Inflation and Price Flexibility*

Ivan Petrella[†]
Collegio Carlo Alberto, University of Turin
University of Warwick & CEPR

Emiliano Santoro[‡]
Catholic University of Milan

Yannik Winkelmann[§]
University of Tübingen

October 2024

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-COVID-19 inflation surge. Tracking real-time movements in price flexibility is crucial for assessing inflation dynamics and has the potential to improve forecasts and inform monetary-policy strategies.

JEL codes: C22, E30, E31, E37.
Keywords: Inflation, price flexibility, Ss models, state-dependence.

*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, and Joseph Vavra.

[†]Collegio Carlo Aberto, University of Turin, 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 nominal 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 measured the response of prices to nominal *stimuli*, little attention has been paid to the sources and characteristics of changes in price flexibility.[1] Most notably, 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 aim to address this gap by demonstrating that state-dependence in price flexibility is crucial for generating accurate inflation projections, enhancing our understanding of inflation dynamics, and supporting more effective inflation-targeting mandates.

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 various 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). Both functions vary substantially over time, displaying some distinct asymmetries that are extremely informative about the microeconomics of price setting. Notably, changes in the price-adjustment cost structure triggered by the Great Recession have led to a pronounced downward shift and flattening of the adjustment hazard, implying greater inaction in price adjustment. During this time span, we observe a significant shift to a price-setting regime characterized by a substantial increase in the dispersion of price changes, along with a concurrent decline in the frequency of price adjustment. In other words, price changes became larger but less frequent, compared to the pre-Great Recession period. This pattern reverts at the onset of the COVID-19 recession.

At the aggregate level, we show that the response of inflation to a one-off nominal shock, *price flexibility*, varies significantly over time, peaking during 2008-2011—almost double the pre-recession level—before halving by 2016 and then climbing steadily to its latest peak during the COVID-19 crisis. Concurrently, inflation has fluctuated sharply since the onset of the Great Recession, being almost twice as volatile, even excluding the post-COVID-19 sample. We suggest that changes in price flexibility shape inflation dynamics. The idea is 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 recent years.

[1] In this regard, Caballero and Engel (1993b) and Berger and Vavra (2017) are notable exceptions.

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*.[2] While the intensive margin was typically the primary driver of price flexibility up until the Great Recession, state dependence 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.

The central message of this paper is that regime shifts in price flexibility are key to understand inflation dynamics. To see this, we estimate a regime-dependent model of inflation, setting price flexibility as a state variable. 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 test if the Bank of England and broader market participants account for such state dependence in their forecasting practices, and examine whether their forecast errors are uncorrelated with flexibility regimes. Inflation forecasts are generally unbiased when aggregate price flexibility is low or average, but evidence suggests a significant negative bias during periods of high price flexibility, even excluding the post-COVID-19 part of the sample. In fact, we devise a counterfactual experiment in which the Bank of England's forecasts at the onset of the last inflation spike are adjusted for the high-flexibility bias, showing how this factor should have suggested to abandon the view of a rapid return of inflation to target.

Our work has important implications for monetary authorities aiming to stabilize inflation. 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. This 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 underscores the practical importance of tracking real-time changes in price flexibility to improve the assessment of inflation persistence and shock passthrough, ultimately supporting more informed monetary policy decisions.

Related literature Our work relates to a number of studies that have examined the connection between micro price changes and aggregate inflation.[3] The paper that connects most

[2] 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.

[3] See, among others, Bils and Klenow (2004), Dotsey and King (2005), Alvarez et al. (2006), Gertler and Leahy (2008), Klenow and Kryvtsov (2008), Nakamura and Steinsson (2008), Gagnon (2009), Costain and Nakov (2011), Midrigan (2011), Nakamura et al. (2011), Alvarez and Lippi (2014), Karadi and Reiff (2019), Berardi et al. (2015), Alvarez et al. (2016), Nakamura et al. (2018).

closely to our analysis is Berger and Vavra (2017), who report that price flexibility is time-varying. Relative to this work, our novel contribution is to show how accounting for time variation in price flexibility improves our understanding of inflation dynamics: in this respect, we document that inflation is less persistent and more volatile in periods of relatively high price flexibility, an show how neglecting this fact can lead to a negative prediction bias. We also relate to Luo and Villar (2021) by highlighting the importance of time variation in the hazard function for quantifying the pass-through of monetary shocks to prices. In fact, movements in firms' incentives to price adjustment are shown to be prominent (in this respect, see also Hobijn et al., 2006).

Our work also builds on a number of papers that devise and estimate specific structural models that connect movements in the distribution of price changes to price flexibility (see, e.g., Midrigan, 2011, Alvarez et al., 2016 and Vavra, 2014, among others). An empirical limitation of these models is their reliance on specific shocks to price-setting units, whereas our approach is more agnostic. This is advantageous, as it allows us to avoid committing to a specific price-adjustment mechanism *ex ante*. In fact, we demonstrate that the observed pattern of time variation in the price change distribution aligns with a different mix of first- and second-moment shocks, as well as shifts in firms' endogenous incentives to adjust their prices.

We also relate to some recent empirical contributions employing individual consumer prices from the UK. In this respect, Bunn and Ellis (2012) have been among the first to investigate the key characteristics of the frequency of price setting and the hazard functions implied by the microdata from the Office for National Statistics (ONS), while Dixon et al. (2020) have focused on the impact of the Great Recession on price setting. The latter, in particular, attributes little importance to endogenous macroeconomics effects on pricing, while our evidence points to a certain prominence of state dependence in price setting, and more so during the sample that is not accounted for in their analysis (i.e., post-2013), during which the extensive margin of price adjustment overcomes the intensive one (i.e., the frequency of adjustment) 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 importance. Our application underlines the role of the selection effect for aggregate inflation (see, on this, Carvalho and Kryvtsov, 2021 and references therein), along with stressing its time variation. Finally, Chu et al. (2018) emphasize that the distribution of price changes can be used to forecast inflation. We extend this by showing that price flexibility—capturing essential information from key micro price statistics—also provides valuable insights for predicting inflation.

Structure The rest of the paper is organized as follows. Section 2 discusses the key characteristics of the ONS microdata on consumer prices. Section 3 reviews the generalized Ss model and takes it to the data. Section 4 examines time variation in the distribution of price gaps and the adjustment hazard, as well as the relative importance of adjustments along the intensive and the extensive margin at different points in 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. Prices are collected on a monthly basis, for more than 700 categories of goods and services, and published with a month lag. Our sample covers the 1996:M2-2024:M8 time window, thus resulting into about 37.4 million observations (see Table 1). 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.[4] 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.[5] 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

[4]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.

[5]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.

Table 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.

sells the item for £22, and one for £26. In February 2013 the price quotes for the same type of good were recorded at £25 and £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.[6] 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 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.[7]

To aggregate the individual price quotes into a single price, we also make use of the following weights produced by the ONS:[8] 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);[9] the *stratification* weights, which reflect the fact that purchasing patterns may

[6]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.

[7]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 ($< 0.02\%$). Appendix A provides additional details on the construction of the dataset.

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

[9]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.

differ markedly by region or type of outlet;[10] finally, the *item* and *COICOP* weights reflect consumers' expenditure shares in the national accounts.

2.1 Variable definition

After deriving our price quotes in line with the criteria set out above, it is important to make a distinction between regular and temporary price changes such as *sales*, which tend to behave significantly differently from that of regular prices (see Eichenbaum et al., 2011 and Kehoe and Midrigan, 2015). To this end, we first exclude all the price quotes to which the ONS attaches a sales indicator.[11] As a second step, we implement a symmetric V-shaped filter, as defined by Nakamura and Steinsson (2010b), for the remaining price quotes. According to the filter, the sale price of item i at time t, $P^s_{i,t}$, is identified as follows: i) it is lower than last period's price (i.e., $P^s_{i,t} < P_{i,t-1}$) and ii) the next period's price is equal to last period's price (i.e., $P_{i,t+1} = P_{i,t-1}$). A recovery price $P^r_{i,t}$, instead, meets the following criteria: i) it is greater than last period's price (i.e., $P^r_{i,t} > P_{i,t-1}$) and ii) it is such that $P^r_{i,t} = P_{i,t-2}$. Once a price quote has been identified as being a sale or a recovery price, we discard it from the sample.[12]

Item substitutions are a further reason of concern when trying to identify price trajectories, as they require a certain degree of judgment to establish what portion of a price change is due to quality adjustment, and which component reflects a pure price adjustment. Product substitutions occur whenever an item in the sample has been discontinued from its outlet, and the ONS identifies a similar replacement item to the price going forward. Therefore, it is reasonable to expect that product turnovers are followed by price changes that either reflect uncaptured quality changes (Bils, 2009), or simply reflect a low-cost opportunity to reset prices that has nothing to do with the underlying sources of price rigidity, as argued by Nakamura and Steinsson (2008). In line with previous contributions, we interrupt a trajectory whenever it encounters a substitution flag, as indicated by the ONS (see, e.g., Berardi et al., 2015, Berger and Vavra, 2017, and Kryvtsov and Vincent, 2021).

Table 1 shows that, after these preliminary steps, we are down to a monthly average of $88,600$ price quotes. As a final step, we define the price change of item i at time t as $\Delta p_{i,t} = \log\left(P_{i,t}/P_{i,t-1}\right)$.[13]

[10] 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).

[11] For a price to be marked as being associated with a sale, the ONS requires the latter to be available to all potential costumers—so as to exclude quantity discounts and membership deals—and that it only entails a temporary or an end-of-season price reduction. This definition excludes clearance sales of products that have reached the end of their life cycle.

[12] See also Nakamura and Steinsson (2008) and Vavra (2014). As an alternative approach, in place of the price associated with a sale, Klenow and Kryvtsov (2008) report the last regular price, until a new regular price is observed. Our results are robust to this approach.

[13] 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 definition has the advantage of being bounded and less sensitive to outliers. The results—virtually unchanged with respect to the ones we report—are available from the authors, upon request.

2.2 Stylized facts

This section presents some key facts about the behavior of the ONS microdata, and their implications for inflation dynamics. We start by re-writing inflation 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 (Δp_t):

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

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 $fr_t^{-1} \sum_i \omega_{i,t} \mathbb{1}_{\{\Delta p_{i,t} \neq 0\}} \Delta p_{i,t}$. All the statistics derived from microdata display a pronounced seasonality (see, e.g., Alvarez et al., 2006), which we remove by computing the 12-month moving average.

The top panels of Figure 1 report fr_t and Δp_t, respectively. As expected, the average price change tracks CPI inflation closely, at least until the end of the Great Recession, to then resume a tight comovement towards the end of 2015. As for the frequency of adjustment, it reflects a contractionary trend beginning with the inflation decline initiated in 2012—falling well below its previous sample average—before showing a significant reversal during the COVID-19 pandemic.[14] As for the frequency of price adjustments, its value 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 "only" around 4%. In the next section, we will show how these most recent figures can be explained by the fact that predetermined price adjustments—as measured by frequency—represent only part of the overall adjustment activity. Indeed, state dependence in price setting will be shown to play a prominent role, with significant implications for inflation volatility and persistence.

[14] The average frequency of price adjustment prior to its drop is slightly below 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. By contrast, Bunn and Ellis (2012) include utility prices and sales, while Dixon and LeBihan (2012), Dixon et al. (2020) and Dixon and Tian (2017) only include sales.

Figure 1: Frequency, Average Price Changes, and Dispersion

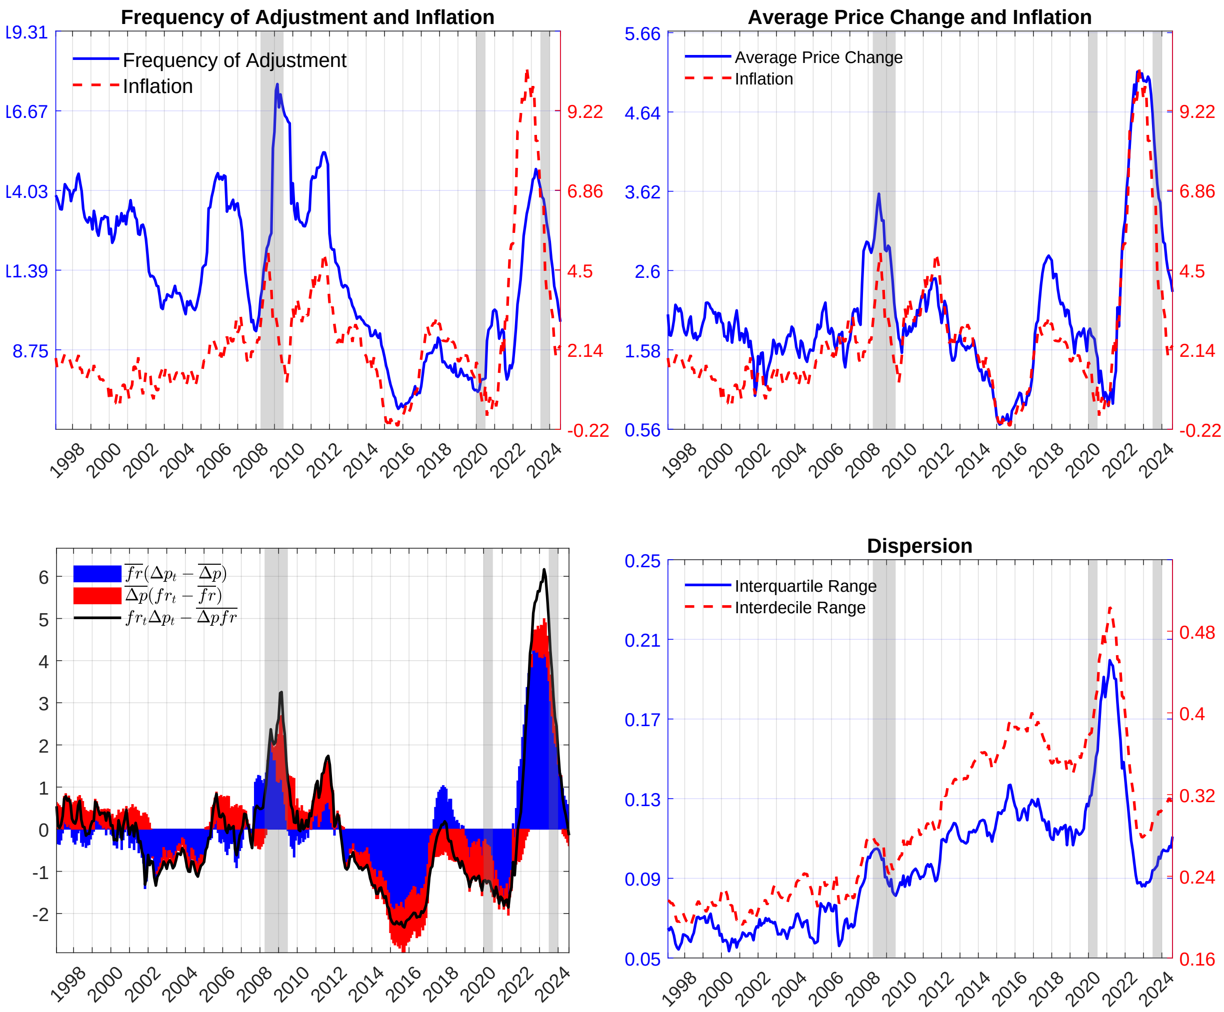

Notes: The frequency of price adjustment, fr_t , measures the share of prices being adjusted in every month, and is computed as $\sum_i \omega_{i,t} \mathbb{1}_{\{\Delta p_{i,t} \neq 0\}}$, where $\omega_{i,t}$ denotes the CPI weight associated to good i at time t, and $\mathbb{1}_{\{\Delta p_{i,t} \neq 0\}} = 1$ if $\Delta p_{i,t} \neq 0$ and zero otherwise. The average price, instead, is denoted by Δp_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: $\pi_t - \overline{fr\Delta p} = \overline{fr}(\Delta p_t - \overline{\Delta p}) + \overline{\Delta p}(fr_t - \overline{fr}) + (\Delta p_t - \overline{\Delta p})(fr_t - \overline{fr})$. 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.

