adjustment—and consequently, a stronger inflation response to shocks—emerged without necessarily requiring significant changes in the hazard function. As a result, the increase in price flexibility contributed to substantial inflationary pressures, and inflation volatility, over this period. This aligns with the insights of Cavallo et al. (2024), who highlight how large shocks and significant price misalignments amplify the pass-through of aggregate shocks, as observed in the post-COVID period.

More generally, movements in price flexibility do not appear to occur randomly: $\mathcal{F}_t$ goes from being positively correlated with output growth in the decade preceding the Great Recession (0.35), to comoving negatively thereafter (-0.14, considering 2019 as the endpoint of the second subsample). As for the correlation with the rate of inflation, it is generally positive, particularly in the 2009-2020 time interval (0.79). It is worth emphasizing how changes in these correlations over the two subsamples are, again, coherent with a shift from a setting where most of the price changes are predetermined to one where the extensive margin gains relevance,[22] representing the main contributor to price flexibility and, consequently, inflation volatility becomes particularly high.

[22]Henkel et al. (2023) report a similar view for selected Eurozone countries, indicating that state dependence in price setting has considerably added to the COVID-19 shock.

Figure 3: EXAMPLES OF SHIFTS IN PRICE GAP DISTRIBUTION AND HAZARD FUNCTIONS

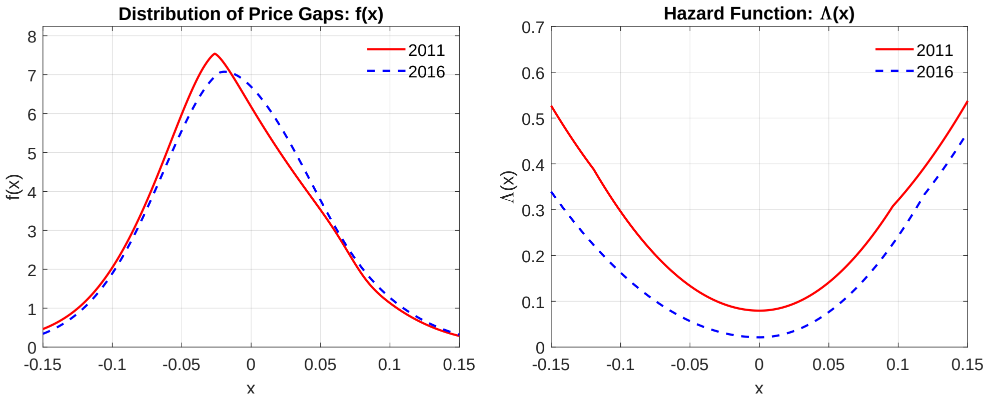


(a) 2011 vs 2016

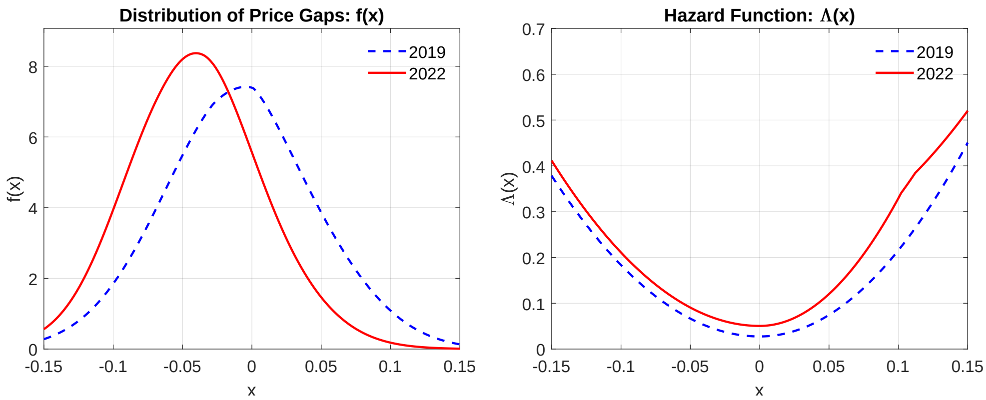


(b) 2019 vs. 2022

Notes: Comparison of the estimated price gap distribution (left panel) and hazard function (right panel) for two cases of shifts in aggregate price flexibility: high-flexibility regime (red), low-flexibility regime (blue). These functions are estimated monthly across different years, with the reported values representing the average estimate for each respective year. Panel (a) illustrates the decline in the hazard function following the Great Recession (comparison of 2011 vs. 2016), while Panel (b) depicts the post-Pandemic inflation shock (comparison of 2019 vs. 2022).

**Price flexibility and money non-neutrality** Our analysis highlights a great deal of variation in aggregate price flexibility. However, we acknowledge the index we use is not the only available measure of price flexibility. Alvarez et al. (2016) put forward a sufficient statistic for money non-neutrality, intended as the cumulative output response to a nominal shock. They prove that, in a variety of sticky-price models, this is proportional to the steady-state ratio of the kurtosis of the size distribution of price changes to the frequency of price adjustments.

Figure 4: COMPARISON WITH ALVAREZ ET AL. (2016)

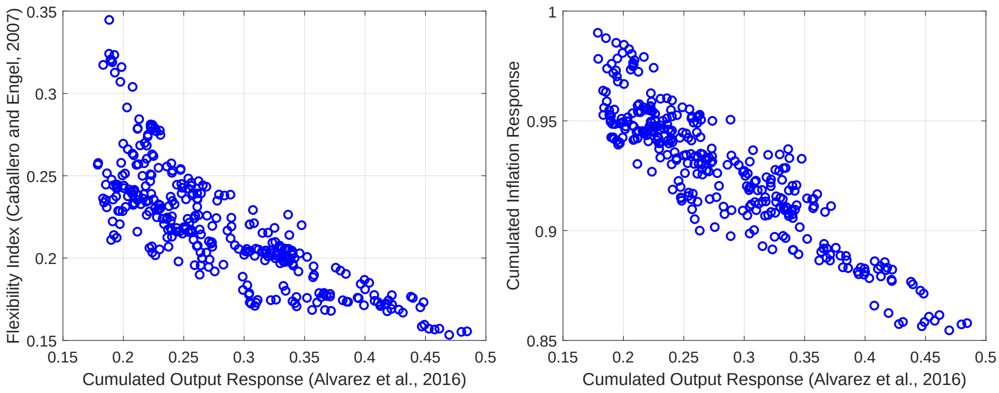


Notes: The left panel of the figure reports a scatter plot of the cumulated output response to a monetary policy shock, as computed by Alvarez et al. (2016), against the index of price flexibility, as computed by Caballero and Engel (2007). The right panel, instead, features a scatter plot of the cumulated output response to a monetary policy shock against the cumulated inflation response to a one-off 1% nominal shock, where we cumulate the inflation response over a 18-month period.

The left panel of Figure 4 proposes a direct comparison between $\mathcal{F}_t$ and $\frac{\text{Kur}_t(\Delta p_i)}{fr_t(\Delta p_i)}$, where the latter is obtained by computing in every month a quantilic version of the kurtosis of price changes, $\text{Kur}_t(\Delta p_i)$, estimating it for each month of the sample (see, e.g., Groeneveld and Meeden, 1984) and the frequency of price adjustment, $fr_t(\Delta p_i)$. A clear (convex) negative relationship emerges, despite the two statistics not being directly comparable, as one measures the *instantaneous* pass-through of nominal shocks to prices, while the other focuses on the *cumulative* impact of nominal shocks on output. In fact, it may well be the case that a shock exerts a relatively low impact on prices, taking a long time to be fully absorbed and leading to a large cumulative output response. To account for this, we compute the cumulative response of inflation over the 18 months following a one-off 1% nominal shock. The right panel of Figure 4 shows a striking (negative) correlation of our cumulative measure of price stickiness with the metric elaborated by Alvarez et al. (2016). This reinforces our confidence in the empirical framework we rely upon to track movements in the price gap distribution and the hazard function.

# 5 State dependence in inflation dynamics

Having established that price flexibility exhibits significant fluctuations throughout the sample under examination, a natural question arises: do these movements matter for our understanding of inflation dynamics? A straightforward exercise may help contextualize our analysis of the connection between price flexibility and inflation dynamics. To this end, we rely upon the $Ss$ model estimates to derive the response of inflation to an aggregate nominal shock across two distinct periods—one characterized by relatively strong and the other by relatively weak pass-through of nominal shocks to inflation, respectively.[23] Figure 5 illustrates that inflation is more responsive and less persistent during periods of relatively high price flexibility. In light of this, price flexibility likely holds valuable information for analyzing inflation dynamics. This insight arises naturally in environments characterized by state-dependent pricing. TWe now examine whether aggregate inflation exhibits non-linearities consistent with these properties, and discuss some key implications for the practice of inflation targeting.

## 5.1 Price flexibility and inflation dynamics

We seek to examine how inflation generally behaves in periods of relatively high and low flexibility. To this end, we employ a regime-switching autoregressive moving average model, where the transition across regimes is a smooth function of the degree of price flexibility. The STARMA(p,q) model is a generalization of the smooth transition autoregression model proposed by Granger

[23] As we only identify the price gap distribution at each point in time, we are not able to disentangle the contribution of the aggregate shock from that of idiosyncratic shocks. Therefore, for purely illustrative purposes, we choose an autoregressive specification for the first-moment shock. More details are available in E.

Figure 5: IMPULSE RESPONSES FROM THE $Ss$ MODEL

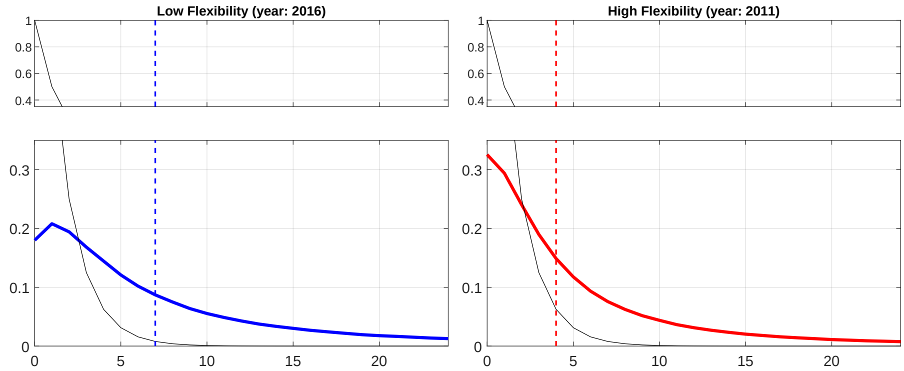


Notes: The graphs display the average inflation response to a $1\%$ aggregate nominal shock, $\mu_t$, in two periods of relatively low and high price flexibility. The shock is assumed to die out with a persistence component of $0.5$ and is depicted by the thin black line (with a negative sign). The left panel (low price flexibility) plots the average inflation response in 2011, while the right panel (high price flexibility) plots the average inflation response in 2016. In each of the two panels the vertical line indicates the half-life of the shock.

and Terasvirta (1993). Estimating a traditional ARMA(p,q) for each regime separately entails a certain disadvantage in that we may end up with relatively few observations in a given regime, which typically renders the estimates unstable and imprecise. By contrast, we can effectively rely upon more information by exploiting variation in the probability of being in a particular regime, so that estimation and inference for each regime are based on a larger set of observations (Auerbach and Gorodnichenko, 2012).[24]

We assume that inflation can be described by the following model:

$$\begin{aligned} \pi_t \;=\; & G\left(\widetilde{\mathcal{F}}_{t-1}, \gamma\right)\left(\phi_0^H + \sum_{j=1}^{p} \phi_j^H \pi_{t-j} + \varepsilon_t^H + \sum_{i=1}^{q} \theta_i^H \varepsilon_{t-i}^H\right) \\ & + \left[1 - G\left(\widetilde{\mathcal{F}}_{t-1}, \gamma\right)\right]\left(\phi_0^L + \sum_{j=1}^{p} \phi_j^L \pi_{t-j} + \varepsilon_t^L + \sum_{i=1}^{q} \theta_i^L \varepsilon_{t-i}^L\right), \end{aligned} \tag{7}$$

[24]Estimating the properties of a given regime by relying on the dynamics of inflation in a different regime would bias our results towards not finding any evidence of non-linearity. In light of this, the asymmetries we will be reporting in the remainder of this section acquire even more statistical relevance.

with $\varepsilon_t^i \sim N(0, \sigma_i^2)$ for $i = \{L, H\}$. Moreover, we set $G\left(\widetilde{\mathcal{F}}, \gamma\right) = (1 + e^{-\gamma\widetilde{\mathcal{F}}})^{-1}$, where $\widetilde{\mathcal{F}}$ denotes the normalized flexibility index and $\gamma$ is the speed of transition across regimes.[25] We allow for different degrees of inflation persistence across the two regimes, as captured by the regime-specific autoregressive and moving average coefficients, as well as for different volatilities of the innovations in either regime. The likelihood of the model can be easily computed by recasting the system in state space (see, e.g., Harvey, 1993, Ch.4). We use Monte Carlo Markov-chain methods developed in Chernozhukov and Hong (2003) for estimation and inference. The parameter estimates, as well as their standard errors, are directly computed from the generated chains (see F for further details).

Focusing on the post-1996 sample, we estimate the model by imposing that, in both regimes, the long-run inflation forecast is $2\%$, consistent with the Bank of England's mandate. The parameter $\gamma$ captures the speed at which we switch between classifying periods as high or low flexibility regimes, and its identification relies on non-linear moments. We estimate this parameter by selecting the value that maximizes the likelihood function. The estimated value of $\gamma$ implies that roughly $20\%$ of the observations are classified in the high-flexibility (low-flexibility) regime, defined by $G\left(\widetilde{\mathcal{F}}_{t-1}; \gamma\right) > 0.8$ ($G\left(\widetilde{\mathcal{F}}_{t-1}; \gamma\right) < 0.2$). The upper-left panel of Figure 6 reports $G\left(\widetilde{\mathcal{F}}_{t-1}; \gamma\right)$. This specification clearly identifies the 2009-2012 and post-2021 periods as characterized by high price flexibility, whereas the 2002-2005 and 2015-2016 periods are marked by low flexibility. Based on the Akaike criterion, we select $p = 2$ and $q = 1$.[26]

The bottom panels of Figure 6 present the impulse-response functions to a one-standard deviation shock to inflation in each of the two regimes, and compares them to the response from an equivalent linear model. Consistent with Figure 5, the inflation response is more muted and significantly more persistent during periods of relatively low price flexibility, with the half-life of the shock being nearly 50% longer, compared to periods of high flexibility. Furthermore,

[25] We compute a backward-looking MA(12) of the flexibility index to smooth out high-frequency variability and get rid of the seasonality in the data. Moreover, we lag the index by one month, in order to avoid potential endogeneity with respect to CPI inflation.

