

Inflation Dynamics and Price Flexibility in the UK*

Ivan Petrella[†]
University of Warwick
& CEPR

Emiliano Santoro[‡]
University of Copenhagen

Lasse de la Porte Simonsen[§]
Birkbeck College

May 2018

Abstract

Using microdata underlying the UK consumer price index we document how the distribution of price changes has evolved over the last two decades and how its variation reflects into the dynamics of price flexibility. In the period after the Great Recession the dispersion of price changes has displayed a sustained upward trend, while the frequency of adjustment has dropped. We rationalize these facts within a stylized menu cost model, showing that a persistent increase in the inaction to nominal shocks may be compatible with the observed diverging trends. We support this prediction by estimating a generalized S_s model. Price flexibility displays pronounced time variation. In fact, over the last decade the capacity of nominal stimulus to generate inflation has decreased substantially. Despite the marked non-linearity of inflation dynamics dictated by changes in the degree of price flexibility, we show that neither the Bank of England nor professional forecasters seem to account for this type of state dependence when forming inflation expectations. In fact, they both tend to overestimate inflation persistence in periods of relatively high price flexibility, especially at medium-term forecast horizons.

JEL codes: E30, E31, E37, C22.

Keywords: Price flexibility, S_s models, inflation.

*We thank—without implicating—Luca Dedola, Federico Di Pace, Ana Galvao, Omar Rachedi, Søren Hove Ravn, Federico Ravenna, and seminar participants at the Bank of Finland. We also thank Rowan Kelsoe, Richard Campbell and Gareth Clews from the Office of National Statistics for their advice and help with the data.

[†]Warwick Business School, University of Warwick and CEPR. *Address:* Scarman Rd, CV4 7AL Coventry, United Kingdom. *E-mail:* Ivan.Petrella@wbs.ac.uk.

[‡]Department of Economics, University of Copenhagen. *Address:* Østerfarimagsgade 5, Building 26, 1353 Copenhagen, Denmark. *E-mail:* emiliano.santoro@econ.ku.dk.

[§]Department of Economics, Mathematics and Statistics, Birkbeck College, University of London. *Address:* Malet Street, London W1E 7HX, UK. *Email:* l.simonsen@mail.bbk.ac.uk.

“(...) I hope that researchers will strive to improve our understanding of inflation dynamics and its interactions with monetary policy.”

Janet Yellen, October 2016

1 Introduction

Over the last decade the increasing availability of disaggregated data has allowed economists to attain a deeper understanding of consumer micro price behavior and its implications for price flexibility at the macroeconomic level. The degree of aggregate price flexibility lies at the core of the monetary policy transmission mechanism, ultimately embodying Central Banks’ capacity to stimulate output growth and inflation. As a result, a wide number of empirical contributions have been concerned with measuring the response of prices to nominal demand shocks. However, less emphasis has been placed on the extent and characteristics of time variation in aggregate price flexibility,¹ and how this information can be usefully employed to study inflation dynamics.

Using microdata underlying the UK consumer price index (CPI) we document how the distribution of price changes has evolved over the last two decades and how that reflects into the behavior of price flexibility. While in the first half of the sample the frequency of adjustment has been roughly stable, during the last decade it has displayed substantial variation, dropping markedly since after the Great Recession. Over the same period the dispersion of price changes denotes a sustained increase. These facts stand in contrast with the behavior of US microdata, where the cross-sectional standard deviation of price changes typically comoves positively with the frequency of adjustment (Vavra, 2014 and Berger and Vavra, 2017).

To frame our empirical facts, we employ the menu cost model popularized by Barro (1972) and Dixit (1991). Within this setting, diverging trends in the dispersion of price changes and the frequency of adjustment may emerge as the result of a persistent increase in the fixed cost of adjustment and/or a drop in the cost of deviating from the optimal price: as long as the resulting expansion in the area where it is not worth adjusting prices (i.e., inaction region) overcomes the effects of low frequency movements in the dispersion of price gaps (i.e., the wedge between the actual and the optimal price), the distribution of price changes becomes more dispersed and firms hit the adjustment bands less frequently. Otherwise, when the opposite holds true, the frequency of adjustment and the dispersion of price changes comove positively (Vavra, 2014). What clearly emerges from our analysis is that, regardless of the time window we examine, the observed price movements cannot simply reflect predetermined price adjustments (i.e., adjustments that are determined ahead of the shock and would materialize into a mere shift of the distribution)—as embodied by time-dependent models of price setting—which would imply a counterfactual comovement between the statistics of interest. We test our predictions by estimating the generalized Ss model developed by Caballero and Engel (2007), fitting the

¹In this respect, Caballero and Engel (1993b) and Berger and Vavra (2017) represent some notable exceptions.

distribution of price gaps and the hazard function (i.e., the probability of individual price adjustment) over the price quotes available in each month. We show that, by the end of the sample, about five times as many firms appear inactive, as compared with the pre-2010 time window. In line with the framework employed to build our comparative-statics analysis, this implies that the expansion in the inaction region dominates the increase in the dispersion of price changes.

We detect sizable changes in the distribution of price gaps and the hazard function. Inevitably, these reflect in the way shocks are propagated to the economy. To dig deeper into the connection between individual price adjustment and the response of aggregate inflation to nominal stimulus, we compute a measure of aggregate price flexibility. This is then decomposed into predetermined price adjustments (i.e., the intensive margin) and adjustments triggered or canceled by the shock (i.e., the extensive margin). The latter appear rather relevant, and more so in periods of particularly volatile inflation. In fact, during these episodes the difference between actual inflation and its ‘Calvo counterfactual’— i.e., the inflation rate obtained by setting the period hazard function to a constant equal to the intensive margin—is particularly large. To quantify the importance of price adjustments along the extensive margin, we compare the observed distribution of price changes to a counterfactual distribution that obtains in the absence of shocks capable of generating episodes of massive repricing, such as those triggered by changes in the value-added tax (VAT). In support of the view that these episodes cannot be seen as mere translations of the distribution of price gaps (see, e.g., Gagnon et al., 2013), we appreciate both a major upward shift and a steepening of the hazard function across all episodes of VAT change, as well as a large increase in the frequency of adjustment.

All in all, price flexibility displays sizable variation over the entire sample. Most remarkably, some large movements can be appreciated since after the Great Recession. The response of aggregate inflation to nominal demand shocks increases substantially during the contraction in real activity—eventually reaching its (sample) peak in 2011—thus reverting and attaining its minimum in the first quarter of 2017. This implies that, over the last decade, the capacity of nominal stimulus to generate inflation has decreased markedly. Along with observing substantial time variation in the pass-through of nominal demand shocks to prices, large changes in price flexibility tend to occur in correspondence of sizable departures of CPI inflation from the Bank of England’s institutional target. In this respect, two facts stand out when examining inflation dynamics in the post-Great Recession sample: i) inflation has been outside the 1%-3% interval for a total of 22 out of 40 quarters, whereas the same happened only during 11 quarters in the previous decade; ii) over the same period, inflation has shot above and below the target, reaching both its maximum (+4.8%) and minimum value (-0.1%) in the overall sample. In light of this, accounting for time variation in price flexibility may help us understand why hitting the inflation target may have proven to be rather difficult for a prolonged period of time. To dig deeper into the responsiveness of prices to nominal stimuli and the dynamics of aggregate price inflation, we examine the latter in response to an inflationary shock over different price-flexibility regimes, and compare the results to those obtained in a linear setting.

The half-life of the inflation response is twice as big in periods of relatively low flexibility, along with appearing remarkably close to the one obtained in the linear setting. In light of this, we posit that neglecting that inflationary shocks are propagated at different speeds depending on the overall degree of price flexibility may lead to overstating inflation persistence. We test this implication, and show that the Bank of England and other market participants do not appear to be taking into account changes in price flexibility when computing their inflation expectations. In fact, price flexibility accounts for roughly 25% of the variability in the absolute forecast error at a four-quarter horizon.

Our work relates to a number of studies that have examined the connection between microprice changes and aggregate inflation.² Among these, Berger and Vavra (2017) represents the contribution that is more in line with the spirit of the present paper. Compared with this study, we highlight the emergence of secular variation in the distribution of UK price changes and, in this respect, we point to some distinctive patterns in the decade following the Great Recession. Moreover, we elaborate on the role of state dependence in price flexibility and its implications for predicting inflation dynamics. Our work also relates to a number of papers that devise and estimate specific structural models that connect changes in the distribution of price changes to price flexibility (see, e.g., Midrigan, 2011, Alvarez et al., 2016 and Vavra, 2014, among others). As discussed by Berger and Vavra (2017), an empirical limitation of this approach is to rely on specific shocks to the price-setting units, while our approach is more agnostic, in this sense. This represents a strategic advantage in the analysis of UK microdata, where the implied pattern of time variation in the distribution of price changes has been somewhat discontinuous, emerging at different points in time as the result of a different mix of first- and second-moment shocks, as well as secular changes in the determinants of the inaction region of price setting. Finally, our work relates to Gagnon et al. (2013) in that we focus on the distinction between price adjustments that are determined ahead of shocks, and those that are triggered or canceled by the shocks. Compared with this study, our empirical model allows us to examine the behavior of the distribution of price gaps and that of the hazard function in connection with different episodes of VAT changes, thus highlighting important asymmetries over different margins of price setting.

Our paper also features some broad connection with recent empirical contributions employing individual UK consumer prices. In this respect, Bunn and Ellis (2012) have been among the first to appreciate 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. (2014) have focused on the impact of the Great Recession on price setting. As compared with these papers, we pose particular emphasis on state dependence in price flexibility, as well as on its role for the transmission of nominal demand shocks. Moreover, our application underlines the importance of the selection effect for aggregate inflation (see, on this, Carvalho and

²See, 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), Berardi et al. (2015), Alvarez et al. (2016), Nakamura et al. (2018).

Kryvtsov, 2017 and references therein). Specifically, we highlight the versatility of the empirical approach proposed by Caballero and Engel (2007), and show how this can be followed to appreciate the importance of the extensive margin of price adjustment for inflation dynamics in the UK.