Table 2: PRICING MOMENTS AND MACROECONOMIC VARIABLES 1997:M1-2024:M8

	Rotemberg Filter			Year-over-Year Filter		
	fr_t	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	fr_t	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$
$\mathbb{I}(\texttt{Rec.})$	0.600	1.562	2.880	0.862	1.114	2.821
y_t	-0.068	-0.278**	-0.267**	-0.201	-0.174	-0.145
π_t	0.292**	-0.424***	-0.444***	0.355***	-0.596***	-0.665***
fr_t	–	-0.377***	-0.431***	–	-0.321**	-0.427***

	Quadratic Detrending			Hamilton Filter		
	fr_t	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	fr_t	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$
$\mathbb{I}(\texttt{Rec.})$	2.586**	1.263	0.959	3.423**	2.055	3.775
y_t	-0.188	-0.292**	-0.213*	-0.107	-0.409***	-0.372***
π_t	0.560***	-0.483***	-0.565***	0.424***	-0.613***	-0.641***
fr_t	–	-0.428***	-0.607***	–	-0.422***	-0.523***

Notes: The first row of each table reports the value of the correlation coefficients associated with a recession dummy. The variables involved are fr_t, the frequency of price adjustment, and two quantilic measures of dispersion of price changes, where $q_{n,t}$ measures the $n-$th quantile of the distribution of price changes. The remaining rows report pairwise correlations with y_t, which denotes detrended GDP, π_t, standing for aggregate CPI inflation, and fr_t. Aside of the inflation rate, all series are obtained by detrending their raw counterparts by means of: i) Rotemberg (1999) version of the HP filter, which sets the smoothing coefficient so as to minimize the correlation between the cycle and the first difference of the trend estimate (top left panel); ii) linear and quadratic detrending of the series (bottom left panel); iii) year-over-year change, as suggested in Stock and Watson (2019) (top right panel); (iv) two-years difference, as suggested by Hamilton (2018) (bottom right panel). ***/**/* indicates statistical significance at the 1/5/10% level, respectively (the standard errors for the cyclicality calculation are adjusted for serial correlation using a Newey-West correction with optimal lag length).

The bottom-left panel of the figure looks deeper into the cyclical behavior of the rate of inflation, decomposing inflation deviations from the sample mean as in, e.g., Gagnon (2009):

$$\pi_t - \overline{fr\Delta p} = \overline{fr}(\Delta p_t - \overline{\Delta p}) + \overline{\Delta p}(fr_t - \overline{fr}) + (\Delta p_t - \overline{\Delta p})(fr_t - \overline{fr}), \tag{2}$$

where i) the first term captures the variation in aggregate inflation associated with changes in the average change of those prices that change from one period to the next one, ii) the second one accounts for variation in the frequency of adjustment, and iii) the last term accounts for joint variation of the two moments around their respective sample means. Notably, only about half of inflation variability is explained by that of average price changes, 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. Focusing on the post-COVID-19 peak in inflation, instead, most of the deviation from the mean can be attributed to the price-change component (in line with evidence documented by Montag and Villar, 2022, for the US). However, changes in the frequency of price adjustment, coupled with the interaction 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.[15] 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, 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.[16]

Notably, also the detrended versions of the two statistics consistently display negative correlation. To see this, Table 2 reports the correlation between different detrended measures of the frequency of adjustment and the dispersion of price changes, as well as the degree of comovement of each of the two with detrended GDP. Potential spurious correlation arising from low-frequency movements in the series of interest is set aside through detrendization—using different filters—with the exception of the inflation rate. Albeit denoting a marked increase in

[15] 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.

[16] Figure B.1 in Appendix B shows 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.

correspondence of the two recessionary episodes, the frequency of price adjustment displays no particular comovement with the business cycle. As for the dispersion of price changes, some degree of countercyclicality is instead detected.

As stressed by Vavra (2014), the cyclical properties of these statistics and 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.[17] In this respect, Berger and Vavra (2017) show how this empirical setting provides a good fit to the data generated by different structural models (e.g., Golosov and Lucas, 2007 and Nakamura and Steinsson, 2010a). Second, as we allow for time variation in the determinants of price adjustment, we can estimate the model over each cross section of price microdata, matching different price-setting statistics.

To contextualize the framework assume that, due to price rigidities, the log of firm i's actual price may deviate from the log of the target or reset price, which is denoted by p^*_{it}. Thus, we define the price gap as $x_{it} \equiv p_{it-1} - p^*_{it}$, implying that a positive (negative) price gap is associated with a falling (increasing) price when the adjustment is actually made. A price is adjusted when the associated price gap is large enough, and $p_{it} = p^*_{it}$ after the adjustment has taken place. Assuming l_{it} periods since the last price change, the adjustment reflects the cumulated shocks: $\Delta p_{it} = \sum_{j=0}^{l_{it}} \Delta p^*_{it-j}$, with $\Delta p^*_{it} = \mu_t + \upsilon_{it}$, where μ_t is a shock to nominal demand and υ_{it} is an idiosyncratic shock.

Caballero and Engel (2007) assume *iid* idiosyncratic shocks to the adjustment cost. Thus, by integrating over their possible realizations, we obtain the adjustment hazard, $\Lambda_t(x)$. This is defined as the (time t) probability of adjusting—prior to knowing the current adjustment cost draw—by a firm that would adjust by x in the absence of adjustment costs (i.e., as if the adjustment cost draw was equal to zero). Caballero and Engel (1993a) prove that the probability of adjusting is non-decreasing in the absolute size of a firm's price gap (i.e., the so-called 'increasing hazard property'). 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

[17] 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.

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., $\Lambda_t(x) = \Lambda_t > 0,\ \forall x$).

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.[18] 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 θ_t and $\psi_t > 0$ capture the location and the scale of the distribution, whereas $0 < \varrho_t < 1$ accounts for its 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 ($\nu_t \to \infty$). Most importantly, it can capture intermediate cases between the Normal and the Laplace distribution, which is 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:[19]

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

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

[18] 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. In fact, they show—for the case of symmetric functional forms—that $\Lambda_t(x)$ and $f_t(x)$ are fully encoded in distribution of price changes and fr_t. While our estimates do not explicitly take into account this mapping, in Section 4 we show that our estimates return measures of the cumulative response of prices to a monetary stimulus which are consistent with alternative measures of money non-neutrality, such as the one proposed by Alvarez et al. (2016).

[19] Unlike Berger and Vavra (2017), we allow for asymmetry in the hazard function, along with the distribution of price gaps. In fact, the distribution of price changes is characterized by some sizeable asymmetry, as well as a marked tendency of skewness to vary over time. Moreover, there is a sizable difference between the frequency of price adjustment associated with positive and negative price changes (both statistics are reported in Figure H.1). By allowing both the distribution of price gaps and the hazard function to be asymmetric, we avoid assuming that the underlying asymmetries in the data are entirely driven by either of the two functions. In support of this choice, the assumption of symmetric branches in the hazard function is generally rejected by the data.

Villar (2021).[20]

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 (see Groeneveld, 1998).[21] 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.[22] All the estimated parameters and the derived statistics inherit some pronounced seasonal variation from the raw data. Thus, we report their 12-month moving-average counterparts.

Identification Appendix C 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 the 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 $\Lambda_t(x)$.

3.2 Estimates

Figure 2 graphs the estimated price gap distributions and the hazard functions.[23] Notably, $f_t(x)$ widens considerably after 2010, with the post-COVID-19 inflationary episode marked by a significant downward shift, followed by a rebound. Since early 2021, a shift in the price

[20] We have also checked that our results are robust to plausible variations to the specification of these functional forms. Using a Pearson Type 7 distribution, 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.

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

[22] 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).

[23] Figures D.1 and D.2 report the estimated parameters, while Figure D.3 reports the fit of selected data moments, and shows that the empirical model is able summarize the main stylized facts.

Figure 2: Estimated Price Gap Distributions and Hazard Functions

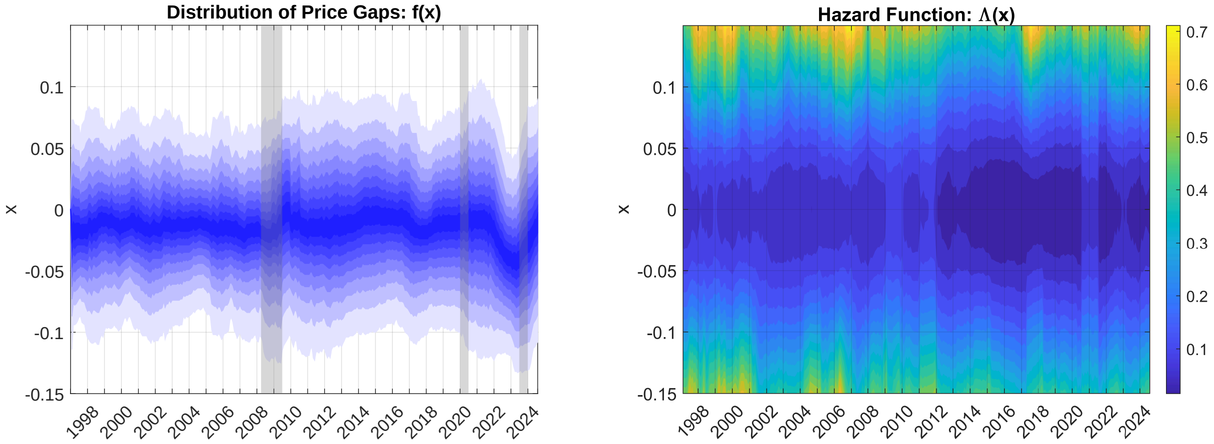

Note: Estimated Price Gap Distributions (left panel) and Hazard Functions (right panel), for each month in the sample. The shaded vertical band in the left panel indicates the duration of recessionary episodes.

gap distribution indicates that a large fraction of price quotes have fallen below their optimal level, suggesting that substantial price increases have become likely in the subsequent period. As for $\Lambda_t(x)$, its shape changes significantly over the entire sample, with periods of high and low probability of adjusting prices that alternate over time. While the 2009-2012 time span is characterized by a relatively high and steep hazard function—most likely as the result of three VAT changes taking place over this short time window—[24] we observe a prolonged period, up until the pandemic, during which it shifts downwards and flattens.

Relative movements in $f_t(x)$ and $\Lambda_t(x)$ may help explain the emergence of diverging patterns in the frequency of adjustment and the dispersion of price changes, as those occurring between the Great Recession and COVID-19. Figure 3 digs into this: in the left panel we report the adjustment thresholds associated with 5% and 7% hazard probabilities, while the right panel reports the share of prices within the corresponding thresholds. Notably, before the Great Recession fewer than 10% of prices had less than a 5% probability of being adjusted. By 2020, this figure had increased to 50%, meaning half of all price quotes exhibited less than a 5% probability of adjustment. In the latter part of the sample, the magnitude of price changes increased significantly, surpassing any previously observed levels: a typical price quote with a 5% probability of adjustment had a price gap of roughly 4%, compared to just 2.5% before the Global Financial Crisis. Eventually, the distribution of price changes is characterized by fewer but larger changes, with a skew towards price increases.

According to conventional menu cost models, a decline in the frequency of adjustment, coupled with a surge in its dispersion, may be rationalized by an expansion in the inaction region that overcomes the effects of a positive shift in the dispersion of price gaps.[25] By the

[24]Namely, a reduction, from 17.5% to 15%, on December 1, 2008, followed by two hikes: one, up to 17.5%, on January 1, 2010, and one, further up to 20%, on January 4, 2011. In particular, the first two VAT changes are associated with a marked shift in the estimated parameters of the hazard function, as it is visible from Figure D.2.

[25]Appendix E reports a stylized menu cost model that stresses how changes in the incentives firms face when

Figure 3: Inaction in price adjustment

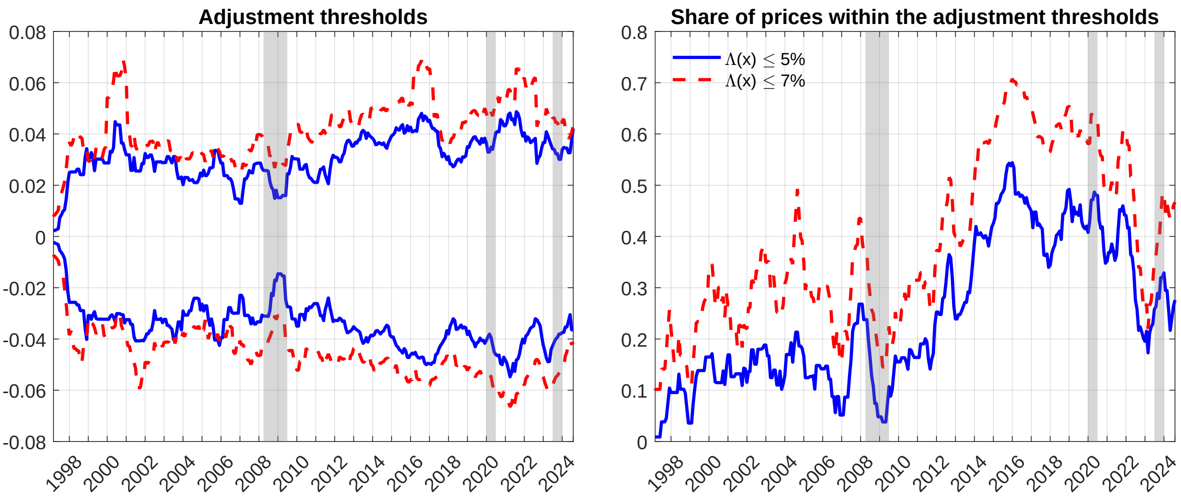

Note: Left panel: adjustment thresholds associated with 5% and 7% hazard probabilities; right-panel: share of prices within the adjustment thresholds. The shaded vertical bands indicate the duration of recessionary episodes.

end of the sample about five times as many firms are likely not to adjust, as compared with the pre-2010 time window. This stands as indirect evidence that increasing price rigidity, as captured by the downward shift in the hazard function, dominates the increase in the dispersion of $f_t(x)$. Note also that nominal rigidity appears more pronounced with positive price gaps than with negative ones, indicating a greater degree of downward price stickiness. On a more general note, changes in the shape of the distribution of price gaps, coupled with a flattening of the hazard function, imply that non-predetermined price adjustments—which are more likely to occur for large price gaps—have played an increasingly important role in the recent past. This explains why, in spite of the average frequency of price changes being much lower over the last 15 years, inflation volatility has been significantly higher. For context, in the post-Great Recession period inflation volatility has been almost 5 times greater than in the previous 10 years. Even excluding post-2020 observations, inflation in the second part of the sample is more than twice as volatile as in the first one. Strikingly, over the past 15 years, (year-on-year) inflation has been outside the 1%-3% range for 34 quarters—i.e., roughly 50% of the time—compared to just 5 quarters in the first decade of the sample.