[26] Our key insights are not affected by the exact specification of the STARMA(p,q) model (see G). The results are also robust to plausible values of $\gamma$.

Figure 6: PRICE FLEXIBILITY AND INFLATION DYNAMICS

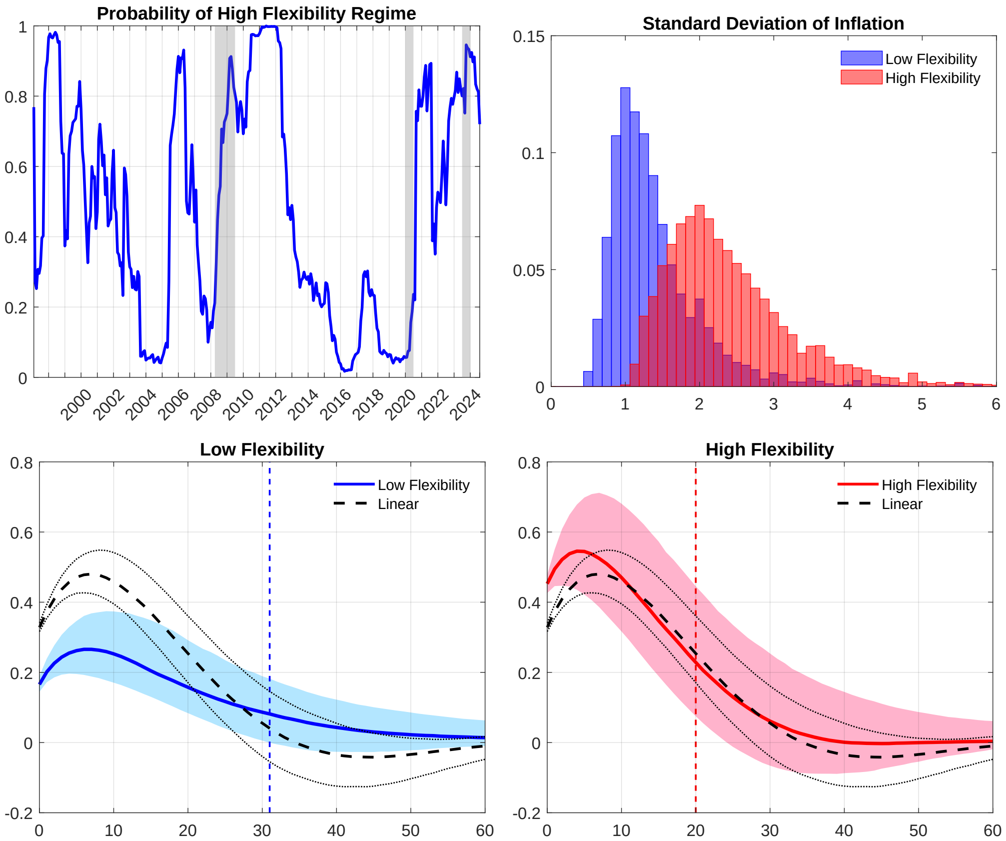


Notes: The upper panels report the probability of being in a high flexibility regime, $G\left(\widetilde{\mathcal{F}}_t, \gamma\right) = (1 + e^{-\gamma \widetilde{\mathcal{F}}_t})^{-1}$, and the distributions of the estimated inflation volatility in the high and low price flexibility regimes. The lower panels report the responses of inflation to a one-standard deviation shock in the STARMA(2,1) model. Specifically, the bottom-left (right) panel graphs the response in the low (high) price flexibility regime. In both cases, we also report the the response from a (linear) ARMA(2,1) model. $68\%$ confidence intervals, and the distribution of inflation volatility, are built based on the Markov Chain Monte Carlo (MCMC) method developed in Chernozhukov and Hong (2003). In each of the two IRFs charts the vertical dashed line indicates the half-life of the shock.

the implied inflation volatility is twice as large in the high-flexibility regime, with the on-impact response 50% higher relative to the linear model. These results are broadly supportive of the basic insights of the $Ss$ model illustrated in the previous section, and highlight the importance of keeping track of the degree of price flexibility.

Neglecting the non-linearities tied to price flexibility may result in a significant underestimation of inflation's response during periods of high flexibility, and an overestimation during periods of low flexibility. This is highly relevant, from the perspective of forecasting and policy making. Before we delve into this, though, we want to establish to which extent inflation volatility and persistence correlate with price flexibility. While the overall level of inflation depends on the blend and magnitude of shocks impacting the economy, the degree of price flexibility is likely to play a key role in shaping their propagation. In this respect, it is useful to recall that Forbes et al. (2018) highlight how UK inflation has shown relatively high volatility and low persistence in the 2008-2012 time-span and, to a lesser extent, around the early 2000s.[27] These periods have been associated with large departures of inflation from the target. Consistently, price flexibility as from our estimates peaks in both periods.

The upper-right panel of Figure 6 reports the distribution of the estimated inflation volatility, in the high and low price flexibility regimes. Periods of high flexibility display significantly greater volatility, whereas inflation volatility is substantially lower under low price flexibility. This suggests that time-varying inflation volatility can, at least in part, be attributed to the changing degree of price flexibility. For instance, during the particularly low-volatility period of 2014–2016, YoY inflation has reached its lowest point, dipping below zero for the first time in the post-WWII period. Analyses from the Bank of England attribute such weak inflation to the decline in oil prices and the depreciation of the Pound.[28] Our analysis suggests that low price flexibility may have extended this period of subdued inflation. By contrast, the post-COVID period and the 2009–2012 period, both identified as high price flexibility phases, are marked by notably higher inflation volatility.

## 5.2 State dependence and inflation projections

An immediate implication of the analysis so far is that inflation volatility and persistence may vary significantly, depending on aggregate price flexibil-

[27]Volatility is measured by standard deviation of the mean-reverting component of their model of inflation.

[28]See, e.g., the Inflation Report published by the Bank of England on February 12, 2015.

ity. Specifically, inflation tends to be more volatile, less persistent, and generally higher when flexibility is high. In this section, we test whether the Bank of England and professional forecasters factor in the state-dependent properties of inflation dynamics related to price flexibility, when forming their inflation expectations. If this properly accounted for, the resulting inflation forecast errors should remain uncorrelated with the flexibility regime.

Each quarter, the Bank of England's Inflation Report publishes YoY inflation forecasts from the Monetary Policy Committee, alongside forecasts from market participants. Both sets of forecasts target the Bank of England's inflation index, which switched from RPIX to CPI inflation in December 2003. We construct quarterly forecast errors as the difference between the appropriate (mean) forecast[29] and realized inflation at a given horizon: $e_{t+h|t} = \pi^{YoY}_{t+h|t} - \pi^{YoY}_{t+h}$. Therefore, positive (negative) errors denote an overprediction (underprediction) of inflation. Forecast errors are then regressed on the logistic transformation of the flexibility index, $G\left(\widetilde{\mathcal{F}}_{t-1};\gamma\right)$. Specifically, we use a quadratic spline function with a knot at 0.5:

$$e_{t+h|t} = a_0 + a_1(G_{t-1} - 0.5) + a_2(G_{t-1} - 0.5)^2 + a_3\mathbb{1}_{\{G_{t-1}>0.5\}}(G_{t-1} - 0.5)^2, \quad (8)$$

where $\mathbb{1}_{\{G_{t-1}>0.5\}}$ is an indicator function equal to 1 when $G_{t-1} > 0.5$ and zero otherwise. This specification allows us to capture various potential relationships between the flexibility regime and the bias in inflation forecasts. The analysis is conducted on a sample of qurterly forecasts produced by the Bank of England and professional forecasters, with $h = 0, ..., 7$ and over the 1998-2024 time interval.

Table 1 summarizes the regression results. The first six columns present the estimated forecast bias (along with associated p-values) for low, average, and high levels of flexibility (i.e., $G = 0.2, 0.5, 0.8$). The last two columns of the table provide the p-value for the null hypothesis that no relationship exists between the forecast error and the flexibility regime (i.e., $H_0 : a_1 = a_2 = a_3 = 0$), as well as the corresponding R-squared (adjusted for the number of regressors).

[29]The results remain virtually unchanged if we use the median or the mode, instead of the mean forecast.

Table 1: INFLATION FORECAST ERRORS AND PRICE FLEXIBILITY

| (a) BoE MPC RPIX/CPI Forecast Error Bias | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| Horizon | $G = 0.2$ | | $G = 0.5$ | | $G = 0.8$ | | F$-stat$ | $\tilde{R}^2$ |
| 0 | 0.04 | [0.19] | -0.02 | [0.58] | -0.05 | [0.15] | 0.22 | 1.38 |
| 1 | 0.12 | [0.23] | 0.02 | [0.86] | -0.13 | [0.34] | 0.38 | 0.06 |
| 2 | 0.24 | [0.19] | -0.02 | [0.94] | -0.37 | [0.20] | 0.12 | 2.73 |
| 3 | 0.34 | [0.20] | -0.16 | [0.62] | -0.77 | [0.10] | 0.02 | 6.29 |
| 4 | 0.43 | [0.19] | -0.37 | [0.38] | -1.21 | [0.06] | 0.00 | 9.63 |
| 5 | 0.55 | [0.11] | -0.61 | [0.21] | -1.70 | [0.03] | 0.00 | 15.33 |
| 6 | 0.62 | [0.07] | -0.72 | [0.20] | -1.92 | [0.02] | 0.00 | 17.34 |
| 7 | 0.60 | [0.08] | -0.64 | [0.26] | -1.86 | [0.03] | 0.00 | 14.96 |
| (b) Market Participants' Forecast Error Bias | | | | | | | | |
| Horizon | $G = 0.2$ | | $G = 0.5$ | | $G = 0.8$ | | F$-stat$ | $\tilde{R}^2$ |
| 0 | 0.02 | [0.65] | -0.09 | [0.23] | -0.10 | [0.10] | 0.25 | 1.09 |
| 1 | 0.02 | [0.87] | -0.05 | [0.71] | -0.11 | [0.44] | 0.88 | -2.28 |
| 2 | 0.19 | [0.26] | -0.05 | [0.82] | -0.37 | [0.21] | 0.18 | 1.85 |
| 3 | 0.28 | [0.26] | -0.23 | [0.50] | -0.79 | [0.10] | 0.04 | 5.47 |
| 4 | 0.33 | [0.28] | -0.47 | [0.28] | -1.24 | [0.06] | 0.01 | 8.57 |
| 5 | 0.44 | [0.16] | -0.74 | [0.16] | -1.76 | [0.03] | 0.00 | 14.67 |
| 6 | 0.50 | [0.09] | -0.88 | [0.14] | -2.02 | [0.02] | 0.00 | 17.30 |
| 7 | 0.47 | [0.12] | -0.82 | [0.18] | -1.99 | [0.02] | 0.00 | 15.04 |

Notes: The table reports the results of a quadratic spline regression of the forecast errors $e_{t+h|t}$ (for different forecast horizons, $h$, measured in quarters) on a quarterly average of an indicator of the normalized price flexibility index, $\widetilde{\mathcal{F}}$: $G_{t-1} = G(\widetilde{\mathcal{F}}_{t-1};\gamma) = (1+e^{-\gamma\widetilde{\mathcal{F}}_{t-1}})^{-1}$. The regression is specified as $e_{t+h|t} = a_0 + a_1\,(G_{t-1}-0.5) + a_2\,(G_{t-1}-0.5)^2 + a_3\,(G_{t-1}-0.5)^2\,\mathbb{1}_{\{G_{t-1}>0.5\}}G_{t-1}^2$, where $\mathbb{1}_{\{G_{t-1}>0.5\}}$ is an indicator function taking value 1 when $G_{t-1} > 0.5$ and zero otherwise. The upper panel refers to the Bank of England MPC's RPIX/CPI forecast errors, while the bottom panel considers market participants' forecast errors. For each value of $G$, the two columns report the fitted $\hat{e}_{t+h|t}$ evaluated at different levels of the indicator and the p-value associated with the null hypothesis that $\hat{e}_{t+h|t}$ is equal to 0 (this is calculated using Newey-West standard errors), respectively. The penultimate column (F-$stat$) reports the p-value of the null hypothesis that all the coefficients associated to the flexibility regime are equal to 0 (i.e., $H_0 : a_1 = a_2 = a_3 = 0$). The last column reports the adjusted R-squared, denoted by $\tilde{R}^2$.

While inflation forecasts tend to be unbiased when aggregate price flexibility is low or average, there is evidence of a significant negative bias during periods of high price flexibility. These findings support the notion that information regarding price flexibility is not fully utilized by either the Central Bank or market participants. In particular, we detect a significant negative bias in inflation forecasts from three quarters ahead ($h > 2$). This bias is not only statistically significant, but also quantitatively relevant: in periods of high flexibility, for $h > 4$, the forecast is 150 basis points lower than it should be when accounting for the level of flexibility. Accounting for this negative bias during periods of high flexibility alone explains above $15\%$ of the variability in the forecast error in our sample.

**Robustness and extensions** A potential concern with the specification in equation (8) is that large forecast errors may simply reflect large inflation realizations, which often follow periods of high price flexibility. To address this, Table F1 in Appendix F presents the forecast error as a percentage of realized inflation, confirming the broad pattern of results. Notably, for $h > 4$, the bias can reach up to 50% of realized inflation. Additionally, we examine whether aggregate price flexibility influences forecast uncertainty, measured as the squared absolute forecast error (see, e.g., Reifschneider and Tulip, 2019). In line with the prediction of a state-dependent model, our findings indicate that high flexibility periods are associated with significantly more volatile inflation outcomes, with absolute forecast errors for $h > 4$ nearly twice as large as in low-flexibility periods (see Table F2, Appendix F). Thus, when projecting inflation, practitioners should account for heightened uncertainty during periods of high price flexibility.