The rest of the paper is organized as follows. In Section 2 we discuss the key characteristics of the ONS microdata on consumer prices. Section 3 discusses the menu cost model that frames our empirical analysis. Section 4 reviews the generalized *Ss* model developed by Caballero and Engel (2007), takes it to the data. Section 5 assesses time variation in price flexibility and identifies the relative contribution of adjustments along the intensive and the extensive margin. Finally, Section 6 discusses the implications of state dependence in price flexibility for inflation dynamics.

2 Microdata on consumer prices

We use ONS microdata that underpin the UK CPI. Prices are collected on a monthly basis, for more than 1,100 categories of goods and services, and published with a month-lag. Our sample covers the 1996:M2-2017:M8 time window, thus resulting into about 27.5 million observations (see Table 1). Each month around 106,000 prices are collected by a market research firm on behalf of the ONS. There are also about 140 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 theoretically allows identification of some price setters, or because of the need to frequently adjust for quality changes.³ The price quotes are recorded on or around the second or third Tuesday of the month, with the exact date being kept secret so as 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 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 our sample. We refer to the remaining subset of prices—which make for approximately 60% of those included in the CPI—as *Classification Of Individual CO*nsumption *by Purpose* (COICOP) approved 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, 146 different locations are visited.⁴ For a given location, the shop code is a unique but anonymized *id*

³This is typically the case for personal computers, whose frequent model upgrades impose the use of hedonic regressions to enhance comparisons across time.

⁴Until 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	27,479,532	27,314,761	23,258,171	19,954,005
Avg. per Month	106,099	105,462	89,800	77,042
Price Trajectories	4,333,302	4,314,903	3,196,697	2,880,332
Avg. CPI Weight	60.73%	60.37%	52.22%	46.48%
Sales and Recoveries				
Avg. per Month (Unweighted)	9.07%	9.10%	8.84%	
Avg. per Month (Weighted)	7.46%	7.49%	7.15%	
Product Substitutions				
Avg. per Month (Unweighted)	6.67%	6.67%	5.30%	
Avg. per Month (Weighted)	5.04%	5.05%	3.91%	

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 for which we can uniquely identify a price trajectory; *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, we compute the total number of price trajectories, 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 have been obtained by accounting for CPI, item-specific, stratum and shop (i.e., elementary aggregate) weights. Sample period: 1996:M2-2017:M8.

associated with the outlet from which the quote is recorded. In turn, each shop is further 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 the same outlet, but for multiple locations within the same region. As an example, in March 2013 we pick an item with the following characteristics: ‘Women’s Long Sleeves Top’ (*id*: 510223) sold in multiple outlets (*shop type*: 1) within the region of London (*region*: 2). With these coordinates at hand we retrieve two different price quotes: one location sells the item for £22 and one for £26. In February 2013 the price quotes for the same goods 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 univocally identified, we look at ‘base prices’, which are intended as the January’s price for each of the goods under scrutiny.⁵ Given this information we are able to uniquely identify the price trajectories for the two types of good. Even after conditioning on base prices, though, a small portion of price trajectories are still not uniquely identified (about 0.1%, on average):

⁵The base price is typically relied upon in order to normalize price quotes and calculate price indices, or to adjust for changes in the quality and/or quantity of a given good.

we opt for discarding these. 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 12,000 quotes per month.⁶

To aggregate the individual price quotes into a single price we also make use of the following weights produced by the ONS:⁷ 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 branch);⁸ the *stratification* weights, which reflect the fact that purchasing patterns may differ markedly by region or type of outlet;⁹ finally, the *item* and *COICOP* weights are used to 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. We start by dealing with sales, whose behavior tends to be significantly different 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. 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.¹⁰ As a second step, we apply a symmetric V-shaped filter, as defined by Nakamura and Steinsson (2010), on the remaining price quotes. According to the filter, the sale price of item i at time t , $P_{i,t}^s$, is identified as follows: i) it is lower than last period’s price (i.e., $P_{i,t}^s < P_{i,t-1}$) 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 $P_{i,t}^r$, instead, meets the following criteria: i) it is greater than last period’s price (i.e., $P_{i,t}^r > P_{i,t-1}$) and ii) it is such that $P_{i,t}^r = P_{i,t-2}$. Once a price quote has been identified as being a sale or a recovery price, we discard it from the sample.¹¹

Item substitutions are a further reason of concern when trying to identify price trajectories, as they require a certain 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

⁶Due to a particularly low coverage, Housing, Water, Electricity, Gas and Other Fuels (COICOP 4) and Education (COICOP 10) are excluded from the sample. We also exclude price changes larger than 300%, which we deem as being due to measurement errors. These take place rarely ($< 0.01\%$). Appendix A provides additional details on the construction of the dataset.

⁷See Chapter 7 of the ONS CPI Manual (ONS, 2014).

⁸In this case the ONS enters a single price of milk, but the weight attached to this is large, so as to reflect that all the branches of Tesco within the region have posted the same price.

⁹In this respect, 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).

¹⁰This definition excludes clearance sales of products that have reached the end of their life cycle.

¹¹See 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.

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 (see Bills, 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 the literature (see, e.g., Berardi et al., 2015, Berger and Vavra, 2017, and Kryvtsov and Vincent, 2017), we interrupt a trajectory whenever it encounters a substitution flag, as indicated by the ONS. Table 1 shows that, after these steps, we are down to a monthly average of 79,000 price quotes. Finally, we define the price change of item i at time t as $\Delta p_{i,t} = \log(P_{i,t}/P_{i,t-1})$.¹²

2.2 Data facts

This section unveils a number of stylized facts about the behavior of the ONS microdata.¹³ The top panels of Figure 1 report the time path of the frequency of adjustment and the average magnitude of price changes: decomposing inflation as the product of these statistics carries important information on the relationship between the distribution of price changes and inflation itself (see, e.g., Gagnon, 2009). As expected, the average price change tends to display a strong degree of positive comovement with CPI inflation, at least until the end of the Great Recession. Thus, in the last part of 2015 the two series are back moving in tandem. As for the frequency of adjustment, it is interesting to notice how this tracks very closely the contraction in the rate of inflation that starts in 2012—going well below its sample average up to that point—while only displaying a weak reversion towards the end of 2015. In the bottom panels of the figure, both statistics are split between positive and negative price changes. The frequency of positive price changes is greater than that associated with negative adjustments throughout the entire sample, while the opposite broadly holds true when comparing the average price changes in either direction. Focusing on the post-recession sample, we appreciate two key aspects: i) the downward trend in the frequency, as observed in the first panel of the figure, is mostly due to the component associated with positive price changes; ii) notwithstanding that the average of positive price changes displays a weak tendency to increase, the (mirror image of the) average of negative price changes denotes a more robust upward trend.¹⁴ Both facts point to a certain degree of asymmetry in price adjustment.

Figure 2 plots higher moments of the distribution of price changes.¹⁵ Notably, the standard deviation displays a very large increase in the aftermath of the Great Recession. In fact, as

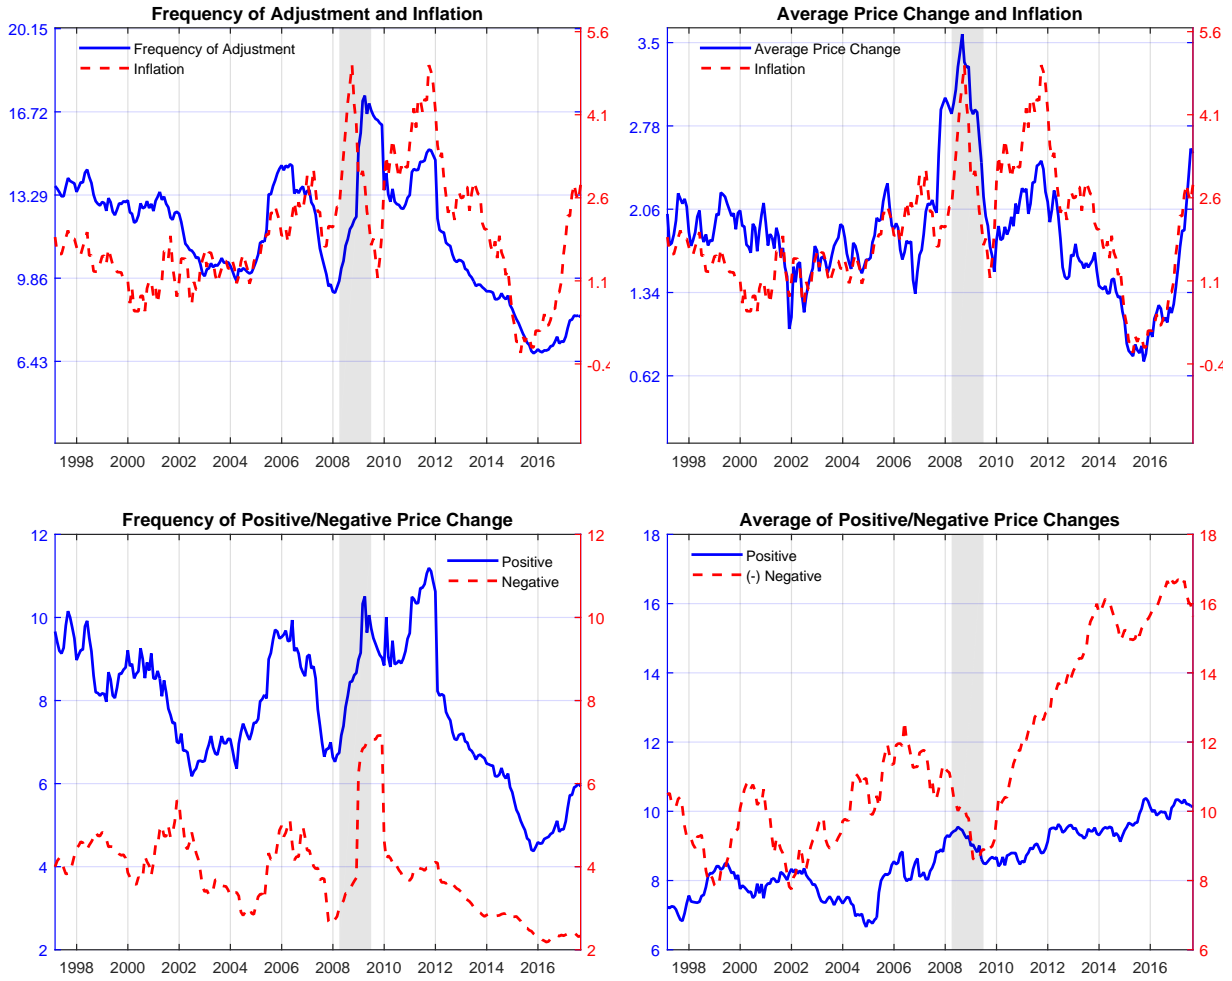
¹²We also compute price changes as $\Delta p_{i,t} = 2 \frac{P_{i,t} - P_{i,t-1}}{P_{i,t} + P_{i,t-1}}$. This definition has the advantage of being bounded and less sensitive to outliers. The results—virtually unchanged with respect to the ones we report—are available from the authors, upon request.