4 Inspecting price setting in a time-varying environment

The estimates of 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

deciding to change prices can provide us with a rationale for diverging movements in the dispersion of price changes and the frequency of adjustment.

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 be naturally decomposed into an intensive and an extensive margin component. On one hand, 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. On the other hand, 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.[26] 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.

Figure 4 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). $\mathcal{F}_t$ displays sizable variation over time, rising substantially during the Great Recession, and showing some secondary peaks during the following recessions. This is consistent with our analysis of the distribution of price gaps. After the Great Recession, both the intensive and the extensive margin of price adjustment contract, though the fall in the former is much more abrupt (in line with the sustained drop in the frequency of adjustment). Concurrently, the extensive margin takes over during this contraction in aggregate price flexibility: this can be explained upon the fact that, over this phase, the expansion in the set of price gaps denoting an extremly low likelihood of reset implies that fewer quotes are pushed near the adjustment boundaries. Even after both margins revert in 2016, the extensive margin remains largely dominant, though.

To see why we observe such a switch in the relative contribution of the two margins, it is useful to recall Caballero and Engel (2007) and their transformation of (6):

$$\mathcal{F}_t = \int \Lambda_t(x) f_t(x) \left[1 + \eta_t(x)\right] dx \tag{7}$$

where $\eta_t(x) = x\Lambda_t'(x) / \Lambda_t(x)$ is the elasticity of the hazard function with respect to the price gap. A downward shift in the hazard function magnifies $\eta_t(x)$ and, as a result, the importance of the extensive margin relative to the intensive one. This is exactly what happens after the Great Recession, as it can be appreciated by inspecting the estimated constant of the hazard function (see Figure D.2 in Appendix D). Alternatively, the same point can be made by approximating the flexibility index as $F_t \cong Int_t + 2\left[Int_t - \Lambda_t(0)\right]$:[27] from this, it is clear how

[26] 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.

[27] For a formal proof, please refer to Caballero and Engel (2007).

Figure 4: PRICE FLEXIBILITY AND DIFFERENT MARGINS OF PRICE ADJUSTMENT

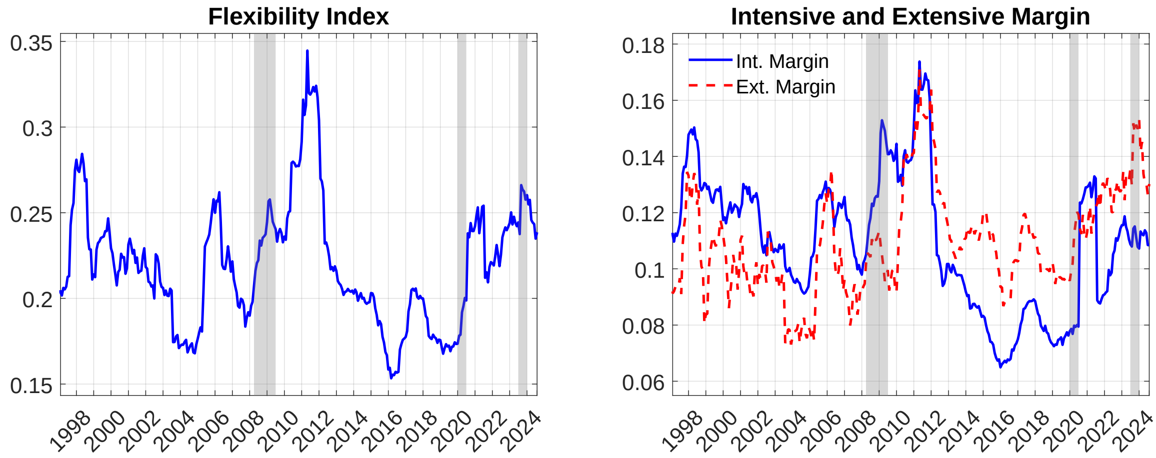

Notes: The left panel reports the estimated index of price flexibility, which is decomposed in the right panel between the intensive and the extensive margin of price adjustment. The shaded vertical bands indicate the duration of recessionary episodes.

a downward shift in a_t—which is equivalent to lowering $\Lambda_t(0)$—translates into an increase in the importance of the extensive margin relative to the intensive one, *ceteris paribus*. It is important to recognize that such a shift in the hazard function is broadly in line with an increase in market power that determines a drop in the cost of being away from the optimal price. In fact, this view is consistent with the sizable increase in the markup that has been observed during the post-recession period, as recently documented by DeLoecker and Eeckhout (2018) and Bell and Tomlinson (2018) for the UK economy.[28]

Turning back to the movements in price flexibility, these 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 the pandemic 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,[29] representing the main contributor to price flexibility and, consequently, inflation volatility becomes particularly high (see Figure 4). This is central to our analysis in the remainder of the paper.

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

[28]Both papers show that the markup has displayed only a modest increase in the 1996-2007 period, while increasing substantially thereafter.

[29]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 5: Comparison with Alvarez et al. (2016)

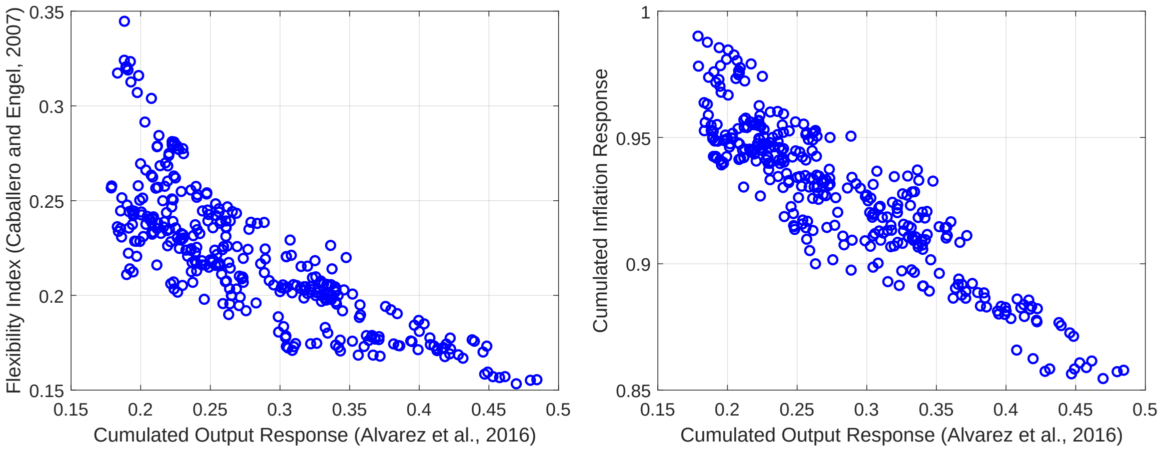

Note: The left panel of the figure reports a scatter plot of the cumulated output response to a monetary policy shock, as computed by Alvarez et al. (2016), against the index of price flexibility, as computed by Caballero and Engel (2007). The right panel, instead, features a scatter plot of the cumulated 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.

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 ($\text{Kur}(\Delta p_i)$) to the frequency of price adjustments ($fr(\Delta p_i)$), $\frac{\delta}{6\epsilon}\frac{\text{Kur}(\Delta p_i)}{fr(\Delta p_i)}$, where δ denotes the size of the monetary shock and ϵ is the elasticity of labor supply.

The left panel of Figure 5 provides a direct comparison between $\mathcal{F}_t$ and $\frac{\text{Kur}_t(\Delta p_i)}{fr_t(\Delta p_i)}$, which is obtained by computing a quantilic version of the kurtosis of price changes, estimating it for each month of the sample.[30] 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 5 shows a striking (negative) correlation of our cumulative measure of price stickiness with the metric elaborated by Alvarez et al. (2016).[31] This reinforces our confidence in the empirical framework we rely upon to track movements in the price gap distribution and the hazard function. The following section provides an in-depth analysis of the relationship between fluctuations in price flexibility and inflation dynamics.

[30]Specifically, this obtained as $\frac{q_{90,t}-q_{62.5,t}+q_{37.5,t}-q_{10,t}}{q_{75,t}-q_{25,t}}$ (see, e.g., Groeneveld, 1998).

[31]The time-series profile of the two measures of money non-neutrality can be observed in Figure H.2.

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 use estimates from the Ss model 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.[32] Figure 6 illustrates that inflation is more responsive and less persistent during periods of relatively high price flexibility. In view of this, price flexibility likely holds valuable information for analyzing inflation dynamics. This insight arises naturally in environments characterized by state-dependent pricing. The remainder of this section examines whether aggregate inflation exhibits non-linearities consistent with these properties and discusses the 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 and Terasvirta (1993).[33] 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).[34]

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

$$\begin{aligned}\pi_t \;&=\; G\left(\widetilde{\mathcal{F}}_{t-1},\gamma\right)\left(\phi_0^H+\sum_{j=1}^{p}\phi_j^H\pi_{t-j}+\varepsilon_t^H+\sum_{i=1}^{q}\theta_i^H\varepsilon_{t-i}^H\right)\\ &\quad+\left[1-G\left(\widetilde{\mathcal{F}}_{t-1},\gamma\right)\right]\left(\phi_0^L+\sum_{j=1}^{p}\phi_j^L\pi_{t-j}+\varepsilon_t^L+\sum_{i=1}^{q}\theta_i^L\varepsilon_{t-i}^L\right),\end{aligned} \tag{8}$$

with $\varepsilon_t^i \sim N\left(0,\sigma_i^2\right)$ for $i=\{L,H\}$. Moreover, we set $G\left(\widetilde{\mathcal{F}},\gamma\right) = (1+e^{-\gamma\widetilde{\mathcal{F}}})^{-1}$, where $\widetilde{\mathcal{F}}$

[32] 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 Appendix F.

[33] The STARMA(p,q) model also generalizes the threshold ARMA(p,q) model in DeGooijer (2017).

[34] 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.

Figure 6: Impulse Responses from the Ss Model

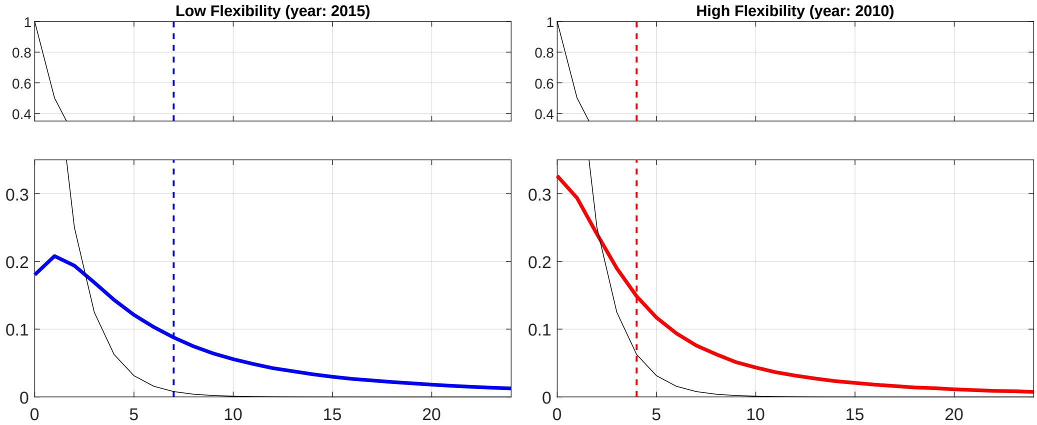

Note: The graphs display the average inflation response to a 1% aggregate nominal shock, μ_t, in two periods of relatively low and high price flexibility. The shock is assumed to die out with a persistence component of 0.5 and is depicted by the thin black line (with a negative sign). The left panel (low price flexibility) plots the average inflation response in 2010, while the right panel (high price flexibility) plots the average inflation response in 2015. In each of the two panels the vertical line indicates the half-life of the shock.

denotes the normalized flexibility index and γ is the speed of transition across regimes.[35] 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, 1990). 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.[36]

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 γ 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. This ensures 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 7 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 $q = 1$.[37]

The bottom panels of Figure 7 presents the impulse-response functions to a one-standard

[35] We employ a backward-looking MA(12) of the flexibility index to get rid of seasonality in the data. Moreover, we lag the index by one month, in order to avoid potential endogeneity with respect to CPI inflation.

[36] See Appendix G for further details.

[37] Figures H.4 and H.5 in Appendix H report the results for two alternative specifications. Our key insights are not affected by the exact specification of the STARMA(p,q) model. The results are also robust to variations in γ.