#### 5.2.1 Price flexibility and post-pandemic inflation forecasts

Failure to predict the scale and persistence of inflation is widely acknowledged, and has drawn criticism towards the Bank of England, ultimately leading to an external review of its inflation forecasts. Bernanke (2024) notes that similar forecast errors also characterize those of professional forecasters and, more generally, Central Banks across G7 countries. Figure 7 compares the estimated bias from the specification in Eq. 8 using the full sample (this corresponds to the results in Table 1) with the bias for the subsample ending in 2020 Q4. While the underprediction of the latest inflationary peak—which coincides with a period of rather elevated price flexibility—affects our estimates, evidence of a significant and substantial bias is already present in data prior to the post-Pandemic inflation surge.

This raises the question of whether earlier recognition of this evidence might have reshaped the prevailing narrative on inflation persistence that dominated global and UK policy discussions from mid-2021 to mid-2022. During this period, the Bank of England consistently predicted a rapid return to target inflation, within a two-year horizon. Figure 8 compares the BoE's inflation forecasts with realized inflation and bias-adjusted forecasts, over four quarters starting in 2020 Q4. To avoid look-ahead bias, we re-estimate the bias-adjustment model in Equation (8) every time we produce a forecast, using real-time data. While the precise magnitude of the inflationary spike was by any means unforeseeable,[30] adjusting for the bias would have consistently resulted in higher inflation forecasts, particularly at longer horizons, over this period. In fact, bias-adjusting the forecast would have signaled that inflationary pressures were already building with the reopening of the economy in late 2020. While inflation remained well below the Central Bank's target at the time, the bias-adjusted forecast predicted a clear overshoot, with inflation exceeding 3% in 2022—contrasting with the Bank of England's projection of a slow convergence to the target. More broadly, these adjusted forecasts suggested a more cautious view on the transitory nature of the inflation spike as early as 2021 Q1, indicating that inflation would remain well above target by the end of the forecast horizon.

While the quantitative improvements in forecast accuracy from incorporating time-varying price flexibility may seem modest, their economic significance should not be overlooked. During the post-COVID inflation surge, the Bank of England and other forecasters predicted a relatively swift return of inflation

[30] Applying Blanchard and Bernanke (2023)'s model to the UK, Haskel et al. (2024) highlight that this is largely attributable to the (unforeseen) rapid rise in energy and food commodity prices during the period under scrutiny.

Figure 7: INFLATION FORECAST ERROR BIAS

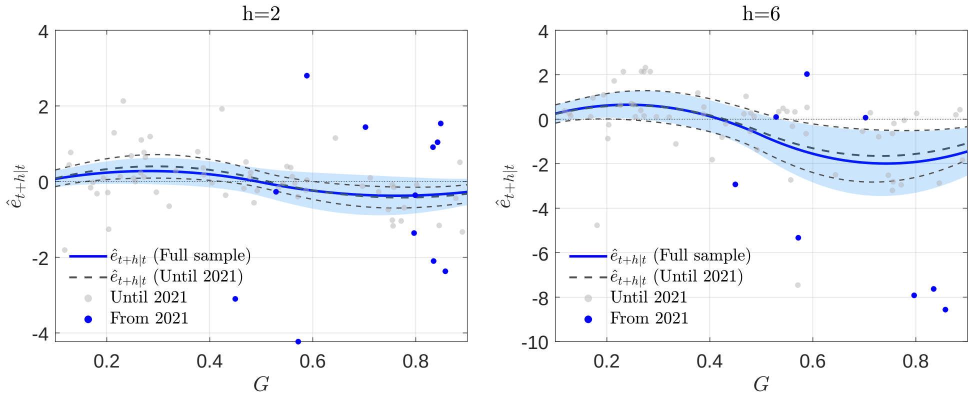


(a) BoE Forecast

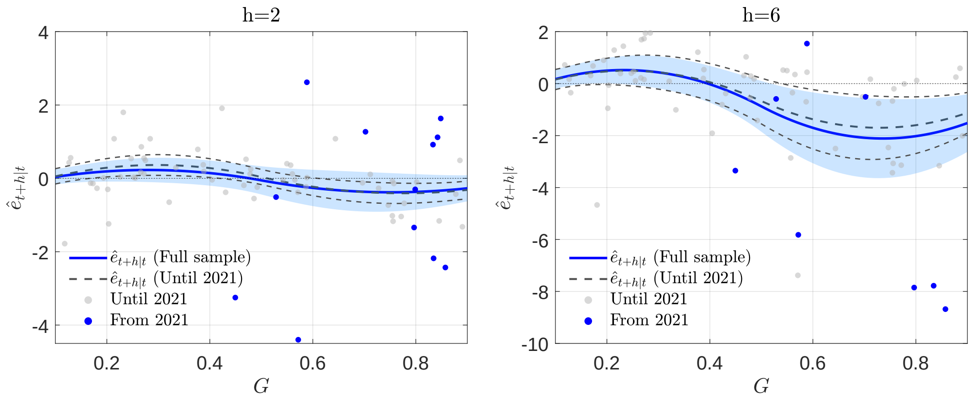


(b) Market Participants' Forecast

Notes: Each panel reports the expected forecast error, $\hat{e}_{t+h|t}$, conditional on values of a the high-flexibility regime probability, $G$. The blue line describes the mean forecast based on the full sample (1997:Q3 - 2024:Q2).The dashed line shows the conditional forecast error based on the sample excluding the post-COVID-19 period (i.e., from 2021:Q1). The bands indicate the 90% confidence interval. The dots plot observed forecast error against the mean of $G$ in the quarter. Panel (a) and (b) refer to the Bank of England MPC's RPIX/CPI forecast errors, while the panel (c) and (d) considers market participants' forecast errors. Negative values indicate a forecast that underpredicts the actual inflation outcome.

to the 2% target—an expectation that ultimately proved incorrect. To put this in context, it is worth recalling that the Bank of England MPC's June 2021 minutes report that the Committee expected "the direct impact of commodity prices on CPI inflation would be transitory". Later that year, Governor Bailey reiterated this view, maintaining that despite early signs of rising prices, "the price

Figure 8: BIAS-ADJUSTED FORECASTS

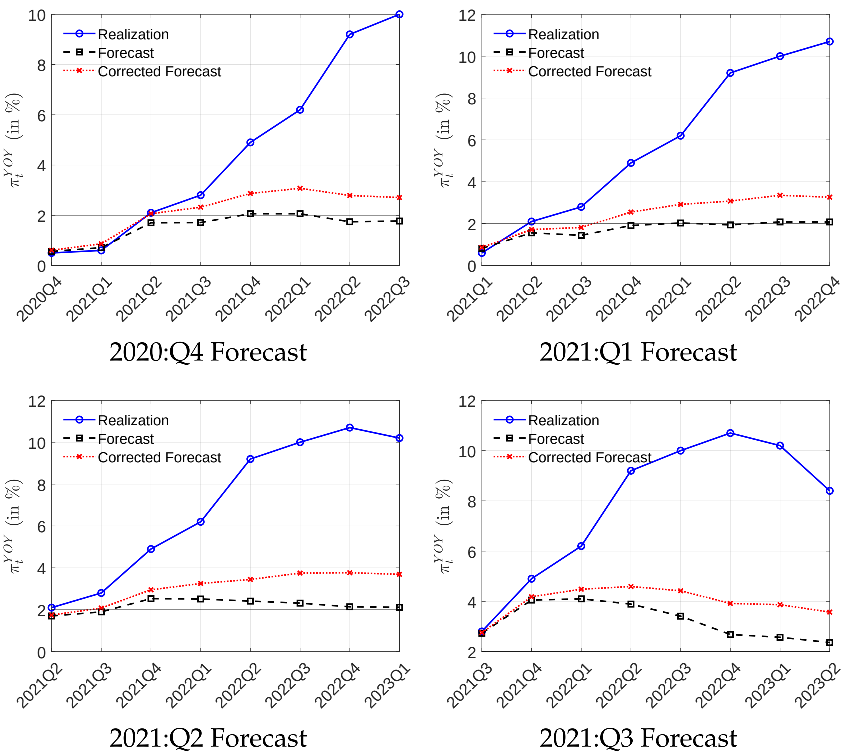


Notes: Each panel reports the CPI inflation forecasts produced by the BoE in the last quarter of 2020 and the first three quarters of 2021, as well as the corresponding bias-corrected forecasts and realizations. The bias adjustment is based on the fitted values from estimating Equation (8) re-estimated with available data in real-time. To get the bias, the high-flexibility regime probabilities $G$ from the first month of the respective quarter is used. Specifically, $G(\widetilde{\mathcal{F}}_t; \gamma) = [0.82, 0.80, 0.76, 0.90]$ for the quarters under consideration.

pressures will be transient" (Bailey, 2021).[31] Our counterfactual exercise suggests that, had forecasters accounted for the prevailing high price flexibility, they would have recognized that inflationary pressures were likely to persist for longer, even without additional shocks. For instance, the forecast produced

[31] Even after a substantial revision to its projections in Q3 2021, the Bank of England continued to emphasize the transitory nature of the inflation surge. This stance was reflected in statements such as, "The Committee's central expectation is that current elevated global and domestic cost pressures will prove transitory," which appeared as a summary view in the MPC minutes of both August and September 2021.

at the time of the June 2021 meeting would have indicated inflation remaining persistently above target, nearing 4% by the end of the forecast horizon, rather than gradually returning to target. This insight could have prompted the Bank of England to revise its policy stance earlier or more decisively. In this sense, the improvement in the forecast performance has not to be merely read from a statistical standpoint. Rather, it highlights a missed opportunity to identify inflation persistence in real time, with potentially significant implications for the timing and calibration of monetary policy.

# 6 Concluding remarks

Analyzing UK price microdata, we document substantial time variation in aggregate price flexibility, emphasizing the critical role of the extensive margin of price adjustment. Relying solely on the frequency of price adjustment significantly understates both the level and variability of price flexibility. Our findings highlight the prominence of state-dependent price setting, particularly in the aftermath of the Great Recession and during the recent post-Pandemic inflation surge.

Most importantly, we identify strong non-linearity in the price response to inflationary shocks, fundamentally driven by shifts in price flexibility. Neither the Bank of England nor professional forecasters fully account for state dependence when projecting future inflation. As a result, they tend to underestimate inflationary pressures during periods of high price flexibility, when inflation is more volatile and less persistent. This oversight was particularly consequential in the early phase of the post-Pandemic inflation surge, when forecasters consistently projected a swift return of inflation to target.

Our findings highlight the crucial role of timely access to the microdata underlying aggregate price indices for Central Banks. In this regard, the UK ONS stands out among national statistical agencies for its commitment to making this data publicly available almost in real-time. These data are not only crucial for refining theoretical models of price adjustment, but also for tracking inflation dynamics and identifying shifts in sectoral pricing behavior. By integrating timely information from price microdata into their forecasting models, policymakers can improve their understanding of inflationary forces and enhance their ability to respond swiftly to changing economic conditions.

# References


ALTONJI, J. G. AND L. M. SEGAL (1996): "Small-Sample Bias in GMM Estimation of Covariance Structures," *Journal of Business & Economic Statistics*, 14, 353–366.

ALVAREZ, F., H. LEBIHAN, AND F. LIPPI (2016): "The Real Effect of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach," *American Economic Review*, 106, 2817–2851.

ALVAREZ, F. AND F. LIPPI (2014): "Price Setting With Menu Cost for Multiproduct Firms," *Econometrica*, 82, 89–135.

ALVAREZ, F., F. LIPPI, AND A. OSKOLKOV (2023): "The Macroeconomics of Sticky Prices with Generalized Hazard Functions," *The Quarterly Journal of Economics*, 137, 989–1038.

ALVAREZ, L. J., E. DHYNE, M. HOEBERICHTS, C. KWAPIL, H. L. BIHAN, P. LÜNNEMANN, F. MARTINS, R. SABBATINI, H. STAHL, P. VERMEULEN, AND J. VILMUNEN (2006): "Sticky Prices in the Euro Area: A Summary of New Micro-Evidence," *Journal of the European Economic Association*, 4, 575–584.

ASCARI, G. AND T. HABER (2022): "Non-Linearities, State-Dependent Prices and the Transmission Mechanism of Monetary Policy," *The Economic Journal*, 132, 37–57.

AUERBACH, A. J. AND Y. GORODNICHENKO (2012): "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, 4, 1–27.

BAILEY, A. (2021): "The Hard Yards. Speech given at the Society of Professional Economists Annual Dinner, London, 27 September 2021. Bank of England," .

BELL, T. AND D. TOMLINSON (2018): "Is everybody concentrating? Recent trends in product and labour market concentration in the UK," Briefing, Resolution Foundation.

BERARDI, N., E. GAUTIER, AND H. L. BIHAN (2015): "More Facts about Prices: France Before and During the Great Recession," *Journal of Money, Credit and Banking*, 47, 1465–1502.

BERGER, D. AND J. VAVRA (2018): "Dynamics of the U.S. price distribution," *European Economic Review*, 103, 60–82.

BERNANKE, B. S. (2024): "Forecasting for Monetary Policy Making and Communication at the Bank of England: A Review," Tech. rep., Bank of England, Independent Evaluation Office.

BILS, M. (2009): "Do Higher Prices for New Goods Reflect Quality Growth or Inflation?" *The Quarterly Journal of Economics*, 124, 947–985.

BILS, M. AND P. J. KLENOW (2004): "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112, 947–985.

BLANCHARD, O. J. AND B. S. BERNANKE (2023): "What Caused the US Pandemic-Era Inflation?" NBER Working Papers 31417, National Bureau of Economic Research, Inc.

BUNN, P. AND C. ELLIS (2012): "Examining The Behaviour Of Individual UK Consumer Prices," *Economic Journal*, 122, 35–55.

CABALLERO, R. J. AND E. M. R. A. ENGEL (2007): "Price stickiness in Ss models: New interpretations of old results," *Journal of Monetary Economics*, 54, 100–121.

CALVO, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 12, 383–398.

CAPLIN, A. S. AND D. F. SPULBER (1987): "Menu Costs and the Neutrality of Money," *The Quarterly Journal of Economics*, 102, 703–725.

CARLSSON, M. AND O. N. SKANS (2012): "Evaluating Microfoundations for Aggregate Price Rigidities: Evidence from Matched Firm-Level Data on Product Prices and Unit Labor Cost," *American Economic Review*, 102, 1571–1595.