¹³Throughout the paper all statistics derived from the microdata on prices are reported as a 12-month moving average, so as to get rid of the seasonality in the data.

¹⁴Figure B.1 in Appendix B shows that composition effects have no role in generating various patterns presented in this subsection. To this end, we compare the moments of the distribution of price changes with their homologues obtained by averaging the corresponding moments of the price quotes obtained for each of the 25 COICOP group categories.

¹⁵To avoid that zero price changes dominate the distribution, we follow Vavra (2014) and much of the literature in that they considers only non-zero price movements.

Figure 1: FREQUENCY OF ADJUSTMENT AND AVERAGE PRICE CHANGES



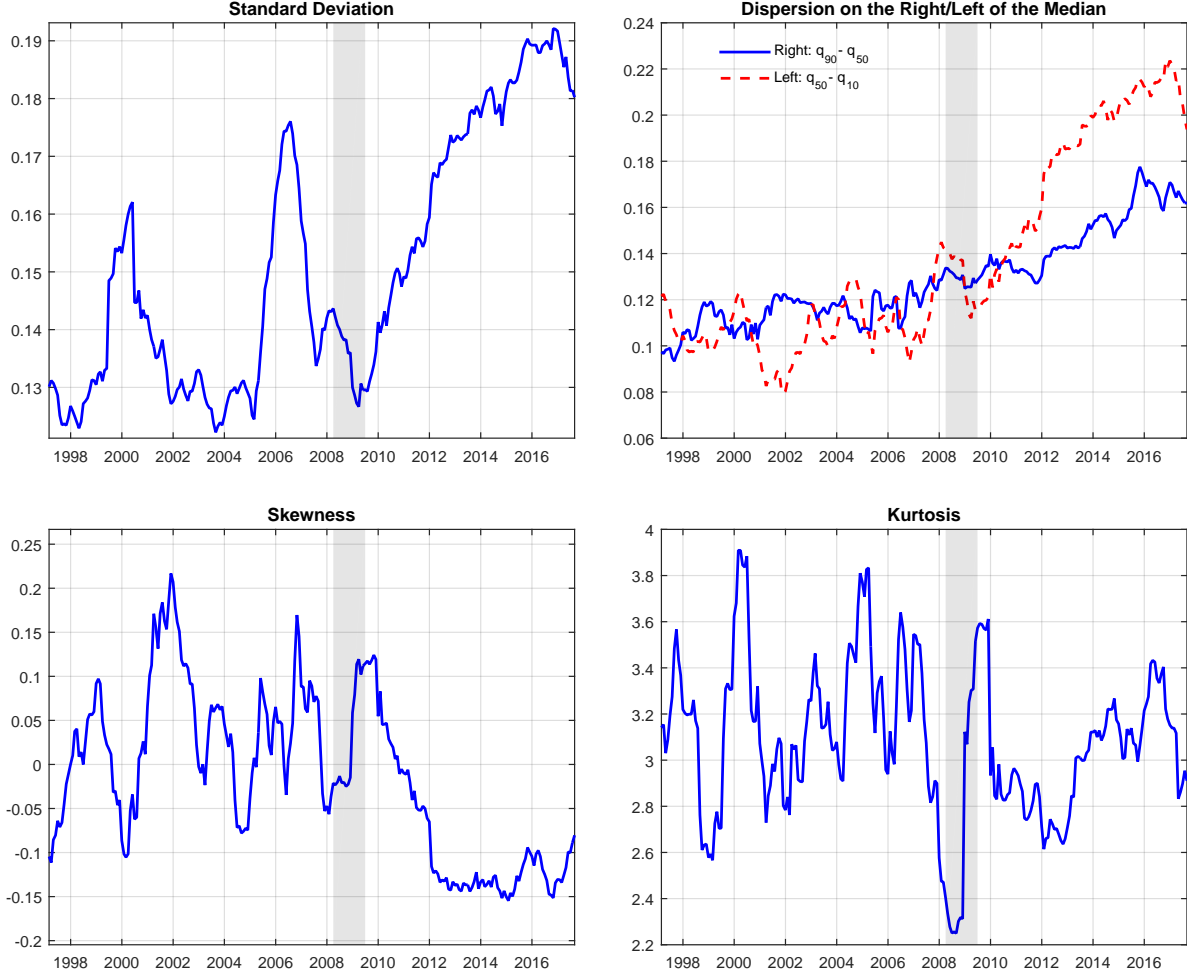
Notes: The shaded vertical band indicates the duration of the Great Recession. The inflation rate graphed in the upper panel of the figure is the official CPI inflation rate published by the ONS. In the bottom-right panel we report the mirror image of the average of the negative price changes.

displayed by the top-right panel of the figure, dispersion increases on either side of the median, though negative price changes denote a stronger acceleration in volatility, as compared with positive price changes. In light of this it should be stressed that the fall in CPI inflation occurring in the post-2010 sample is to a large extent a manifestation of the trend in the dispersion of negative price changes—relative to that of positive ones—rather than reflecting a mere shift in the mode of the density. This fact, coupled with the observation of diverging trends in the relative size of average positive/negative price changes, inevitably reflects into the dynamics of the skewness, which fluctuates around a positive mean in the pre-2010 sample, becoming persistently negative thereafter.

Table 2 reports the correlation between some of the key moments of the distribution of price changes, CPI inflation and a business cycle indicator.¹⁶ To set aside potential spurious correlation emanating from the low-frequency behavior of the series under examination, we

¹⁶Appendix C contains more details on the derivation of the monthly coincident indicator of economic activity.

Figure 2: MOMENTS OF THE DISTRIBUTION OF PRICE CHANGES



Notes: Price dispersion on the right (left) side of the median price quote is computed as $q_{50} - q_{10}$ ($q_{90} - q_{50}$). The skewness and kurtosis of the distribution of price changes are measured as $\frac{q_{90,t} + q_{10,t} - 2q_{50,t}}{q_{90,t} - q_{10,t}}$ and $\frac{q_{90,t} - q_{62.5,t} + q_{37.5,t} - q_{10,t}}{q_{75,t} - q_{25,t}}$, respectively. The shaded vertical band indicates the duration of the Great Recession.

detrend all of them aside of the inflation rate.¹⁷ Turning our attention to the frequency of adjustment and the dispersion of price changes, it is important to stress that they also display somewhat different cyclical behaviors. Looking at the entire sample, the frequency moves countercyclically, while dispersion is procyclical. However, the sign of these correlations is only preserved in the post-recession sample, while during the previous decade both statistics have behaved procyclically. Also their pairwise correlation seems to vary substantially across the two subsamples—going from being positive in the first decade to negative thereafter—though measuring dispersion through inter-quantile differences, so as to account for asymmetry, points to a negative correlation. As for the higher moments of the distribution, the skewness signals a marked countercyclical behavior, while the correlation between kurtosis and the cyclical indi-

¹⁷Moreover, when splitting the sample we exclude the period around the Great Recession (2007:M3-2010:M6), so as to avoid that the correlations among the key variables are dominated by the major macroeconomic turmoil in that period. In light of this, when interpreting the cyclical properties of the data in the two subsamples it is worth stressing that the correlations are most likely picked up by the behavior of the series in periods of relatively stronger/weaker expansion, rather than by alternative cyclical phases.

Table 2: CORRELATIONS OF PRICING MOMENTS WITH MACROECONOMIC VARIABLES

Full Sample						
	fr_t	σ_t^2	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	$Skew_t$	$Kurt_t$
y_t	-0.569***	0.264***	0.334***	0.422***	-0.363***	-0.322***
π_t	0.169***	0.000	-0.016	-0.147**	-0.024	-0.281***
fr_t	—	0.162**	-0.510***	-0.737***	0.470***	0.286***
Pre-Recession						
	fr_t	σ_t^2	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	$Skew_t$	$Kurt_t$
y_t	0.455***	0.612***	-0.121	-0.092	-0.015	0.171*
π_t	0.387***	0.213**	-0.416***	-0.410***	0.177*	0.181**
fr_t	—	0.569***	-0.120	-0.511***	0.356***	-0.055
Post-Recession						
	fr_t	σ_t^2	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	$Skew_t$	$Kurt_t$
y_t	-0.399***	0.221**	0.137	0.428***	-0.244**	0.291***
π_t	0.467***	0.077	-0.275***	-0.303***	-0.216**	-0.530***
fr_t	—	-0.475***	-0.646***	-0.854***	0.383***	-0.292***

Notes: fr_t denotes the frequency of adjustment; σ_t^2 stands for the volatility of the distribution of price changes; $q_{n,t}$ measures the n -th quantile of the distribution of price changes; $Skew_t$ denotes the skewness of the distribution of price changes and is measured as $\frac{q_{90,t} + q_{10,t} - 2q_{50,t}}{q_{90,t} - q_{10,t}}$; $Kurt_t$ denotes the kurtosis of the distribution of price changes and is measured as $\frac{q_{90,t} - q_{62.5,t} + q_{37.5,t} - q_{10,t}}{q_{75,t} - q_{25,t}}$; y_t is a business cycle indicator; π_t indicates aggregate CPI inflation. Aside of the inflation rate, all series are obtained by detrending their raw counterparts by means of Rotemberg's (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. ***/**/* indicates statistical significance at the 1/5/10% level, respectively.

cator is heavily influenced by the only recession in the time window considered, being negative in the whole sample, while turning positive in the subsamples that exclude the Great Recession.¹⁸ Table 3 broadly confirms these tendencies, while showing that the frequency of negative price changes denotes stronger countercyclicality—as compared with its counterpart computed for positive adjustments—both in the full sample and in the last decade. Concurrently, the procyclicality of the dispersion is a phenomenon that tends to characterize price changes taking place on the left side of the median—mostly in the post-recession sample—while the dispersion of negative price changes varies substantially depending on both the specific subsample and the way dispersion is measured.