Figure 7: Price Flexibility and Inflation Dynamics

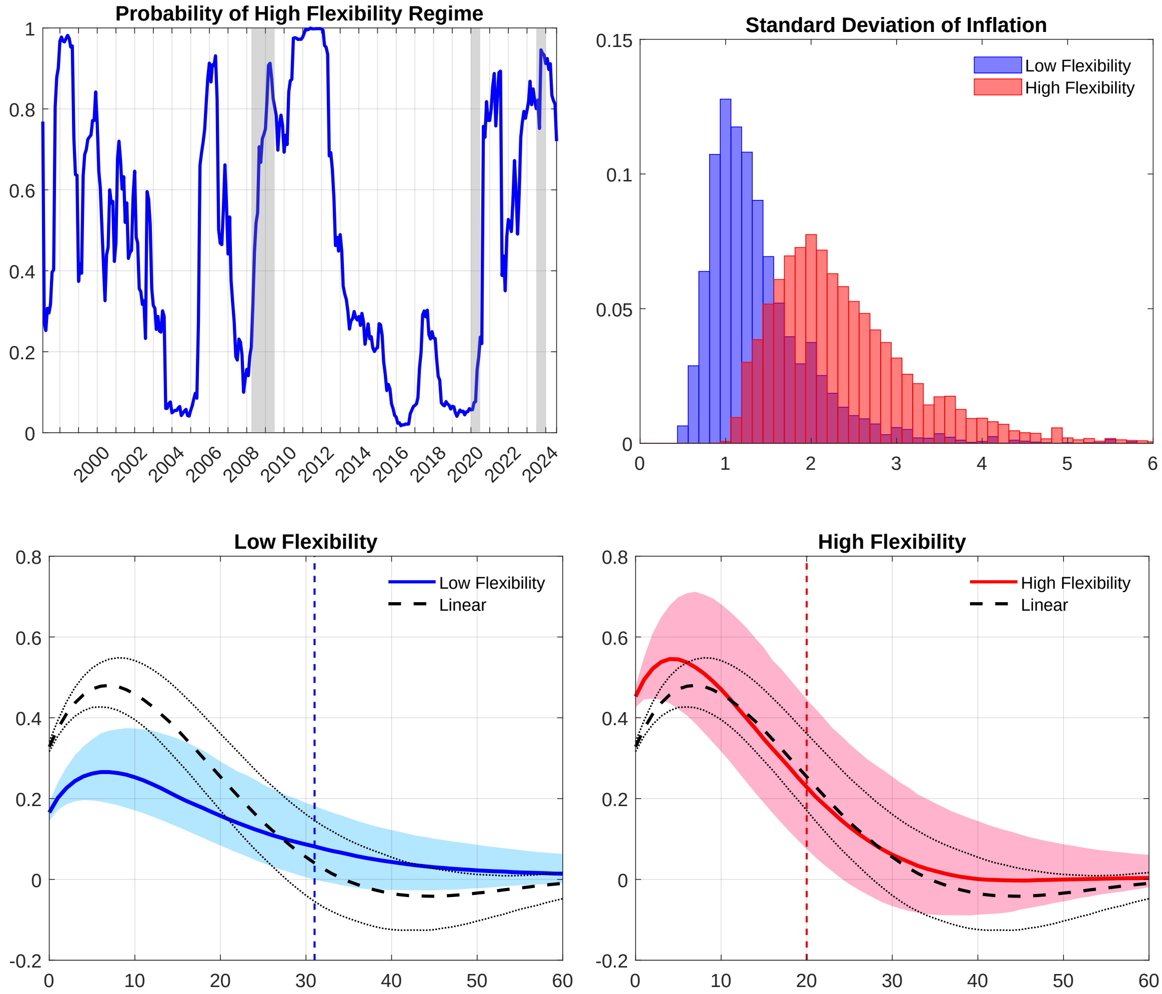

Note: The upper panels report the probability of being in a high flexibility regime, $G\left(\widetilde{\mathcal{F}}_t, \gamma\right) = (1+e^{-\gamma\widetilde{\mathcal{F}}_t})^{-1}$, and the distributions of the estimated inflation volatility in the high and low price flexibility regimes. The lower 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.

deviation shock to inflation in each of the two regimes, and compares them to the response from an equivalent linear model. Consistent with Figure 6, 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, 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 nonlinearities 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.[38] 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 7 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, year-on-year inflation 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.[39] 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. In the next subsection, we explore how non-linearities stemming from time variation in price flexibility may explain why the Bank of England and professional forecasters at large have assumed that the impact of large inflationary shocks would have been shorter-lived in the post-pandemic period, in spite of elevated price flexibility.

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 flexibility. 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.

[38] Volatility is measured by standard deviation of the mean reverting component of their model of inflation.
[39] See, e.g., the Inflation Report published on February 12, 2015.

Table 3: 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	[0.20]	-0.02	[0.59]	-0.05	[0.16]	0.22	1.32
1	0.06	[0.23]	0.01	[0.87]	-0.06	[0.33]	0.36	0.18
2	0.08	[0.19]	-0.01	[0.93]	-0.12	[0.20]	0.11	2.90
3	0.09	[0.19]	-0.04	[0.61]	-0.20	[0.10]	0.02	6.50
4	0.09	[0.18]	-0.08	[0.37]	-0.24	[0.06]	0.00	9.87
5	0.09	[0.10]	-0.10	[0.21]	-0.28	[0.03]	0.00	15.58
6	0.09	[0.06]	-0.10	[0.20]	-0.27	[0.02]	0.00	17.53
7	0.07	[0.08]	-0.08	[0.26]	-0.23	[0.03]	0.00	15.03

(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	[0.24]	-0.09	[0.11]	0.26	0.95
1	0.01	[0.86]	-0.03	[0.69]	-0.06	[0.41]	0.86	-2.20
2	0.06	[0.26]	-0.02	[0.81]	-0.12	[0.20]	0.17	2.01
3	0.07	[0.25]	-0.06	[0.49]	-0.20	[0.10]	0.03	5.67
4	0.07	[0.26]	-0.10	[0.28]	-0.25	[0.06]	0.01	8.79
5	0.07	[0.14]	-0.12	[0.15]	-0.29	[0.02]	0.00	14.92
6	0.07	[0.09]	-0.13	[0.14]	-0.29	[0.02]	0.00	17.49
7	0.06	[0.13]	-0.10	[0.18]	-0.25	[0.02]	0.00	15.12

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 $e_{t+h|t}/h = a_0 + a_1\,(G_{t-1} - 0.5) + a_2\,(G_{t-1} - 0.5)^2 + a_3\,(G_{t-1} - 0.5)^2\,\mathbb{1}_{\{G_{t-1}>0.5\}}G^2_{t-1}$, 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 $\hat{e}_{t+h|t}$ 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 $\tilde{R}^2$.

Each quarter, the Bank of England's Inflation Report publishes year-on-year 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[40] and realized inflation at a given horizon: $e_{t+h|t} = \pi_{t+h|t} - \pi_{t+h}$. Therefore, positive (negative) errors denote an overprediction (underprediction) of inflation. Forecast errors, normalized by the forecast horizon, 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:

$$e_{t+h|t}/h = a_0 + a_1(G_{t-1} - 0.5) + a_2(G_{t-1} - 0.5)^2 + a_3\mathbb{1}_{\{G_{t-1}>0.5\}}(G_{t-1} - 0.5)^2, \tag{9}$$

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 3 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, as well as the corresponding R-squared (adjusted for the number of regressors).

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 economically relevant. At the peak forecasts are, on average, more than 150 basis points below actual inflation outturns. Accounting for this negative bias during periods of high flexibility alone explains nearly 20% of the variability in the forecast error in our sample.

It is natural to question to what extent our results are influenced by the post-COVID-19 experience. 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 8 compares the estimated bias from the specification in Equation 9 using the full sample (this corresponds to the results in Table 3) with the bias for the subsample ending in 2020. 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

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

Figure 8: Inflation Forecast Error Bias

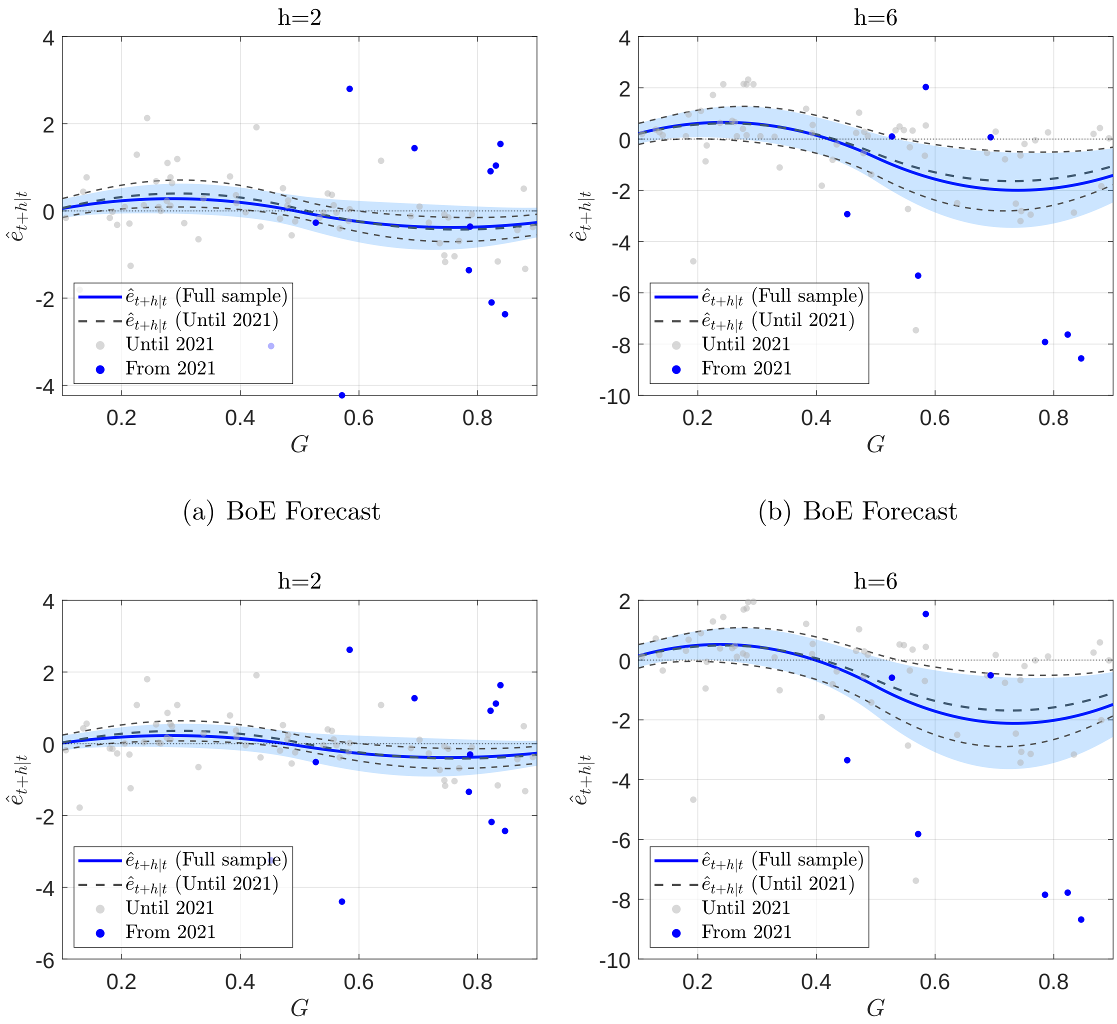

(a) BoE Forecast

(b) BoE Forecast

(c) Market Participants' Forecast

(d) 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.

is already present in data prior to 2020.

This prompts the question of whether earlier recognition of this evidence might have reshaped the 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 9 compares the BoE's inflation forecasts with realized inflation and bias-adjusted forecasts. To avoid look-ahead bias, the latter rely only on pre-2020:Q4 data. While the precise magnitude of the inflationary spike was by any means unforeseeable,[41] adjusting for the bias would have markedly reduced fore-

[41] Applying Blanchard and Bernanke (2023)'s model to the UK, Haskel et al. (2024) highlight that this is

Figure 9: Bias-adjusted Forecasts

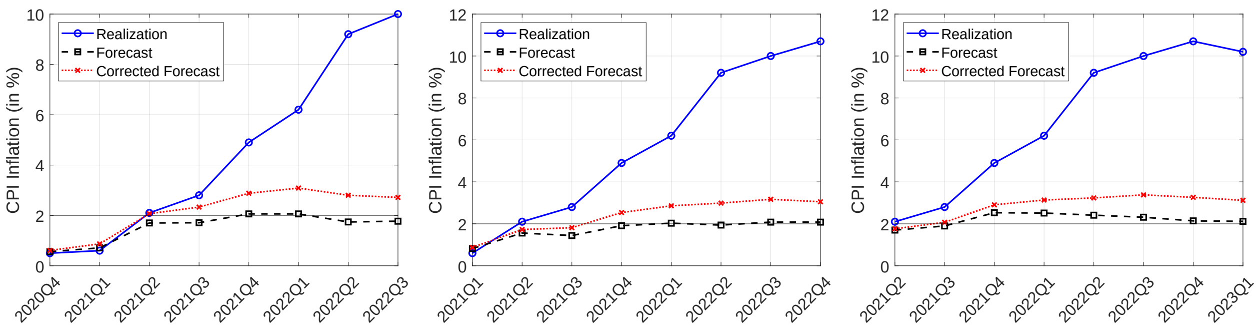

(a) Forecast made in 2020:Q4 (b) Forecast made in 2021:Q1 (c) Forecast made in 2021:Q2

Notes: Each panel reports the CPI inflation forecasts produced by the BoE in the last quarter of 2020 and the first two 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 (9) on the sample up to 2020:Q4. To get the bias, the high-flexibility regime probabilities G from the first month of the respective quarter is used. Specifically, $G(\widetilde{\mathcal{F}}_t; \gamma) = [0.8164, 0.7991, 0.7560]$ for $t = [2020{:}10, 2021{:}01, 2021{:}04]$.

cast errors. In fact, bias-adjusted forecasts suggest a more cautious view about the transitory nature of the inflation spike as early as 2021:Q1, indicating that inflation would have remained well above target by the end of the forecast horizon.

6 Concluding remarks

Looking at UK price microdata, we document distinctive patterns of time variation in some key moments of the underlying process of price adjustment. A key implication of our analysis is that employing time-dependent price-setting protocols to match the frequency of adjustment would understate time variation in price flexibility, which is heavily influenced by the extensive margin of price setting. In fact, our evidence assigns a prominent role to state-dependent price setting, especially in the post-Great Recession sample. In doing so, we underscore the importance of accommodating asymmetry and time variation in the hazard function and the distribution of price gaps. In this regard, further research should investigate the sources of such time variation and explore its connections to firm dynamics, market concentration, and other pertinent micro- and macroeconomic factors.

We highlight a pronounced non-linearity in the price response to inflationary shocks, which is fundamentally driven by the degree of price flexibility. Neither the Bank of England nor professional forecasters appear consider state dependence when projecting CPI inflation. Both sets of forecasters tend to underestimate the impact of inflationary shocks during periods of relatively high price flexibility, when inflation is more volatile and less persistent. Indeed, during the recent inflation surge, recognizing the relatively high level of price flexibility could have prompted the Bank of England to exercise greater caution in projecting a swift return of

largely attributable to the (unforeseen) rapid rise in energy and food commodity prices during the period under scrutiny.

inflation to target in the early stages of the increase.

Employing timely price microdata to develop proxies for aggregate price flexibility is valuable not only for tracking overall inflation dynamics, but also for detecting shifts in how sectoral prices adjust to demand and supply shocks, deepening our understanding of relative-price fluctuations. Recognizing price flexibility as a key state variable in assessing inflation dynamics is therefore essential for policymakers seeking to respond effectively 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.

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

Barro, R. J. (1972): "A Theory of Monopolistic Price Adjustment," *Review of Economic Studies*, 39, 17–26.

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 (2017): "Dynamics of the U.S. Price Distribution," *European Economic Review*, Forthcoming.

Bernanke, B. S. (2024): "Forecasting for Monetary Policy Making and Communication at the Bank of England: A Review," .

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 (1993a): "Microeconomic Adjustment Hazards and Aggregate Dynamics," *The Quarterly Journal of Economics*, 108, 359–383.

——— (1993b): "Microeconomic rigidities and aggregate price dynamics," *European Economic Review*, 37, 697–711.

——— (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.

Carvalho, C. and O. Kryvtsov (2021): "Price Selection," *Journal of Monetary Economics*, 122, 56–75.

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.