CARVALHO, C. AND O. KRYVTSOV (2021): "Price Selection," *Journal of Monetary Economics*, 122, 56–75.

CAVALLO, A., F. LIPPI, AND K. MIYAHARA (2024): "Large Shocks Travel Fast," *American Economic Review: Insights*, 6, 558–574.

CHERNOZHUKOV, V. AND H. HONG (2003): "An MCMC approach to classical estimation," *Journal of Econometrics*, 115, 293–346.

CHU, B. M., K. P. HUYNH, D. T. JACHO-CHÁEZ, AND O. KRYVTSOV (2018): "On the Evolution of the United Kingdom Price Distributions," *Annals of Applied Statistics*.

COSTAIN, J. AND A. NAKOV (2011): "Distributional dynamics under smoothly state-dependent pricing," *Journal of Monetary Economics*, 58, 646–665.

DEDOLA, L., L. HENKEL, C. HÖYNCK, C. OSBAT, AND S. SANTORO (2024): "What does new micro price evidence tell us about inflation dynamics and monetary policy transmission?" *Economic Bulletin Articles*, 3.

DELOECKER, J. AND J. EECKHOUT (2018): "Global Market Power," NBER Working Papers 24768, National Bureau of Economic Research.

DIXON, H. D. AND C. GRIMME (2022): "State-dependent or time-dependent pricing? New evidence from a monthly firm-level survey: 1980–2017," *European Economic Review*, 150.

DIXON, H. D., K. LUINTEL, AND K. TIAN (2020): "The Impact of the 2008 Crisis on UK Prices: What We Can Learn from the CPI Microdata," *Oxford Bulletin of Economics and Statistics*, 82, 1322–1341.

DIXON, H. D. AND K. TIAN (2017): "What We can Learn About the Behaviour of Firms from the Average Monthly Frequency of Price-Changes: An Application to the UK CPI Data," *Oxford Bulletin of Economics and Statistics*, 79, 907–932.

DORSEY, R. E. AND W. J. MAYER (1995): "Genetic Algorithms for Estimation Problems with Multiple Optima, Nondifferentiability, and Other Irregular Features," *Journal of Business & Economic Statistics*, 13, 53–66.

EICHENBAUM, M., N. JAIMOVICH, AND S. REBELO (2011): "Reference Prices, Costs, and Nominal Rigidities," *American Economic Review*, 101, 234–262.

FORBES, K., L. KIRKHAM, AND K. THEODORIDIS (2018): "A Trendy Approach to UK Inflation Dynamics," CEPR Discussion Papers 12652.

GAGNON, E. (2009): "Price Setting during Low and High Inflation: Evidence from Mexico," *The Quarterly Journal of Economics*, 124, 1221–1263.

GAUTIER, E., C. CONFLITTI, R. P. FABER, B. FABO, L. FADEJEVA, V. JOUVANCEAU, J.-O. MENZ, T. MESSNER, P. PETROULAS, P. ROLDAN-BLANCO, ET AL. (2024): "New facts on consumer price rigidity in the euro area," *American Economic Journal: Macroeconomics*, 16, 386–431.

GELMAN, A., D. DUNSON, D. RUBIN, H. S. STERN, AND J. B. CARLIN (2004): *Bayesian Data Analysis*, Chapman & Hall/CRC Texts in Statistical Science.

GHASSIBE, M. AND A. NAKOV (2025): "Business Cycles with Pricing Cascades," Unpublished Manuscript, Universitat Pompeu Fabra.

GOLOSOV, M. AND R. E. LUCAS (2007): "Menu Costs and Phillips Curves," *Journal of Political Economy*, 115, 171–199.

GRANGER, C. W. J. AND T. TERASVIRTA (1993): *Modelling Non-Linear Economic Relationships*, Oxford University Press.

GROENEVELD, R. A. AND G. MEEDEN (1984): "Measuring skewness and kurtosis," *Journal of the Royal Statistical Society: Series D*, 33, 391–399.

GUIMARAES, B. AND K. D. SHEEDY (2011): "Sales and Monetary Policy," *American Economic Review*, 101, 844–876.

HARVEY, A. C. (1993): *Time Series Models. Second Edition*, Harvester Wheatsheaf.

HASKEL, J., J. MARTIN, AND L. BRANDT (2024): "Recent UK Inflation: An Application of the Bernanke-Blanchard Model," Unpublished Manuscript, Bank of England.

HENKEL, L., E. WIELAND, A. BŁAŻEJOWSKA, C. CONFLITTI, B. FABO, L. FADEJEVA, J. JONCKHEERE, P. KARADI, P. MACIAS, J.-O. MENZ, AND P. A. SEILER (2023): "Price setting during the Coronavirus (COVID-19) pandemic," Occasional Paper Series 324, European Central Bank.

KARADI, P., J. AMANN, J. S. BACHILLER, P. SEILER, AND J. WURSTEN (2023): "Price setting on the two sides of the Atlantic - Evidence from supermarket scanner data," *Journal of Monetary Economics*, 140, 1–17.

KARADI, P. AND A. REIFF (2019): "Menu Costs, Aggregate Fluctuations, and Large Shocks," *American Economic Journal: Macroeconomics*, 11, 111–146.

KARADI, P., R. SCHOENLE, AND J. WURSTEN (2024): "Price Selection in the Microdata," *Journal of Political Economy Macroeconomics*, 2, 228–271.

KEHOE, P. AND V. MIDRIGAN (2015): "Prices are sticky after all," *Journal of Monetary Economics*, 75, 35–53.

Klenow, P. J. and O. Kryvtsov (2008): "State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?" *The Quarterly Journal of Economics*, 123, 863–904.

Komunjer, I. (2007): "Asymmetric power distribution: Theory and applications to risk measurement," *Journal of Applied Econometrics*, 22, 891–921.

Kryvtsov, O. and N. Vincent (2021): "The Cyclicality of Sales and Aggregate Price Flexibility," *The Review of Economic Studies*, 88, 334–377.

Lein, S. M. (2010): "When do firms adjust prices? Evidence from micro panel data," *Journal of Monetary Economics*, 57, 696–715.

Luo, S. and D. Villar (2021): "The price adjustment hazard function: Evidence from high inflation periods," *Journal of Economic Dynamics and Control*, 130.

Midrigan, V. (2011): "Menu Costs, Multiproduct Firms, and Aggregate Fluctuations," *Econometrica*, 79, 1139–1180.

Montag, H. and D. Villar (2022): "Price-Setting During the Covid Era," Economic Working Papers 547, Bureau of Labor Statistics.

Nakamura, E. and J. Steinsson (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *The Quarterly Journal of Economics*, 123, 1415–1464.

——— (2010a): "Monetary Non-neutrality in a Multisector Menu Cost Model," *The Quarterly Journal of Economics*, 125, 961–1013.

——— (2010b): "More Facts About Prices – Supplement to "Five Facts About Prices: A Reevaluation of Menu Cost Models"," Online Supplement.

Nakamura, E., J. Steinsson, P. Sun, and D. Villar (2018): "The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation," *Quarterly Journal of Economics*, Forthcoming.

ONS (2019): *Consumer Price Indices: Technical Manual*, Office for National Statistics, 2014 ed.

Reifschneider, D. and P. Tulip (2019): "Gauging the uncertainty of the economic outlook using historical forecasting errors: The Federal Reserve's approach," *International Journal of Forecasting*, 35, 1564–1582.

Vavra, J. (2014): "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation," *The Quarterly Journal of Economics*, 129, 215–258.

# A ONS Microdata on consumer price

Our sample covers the 1996:M2-2024:M8 time window, thus resulting into about $37.4$ million observations. Each month around $109,000$ prices are collected by a market research firm on behalf of the ONS. There are also about $150$ items for which the corresponding price quotes are centrally collected. These are excluded from the publicly available dataset, as the structure of their market segment might allow the identification of some price setters, or because of the need to frequently adjust for quality changes.[1] Price quotes are recorded on or around the second or third Tuesday of the month (*Index day*), with the exact date being kept secret to avoid abnormal prices that, among other things, may be due to the collection of prices during bank-holiday weeks, or to price manipulations by service providers and retailers. Furthermore, to make sure the collected price quotes are valid prices, the ONS has set various checks in place, both at the collection point and at later stages in the process. As a preliminary step in handling the dataset, we only employ price quotes that have been marked as being validated by the system or accepted by the ONS. Thus, any price quote that has been marked as missing, non-comparable, or temporarily out of stock is excluded from the sample. We refer to the remaining subset of prices—which make for approximately 60% of those included in the CPI—as *Classification Of Individual COnsumption by Purpose* (COICOP) price quotes.

Each price quote is classified by region, location, outlet and item. The region refers to the geographical entity within the UK from which a given price quote is recorded. The location is intended as a shopping district within a given region: on price-collection days, $141$ different locations are visited.[2] For a given location, the shop code is a unique but anonymized *id* associated with the outlet from which the quote is recorded. In turn, each shop is classified according to whether it is independent (i.e., part of a group comprising less than $10$ outlets at the national level) or part of a chain (i.e., more than 10 outlets). Due to a confidentiality agreement between the ONS and the individual shops, for each price quote only the region, outlet and item classifications are published. In light of this, some of the price quotes may not be uniquely identified. This is typically the case when the ONS samples the same item, in outlets that are part of a chain, but for multiple locations within the same region. As an example, in March 2013 we pick an item with the following characteristics: 'Women's Long Sleeves Top' (*id*: 510223) sold in multiple outlets (*shop type*: 1) within the region of London (*region*: 2). With these coordinates at hand we retrieve two different price quotes: one location sells the item for $\pounds 22$, and one for $\pounds 26$. In February 2013 the price quotes for the same type of good were recorded at $\pounds 25$ and $\pounds 26$, respectively. The price quotes are so close that telling the two price trajectories apart may be challenging. To make sure that price trajectories can be uniquely identified, we look at 'base prices', which are intended as the January's price for each of the items under scrutiny.[3] Even after conditioning on base prices, though, a small portion of price trajectories are still not uniquely identified (about $0.6\%$, on average): we opt for discarding them. In Table

[1]This is typically the case for personal computers, whose frequent model upgrades impose the use of hedonic regressions, so as to enhance comparisons across time.

[2]Until August 1996, 180 different locations were being sampled. New locations are chosen every year, with about 20% of them being replaced. As a result, a location is expected to survive an average of about four years in the sample.

[3]The base price is typically relied upon to normalize price quotes and calculate price indices, or to adjust for changes in the quality and/or quantity of a given good.

Table A.1: SUMMARY STATISTICS

| | Categories | | | |
|---|---|---|---|---|
| | COICOP | Unique | History | Regular |
| Price Quotes | | | | |
| Total | $37,390,169$ | $37,171,595$ | $34,063,217$ | $30,401,232$ |
| Avg. per Month | $109,009$ | $108,372$ | $99,309$ | $88,633$ |
| Avg. CPI Weight | $59.81\%$ | $59.51\%$ | $54.89\%$ | $50.07\%$ |
| Sales and Recoveries | | | | |
| Avg. per Month (Unweighted) | $10.44\%$ | $10.46\%$ | $10.69\%$ | |
| Avg. per Month (Weighted) | $5.14\%$ | $5.14\%$ | $4.82\%$ | |
| Product Substitutions | | | | |
| Avg. per Month (Unweighted) | $1.02\%$ | $1.02\%$ | $0.58\%$ | |
| Avg. per Month (Weighted) | $0.48\%$ | $0.48\%$ | $0.25\%$ | |

Notes: *COICOP* stands for the *Classification Of Individual COnsumption by Purpose* price quotes used to calculate the CPI index; *Unique* indicates the COICOP price quotes that are uniquely identified; *History* refers to the subset of price quotes in the Unique category for which we can identify at least two consecutive price quotes; *Regular* refers to the price quotes in the History category that do not correspond to sales, product substitutions, or recovery prices. For each of these categories, we compute the weighted contribution of each category's price quotes to the CPI index, as well as the relative number of price quotes corresponding to sales, recovery prices, and product substitutions. Whenever weighted, these statistics are obtained by accounting for CPI, item-specific, stratum and shop (i.e., elementary aggregate) weights. Sample period: 1996:M2-2024:M8.

A.1 the column labeled 'History' refers to the price quotes with an identifiable history that spans at least two consecutive periods. Following the criteria outlined above, we drop about $9,000$ quotes per month.[4]

To aggregate the individual price quotes into a single price, we also make use of the following weights produced by the ONS:[5] the *shop* weights, which are employed to account for the fact that a single item's price is the same in different shops of the same chain (e.g., a pint of milk at a Tesco store);[6] the *stratification* weights, which reflect the fact that purchasing patterns may differ markedly by region or type of outlet;[7] finally, the *item* and *COICOP* weights reflect consumers' expenditure shares in the national accounts.

[4]Due to a particularly low coverage, Housing and Housing Services(COICOP 4) and Education (COICOP 10) are excluded from the sample. We also exclude price changes larger than $300\%$, which we deem to be due to measurement errors. These take place rarely ($< 0.02\%$).

[5]See Chapter 7 of the ONS CPI Manual (ONS, 2019).

[6]In this case the ONS enters a single price for a pint of milk, but the weight attached to this is 'large', so as to reflect that all Tesco stores within the region have posted the same price.

[7]Four levels of sampling are considered for local price collection: locations, outlets within location, items within location-outlet section and individual product varieties. For each geographical region, locations and outlets are based on a probability-proportional-to-size systematic sampling, where size accounts for the number of employees in the retail sector (locations) and the net retail floor space (outlets).

## A.1 On the representativeness of the data

This section provides additional details on the construction of the dataset. The ONS data have a good coverage of all COICOP sectors, with the exception of Housing and Housing Services (COICOP 4), Communication (COICOP 8) and Education (COICOP 10), whose coverage are less than 21%, 6%, and 3%, respectively. Given the extremely low coveage, we exclude COICOP 4 and 10. We keep COICOP 8, as the available price quotes are clustered in a small subset of items, such as Flower Delivery, Telephone for home use and Phone Accessories.[8]

The left panel of Figure A.1 contrasts the weights assigned to each of the COICOP sectors to those employed to build the CPI (re-normalized to exclude COICOP 4 and 10). Overall, we observe that using the available price quotes results into relatively larger weights for COICOP 1 and 11, whereas sectors 7 and 9 are underweighed. The right panel of Figure A.1 reports the official CPI inflation together with the inflation series retrieved from all the available price quotes (labeled *COICOP*) and the inflation obtained once all filters described in Section 2 are applied (labeled *Regular*). Unfiltered data track quite closely the official numbers, whereas the 'regular' series displays a robust correlation with the official data (roughly 0.84), and shows a positive bias. The latter mainly emerges from the exclusion of sales from the sample.