To summarize the most consistent patterns of the frequency of adjustment and the dis-

¹⁸Villar and Luo (2017) show how different models of price setting account for different signs of the correlation between inflation and the skewness of price changes. In this respect, menu cost models—which feature the price change distribution becoming less skewed as inflation rises—could well rationalize our data in the second subsample, while the Calvo model—which features a positive correlation—could account for the first subsample. In the remainder of the analysis we will show how such characterization is also supported by the behavior of the extensive margin of price adjustment—a hallmark of menu cost models—assumes a prominent role in the aftermath of the Great Recession.

Table 3: CORRELATIONS OF PRICING MOMENTS WITH MACROECONOMIC VARIABLES: THE ROLE OF ASYMMETRY

Full Sample								
	fr_t^+	fr_t^-	dp_t^+	$-dp_t^-$	$q_{75,t} - q_{50,t}$	$q_{50,t} - q_{25,t}$	$q_{90,t} - q_{50,t}$	$q_{50,t} - q_{10,t}$
y_t	-0.330***	-0.636***	0.306***	0.388***	0.188***	0.295***	0.127**	0.417***
π_t	0.529***	-0.110*	0.031	0.285***	0.253***	-0.370***	-0.366***	-0.203***
Pre-Recession								
	fr_t^+	fr_t^-	dp_t^+	$-dp_t^-$	$q_{75,t} - q_{50,t}$	$q_{50,t} - q_{25,t}$	$q_{90,t} - q_{50,t}$	$q_{50,t} - q_{10,t}$
y_t	0.466***	0.210**	0.213**	0.427***	-0.162*	-0.018	-0.037	-0.059
π_t	0.154*	0.173*	-0.001	-0.018	0.057	-0.231***	0.045	-0.406***
Post-Recession								
	fr_t^+	fr_t^-	dp_t^+	$-dp_t^-$	$q_{75,t} - q_{50,t}$	$q_{50,t} - q_{25,t}$	$q_{90,t} - q_{50,t}$	$q_{50,t} - q_{10,t}$
y_t	-0.373***	-0.415***	-0.117	-0.489***	-0.696***	0.479***	0.325***	0.415***
π_t	0.858***	0.556***	-0.171	0.606***	0.535***	-0.760***	-0.702***	-0.619***

Notes: fr_t^+/fr_t^- stands for the frequency of positive/negative price changes; dp_t^+/dp_t^- indicates the average size of positive/negative price changes; $q_{n,t}$ measures the n -th quantile of the distribution of price changes; y_t is a (monthly) business cycle indicator; π_t indicates aggregate CPI inflation. Aside of the inflation rate, all series are obtained by detrending their raw counterparts by means of Rotemberg's (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. ***/**/* indicates statistical significance at the 1/5/10% level, respectively.

persion of price changes: after the Great Recession, the former has displayed a pronounced countercyclicality, while dispersion has been markedly procyclical throughout the entire sample, with both comovements appearing more marked in the case of negative price changes. Otherwise, the pairwise correlation between these statistics has turned deeply negative after the Great Recession. Notably, this picture stands in contrast with the analysis on US micro-data by Vavra (2014), who reports that the cross-sectional standard deviation of price changes is strongly countercyclical and comoves with the frequency of adjustment. To rationalize these facts, he employs a stylized menu cost model, showing how shocks to the dispersion of price gaps may play an important role. In the next section we review the same framework and show how it can also provide us with a rationale for diverging trends in the dispersion of price adjustments and their frequency.

3 Analytical framework

To frame the empirical analysis, 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

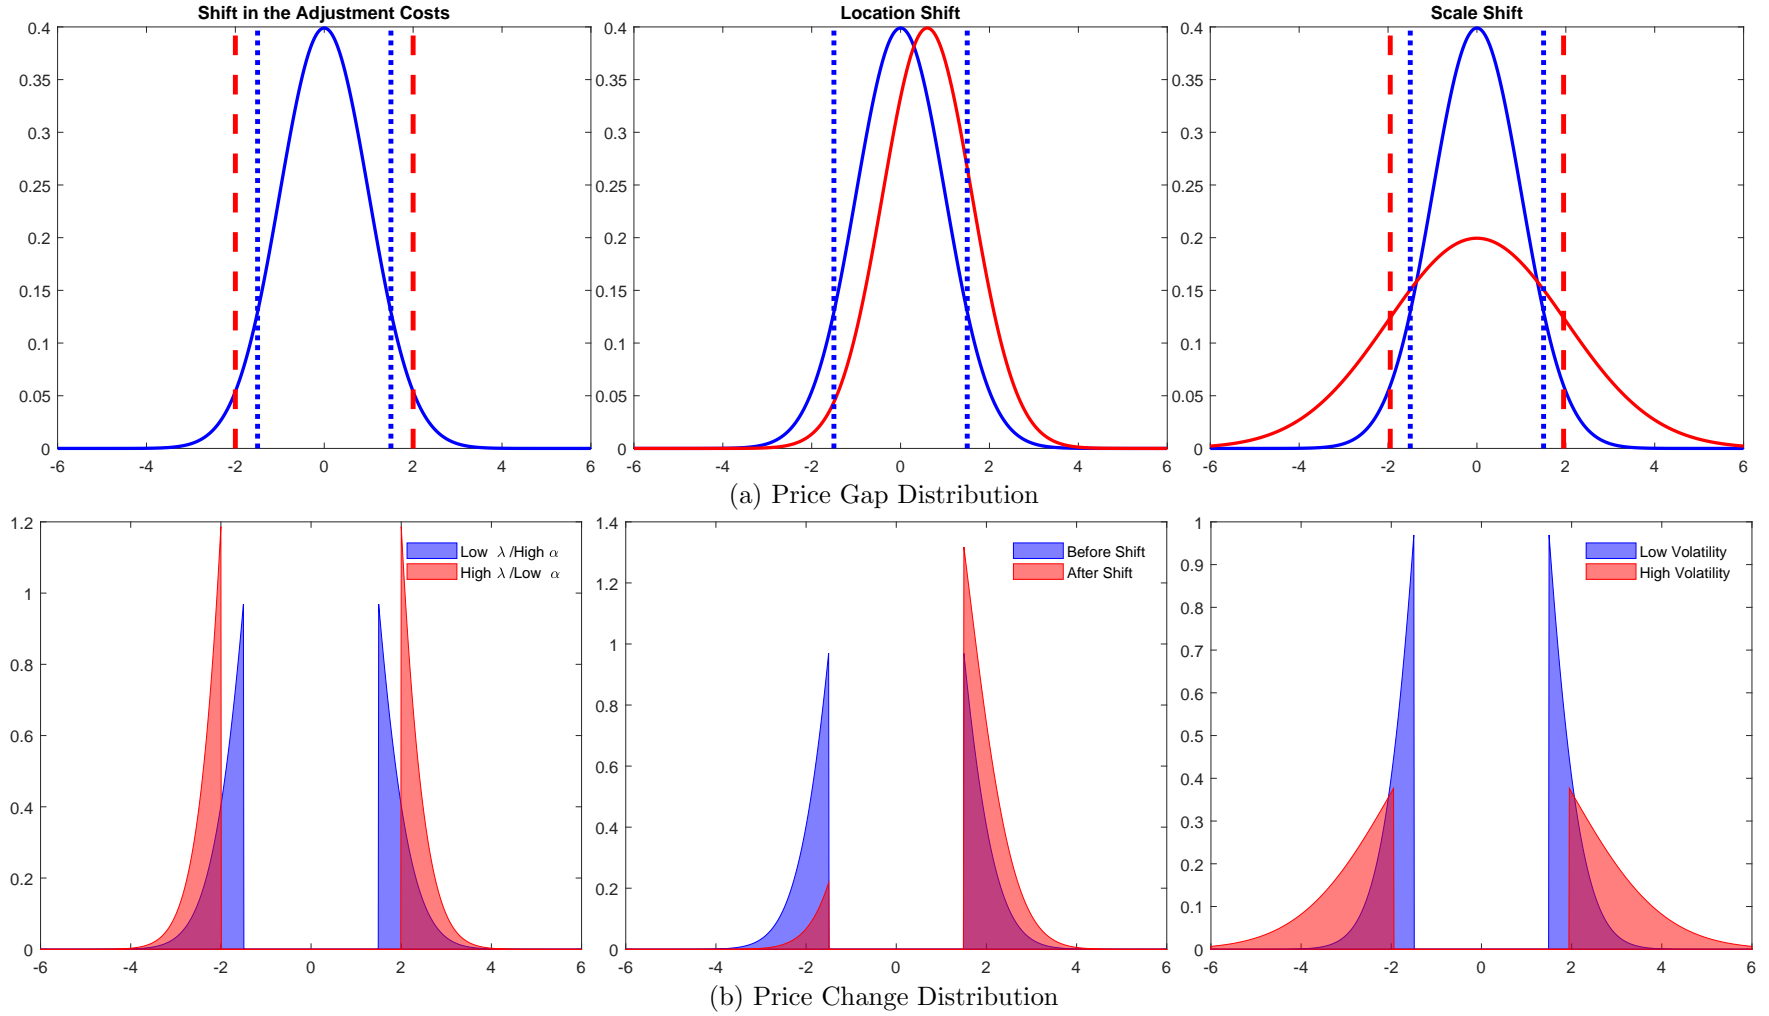
difference between the actual price of a given good and its reset price (i.e., the price that would have prevailed in the absence of price-setting frictions).