DeGooijer, J. G. (2017): *Elements of Nonlinear Time Series Analysis and Forecasting*, Springer.

DeLoecker, J. and J. Eeckhout (2018): "Global Market Power," NBER Working Papers 24768, National Bureau of Economic Research.

Dixit, A. (1991): "Analytical Approximations in Models of Hysteresis," *Review of Economic Studies*, 58, 141–151.

Dixon, H., 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 H. LeBihan (2012): "Generalised Taylor and Generalised Calvo Price and Wage Setting: Micro-evidence with Macro Implications," *Economic Journal*, 122, 532–554.

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.

DOTSEY, M. AND R. G. KING (2005): "Implications of state-dependent pricing for dynamic macroeconomic models," *Journal of Monetary Economics*, 52, 213–242.

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.

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

GERTLER, M. AND J. LEAHY (2008): "A Phillips Curve with an Ss Foundation," *Journal of Political Economy*, 116, 533–572.

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. (1998): "A Class of Quantile Measures for Kurtosis," *The American Statistician*, 52, 325–329.

HAMILTON, J. D. (2018): "Why You Should Never Use the Hodrick-Prescott Filter," *The Review of Economics and Statistics*, 100, 831–843.

HARVEY, A. C. (1990): *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press.

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.

Hobijn, B., F. Ravenna, and A. Tambalotti (2006): "Menu Costs at Work: Restaurant Prices and the Introduction of the Euro," *The Quarterly Journal of Economics*, 121, 1103–1131.

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

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.

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, A. O., E. Nakamura, and L. I. Nakamura (2011): "Price dynamics, retail chains and inflation measurement," *Journal of Econometrics*, 161, 47–55.

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.

Rotemberg, J. J. (1999): "A Heuristic Method for Extracting Smooth Trends from Economic Time Series," NBER Working Papers 7439, National Bureau of Economic Research, Inc.

Stock, J. H. and M. W. Watson (2019): "Slack and Cyclically Sensitive Inflation," NBER Working Papers 25987.

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

Inflation and Price Flexibility

Supplementary Material

Ivan Petrella
Collegio Carlo Alberto, University of Turin, University of Warwick and CEPR.
Email: ivan.petrella@carloalberto.org

Emiliano Santoro
Department of Economics and Finance, Catholic University of Milan.
Email: emiliano.santoro@unicatt.it

Yannik Winkelmann
Department of Economics, University of Tübingen.
Email: yannik.winkelmann@uni-tuebingen.de

A 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.[1]

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.1: Representativeness

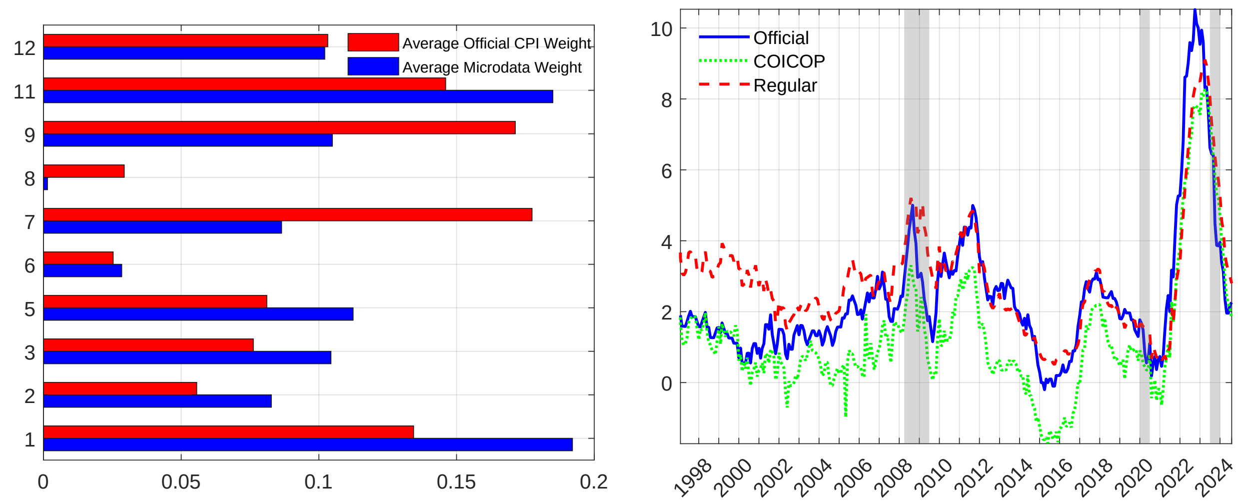

Notes: The left panel contrasts the weights assigned to each of the COICOP sectors to those assigned to build the CPI (re-normalized to exclude COICOP 4 and 10). The right panel reports the official CPI inflation, together with the inflation series retrieved from all the available price quotes (labeled *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.

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

B Aggregation and composition effects

Figure B.1: Aggregate vs Disaggregated Moments

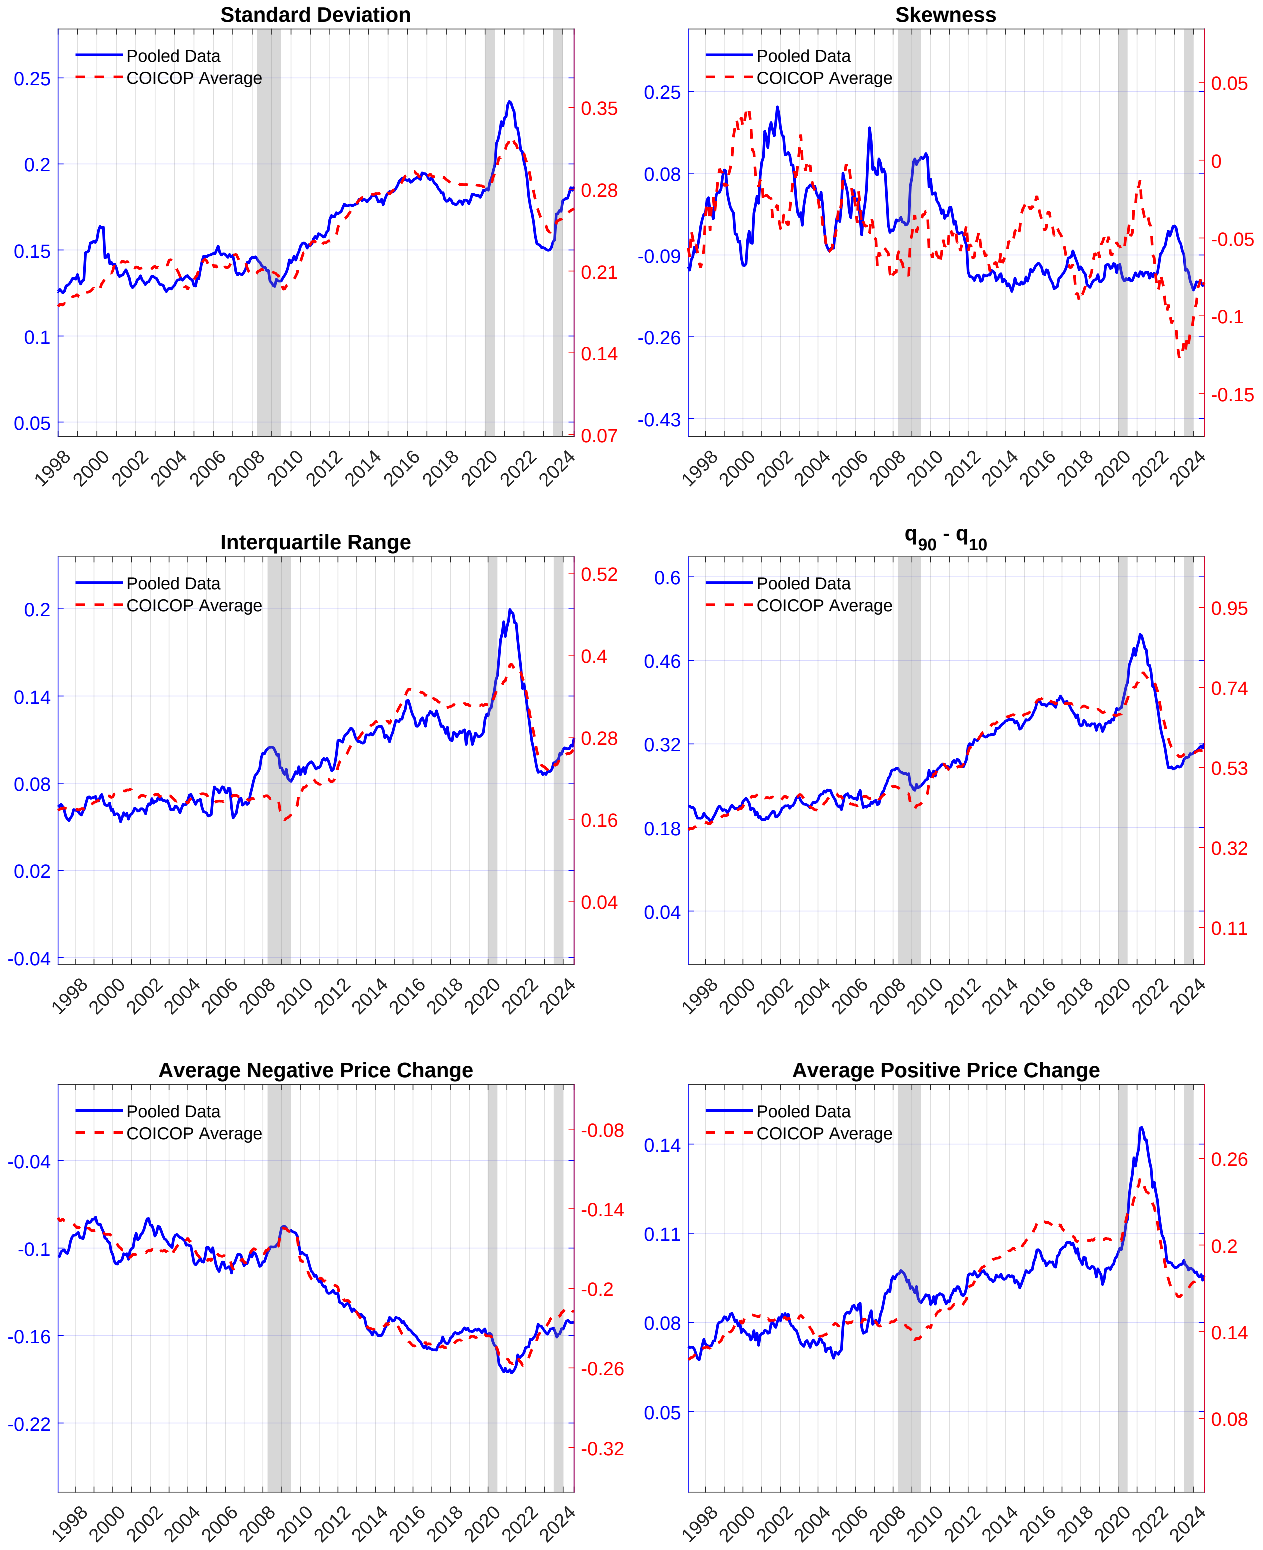

Notes: The figure compares various moments of the distribution of price changes with their counterparts obtained by averaging the corresponding moments of the price quotes obtained for each of the 25 COICOP group categories. The shaded vertical band indicates the duration of recessionary episodes.

C 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 C.1 and Figure C.2 reports 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 (2017) 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 C.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 (2017) 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 C.1: IDENTIFICATION AND THE PARAMETERS OF $f_t(x)$

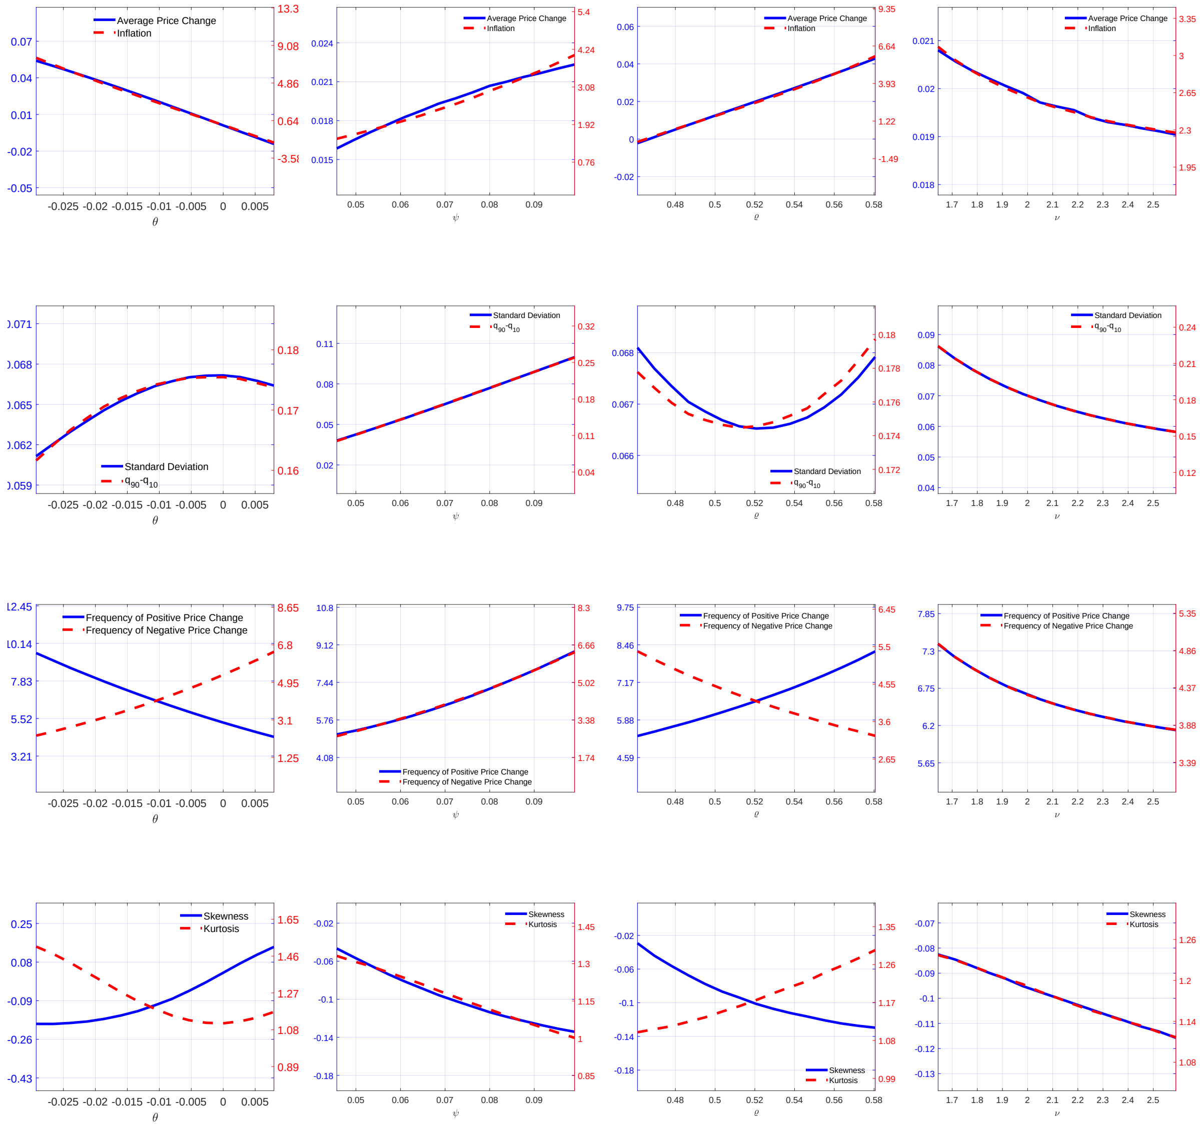

Notes: In each panel, we vary one of the parameters of $f_t(x)$ at the time—while keeping the other coefficients at their baseline estimate—and report its effect on key moments of the price change distribution, as well as the resulting rate of inflation.