Figure A.2 provides a number of additional statistics calculated on price microdata. Figure A.3 reproduces some of the moments used in the model estimation, where the latter are obtained by aggregating sectoral estimates (using 25 COICOP groups). The moments track each other closely, therefore demonstrating that sectoral composition has a limited effect on the aggregate moment statistics.

[8]Due to the small number of price quotes in this sector, the results would be little affected by its exclusion from the analysis.

Figure A.1: REPRESENTATIVENESS

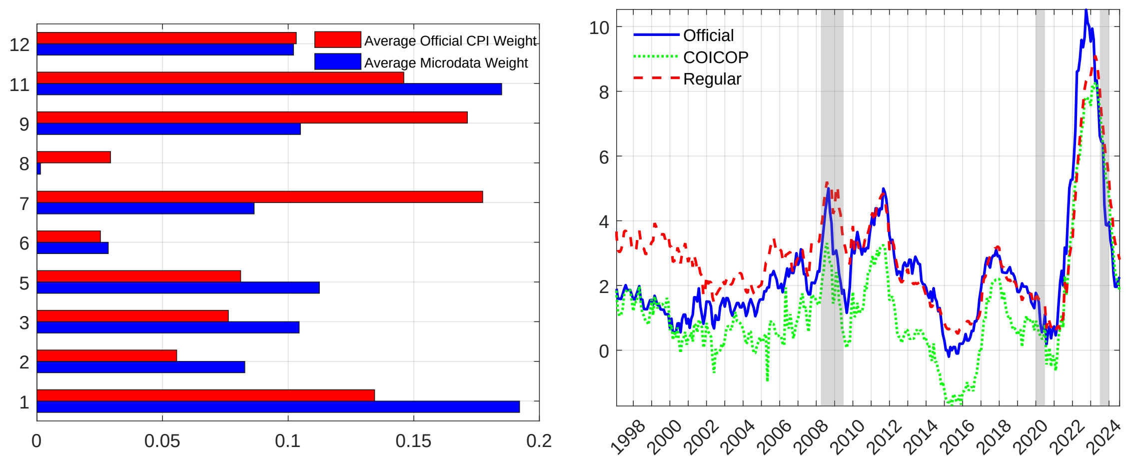


Notes: The left panel contrasts the weights assigned to each of the COICOP sectors to those assigned to build the CPI (re-normalized to exclude COICOP 4 and 10). The right panel reports the official CPI inflation, together with the inflation series retrieved from all the available price quotes (labeled *COICOP*) and the inflation obtained once all filters described in Section 2 are applied (labeled *Regular*). The COICOP codes are (1) Food And Non-Alcoholic Beverages, (2) Alcoholic Beverages, Tobacco And Narcotics, (3) Clothing And Footwear, (5) Furnishings, Household Equipment And Routine Household Maintenance, (6) Health, (7) Transport, (8) Communication, (9) Recreation And Culture, (11) Hotels, Cafes And Restaurants, (12) Miscellaneous Goods And Services.

Figure A.2: ADDITIONAL STATISTICS FROM PRICE MICRODATA

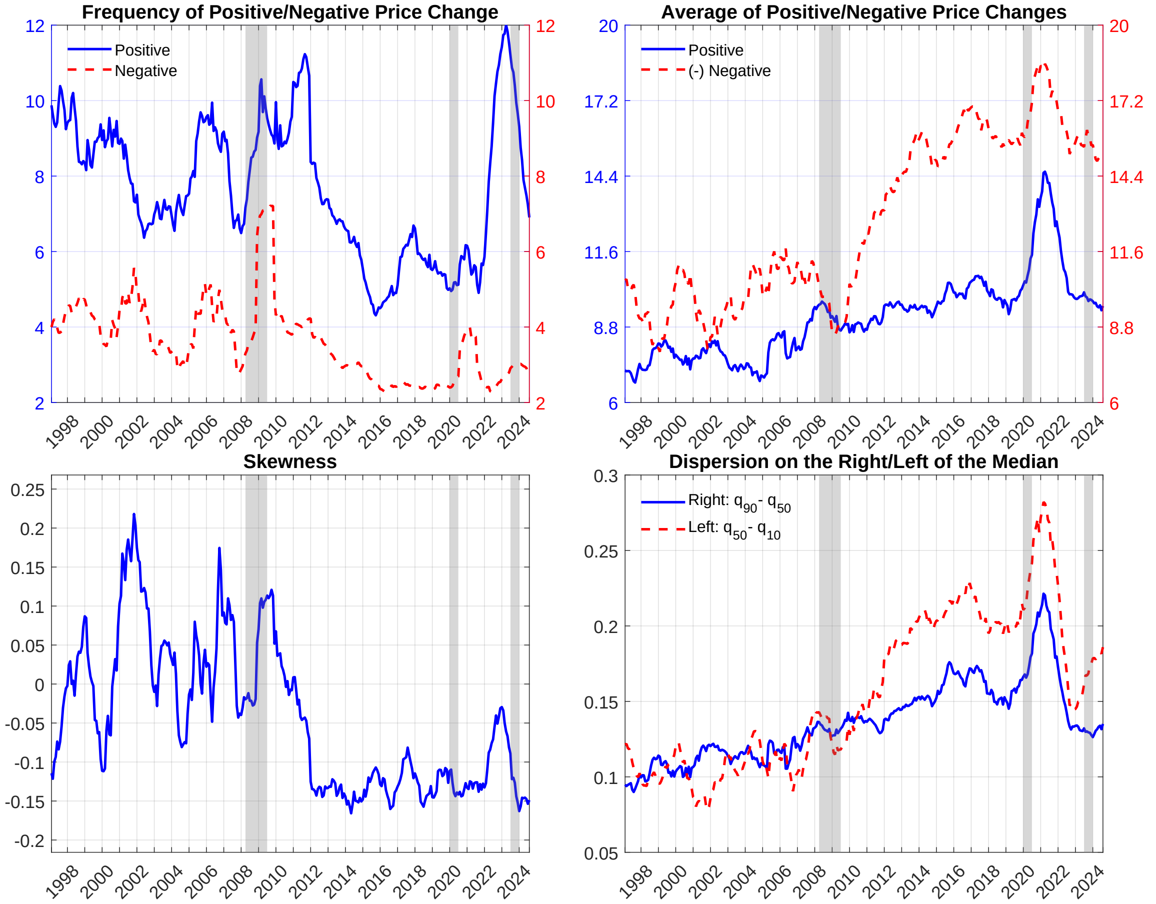


Notes: The frequency of adjustment, $fr_t$ , is computed as $\sum_i \omega_{i,t} \mathbb{1}_{\{\Delta p_{i,t} \neq 0\}}$, where $\omega_{i,t}$ denotes the CPI weight associated to good $i$ at time $t$, and $\mathbb{1}_{\{\Delta p_{i,t} \neq 0\}} = 1$ if $\Delta p_{i,t} \neq 0$ and zero otherwise. The average price change, instead, is computed as $fr_t^{-1} \sum_i \omega_{i,t} \mathbb{1}_{\{\Delta p_{i,t} \neq 0\}} \Delta p_{i,t}$. The positive and negative counterparts of these statistics are obtained by conditioning them on positive and negative price changes, respectively. All series are in percentage terms. In the upper-right panel we report the mirror image of the average of negative price changes. The skewness of the distribution of price changes is calculated as $\frac{q_{90,t}+q_{10,t}-2q_{50,t}}{q_{90,t}-q_{10,t}}$. The lower-right panel reports the price dispersion on the right (left) side of the median price change computed as $q_{50} - q_{10}$ ($q_{90} - q_{50}$). The shaded vertical bands indicate the duration of recessionary episodes.

Figure A.3: Aggregate vs Disaggregated Moments

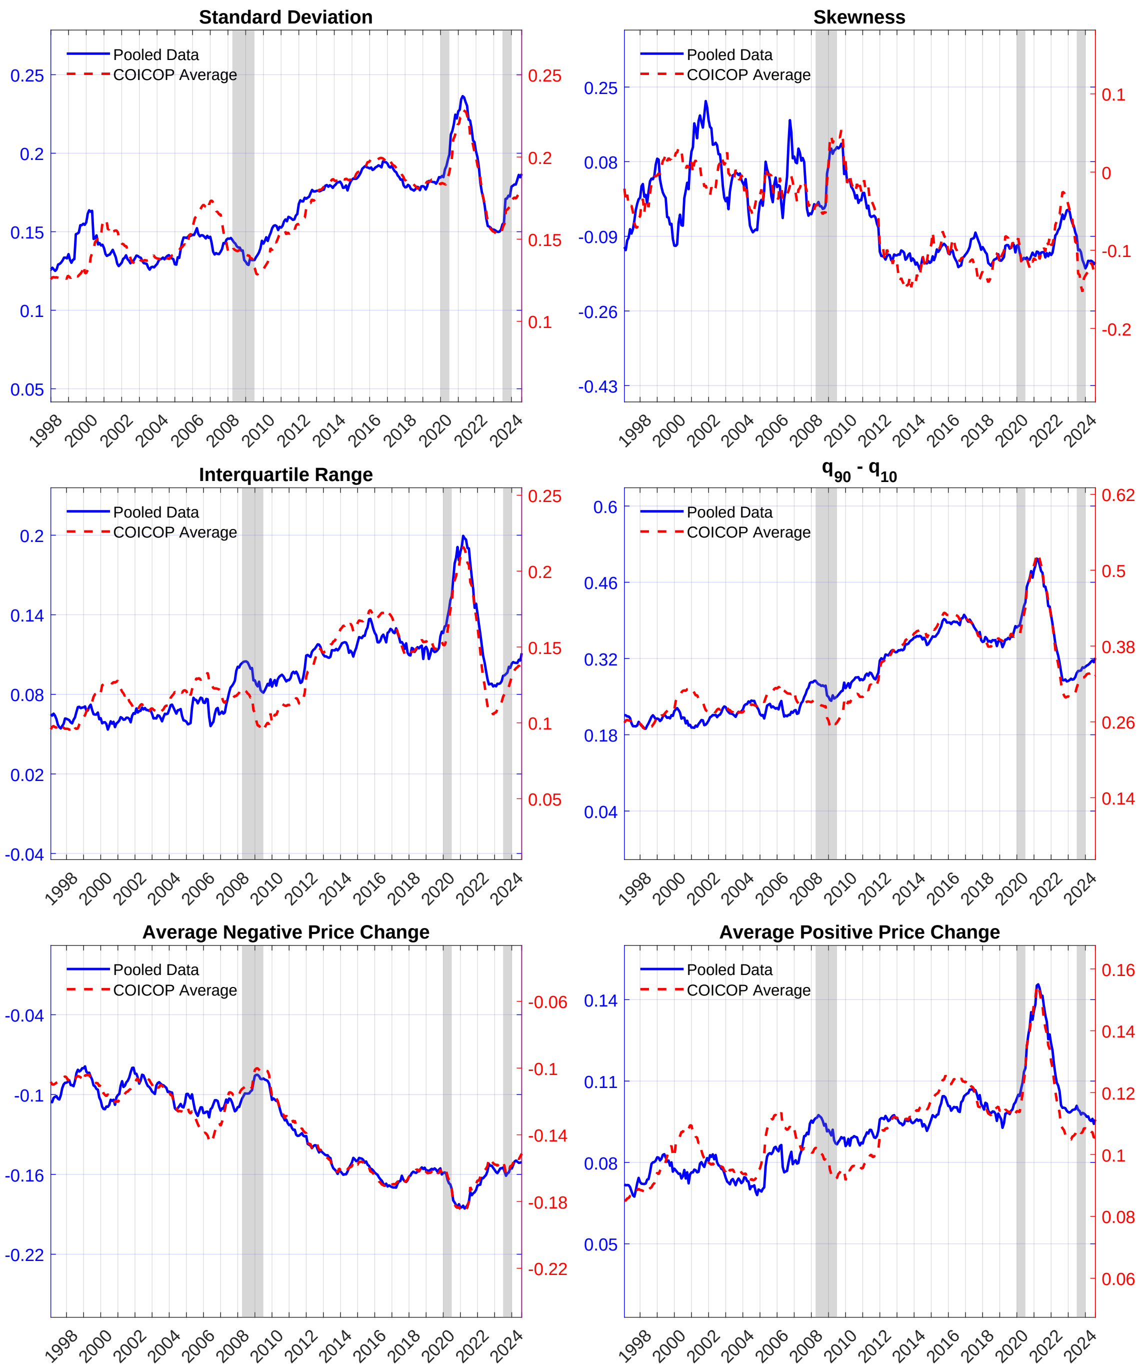


Notes: The figure compares various moments of the distribution of price changes with their counterparts obtained by averaging the corresponding moments of the price quotes obtained for each of the 25 COICOP group categories. The shaded vertical band indicates the duration of recessionary episodes.

# B A simple decomposition of YoY inflation variation

In this section, we provide a step-by-step derivation of how the variability of year-on-year (YoY) inflation can be decomposed into the contributions of frequency and average price changes, as reported in Eq. (2).

Denote month-on-month (MoM) inflation as $\pi_t$. As stated in Eq. (1), we have $\pi_t = fr_t \times \Delta p_t$. Taking a first-order approximation, it follows that

$$\pi_t - \overline{\pi} \approx \overline{fr}(\Delta p_t - \overline{\Delta p}) + \overline{\Delta p}(fr_t - \overline{fr}), \tag{9}$$

where variables with bars denote full-sample averages. Higher-order terms account for the variation in inflation due to the covariation between the frequency of price adjustments and the average price change.

Define the 12-month moving average operator applied to a generic variable $\kappa_t$ as $\widetilde{\kappa}_t \equiv \frac{1}{12}\sum_{j=0}^{11} \kappa_{t-j}$. Year-on-year (YoY) inflation can then be recovered as the (backward-looking) 12-month moving average of the annualized MoM inflation, i.e., $\pi_t^{YoY} = \sum_{j=0}^{11} \pi_{t-j} = \widetilde{\pi}_t^A$, where $\pi_t^A \equiv 12\pi_t$ denotes the annualized MoM inflation.