Firms face a dynamic control problem where x —the deviation of the current price from the optimal price—is defined as the 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 $\sigma = (6\lambda\phi^2/\alpha)^{1/4}$ denotes the standard deviation of price changes. Moreover, $fr = (\alpha/6\lambda)^{1/4} \phi$ is the frequency of adjustment.¹⁹

To provide an overview of the different determinants of the distribution of price gaps and the associated distribution of price changes, Figure 3 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 increases the inaction region, translating into a compression in the frequency of adjustment and an increase in the dispersion of price changes. 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. 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). Finally, a rise in ϕ , as sketched in the last column of the figure, induces both fr and σ to increase.

¹⁹For analytical details and proofs, see Barro (1972) and Vavra (2014).

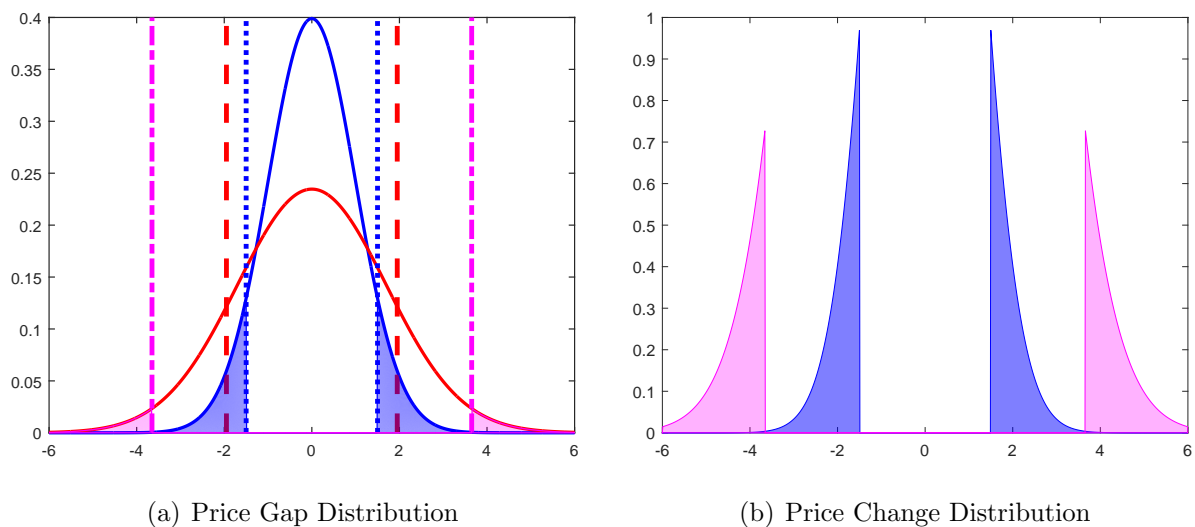
Figure 3: ANALYTICAL FRAMEWORK



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.

In light of the properties we have outlined, Vavra (2014) points to second-moment shocks as potential drivers of the positive comovement between the frequency of adjustment and the price-change dispersion that characterize U.S. CPI data throughout the time window he considers. By contrast, in the ONS microdata under examination the comovement between these two statistics is positive only in the first part of the sample, while turning negative in the following decade, when the two series display diverging trending behaviors. Therefore, as in the analysis of Vavra (2014), second-moment shocks would provide a good account of what might have happened up to the Great Recession. By contrast, shocks to x of either sign would determine relative movements in the dispersion of price changes and the frequency of adjustment which do not square with the data, regardless of the time window we consider. In fact, Section 4 will show that episodes of major repricing—such as those occurring due to changes in the VAT—do not only reflect into pre-determined price adjustments (i.e., adjustments that are determined ahead of the shock and would materialize into a mere shift of the distribution). In light of this, the so-called extensive margin of price flexibility, which accounts for adjustments that are either triggered or cancelled by the VAT change, is shown to play an important role.

Figure 4: A COMBINED INCREASE IN ϕ AND λ



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 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 increase in λ . The right panel reports the distribution of price changes, both in the ex-ante situation and in the case of a combined increase in ϕ and λ .

When looking at the post-recession experience, among the free parameters of the model only a persistent increase in the fixed cost of adjustment and/or a drop in the cost of deviating from the optimal price may account for the diverging trends we observe, conditional on the resulting expansion of the inaction region dominating the effects of positive shifts in the dispersion of price gaps. A *caveat* is in order at this stage: the menu cost model we are employing has been explicitly envisaged to investigate the effects of shocks to x in the neighborhood of

the steady state. Therefore, within this framework secular movements in the frequency of adjustment and the price-change dispersion—as those observed in the post-recession period—can be thought of as resulting from a sequence of persistent changes in the volatility and the cost parameters. Figure 4 considers a situation in which both ϕ and λ increase:²⁰ 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. As a result, we observe a negative comovement between the cross-sectional dispersion of prices and the frequency of adjustment, which is consistent with what we observe in the post-recession period. To dig deeper into these aspects, the next section introduces an accounting framework that proves to be particularly useful at quantifying the link between changes in the timing of individual price adjustments and macro price flexibility, along with formalizing the distinction between predetermined price adjustments and those which are triggered or cancelled by shocks.

4 A generalized Ss model

To account for the connection between price setting at the micro level and price flexibility at the aggregate level we use the generalized Ss model developed by Caballero and Engel (2007). This framework is consistent with lumpy and infrequent price adjustments—which are typically perceived as distinctive traits of price setting—along with encompassing several pricing protocols.²¹ Berger and Vavra (2017) also show that such an accounting approach is capable of providing a good fit to the data generated by different structural models, such as Golosov and Lucas (2007) and Nakamura and Steinsson (2008). To allow for time variation in price flexibility, we estimate the model over each cross section of micro price data, matching different price-setting statistics. More details on the estimation are reported in Section 4.1. In the remainder of this section, instead, we discuss the analytical details of the accounting framework.

Assume that, due to price rigidities, firm i 's (log of) the 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. In a simple Ss model, as the one detailed in the previous section, the price is adjusted when the price gap is large enough, and $p_{it} = p_{it}^*$ after the adjustment has taken place. Assuming l_i periods since the last price change, the adjustment reflects the cumulated shocks: $\Delta p_{it} = \sum_{j=0}^{l_i} \Delta p_{it-j}^*$, with $\Delta p_{it}^* = \mu_t + v_{it}$, where μ_t is a shock to nominal demand and v_{it} is an idiosyncratic shock.

²⁰Once again, a drop in α would lead to qualitatively similar results.

²¹To mention two 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) model—where adjusting firms change prices by such large amounts that the aggregate price is fully flexible, regardless of the frequency of adjustment.

As discussed by Caballero and Engel (2007), the basic *Ss* setting of the previous section can be generalized by assuming *iid* idiosyncratic shocks to the adjustment costs. Thus, by integrating over their possible realizations, we obtain an 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 recovered as

$$\pi_t = - \int x \Lambda_t(x) f_t(x) dx. \quad (1)$$

Notice that the Calvo pricing protocol implies the same hazard across x ’s (i.e., $\Lambda_t(x) = \Lambda_t > 0, \forall x$).

4.1 Taking the model to the data

In order to take the model to the data we need to specify generic functional forms for the distribution of price gaps and the hazard function. Specifically, 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

$$f(x) = \begin{cases} \frac{\delta(\varrho, \nu)^{1/\nu}}{\Gamma(1+1/\nu)} \exp \left[-\frac{\delta(\varrho, \nu)}{\varrho^\nu} \left| \frac{x-\theta}{\phi} \right|^\nu \right] & \text{if } x \leq \theta \\ \frac{\delta(\varrho, \nu)^{1/\nu}}{\Gamma(1+1/\nu)} \exp \left[-\frac{\delta(\varrho, \nu)}{(1-\varrho)^\nu} \left| \frac{x-\theta}{\phi} \right|^\nu \right] & \text{if } x > \theta \end{cases}, \quad (2)$$

with $\delta(\varrho, \nu) = \frac{2\varrho^\nu(1-\varrho)^\nu}{\varrho^\nu + (1-\varrho)^\nu}$. The parameters θ and $\phi > 0$ capture the location and the scale of the distribution, whereas $0 < \varrho < 1$ accounts for its degree of asymmetry. Last, the parameter $\nu > 0$ measures the degree of tail decay: for $\infty > \nu \geq 2$ the distribution is characterized by short tails, whereas it features fat tails when $2 > \nu > 0$. This functional form nests a number of standard specifications, such as the Normal ($\nu = 2$), the Laplace ($\nu = 1$) and the Uniform ($\nu \rightarrow \infty$). 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:

$$\Lambda_t(x) = \min \{ a_t + b_t x^2 \mathbf{I}(x > 0) + c_t x^2 \mathbf{I}(x < 0), 1 \}, \quad (3)$$

where $\mathbf{I}(z)$ 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 otherwise allowing for asymmetric costs of adjustment (so as to be able to capture, for instance,

stickiness, which is implied by $b_t > c_t$).²²

Given the parametric specifications of $f_t(x)$ and $\Lambda_t(x)$ we estimate seven parameters for each cross section of micro price data, 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.²³ We also match the frequency and the average size of prices movements, after distinguishing between 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.²⁴

4.2 Making sense of diverging trends in the frequency and dispersion of price changes

The first two panels of Figure 5 report the estimated scale parameter of $f(x)$ and the inaction region associated with two hazard probabilities (5% and 7%). Both statistics increase substantially in the second decade of the sample, thus implying that—at least over this period—first-moment shocks do not appear as the main determinant of price adjustment. According to our comparative statics analysis in Section 3, a downward trend in the frequency of adjustment, coupled with a surge in its dispersion, may be rationalized by an expansion in the inaction region—as prompted by an increase in the fixed cost of adjustment and/or a drop in the cost of deviating from the optimal price, for instance—that overcomes the effects of a positive shift in the dispersion of price gaps. To verify this is indeed the case, the last panel of Figure 5 reports the share of prices in the inaction region, obtained as the proportion of prices whose $\Lambda(x)$ is lower than a given hazard rate. Notably, by the end of the sample about five times as many firms are inactive, as compared with the pre-2010 time window. This stands as indirect evidence that the expansion in the inaction region, as captured by the downward shift in the hazard function, dominates the increase in the dispersion of $f(x)$.²⁵ As we will discuss in the next section, changes in the shape of the distribution of price gaps, coupled with the expansion of the inaction region, 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.