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

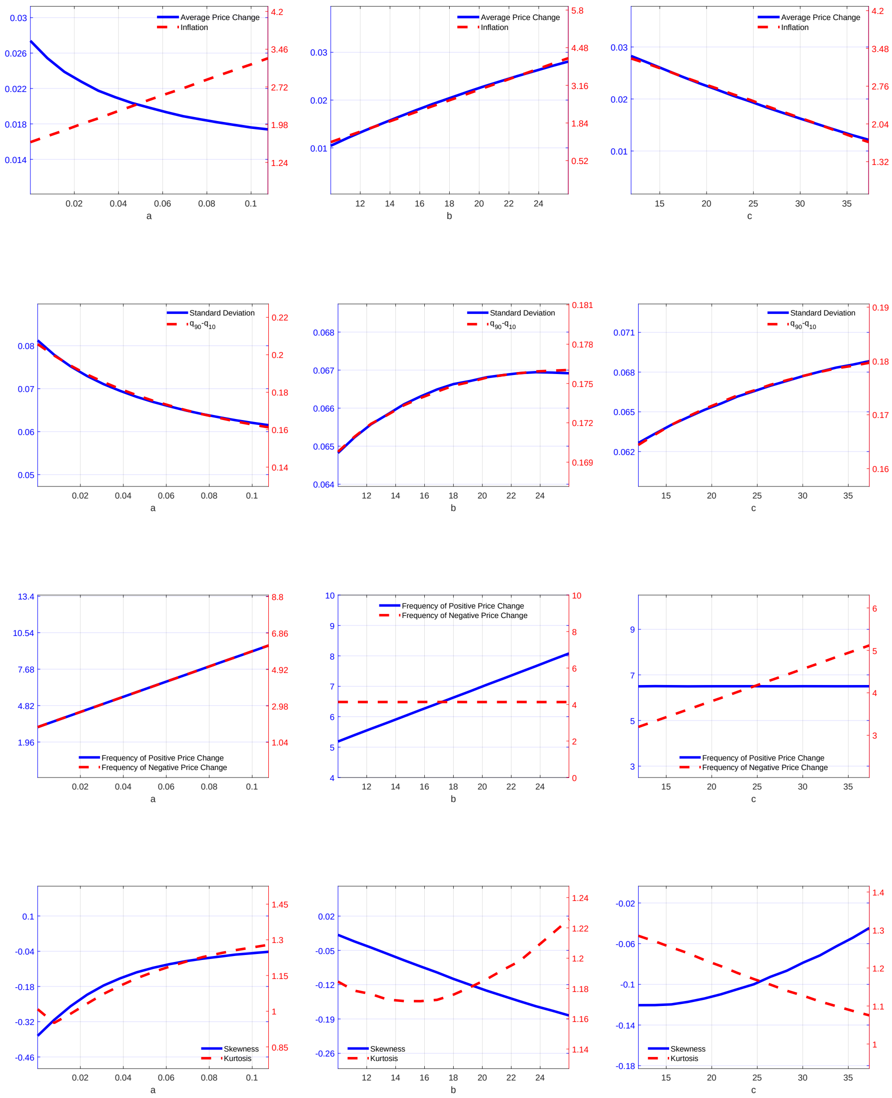

Notes: In each panel, we vary one of the parameters of $\Lambda_t(x)$ at the time—while keeping the other coefficients at their baseline estimate—and report its effect on key moments of the price change distribution, as well as the resulting rate of inflation.

Figure C.3: Model Simulations and Empirical Fit

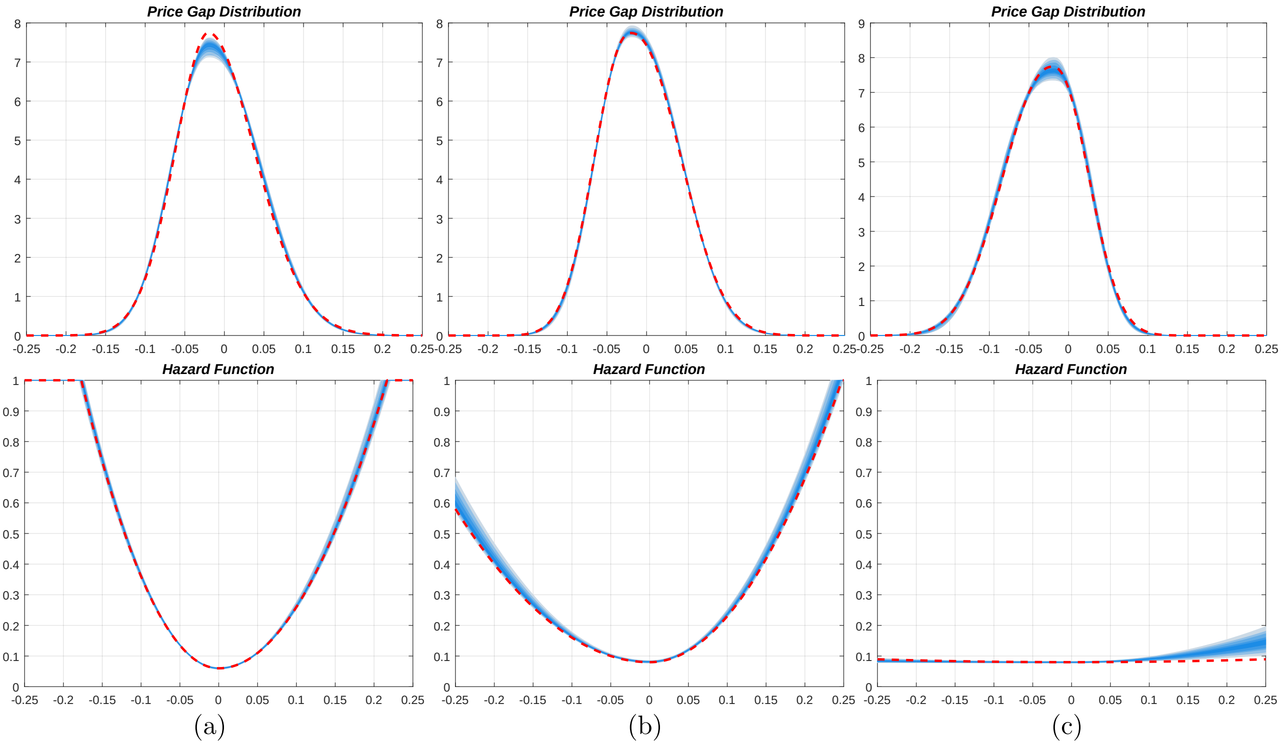

Note: The red line corresponds to the 'true' DGP, while the blue shades correspond to the [5,10,20,. . . 90,95]-th quantile of the estimated price gap distribution (upper panel) and hazard function (lower panel). The following parameterizations are considered: Panel (a): $\theta = -0.02, \psi = 0.07,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

D Model estimates

Figure D.1: Parameters of the Price Gap Distribution

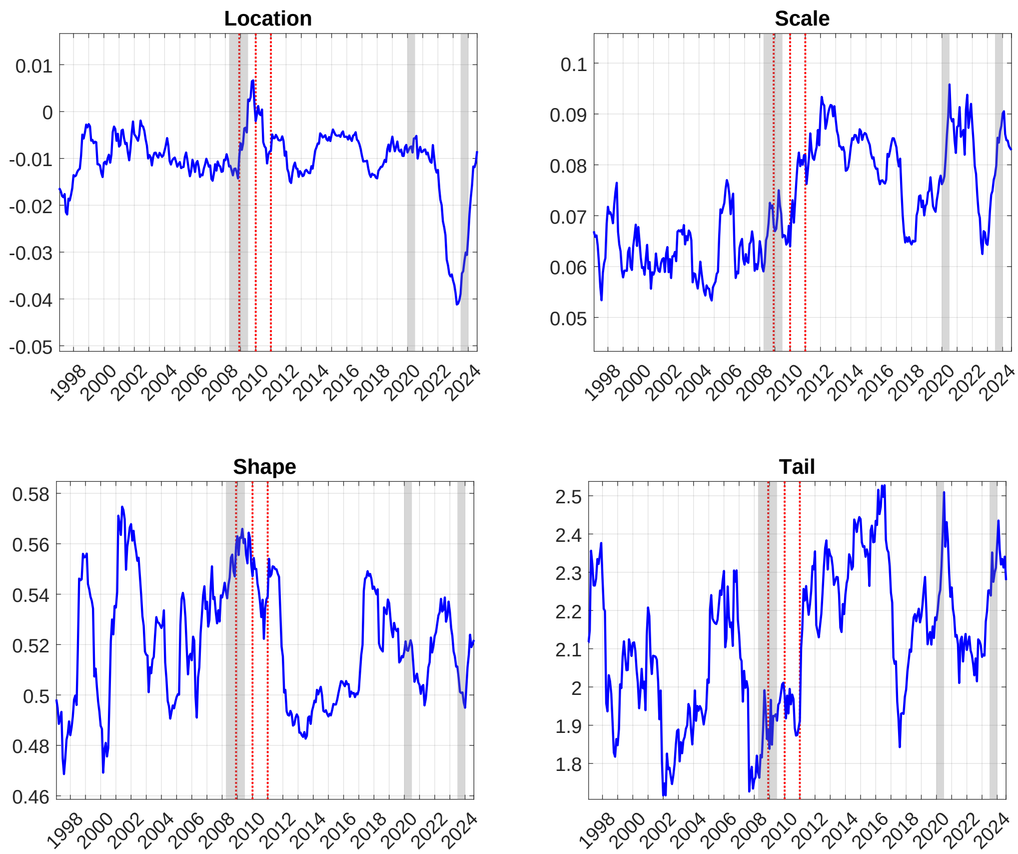

Note: The red lines denote the three VAT changes in the sample. The shaded vertical bands indicate the duration of recessionary episodes.

Figure D.2: Parameters of the Hazard Function

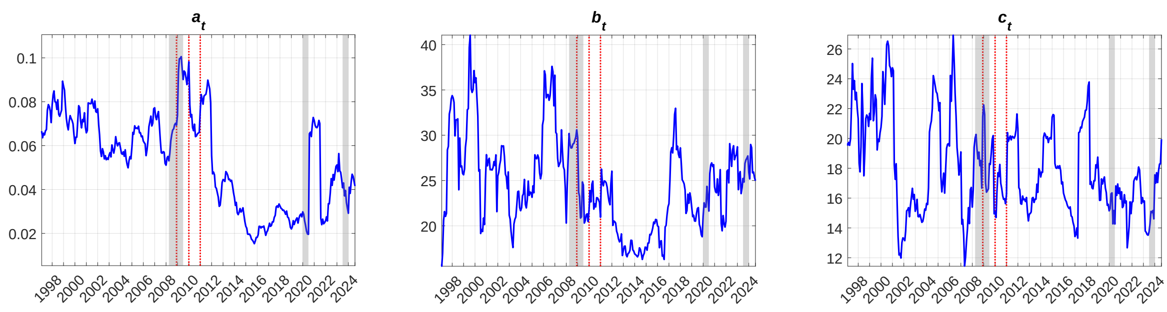

Note: The red lines denote the three VAT changes in the sample. The shaded vertical bands indicate the duration of recessionary episodes.

Figure D.3: Fit of the Ss Model (Selected Moments)

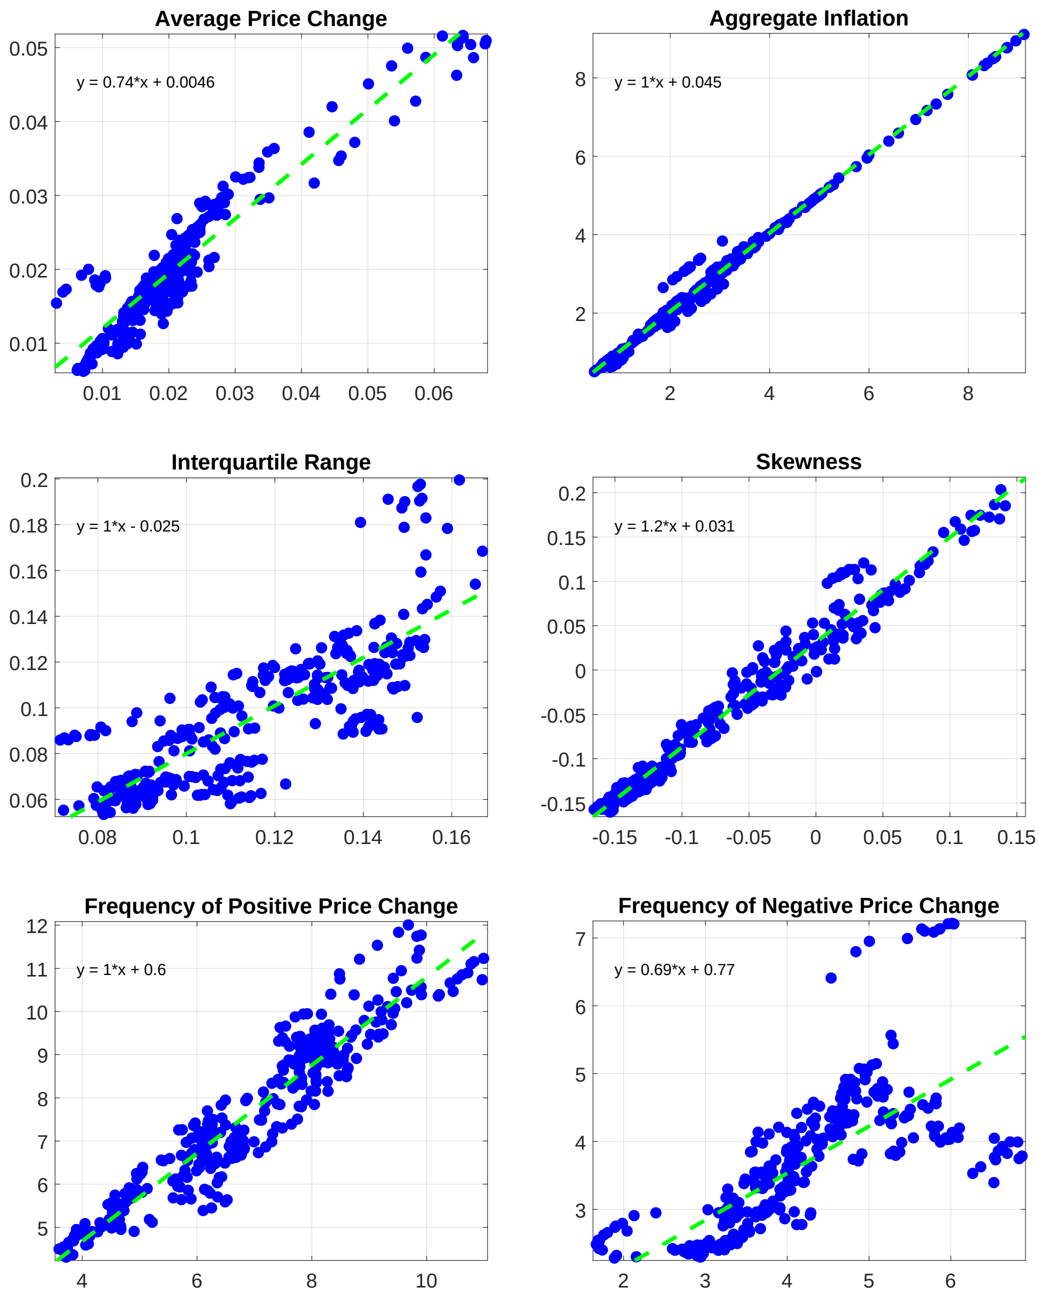

Notes: The figure compares the estimated moments from the Ss model in Section 3 (x-axis) to the moments estimated from the raw data (y-axis). Each chart also reports the linear fit (green/broken) line.