Substituting the first-order approximation into this relationship, we can express the variation of YoY inflation as a function of the variability of the (12-month moving average of the) frequency of price adjustment and price change:

$$\pi_t^{YoY} - \overline{\pi}^{YoY} \approx 12 \times \left\{ \overline{fr}(\widetilde{\Delta p_t} - \overline{\Delta p}) + \overline{\Delta p}(\widetilde{fr}_t - \overline{fr}) \right\}. \tag{10}$$

Finally, defining the annualized average price change as $\Delta p_t^A \equiv 12\Delta p_t$, we obtain

$$\pi_t^{YoY} - \overline{\pi}^{YoY} \approx \overline{fr}(\widetilde{\Delta p_t}^A - \overline{\Delta p}^A) + \overline{\Delta p}^A(\widetilde{fr}_t - \overline{fr}), \tag{11}$$

where the long-run average YoY inflation is given by $\overline{\pi}^{YoY} = \overline{fr} \times \overline{\Delta p}^A$.

# C Model estimates

Figure C.1: Parameters of the Price Gap Distribution

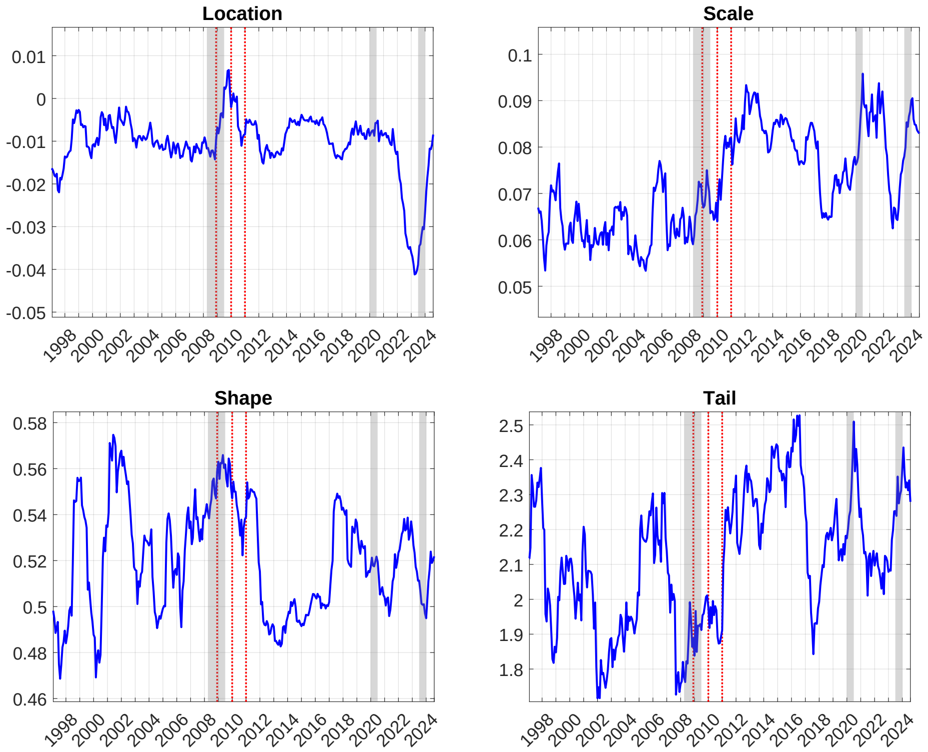


Notes: The red lines denote the three VAT changes in the sample. The shaded vertical bands indicate the duration of recessionary episodes.

Figure C.2: Parameters of the Hazard Function

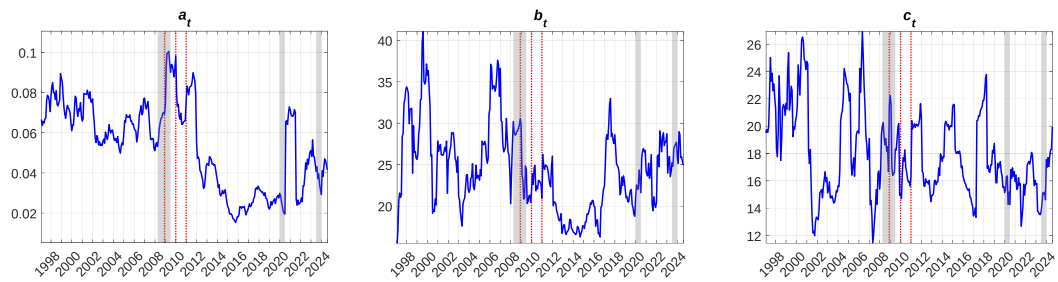


Notes: The red lines denote the three VAT changes in the sample. The shaded vertical bands indicate the duration of recessionary episodes.

Figure C.3: FIT OF THE $Ss$ MODEL (SELECTED MOMENTS)

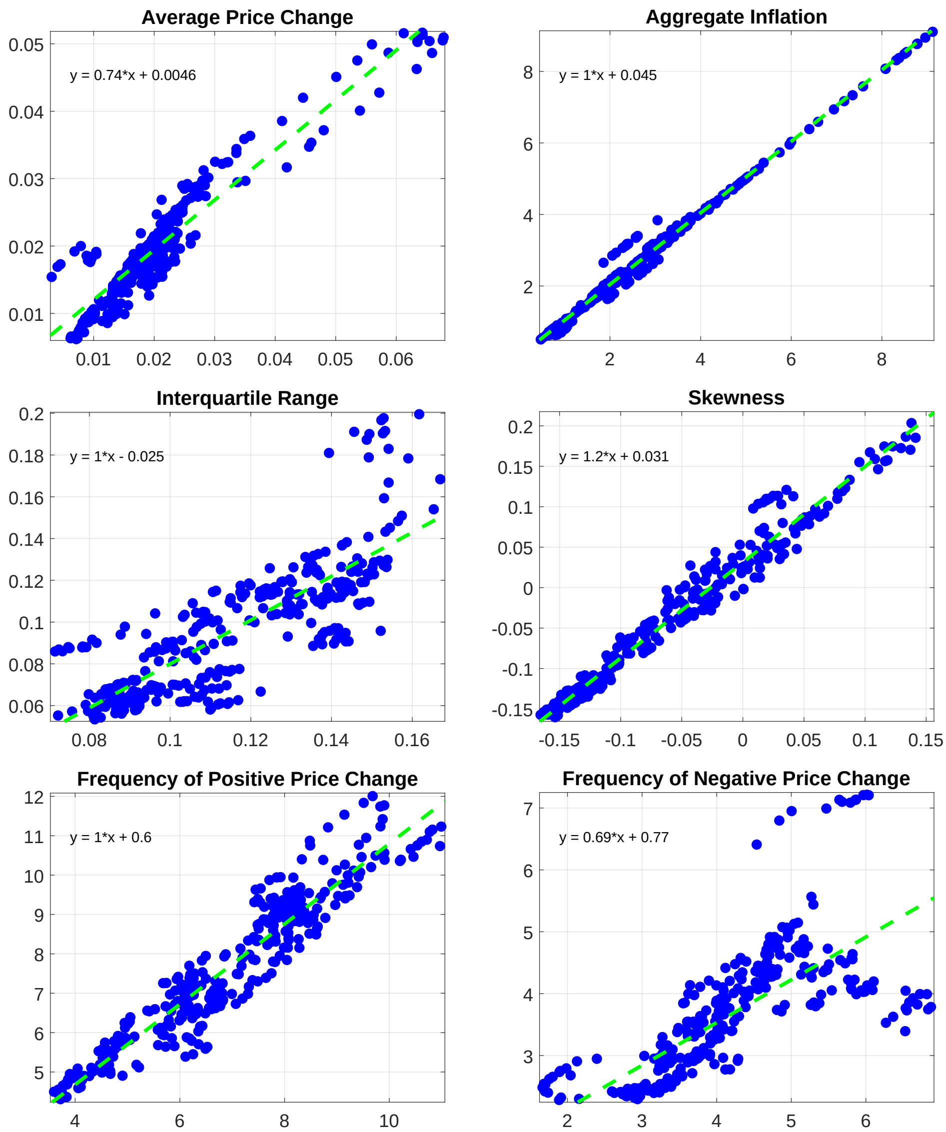


Notes: The figure compares the estimated moments from the $Ss$ model in Section 3 (x-axis) to the moments estimated from the raw data (y-axis). Each chart also reports the linear fit (green/broken) line.

# D Model identification

In this appendix we check whether the SMM estimation strategy we employ for the estimation of the generalized $Ss$ model is able to separately identify the shape of the price gap distribution and the hazard function.

The parameters of the model are identified through their ability to match the selected moments. As noted in Section 3.1, we match the following moments of the distribution of price changes: mean, median, standard deviation, interquartile range, difference between the 90th and 10th quantile of the distribution, as well as (quantile-based) skewness and kurtosis. We also match the frequency and the average size of prices movements, after distinguishing between positive and negative price changes, as well as the observed rate of inflation.

We evaluate the systematic impact of each parameter on the moments that we are matching. To this end, the first exercise we perform consists of investigating whether marginal variation in each of the parameters of the model affects the moments that we are matching. Figure D.1 and D.2 report the results of this exercise. We fix all the parameters at their median estimates, and for each column we vary one of of them at the time (within the range of values that the parameters assume in our estimation) and report the impact of these changes for some selected moments.

All parameters have an impact on a number of moments, and in the expected direction. For instance, increasing the scale (tail) parameter of the price gap distribution increases (decreases) monotonically the implied dispersion of the distribution of (non-zero) price changes, and in both cases decreases the skewness and the kurtosis. Instead, changing the location or the shape parameter has an opposite impact on skewness and kurtosis, and affects non-monotonically the dispersion (with higher dispersion obtained for a more skewed distribution, regardless of the sign of the skewness). As for the parameters of the hazard function, changing the constant term affects equally the frequency of price adjustment, whereas changes in the slope for positive (negative) price gaps impacts the frequency of negative (positive) price changes and the average negative (positive) price changes, leaving invariate the positive (negative) side. These results confirm the observation of Berger and Vavra (2018) for the specific functional forms of the price gap distribution and the hazard function we employ.

Having established that all the parameters have an impact on the moments we attempt to match, a fair question is whether moment matching allows us to appropriately identify/distinguish the shape of the price gap distribution from the shape of the hazard function. In fact, one might question whether the specific model we choose is able to identify a fatter price gap distribution from a steeper hazard function, or a skewed price gap distribution from an asymmetric hazard function. To this end, we simulate samples of 100,000 price changes from the model, and then fit the model on each of these synthetic samples by SMM, matching the same moments we use in the baseline estimation (see Section 3.1). Figure D.3 contrasts the true price gap distribution (upper panel) and hazard function (lower panel) to the estimated counterparts. We look at three possible different parameterizations of the model, and report the 'fan charts' of the estimated functions. The specific parameterizations are merely meant to serve for illustrative purposes: we would obtain very similar evidence by imposing alternative specifications. Finally, for each set of calibrations, we simulate and estimate the model over 200 different samples.

The charts highlight that the model is able to separately identify the shape of the price gap and hazard function in all the settings we consider. The discrepancy between the true parametrization and the estimate is minimal, and the resulting match of the flexibility index and its decomposition is very close to the true one.

It is also important to stress that Berger and Vavra (2018) produce a battery of exercises in support of our approach. Most importantly, they address how well the resulting measure of price flexibility—which only captures the impact response of prices to a nominal shock—reflects overall non-neutrality. To this end, they estimate simulated data from the CalvoPlus model of Nakamura and Steinsson (2008), and report close comovement between the impact response from the structural model and the estimated index of price flexibility from the accounting framework. Notably, this exercise also addresses the criticism towards estimating the generalized $Ss$ model in every period, as if observations were independent across time. In this respect, we should stress that standard structural frameworks tend to impose a rather tight relationship between distributions at a given point in time and how they evolve. In line with our predecessors, we claim that imposing flexible functional forms within a period—in a way that represents an intermediate step between a fully structural approach and a non parametric one—allows us to exploit valuable information, in the perspective of studying time variation in aggregate price flexibility.

Figure D.1: IDENTIFICATION AND THE PARAMETERS OF $f_t\,(x)$

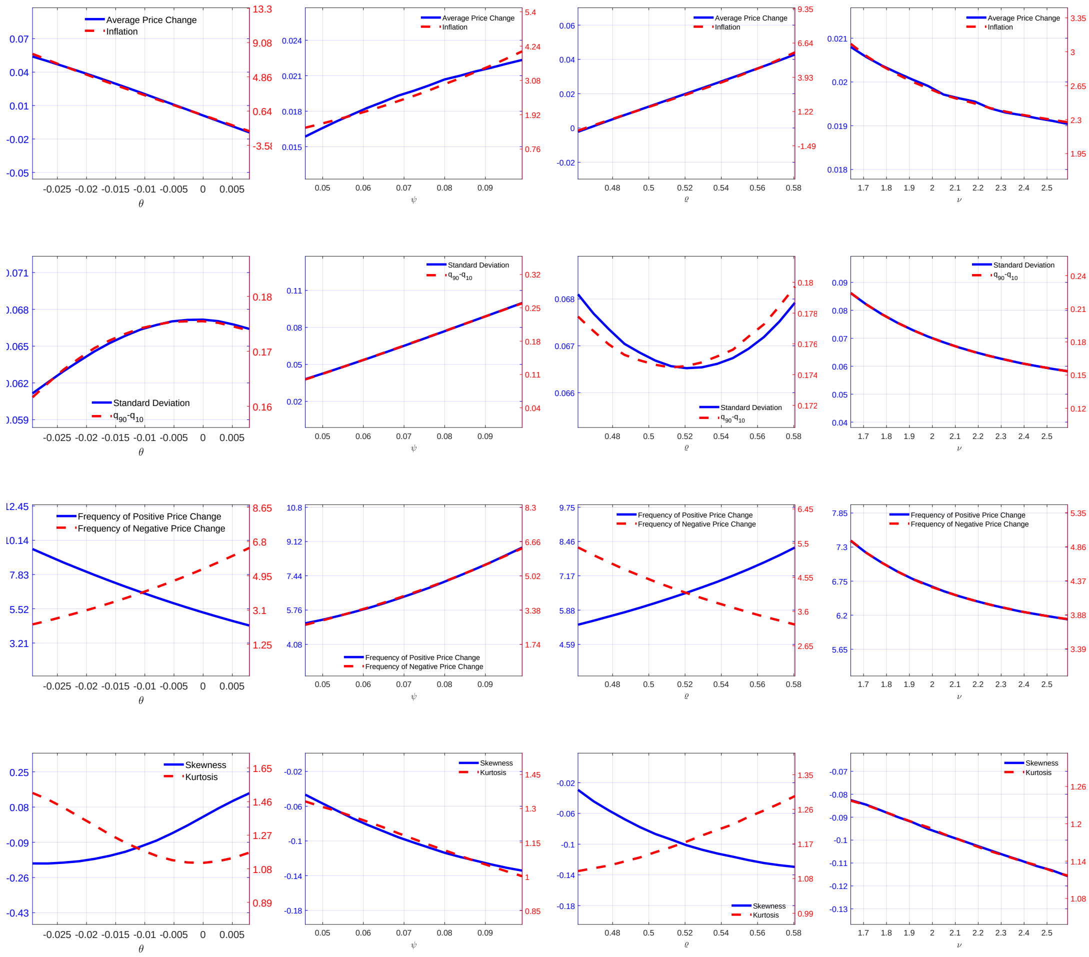


Notes: In each panel, we vary one of the parameters of $f_t\,(x)$ at the time—while keeping the other coefficients at their baseline estimate—and report its effect on key moments of the price change distribution, as well as the resulting rate of inflation.