²²We have checked that the results are robust to plausible variations to this specification. Specifically, using a mixture of two Normal distributions for the price gap and/or the asymmetric inverted normal function for the hazard function delivers results that are qualitatively similar to those reported in the next section.

²³We match quantilic moments as the 3rd and 4th moments of the cross-sectional distribution are quite sensitive to outliers.

²⁴Altonji 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).

²⁵Figure F.2 reports the estimated parameters of the APD, while Figure F.1 graphs the dynamics of both $f(x)$ and $\Lambda(x)$.

Figure 5: DISPERSION OF PRICE GAPS AND THE INACTION REGION



Note: The three panels of the figure report the estimated scale parameter of $f(x)$, the inaction region (for two different hazard rates), and the corresponding share of prices within the inaction region, respectively. The shaded vertical band indicates the duration of the Great Recession.

5 Implications for aggregate price adjustment

The estimation of the generalized Ss model highlights the importance of changes in the distribution of price gaps and the hazard function. These inevitably reflect in the way shocks are propagated to the economy. To dig deeper into the connection between individual price adjustment and the dynamics of the response of aggregate inflation to nominal demand, Caballero and Engel (2007) highlight that, within their accounting framework, one can derive a measure of aggregate price flexibility that accounts for the impact response of realized inflation to a one-off aggregate nominal shock:

$$\mathcal{F}_t = \lim_{\mu_t \rightarrow 0} \frac{\partial \pi_t}{\partial \mu_t} = \underbrace{\int \Lambda_t(x) f_t(x) dx}_{\text{Intensive Margin}} + \underbrace{\int x \Lambda'_t(x) f_t(x) dx}_{\text{Extensive Margin}}. \quad (4)$$

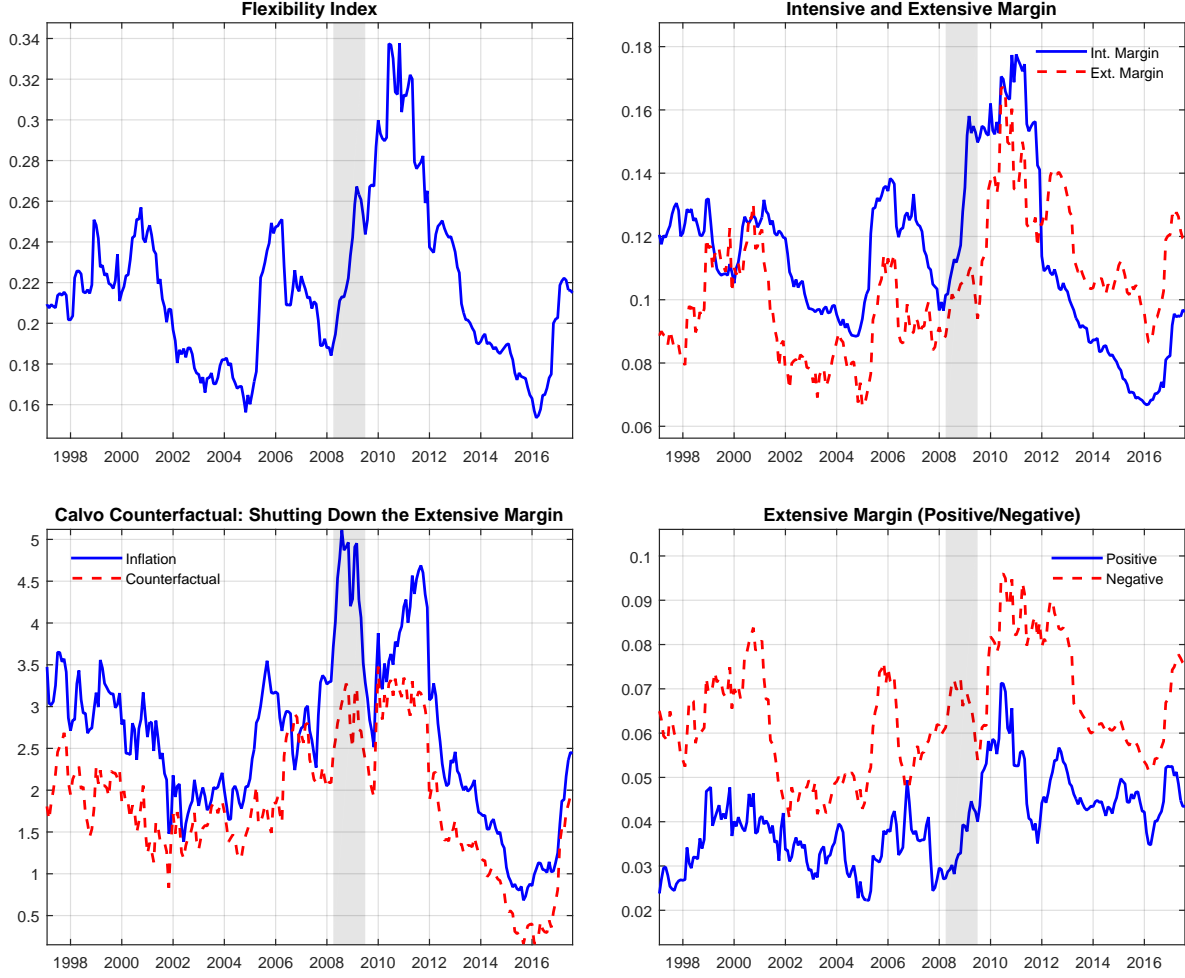
Since this flexibility index is simply derived from the accounting identity (1), its validity as a measure of aggregate flexibility does not require that we take a stand on a specific model of price setting.²⁶

The flexibility index 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 who would have kept their price constant and instead change it, as well as firms who would have adjusted their price but choose not to do it. 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.

The top panels of Figure 6 report the estimated price flexibility index and its decomposition into the intensive and the extensive margin of price adjustment for the period under investigation. Aggregate price flexibility displays sizable variation over time, and even more so in the last part of the sample, rising substantially during the Great Recession, and declining thereafter. This is consistent with our analysis of the distribution of price gaps. In fact, after the Great Recession both the intensive and the extensive margin of price adjustment display similar trending behaviors, though the fall in the former is much more abrupt, in line with the sustained downward trend in the frequency of adjustment. As for the extensive margin, the expansion in the inaction region implies that fewer firms are pushed near the adjustment boundaries. Moreover, it should be stressed that, over most of the decline, the extensive margin

²⁶In this respect, Alvarez et al. (2016) show that the steady-state ratio of kurtosis to frequency is a sufficient statistic for monetary non-neutrality in a wide variety of frameworks. However, as highlighted by Berger and Vavra (2017), while their characterization provides us with a measure of cumulative output response, it does not apply to settings that allow for large shocks to the price gap distribution. Despite these fundamental differences, when comparing the two measures obtained from our data, they display a strong negative correlation, as one would expect on theoretical grounds.

Figure 6: PRICE FLEXIBILITY AND DIFFERENT MARGINS OF PRICE ADJUSTMENT



Notes: The bottom-left panel reports both the rate of inflation obtained from our sample of ONS price quotes and its counterfactual, obtained by setting the period hazard function to a constant equal to the intensive margin. The shaded vertical band indicates the duration of the Great Recession.

tends to contribute more to price flexibility, as compared with the intensive one, even after they both revert in 2016. Otherwise, the relative importance of the frequency of adjustment has generally been higher prior to 2012, with few short lived exceptions. 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 (4):

$$\mathcal{F}_t = \int \Lambda_t(x) f_t(x) [1 + \eta_t(x)] dx \quad (5)$$

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 consequence, the importance of the extensive margin relative to the intensive one. This is exactly what happens in the period under examination, as it can be appreciated by inspecting the estimated constant of the hazard function (see, Figure F.3 in Appendix F). Alternatively, the same point can

be made by approximating the flexibility index as $F_t \cong Int_t + 2[Int_t - \Lambda_t(0)]$:²⁷ from this approximation it is clear how a downward shift in a_t —which is equivalent to $\Lambda_t(0)$ —translates into an increase in the importance of the extensive margin relative to the intensive one, *ceteris paribus*.

Table 4: FLEXIBILITY IN PRICE ADJUSTMENT: CORRELATION WITH REAL ACTIVITY AND INFLATION

Full Sample							
	\mathcal{F}_t	Int_t	Ext_t	Int_t^+	Int_t^-	Ext_t^+	Ext_t^-
y_t	-0.233***	-0.352***	-0.060	-0.532***	-0.190***	-0.210***	0.044
π_t	0.380***	0.398***	0.281***	0.005	0.565***	-0.061	0.467***
Pre-Recession							
	\mathcal{F}_t	Int_t	Ext_t	Int_t^+	Int_t^-	Ext_t^+	Ext_t^-
y_t	0.456***	0.368***	0.412***	0.223**	0.395***	0.403***	0.331***
π_t	-0.012	0.269***	-0.279***	0.062	0.345***	-0.311***	-0.221**
Post-Recession							
	\mathcal{F}_t	Int_t	Ext_t	Int_t^+	Int_t^-	Ext_t^+	Ext_t^-
y_t	-0.527***	-0.428***	-0.559***	-0.363***	-0.416***	-0.289***	-0.632***
π_t	0.678***	0.718***	0.512***	0.372***	0.787***	0.084	0.721***

Notes: The table reports pairwise correlations of output and inflation with the flexibility index, as well as the intensive margin and the extensive margin of price adjustment (together with their counterparts corresponding to positive and negative price gaps). Aside of the inflation rate, all series are obtained by detrending their raw counterparts by means of 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. ***/**/* indicates statistical significance at the 1/5/10% level, respectively.