E A simple analytical setting to frame the stylized facts

We consider the menu cost model popularized by Barro (1972) and Dixit (1991). As illustrated by Vavra (2014), the advantage of this framework is to provide us with a simple analytical setting to keep track of the determinants of the frequency and the dispersion of price changes, as well as the dispersion of price gaps, intended as the difference between the actual price of a given good and its reset price (i.e., the price that would have prevailed in the absence of price-setting frictions). For the sake of our analysis, we will use this model as a prism through which interpreting diverging movements in the frequency of price adjustment and the dispersion of price changes. Othwerise, the model has no presumption to map into the statistical framework employed in the empirical analysis.

Firms face a dynamic control problem where x—the deviation of the current price from the optimal price—is a state variable. A wedge between the state variable and zero entails an out-of-equilibrium cost αx^2, where α can be inversely related to market power. When not adjusting, x follows a Brownian motion $dx = \phi dW$, where W is the increment to the Wiener process. It is possible to change the value of x by applying an instantly effective control at a lump-sum cost λ. From this environment, a simple Ss rule emerges, according to which the optimal policy is 'do not adjust' when $|x| < \sigma$ and 'adjust to zero' when $|x| \geq \sigma$, where $\sigma = \left(6\lambda\phi^2/\alpha\right)^{1/4}$ denotes the standard deviation of price changes. Moreover, $fr = (\alpha/6\lambda)^{1/4}\,\phi$ is the frequency of adjustment.[2]

To provide an overview of different determinants of the distribution of price gaps and the associated distribution of price changes, Figure E.1 considers three possible scenarios: i) a positive shift in the cost of adjustment λ (or, equivalently, a negative shift in α) that affects the inaction region, while leaving the distribution of price gaps unaffected; ii) a first-moment shock that causes a shift in the distribution of price gaps, affecting all x's in the same manner; iii) an increase in the dispersion of the distribution of price gaps (i.e., a rise in ϕ).

As for i), a positive change in λ widens the inaction region, translating automatically into a reduction in the frequency of adjustment and an increase in the dispersion of price changes, which is in line with the behavior of the two statistics in the post-recession sample. As for ii), the immediate effect of a shift in the distribution of price gaps is to push more firms out of the inaction region, thus inducing an increase in the frequency of adjustment. Importantly, this result does not depend on the specific sign of the shock, as all firms' desired price changes will be affected in the same way. Thus, all firms pushed out of the inaction region will denote price changes of the same sign, implying a decrease in their dispersion.[3] Thus, while negative comovement would emerge in this case, it is important to recognize that first-moment shocks would not be suitable to characterize the (diverging) movements in the frequency and the dispersion that have occurred over the post-recession sample.[4] Finally, a rise in ϕ, as sketched in the last column of the figure (iii), induces increased dispersion in the price gap distribution and an expansion in the inaction region. As a result, both fr and σ increase.

Vavra (2014) points to second-moment shocks as potential drivers of the positive comovement between the frequency of adjustment and price-change dispersion in U.S. CPI data. It is clear how this type of shock would not be suitable to rationalize negative comovement. In fact, only an increase in the fixed cost of adjustment and/or a drop in the cost of deviating from the optimal price may account for a concurrent drop (increase) in the frequency of adjustment (dispersion of price changes), as observed after the Great Recession.

[2]For analytical details and proofs, see Barro (1972) and Vavra (2014).

[3]In fact, Vavra (2014) shows that, while in environments with zero inflation small shocks to x do not produce any effect on the frequency of adjustment and the dispersion of price changes, in the presence of positive trend inflation the frequency (dispersion) increases (decreases).

[4]One should note that such movements could also be rationalized in the occurrence of a first-moment shock, whenever the latter hits outside the steady state and shifts the distribution towards its ergodic counterpart. However, our empirical evidence indicates that changes in the price-adjustment cost structure, as reflected in upward trends in the markup associated with several industries/goods, are of primary importance, as opposed to first- or second-moment shocks. In fact, first-moment shocks seem to account only for a small part of the persistent increase in the dispersion of price changes, and mainly when aggregate inflation has come close to zero, towards the end of the sample. In ongoing work we examine these issues in depth.

Figure E.1: Analytical Framework

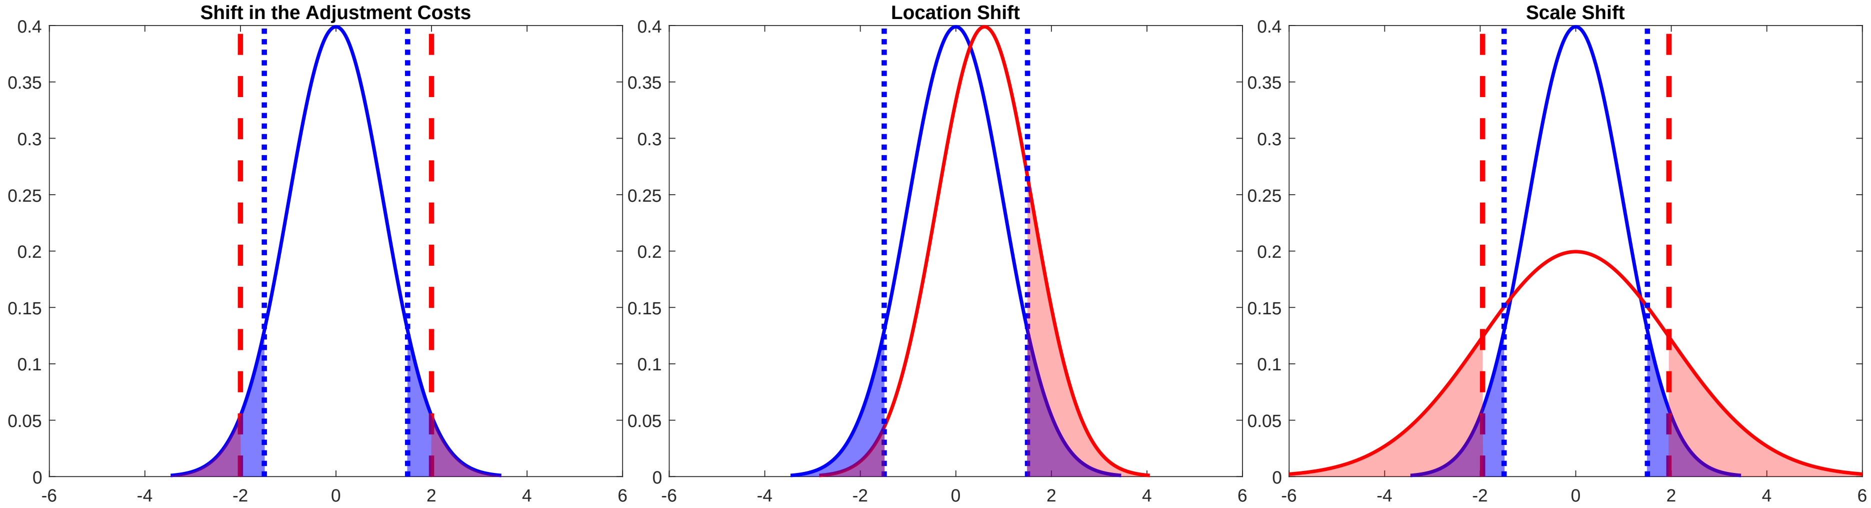

(a) Price Gap Distribution

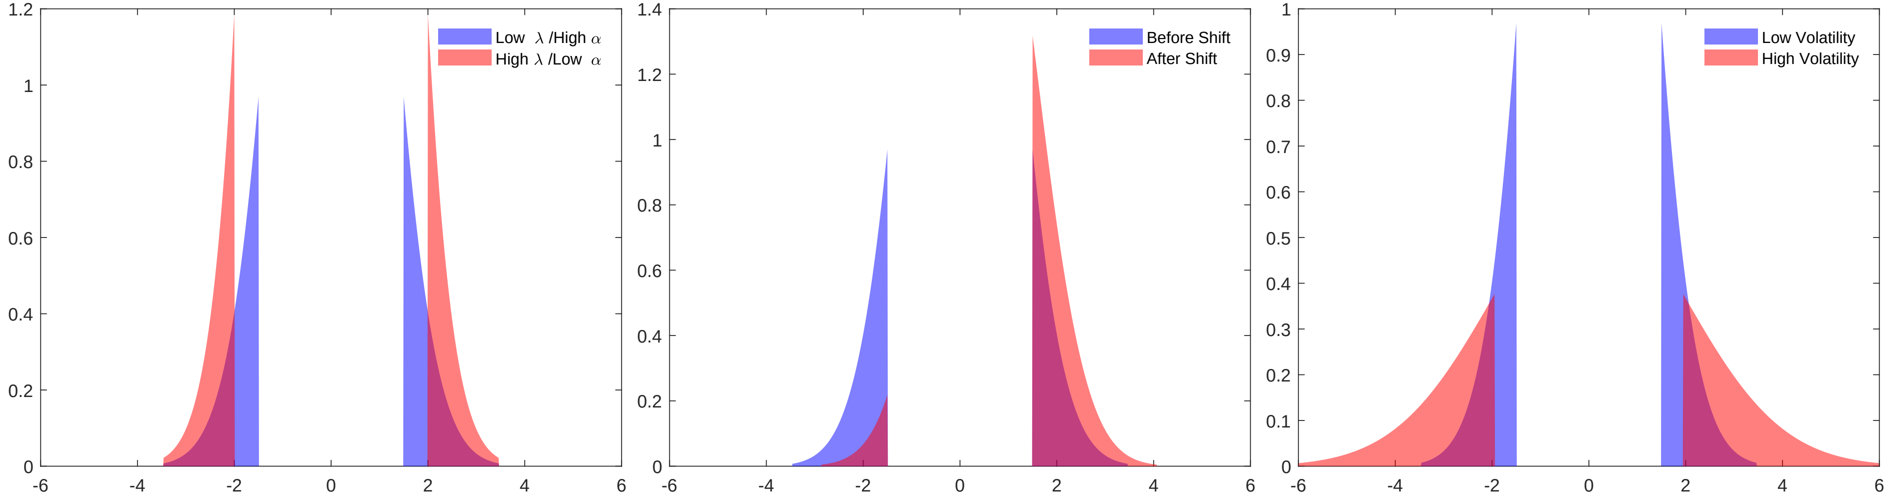

(b) Price Change Distribution

Note: The first column considers a positive shift in λ (or a negative shift in α) that affects the inaction region, while leaving the distribution of price gaps unaffected. The second column considers the effects of a first-moment shock that affects all x's in the same direction. The last column depicts the effects of an increase in ϕ. The upper panels report the ex-ante distribution of price gaps and the corresponding bands delimiting the inaction region (dotted-blue lines), together with their ex-post counterparts (dashed-red lines). The bottom panels report the corresponding distributions of price changes.

It is important to stress that shifts in λ and α would immediately reflect into a change in the inaction region, while leaving the price gap distribution unaffected. In this respect, it is possible to show show how large diverging movements in the dispersion of price changes and the frequency of adjustment in the post-recession period may be rationalized by an expansion of the inaction region that dominates the effects of positive shifts in the dispersion of price gaps. Figure E.2 considers a situation in which both ϕ and λ increase.[5] The rise in the dispersion of price changes determines an expansion in the inaction region, thus increasing the density outside the adjustment bands and, in turn, the frequency of adjustment. This effect is counteracted by the rise in λ, which widens the inaction region further and restricts the density outside the adjustment bands beyond the initial situation. If the expansion in the inaction region is large enough to overcome the increase in dispersion, we observe opposite movements in the cross-sectional dispersion of prices and the frequency of adjustment. This is in line with what we observe in the post-recession period.

Figure E.2: A combined increase in ϕ and λ

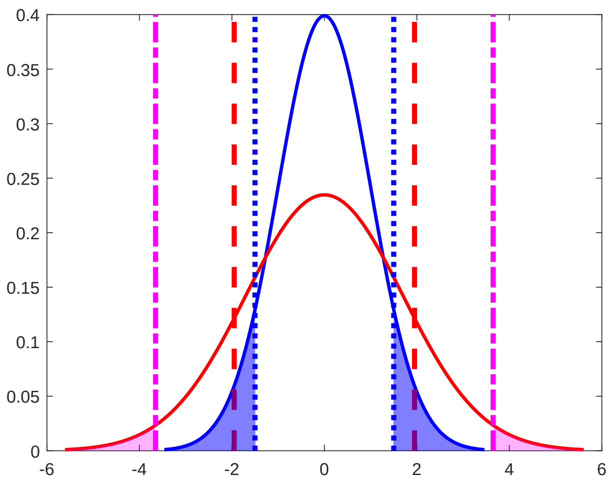

(a) Price Gap Distribution

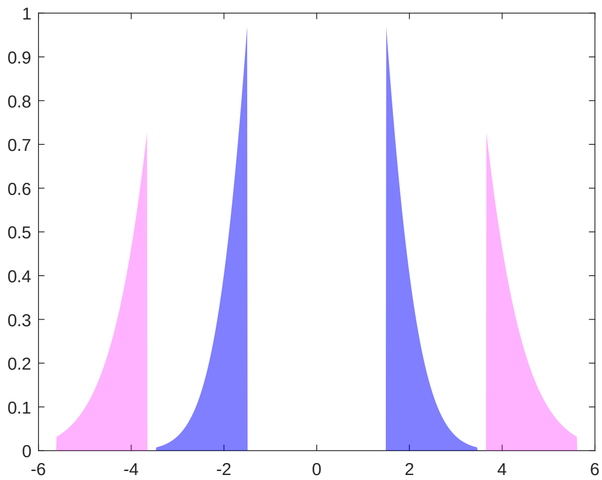

(b) Price Change Distribution

Note: We consider a positive shift in λ that affects the inaction region (while leaving the distribution of price gaps unaffected), combined with an increase in the dispersion of the distribution of price gaps, ϕ. The left panel reports the transformations occurring to the distribution of price gaps and the corresponding bands delimiting the inaction region: the dotted (blue) line refers to the ex-ante situation, the dashed (red) line denotes the effects of the volatility shift, while the dashed-dotted (magenta) line refers to the effects produced by the joint increase in ϕ and λ. The right panel reports the distributions of price changes, both in the ex-ante situation and in the case of a combined increase in ϕ and λ.