Figure D.2: IDENTIFICATION AND THE PARAMETERS OF $\Lambda_t(x)$

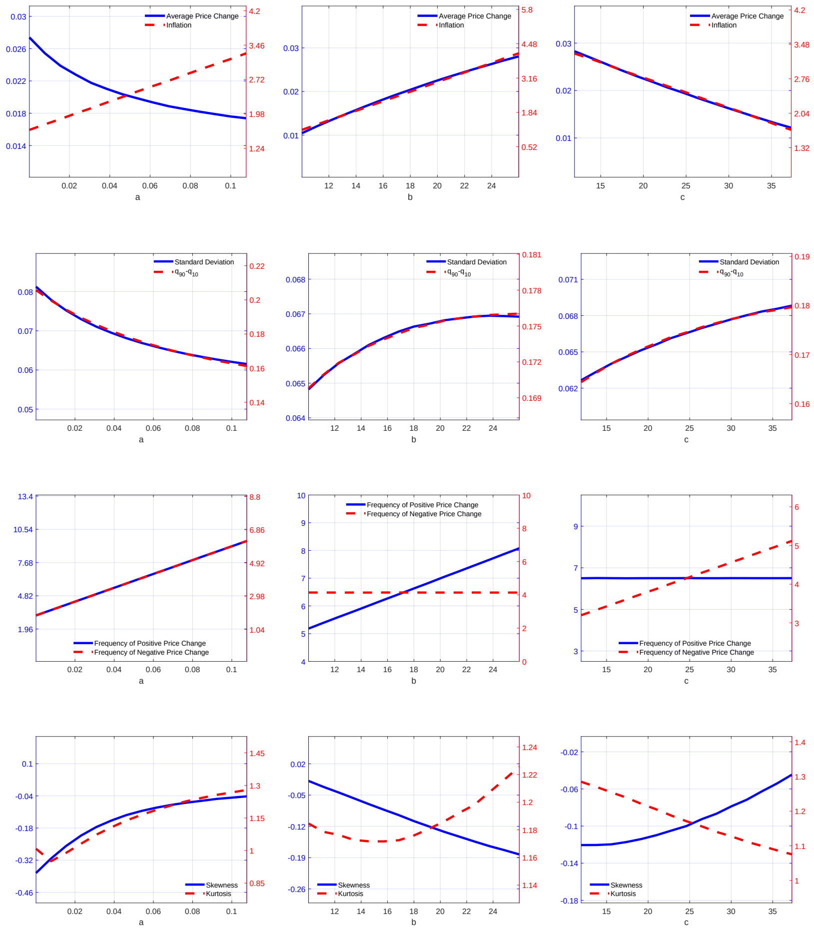

Notes: In each panel, we vary one of the parameters of $\Lambda_t(x)$ at the time—while keeping the other coefficients at their baseline estimate—and report its effect on key moments of the price change distribution, as well as the resulting rate of inflation.

Figure D.3: MODEL SIMULATIONS AND EMPIRICAL FIT

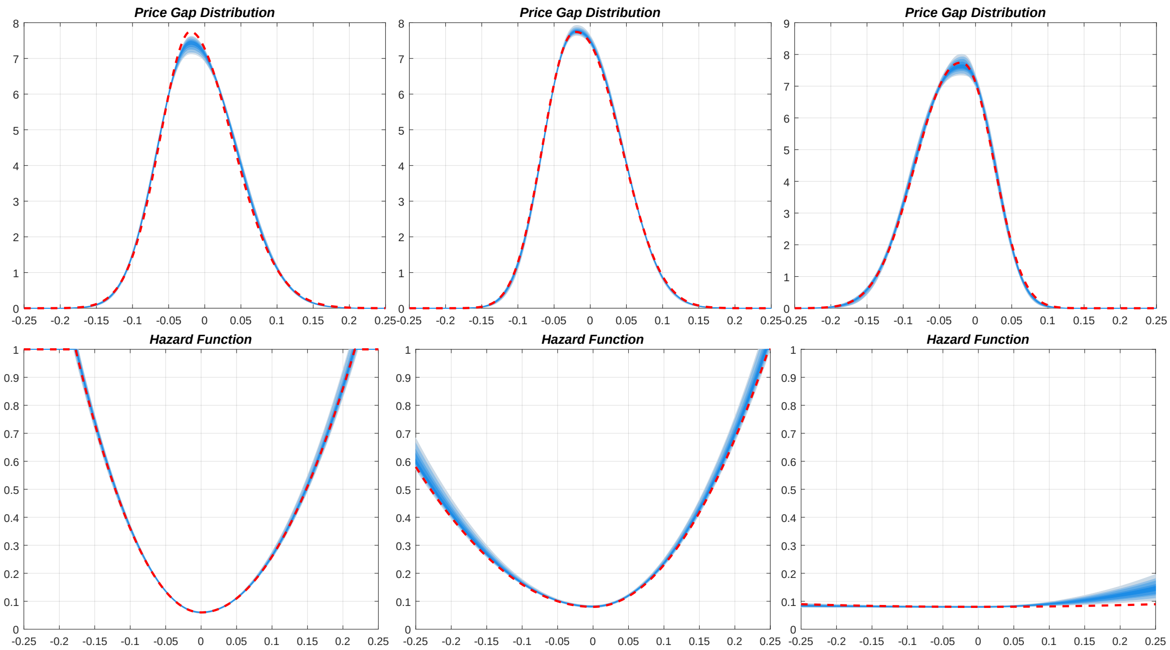


(a) (b) (c)

Notes: The red line corresponds to the 'true' DGP, while the blue shades correspond to the [5,10,20,. . . 90,95]-th quantile of the estimated price gap distribution (upper panel) and hazard function (lower panel). The following parameterizations are considered: Panel (a): $\theta = -0.02, \psi = 0.07, \varrho = 0.42, \nu = 1.9, a = 0.06, b = 20, c = 30$; Panel (b): $\theta = -0.02, \psi = 0.07, \varrho = 0.42, \nu = 2.2, a = 0.08, b = 15, c = 8$; Panel (c): $\theta = -0.02, \psi = 0.07, \varrho = 0.58, \nu = 2.2, a = 0.08, b = 0.15, c = 0.15$.

# E Details on the computation of the impulse response function from the $Ss$ model

This appendix gives a brief account of how we compute the impulse response functions from the generalized $Ss$ model presented in Section 3. We start by specifying a process for the exogenous (first-moment) shock.[9] Specifically, we assume that:

$$\mu_t = \rho\mu_{t-1} + \eta_t.$$

Thus, we fix $\rho = 0.5$ and select a shock $\eta_0 = -1\%$. In light of this, should prices be fully flexible, we would observe a $1\%$ increase of inflation that dies out relatively quickly.

The impulse responses are then calculated as:

$$\begin{aligned} IRF_j &= \mathrm{E}(\pi_{t+j}|\mu_{t+j} = \hat{\mu}_{t+j}) - \mathrm{E}(\pi_{t+j}|\mu_{t+j} = 0) \\ &= -\int z_j \Lambda_t(z) f_t(z)\, dz + \int x_j \Lambda_t(x) f_t(x)\, dx, \end{aligned}$$

where $z_j = x_j + \hat{\mu}_{t+j}$. Note that, by definition, the cumulative impact of the shock equals the sum of the $\mu_t$'s.

# F Estimation of the STARMA (p,q) model

Recall the smooth transition ARMA model, STARMA(p,q), in Section 5.1:

$$\begin{aligned} \pi_t &= G\left(\widetilde{\mathcal{F}}_{t-1};\gamma\right)\left(\phi_0^H + \sum_{j=1}^{p}\phi_i^H\pi_{t-j} + \varepsilon_t^H + \sum_{i=1}^{q}\theta_i^H\varepsilon_{t-i}^H\right) \\ &\quad + \left[1 - G\left(\widetilde{\mathcal{F}}_{t-1};\gamma\right)\right]\left(\phi_0^L + \sum_{j=1}^{p}\phi_i^L\pi_{t-j} + \varepsilon_t^L + \sum_{i=1}^{q}\theta_i^L\varepsilon_{t-i}^L\right). \end{aligned} \tag{12}$$

This can be easily casted in state space. Therefore the likelihood can be calculated recursively using the Kalman filter (see, e.g., Harvey, 1993, Ch.4). Since the model is highly non-linear in the parameters, it is possible to have several local optima and one must try different starting values of the parameters. Furthermore, given the nonlinearity of the problem, it may be difficult to construct confidence intervals for parameter estimates, as well as impulse responses. To address these issues, we use a Markov Chain Monte Carlo (MCMC) method developed in (Chernozhukov and Hong, 2003, henceforth CH). This method delivers not only a global optimum but also distributions of parameter estimates.

Denote with $\theta$ the vector of parameters. We employ the Hastings-Metropolis algorithm to implement CH's estimation method. Specifically, our procedure to construct chains of length $N$ can be summarized as follows:

[9]Since we assume that the shock has the same impact on all price quotes, the shock acts as a location shifter of the price gap distribution.

- *Step 1*: Draw $\vartheta^{(n+1)}$, a candidate vector of parameter values for the chain's $n+1$ state, as $\vartheta^{(n+1)} = \theta^{(n)} + \mathbf{u}_n$ where $\mathbf{u}_n$ is a vector of *iid* shocks taken from a student-t distribution with zero mean, $\nu = 5$ degrees of freedom and variance $\Omega$.

- *Step 2*: Take the $n+1$ state of the chain as

$$\theta^{(n+1)} = \begin{cases} \vartheta^{(n+1)} & \text{with probability } \min\left\{1, \frac{L\left(\vartheta^{(n+1)}\right)}{L\left(\theta^{(n)}\right)}\right\} \\ \theta^{(n)} & \text{otherwise} \end{cases}$$

  where $L(\theta)$ denotes the value of the likelihood of the model evaluated at the parameters values $\theta$.

Specifically, we use an adaptive step for the value of $\Omega$, i.e. this is recalibrated using the accepted draws in the initial part of the chain and then adjusted on the fly to generate $25 - 35\%$ acceptance rates of candidate draws, as proposed in Gelman et al. (2004). We use a total of 50,000 draws, and drop the first 25,000 draws (i.e., the 'burn-in' period). We then pick the 1-every-5 accepted draws to mitigate the possible autocorrelations in the draws. We run a series of diagnostics to check the properties of the resulting distributions from the generated chains. We find that the simulated chains converge to stationary distributions and that simulated parameter values are consistent with good identification of parameters.

CH show that $\overline{\theta} = \frac{1}{N}\sum_{i=1}^{N} \theta^{(i)}$ is a consistent estimate of $\theta$ under standard regularity assumptions of maximum likelihood estimators. CH also prove that the covariance matrix of the estimate of $\theta$ is given by the variance of the estimates in the generated chain. Furthermore, we can use the generated chain of parameter values $\theta^{(i)}$ to construct confidence intervals for the impulse responses.

# G Additional Results and Robustness

Figure G.1: PRICE FLEXIBILITY AND INFLATION PERSISTENCE STARMA(2,4)

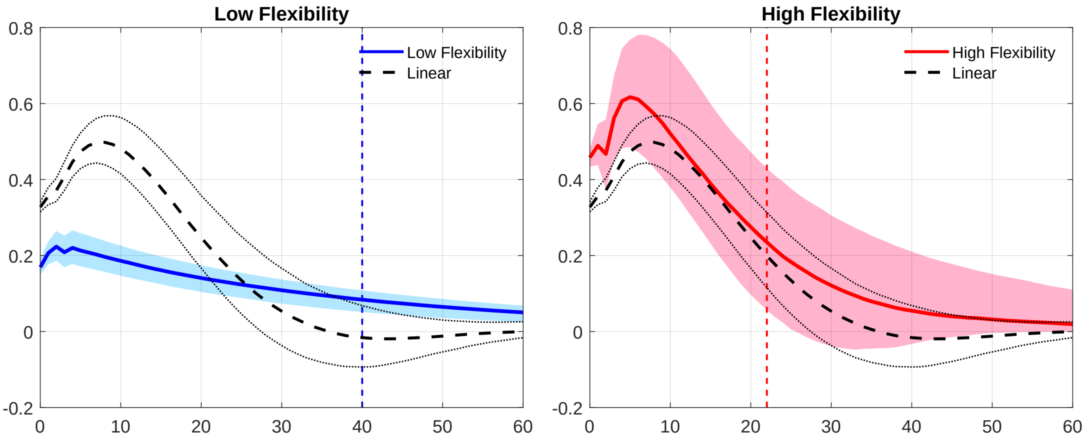


Notes: Figure G.1 reports the responses of inflation to a 1% shock in a STARMA(2,4) model. The left (right) panel graphs the response in the low (high) price flexibility regime. In both cases we also report the response from a (linear) ARMA(2,4) model. 68% confidence intervals are built based on the Markov Chain Monte Carlo (MCMC) method developed in Chernozhukov and Hong (2003). In each of the two charts the vertical line indicates the half-life of the shock.