To gauge the actual contribution of the extensive margin to inflation dynamics, we can take a step further: the bottom-left panel of Figure 6 reports both the overall rate of inflation and its counterfactual, obtained by setting the period hazard function to a constant equal to the intensive margin. As pointed out by Gagnon et al. (2013), this is equivalent to calibrating the Calvo model to match the intensive margin of price adjustment by assuming that the probability of price adjustment, while exogenous to the firm, can vary with the state of the economy (i.e., $\pi_t^{Calvo} = -f r_t^{Calvo} \int x f_t(x) dx$, where $f r_t^{Calvo} = \int \Lambda_t(x) f_t(x) dx$). The presence of an increasing hazard function tends to exacerbate the impact of large shocks (Caballero and Engel, 1991). In fact, the extensive margin proves to be rather important in periods of particularly volatile inflation, so that the difference between the latter and its ‘Calvo counterfactual’ is sizable. In this respect, it is important to appreciate how movements along the extensive margin may reflect some asymmetries in the adjustment of prices in either direction. To this end, the last panel of the figure reports the extensive margin associated with positive and negative price gaps (Ext^+ and Ext^- , respectively).²⁸ Both statistics denote a shift during the

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

²⁸To this end, we simply rely on the following decomposition of the extensive margin: $\int_{-\infty}^0 x \Lambda_t'(x) f_t(x) dx +$

last part of the sample, with Ext^+ leading the increase in the wake of the Great Recession, and Ext^- reflecting the two hikes in the VAT at the beginning of 2010 and 2011. This aspect will be examined in further detail in the next subsection.

From a business cycle perspective, variations in price flexibility do not seem to occur at random: in fact, \mathcal{F}_t goes from being markedly procyclical in the first part of the sample to invert its cyclical in the last decade (see Table 4). As for the correlation with the rate of inflation, this is generally positive, and more so in the post-recession sample, while it is not statistically different from zero in the previous decade. Analogous changes in the correlation with real activity occur when looking at Ext and Int over the two subsamples, while the strong positive correlation with the rate of inflation is even more pronounced in the last decade. It is interesting to notice that both margins denote a negative correlation with the cyclical indicator in the overall sample, which is particularly pronounced for Ext^+ and Int^+ .²⁹ This fact, in conjunction with a correlation with the rate of inflation that is not statistically different from zero, might indicate a certain degree of downward rigidity, given that nominal shocks appear not to be able to stimulate price cuts along both margins. However, notice that the correlation structure of both margins of (negative) price adjustment changes markedly across the two subsamples, indicating that price cuts might have been particularly sticky during the Great Recession.³⁰

As a final note on the change in correlation we observe over the two subsamples, it is worth emphasizing how this is consistent with a shift from an environment where the intensive margin dominates the extensive one, to an environment where the extensive margin assumes a prominent role and inflation volatility is particularly marked (see Figure 6).

5.1 Price adjustment and the importance of state-dependent pricing: a VAT event study

Examining the relative importance of price adjustment along the extensive margin is of crucial importance to contrast time-dependent models that are widely employed in quantitative macroeconomic frameworks, with state-dependent models. In this respect Gagnon et al. (2013) suggest that, if the timing of all price changes was predetermined, following a nominal shock we should observe a shift in the gap distribution, with the shape of the distribution being preserved (see, e.g., the middle panel of Figure 3). Thus, one can measure the importance of adjustment along the extensive margin by comparing the observed distribution of price changes to a counterfactual distribution that obtains in the absence of the shock. Any evidence that the two distributions differ by more than a shift can be attributed to the extensive margin. To

$\int_0^\infty x \Lambda'_t(x) f_t(x) dx$, where Ext_t^- (Ext_t^+) is the first (second) term on the right side of the equality. To see a similar split for the intensive margin, recall that the bottom-left panel Figure 1 reports the frequency of positive and negative price changes.

²⁹When looking at the two separate subsamples, we notice that the countercyclical is a hallmark of the last decade, while the cyclical comovement is positive in the pre-recession period.

³⁰In this respect, Gilchrist et al. (2017) have shown how the interplay between price stickiness and financial frictions faced by firms operating in customer markets might have acted as sources of downward price rigidity during the Great Recession.

this end, we can usefully exploit episodes of massive repricing triggered by changes in the VAT. These are relatively simple to study, because their timing and size are directly observable. The recent UK history has been characterized by three episodes of changes in the VAT: 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. To examine the contribution of VAT changes to the overall degree of price flexibility, Figure 7 reports the distribution of price gaps and that of price changes, together with the corresponding hazard function. Moreover, we report their counterfactuals, obtained by averaging the same function, for the same month of the year, in the previous six years.³¹

Table 5: VAT CHANGES: ACTUAL AND COUNTERFACTUAL STATISTICS

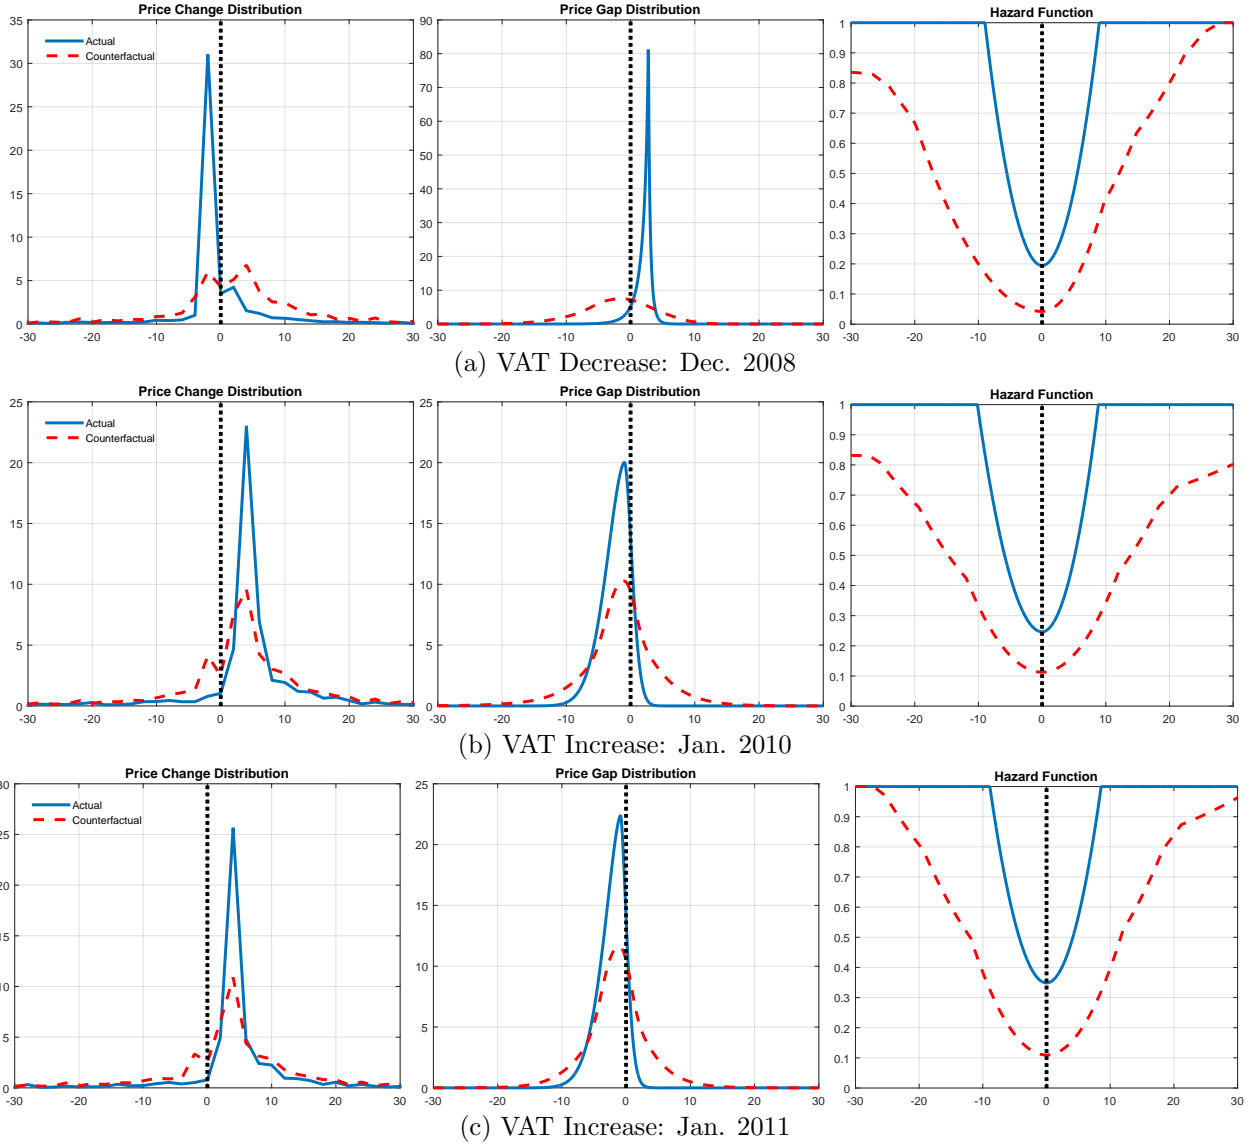
	VAT 1							
	π	\mathcal{F}	Int	Ext	Int^+	Int^-	Ext^+	Ext^-
Actual	-5.941	0.346	0.235	0.111	0.211	0.023	0.105	0.006
Scenario 1	-1.604	0.101	0.060	0.041	0.055	0.005	0.040	0.001
Scenario 2	1.863	0.200	0.096	0.104	0.038	0.058	0.048	0.056
	VAT 2							
	π	\mathcal{F}	Int	Ext	Int^+	Int^-	Ext^+	Ext^-
Actual	11.631	0.471	0.322	0.149	0.019	0.304	0.003	0.146
Scenario 1	4.580	0.181	0.135	0.045	0.008	0.127	0.001	0.045
Scenario 2	4.111	0.218	0.148	0.070	0.043	0.105	0.016	0.054
	VAT 3							
	π	\mathcal{F}	Int	Ext	Int^+	Int^-	Ext^+	Ext^-
Actual	14.487	0.573	0.428	0.145	0.019	0.409	0.002	0.143
Scenario 1	4.708	0.190	0.136	0.053	0.006	0.130	0.001	0.053
Scenario 2	4.258	0.239	0.154	0.086	0.041	0.113	0.020	0.066

Notes: The table reports the inflation rate, the inflation rate that would have been observed had there not been any extensive margin, the flexibility index, the intensive and extensive margins of price adjustment (as well as their counterparts computed for positive and negative price gaps), in the month of a VAT change. Three recent episodes of changes in the VAT are considered: a reduction from 17.5% to 15% on December 1, 2008 (indicated by VAT 1), followed by two hikes, on up to 17.5% on January 1, 2010 and then up to 20% on January 4, 2011 (indicated by VAT 2 and VAT 3, respectively). For every episode we contrast the actual numbers with two alternative scenarios. Scenario 1 considers a case in which the VAT change only impacts on the distribution of price gaps, while keeping the hazard function at its counterfactual in Figure 7. Scenario 2, instead, consider an alternative where neither the hazard function nor the price gap distribution change.