[5]Once again, a drop in α would lead to qualitatively similar results.

F 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.[6] 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:

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

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

G Estimation of the STARMA (p,q) model

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

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

This can be easily casted in state space. Therefore the likelihood can be calculated recursively using the Kalman filter (see Harvey, 1990). 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 non-linearity 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 θ 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:

- *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 Ω.

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

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

[6] 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.

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

Specifically, we use an adaptive step for the value of Ω, 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 θ under standard regularity assumptions of maximum likelihood estimators. CH also prove that the covariance matrix of the estimate of θ 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.

H Additional Results and Robustness

Figure H.1: ADDITIONAL STATISTICS FROM PRICE MICRODATA

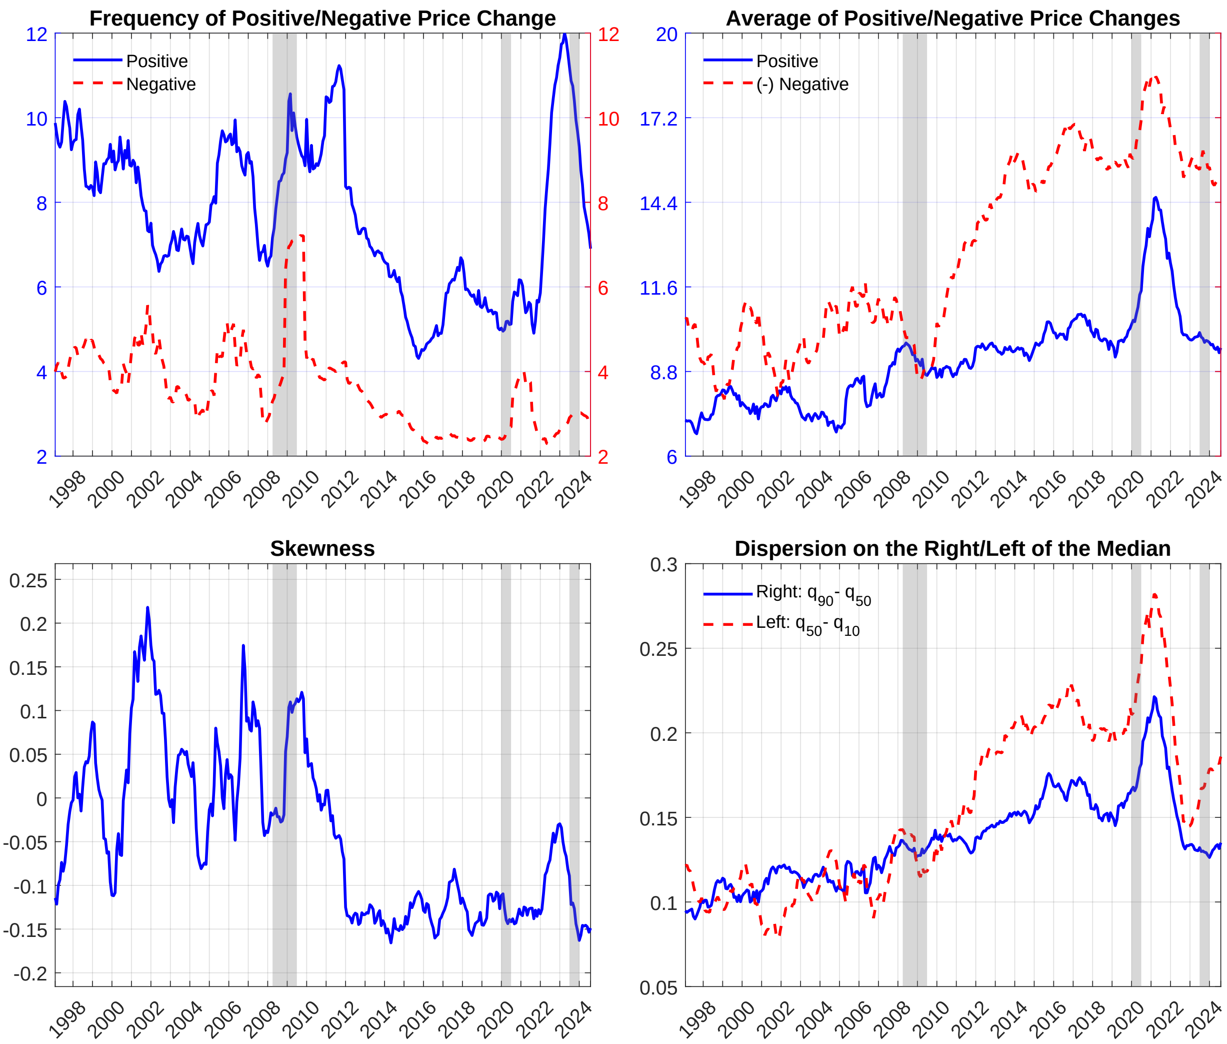

Notes: The frequency of adjustment, fr_t , is computed as $\sum_i \omega_{i,t} \mathbb{1}_{\{\Delta p_{i,t} \neq 0\}}$, where $\omega_{i,t}$ denotes the CPI weight associated to good i at time t, and $\mathbb{1}_{\{\Delta p_{i,t} \neq 0\}} = 1$ if $\Delta p_{i,t} \neq 0$ and zero otherwise. The average 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 H.2: Comparison with Alvarez et al. (2016)

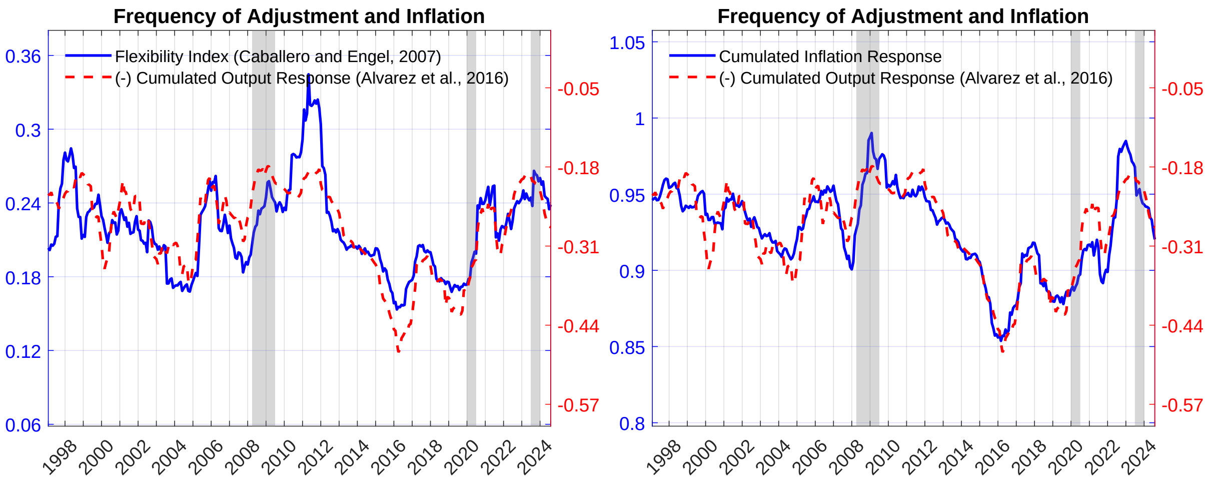

Note: The left panel of the figure reports the cumulated output response to a monetary policy shock (solid-blue line), as computed by Alvarez et al. (2016), as well as the (negative of the) index of price flexibility, as computed by Caballero and Engel (2007) (dashed-red line). The right panel features the cumulated output response to a monetary policy shock (solid-blue line) against the (negative of the) cumulated inflation response (dashed-red line), the latter being cumulated over a 18-month period.

Figure H.3: Price Flexibility and Inflation Persistence STARMA(2,4)

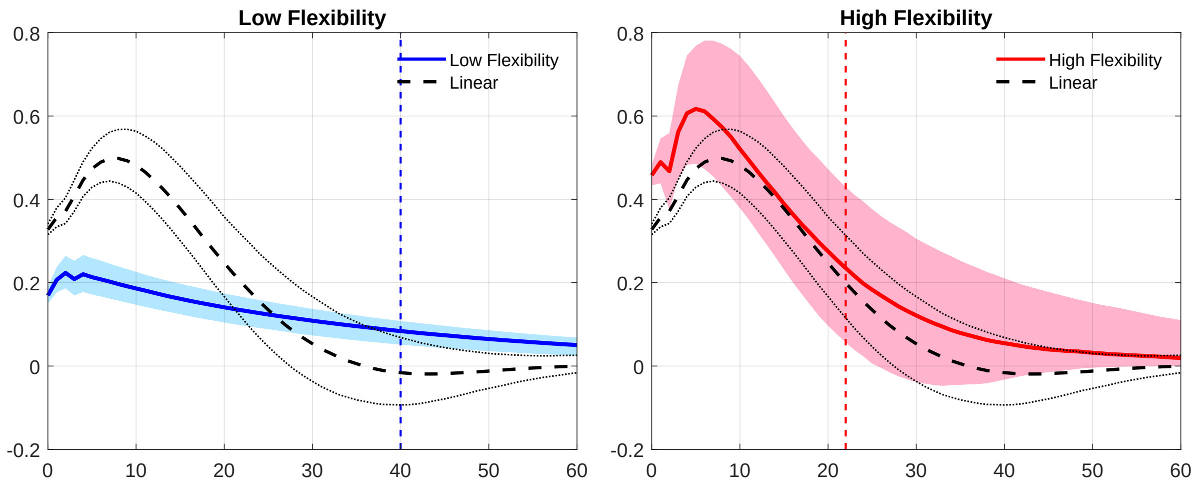

Note: Figure H.4 reports the responses of inflation to a 1% shock in a STARMA(2,4) model. The left (right) panel graphs the response in the low (high) price flexibility regime. In both cases we also report the response from a (linear) ARMA(2,4) model. 68% confidence intervals are built based on the Markov 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 H.4: Price Flexibility and Inflation Persistence STARMA(1,1)

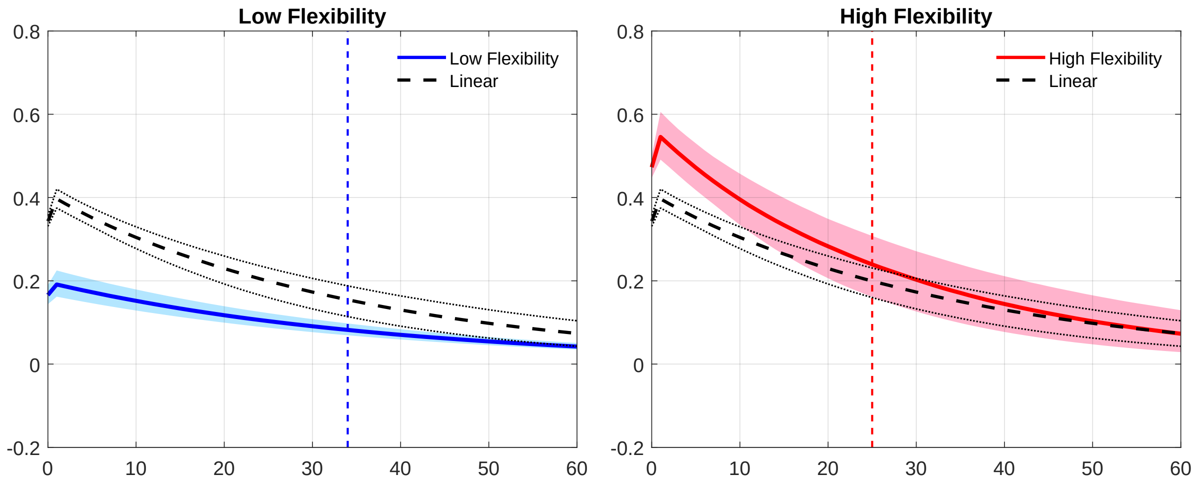

Note: Figure H.4 reports the responses of inflation to a 1% shock in a STARMA(1,1) model. The left (right) panel graphs the response in the low (high) price flexibility regime. In both cases we also report the response from a (linear) ARMA(1,1) model. 68% confidence intervals are built based on the Markov 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 H.5: Price Flexibility and Inflation Volatility

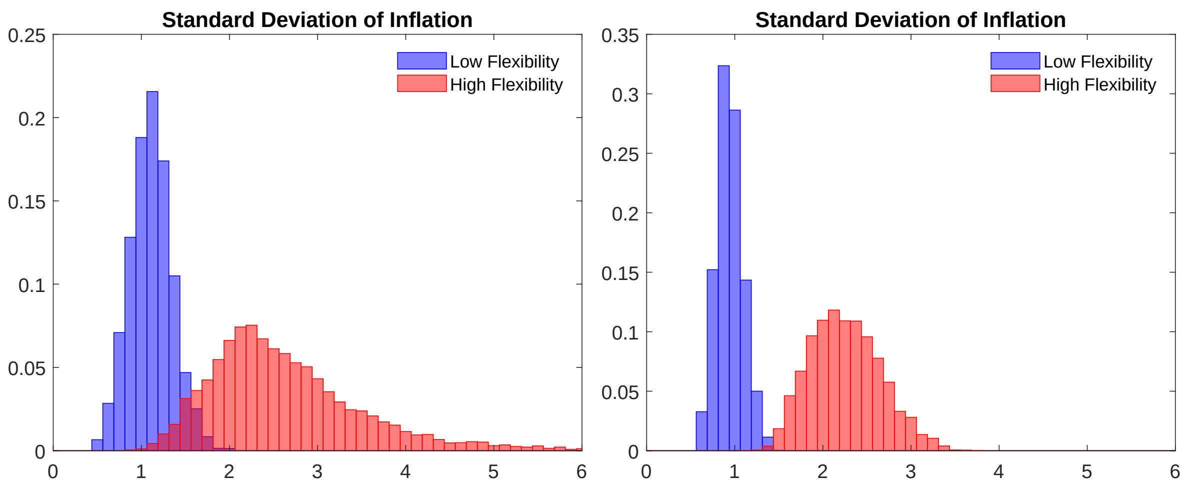

(a) STARMA(2,4)

(b) 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).