Figure G.2: PRICE FLEXIBILITY AND INFLATION PERSISTENCE STARMA(1,1)

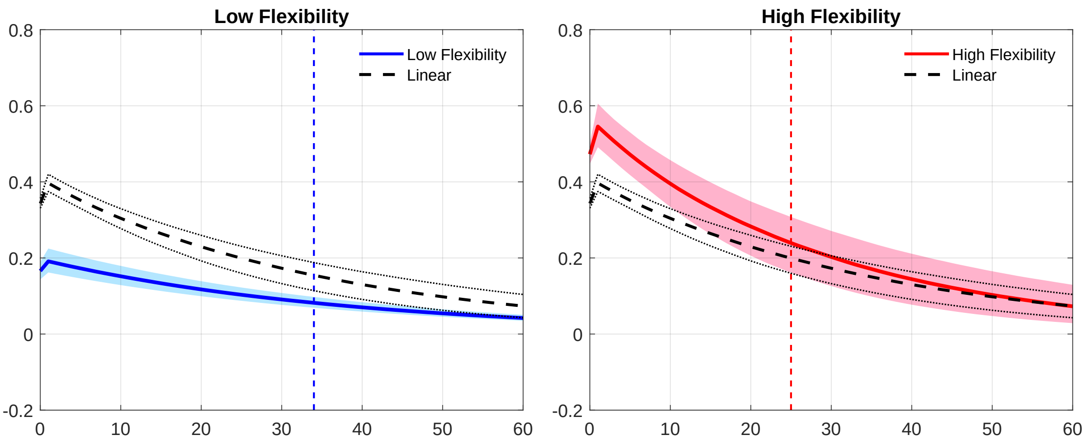


Notes: Figure G.2 reports the responses of inflation to a 1% shock in a STARMA(1,1) model. The left (right) panel graphs the response in the low (high) price flexibility regime. In both cases we also report the response from a (linear) ARMA(1,1) model. 68% confidence intervals are built based on the Markov Chain Monte Carlo (MCMC) method developed in Chernozhukov and Hong (2003). In each of the two charts the vertical line indicates the half-life of the shock.

Figure G.3: PRICE FLEXIBILITY AND INFLATION VOLATILITY

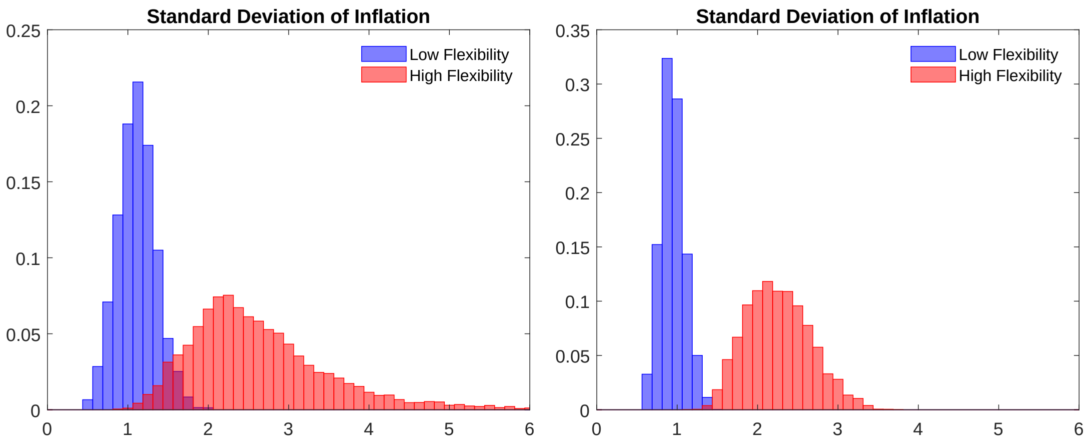


STARMA(2,4) STARMA(1,1)

Notes: Each panel reports the distribution of the estimated inflation volatility in the two regimes. The left panel refers to a STARMA(2,4) model, while the right panel refers to a STARMA(1,1).

Table G.1: INFLATION FORECAST ERRORS AND PRICE FLEXIBILITY

(a) BoE MPC RPIX/CPI Forecast Error Bias: Alternative specification

| Horizon | $G = 0.2$ | | $G = 0.5$ | | $G = 0.8$ | | F$-stat$ | $\tilde{R}^2$ |
|---|---|---|---|---|---|---|---|---|
| 0 | 3.93 | [0.25] | 0.65 | [0.71] | -1.34 | [0.51] | 0.48 | -0.54 |
| 1 | 47.56 | [0.20] | 17.47 | [0.27] | -17.46 | [0.16] | 0.19 | 1.68 |
| 2 | 102.88 | [0.20] | 44.77 | [0.26] | -39.35 | [0.15] | 0.13 | 2.59 |
| 3 | 109.27 | [0.17] | 46.67 | [0.27] | -45.56 | [0.10] | 0.06 | 4.28 |
| 4 | 154.30 | [0.17] | 63.71 | [0.28] | -64.01 | [0.11] | 0.04 | 5.44 |
| 5 | 190.03 | [0.17] | 76.38 | [0.30] | -84.81 | [0.08] | 0.03 | 6.03 |
| 6 | 194.77 | [0.17] | 81.92 | [0.31] | -87.86 | [0.08] | 0.02 | 7.20 |
| 7 | 187.19 | [0.15] | 93.80 | [0.30] | -76.39 | [0.05] | 0.02 | 7.37 |

(b) Market Participants' Forecast Error Bias

| Horizon | $G = 0.2$ | | $G = 0.5$ | | $G = 0.8$ | | F$-stat$ | $\tilde{R}^2$ |
|---|---|---|---|---|---|---|---|---|
| 0 | 1.54 | [0.59] | -2.20 | [0.46] | -2.55 | [0.33] | 0.57 | -0.97 |
| 1 | 37.15 | [0.24] | 13.02 | [0.36] | -14.54 | [0.17] | 0.28 | 0.81 |
| 2 | 88.29 | [0.20] | 39.01 | [0.26] | -34.35 | [0.14] | 0.13 | 2.70 |
| 3 | 91.58 | [0.16] | 38.45 | [0.29] | -40.00 | [0.08] | 0.06 | 4.53 |
| 4 | 129.63 | [0.17] | 52.11 | [0.30] | -55.81 | [0.09] | 0.03 | 5.85 |
| 5 | 153.94 | [0.16] | 58.79 | [0.34] | -73.26 | [0.06] | 0.02 | 6.82 |
| 6 | 166.39 | [0.17] | 66.96 | [0.35] | -80.60 | [0.05] | 0.02 | 7.42 |
| 7 | 164.33 | [0.16] | 82.24 | [0.33] | -71.73 | [0.03] | 0.02 | 7.09 |

Notes: The table reports the results of a quadratic spline regression of the forecast errors as a percentage of realized inflation, $\widetilde{e}_{t+h|t} = 100 \times \frac{e_{t+h|h}}{\pi^{YoY}_{t+h}}$ (for different forecast horizons, $h$, measured in quarters). The regression is specified as $\widetilde{e}_{t+h|t} = a_0 + a_1 (G_{t-1} - 0.5) + a_2 (G_{t-1} - 0.5)^2 + a_3 (G_{t-1} - 0.5)^2 \mathbb{1}_{\{G_{t-1}>0.5\}} G^2_{t-1}$, where $\mathbb{1}_{\{G_{t-1}>0.5\}}$ is an indicator function taking value 1 when $G_{t-1} > 0.5$ and zero otherwise. The upper panel refers to the Bank of England MPC's RPIX/CPI forecast errors, while the bottom panel considers market participants' forecast errors. For each value of $G$, the two columns report the fitted $\hat{\widetilde{e}}_{t+h|t}$ evaluated at different levels of the indicator and the p-value associated with the null hypothesis that $\hat{\widetilde{e}}_{t+h|t}$ is equal to 0 (this is calculated using Newey-West standard errors), respectively. The penultimate column (F-$stat$) reports the p-value of the null hypothesis that all the coefficients associated to the flexibility regime are equal to 0 (i.e., $H_0 : a_1 = a_2 = a_3 = 0$). The last column reports the adjusted R-squared, denoted by $\tilde{R}^2$.

Table G.2: FORECAST UNCERTAINTY AND PRICE FLEXIBILITY

(a) BoE MPC RPIX/CPI (Absolute) Forecast Errors

| Horizon | $\hat{z}_t$ at $G=0.2$ | | $\hat{z}_t$ at $G=0.8$ | | F-$stat$ | $\tilde{R}^2$ |
|---|---|---|---|---|---|---|
| 0 | 0.14 | [0.00] | 0.16 | [0.00] | 0.15 | 2.24 |
| 1 | 0.36 | [0.00] | 0.43 | [0.00] | 0.14 | 2.30 |
| 2 | 0.49 | [0.00] | 0.79 | [0.00] | 0.14 | 2.41 |
| 3 | 0.70 | [0.00] | 1.13 | [0.00] | 0.23 | 1.26 |
| 4 | 0.79 | [0.00] | 1.56 | [0.01] | 0.10 | 3.29 |
| 5 | 0.81 | [0.00] | 1.80 | [0.02] | 0.08 | 3.65 |
| 6 | 0.74 | [0.01] | 1.94 | [0.02] | 0.04 | 5.00 |
| 7 | 0.83 | [0.01] | 1.83 | [0.02] | 0.15 | 2.40 |

(b) BoE MPC RPIX/CPI (Squared) Forecast Errors

| Horizon | $\hat{z}_t$ at $G=0.2$ | | $\hat{z}_t$ at $G=0.8$ | | F-$stat$ | $\tilde{R}^2$ |
|---|---|---|---|---|---|---|
| 0 | 0.03 | [0.00] | 0.04 | [0.00] | 0.23 | 1.28 |
| 1 | 0.26 | [0.01] | 0.40 | [0.01] | 0.17 | 1.96 |
| 2 | 0.58 | [0.03] | 1.33 | [0.04] | 0.14 | 2.34 |
| 3 | 0.88 | [0.08] | 3.20 | [0.08] | 0.12 | 2.69 |
| 4 | 0.67 | [0.37] | 6.07 | [0.10] | 0.06 | 4.31 |
| 5 | 0.30 | [0.78] | 8.47 | [0.13] | 0.04 | 5.25 |
| 6 | 0.06 | [0.96] | 9.90 | [0.13] | 0.03 | 5.68 |
| 7 | 1.05 | [0.41] | 8.98 | [0.15] | 0.18 | 1.92 |

(c) Market Participants' (Absolute) Forecast Errors

| Horizon | $\hat{z}_t$ at $G=0.2$ | | $\hat{z}_t$ at $G=0.8$ | | F-$stat$ | $\tilde{R}^2$ |
|---|---|---|---|---|---|---|
| 0 | 0.15 | [0.00] | 0.21 | [0.00] | 0.18 | 1.84 |
| 1 | 0.40 | [0.00] | 0.43 | [0.00] | 0.11 | 2.87 |
| 2 | 0.46 | [0.00] | 0.81 | [0.00] | 0.08 | 3.56 |
| 3 | 0.63 | [0.00] | 1.18 | [0.00] | 0.13 | 2.58 |
| 4 | 0.70 | [0.00] | 1.63 | [0.01] | 0.04 | 5.12 |
| 5 | 0.73 | [0.00] | 1.87 | [0.02] | 0.04 | 5.30 |
| 6 | 0.67 | [0.01] | 2.08 | [0.01] | 0.02 | 7.35 |
| 7 | 0.73 | [0.01] | 1.99 | [0.02] | 0.06 | 4.64 |

(d) Market Participants' (Squared) Forecast Errors

| Horizon | $\hat{z}_t$ at $G=0.2$ | | $\hat{z}_t$ at $G=0.8$ | | F-$stat$ | $\tilde{R}^2$ |
|---|---|---|---|---|---|---|
| 0 | 0.08 | [0.15] | 0.18 | [0.19] | 0.53 | -0.79 |
| 1 | 0.46 | [0.06] | 0.136 | [0.04] | 0.18 | 1.77 |
| 2 | 0.53 | [0.04] | 1.40 | [0.04] | 0.13 | 2.759 |
| 3 | 0.76 | [0.13] | 3.37 | [0.08] | 0.11 | 3.00 |
| 4 | 0.51 | [0.48] | 6.40 | [0.10] | 0.05 | 4.61 |
| 5 | 0.08 | [0.94] | 8.94 | [0.12] | 0.03 | 5.79 |
| 6 | -0.24 | [0.83] | 10.58 | [0.12] | 0.02 | 6.78 |
| 7 | 0.68 | [0.58] | 9.80 | [0.13] | 0.10 | 3.30 |

Notes: The table reports the results of a quadratic spline regression of the absolute (LHS) and squared (RHS) forecast errors (for different forecast horizons, $h$, measured in quarters) on a quarterly average of an indicator of the normalized price flexibility index, $\widetilde{\mathcal{F}}$: $G_{t-1} = G(\widetilde{\mathcal{F}}_{t-1};\gamma) = (1+e^{-\gamma\widetilde{\mathcal{F}}_{t-1}})^{-1}$. The regression is specified as $z_t = a_0 + a_1(G_{t-1}-0.5) + a_2(G_{t-1}-0.5)^2 + a_3\mathbb{1}_{\{G_{t-1}>0.5\}}(G_{t-1}-0.5)^2$, where $\mathbb{1}_{\{G_{t-1}>0.5\}}$ is an indicator function taking value 1 when $G_{t-1} > 0.5$ and zero otherwise, $z_t = |e_{t+h|t}|$ (tables (a) and (c)) and $z_t = e^2_{t+h|t}$ (tables (b) and (d)). The upper panels refer to the Bank of England MPC's RPIX/CPI forecast errors, while the bottom panels consider market participants' forecast errors. For each value of $G$, the two columns report the fitted value evaluated at different levels of the indicator and together the p-value associated with the null hypothesis that this value is equal to 0 (this is calculated using Newey-West standard errors), respectively. The penultimate column (F-$stat$) reports the p-value of the null hypothesis that all the coefficients associated to the flexibility regime are equal to 0 (i.e., $H_0: a_1 = a_2 = a_3 = 0$). The last column reports the adjusted R-squared, denoted by $\tilde{R}^2$.