Looking at the inflation rate in the month corresponding to a VAT change, we notice that shifts in the distribution of price changes are such that many firms seize the opportunity to adjust prices by more than the VAT change, thus implying that inflationary/deflationary pressures from other sources have been released in the process. In support of the view that

³¹January 2010 has not been included when computing the counterfactual distribution for January 2011, so as to avoid that the second VAT change affects the counterfactual distribution corresponding to the last episode.

Figure 7: EVENT STUDY: VAT CHANGES



Notes: Each line of the figure reports the distribution of price changes, the distribution of price gaps, and the hazard function in the month corresponding to a VAT change. The distribution of price changes is computed by grouping observations into bins of 2% (excluding zeros), and weighting them by their relative importance in the CPI. In all cases, the counterfactuals are computed by averaging the same function, for the same month of the year in the previous 6 years. Three recent episodes of changes in the VAT are considered: a reduction from 17.5% to 15% on December 1, 2008, followed by two hikes, on up to 17.5% on January 1, 2010 and then up to 20% on January 4, 2011.

episodes of massive repricing cannot be seen as mere translations of the distribution of price gaps, we appreciate both a major upward shift and a steepening of the hazard function across all the three episodes of VAT change: in fact, these are associated with a large rise in the frequency of adjustment.

To dig deeper into the role of state-dependent pricing, Table 5 reports some statistics in coincidence with the three VAT changes, as well as two alternative scenarios.³² In the first scenario, we keep the hazard function as that computed in the counterfactual exercise, but let

³²More details on the computation of two alternative scenarios are provided in Appendix D.

the price gap distribution vary as a result of the VAT change. Therefore, we abstract from any amplification that could be potentially induced by state-dependent pricing through upward shifts of the hazard function. The second scenario, instead, considers a hypothetical case in which neither the price gap distribution nor the hazard function are affected by the VAT change. Effectively, this case describes what were likely to be observe had the VAT not changed. From the comparison between actual inflation in the occurrence of a VAT change and the alternative scenarios, two observations are worth emphasizing. First, state-dependent pricing has a major role in amplifying the impact of a VAT change on the rate of inflation, accounting for most of the change in the rate of inflation. Second, in the absence of state-dependent pricing, shifts of the price gap distribution and drops in its dispersion would result in a substantial drop of price flexibility, with both the intensive and the extensive margin decreasing.

Importantly, when comparing the two margins of adjustment, the intensive one is typically twice as large as its counterfactual—indicating that upward shifts in $\Lambda(0)$ are the dominant feature in the occurrence of changes in the VAT—while movements along the extensive margin appear less evident. However, this conclusion is not warranted after conditioning both margins to positive and negative price changes. In this case, substantial variation takes place along the extensive margin coherent with the sign of the underlying price change. For instance, in the occurrence of the drop in the VAT from 17.5% to 15% (December 2008), Ext^+ is 0.048 in the counterfactual, while actually being more than twice as big. The same order of magnitude can be appreciated when making the same comparison for two VAT hikes (in this case we need to focus on Ext^-). Movements in the extensive margin are a reflection of the interplay between the hazard function and the distribution of price gaps. In this respect, Figure 7 shows that all the episodes of VAT change are associated with a close-to-symmetric increase in the steepness of the hazard function, as well as with a shift in the distribution of price gaps in the direction opposite to the sign of a given VAT change. On one hand, this necessarily implies that the extensive margin associated with price gaps coherent with the sign of the adjustment is substantial. On the other hand, the extensive margin associated with price gaps of the opposite sign is very low, as consequence of the hazard function being weighed by a very small probability mass, after the shift in the distribution of price gaps.

On a slightly different note, it should be stressed that price flexibility reaches its maximum over roughly the same period we observe the most recent VAT price changes. In light of our analysis, this comes as little surprise, given that this type of events typically offer price setters with some windows of opportunity to release at least some of the accumulated (positive/negative) price pressure. However, from a normative perspective the opportunity to enhance coordination between fiscal and monetary policy should be carefully considered. Such a prescription might be particularly relevant in contexts such as the one examined, where the potential real effects of the accommodative monetary-policy stance might have been baffled by VAT changes that were mainly inspired by stimulus- or revenue-based considerations.

6 Inflation dynamics and state dependence in price flexibility

The estimation of the generalized Ss model shows that the pass-through of nominal shocks to inflation is highly variable. We also show that—while not hinging on a specific margin of adjustment—flexibility is higher in connection with positive price changes, while downward price adjustments are typically stickier. These properties bear major implications for evaluating the transmission of shocks to nominal demand. In fact, at the eyes of a hypothetical Central Banker, aggregate price flexibility should be regarded as a key state variable to predict the influence of a given monetary policy stance on prices and quantities.

While aggregate price flexibility only accounts for the impact response of inflation to a nominal shock, one would expect it to contain valuable information to address state dependence in inflation dynamics. 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 key disadvantage in that we may end up with relatively few observations in a certain 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:

$$\begin{aligned} \pi_t = & G\left(\tilde{\mathcal{F}}_{t-1}, \gamma\right) \left(\phi_0^H + \sum_{j=1}^p \phi_j^H \pi_{t-j} + \varepsilon_t^H + \sum_{i=1}^q \theta_i^H \varepsilon_{t-i}^H \right) \\ & + \left[1 - G\left(\tilde{\mathcal{F}}_{t-1}, \gamma\right) \right] \left(\phi_0^L + \sum_{j=1}^p \phi_j^L \pi_{t-j} + \varepsilon_t^L + \sum_{i=1}^q \theta_i^L \varepsilon_{t-i}^L \right), \end{aligned} \quad (6)$$

with $\varepsilon_t^i \sim N(0, \sigma_i^2)$ for $i = \{L, H\}$. Moreover, we set $G\left(\tilde{\mathcal{F}}, \gamma\right) = (1 + e^{-\gamma\tilde{\mathcal{F}}})^{-1}$, where $\tilde{\mathcal{F}}$ 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. The likelihood of the model can be easily computed by recasting the system in state space (see, e.g., Harvey, 1990). Given the highly nonlinear nature of the framework, we use Monte Carlo Markov-chain methods developed in Chernozhukov and Hong (2003) for

³³In this respect, the STARMA(p,q) model also generalizes the threshold ARMA(p,q) model (DeGooijer, 2017).

³⁴Note also that 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.

³⁵We 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.

estimation and inference. The parameter estimates, as well as their standard errors, are directly computed from the generated chains.³⁶

As we focus on the post-1996 sample, we calibrate the constant terms ϕ_0^H and ϕ_0^L so that in both regimes the long-run inflation forecast is 2%, consistent with the mandate of the Bank of England. Whereas one can potentially estimate the speed of transition between regimes, the identification of γ relies on nonlinear moments. Moreover, in short samples the estimates may be sensitive to a handful of observations. Therefore, we decide to calibrate γ so that roughly 25% (30%) of the observations are classified to be in the high-flexibility (low-flexibility) regime, where this is defined by $G(\tilde{\mathcal{F}}_{t-1}; \gamma) > 0.8$ ($1 - G(\tilde{\mathcal{F}}_{t-1}; \gamma) > 0.8$).³⁷ Thus, based on the Akaike criterion, we choose $p = 1$ and $q = 7$.³⁸

Figure 8 reports the impulse-response functions to a 1% shock to inflation in each of the two regimes, and compares them to the response from an equivalent linear model. Inflation is much more persistent in periods characterized by a relatively low price flexibility, with the half-life of the shock being almost twice as large, as compared with periods of high flexibility. In fact, the estimated inflation volatility is 1.44 in the high-flexibility regime and 0.91 in the low-flexibility regime. These results are broadly supportive of the basic insights of the *Ss* model illustrated in the previous sections, and highlight the importance of keeping track of the degree of price flexibility.

Notably, the impulse-response function from the linear model is consistent with the behavior of inflation in the low-flexibility regime. A direct implication of this is that neglecting that shocks are propagated at different speeds—depending on the overall degree of price flexibility—would entail an overestimation of their inflationary impact during windows of relatively high flexibility. This should be particularly evident at medium-term forecast horizons, i.e. when the difference between the responses from the linear and the nonlinear model is somewhat larger. This begs the following question: do the Bank of England or other market participants take price flexibility into account when computing their inflation expectations? In the remainder of this section we turn our attention to addressing this issue. In this respect, our premise delivers a key testable implication: if state dependence in price flexibility is accounted for by the forecaster, the resulting inflation forecast errors should be orthogonal to the flexibility regime.


In every quarter, the Inflation Report of the Bank of England publishes (year-on-year) Monetary Policy Committee’s inflation forecasts, along with market participants’ forecasts. Both types of forecasts refer to the Bank of England’s inflation target, which has switched from RPIX inflation to CPI inflation in 2004:Q1. Thus, we construct quarterly (absolute)

³⁶See Appendix E for further details.

³⁷Figure G.1 in Appendix G reports the dynamics of $G(\tilde{\mathcal{F}}_{t-1}; \gamma)$. Clearly, this specification identifies the 2009-2012 period as being characterized by a high-flexibility regime, whereas the 2002-2005 and 2015-2016 periods are marked by low price flexibility. The qualitative results are robust to variations in γ .

³⁸Note that the modified AIC information criterion indicates a STARMA(1,3). Figures G.2 and G.3 in Appendix G report the results for this alternative setting. Our key insights are not affected by the exact specification of the STARMA(p,q) model.

Figure 8: PRICE FLEXIBILITY AND INFLATION PERSISTENCE



Note: This figure reports the responses of inflation to a 1% shock in the STARMA(1,7) model. The 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(1,7) model. 68% confidence intervals are built based on the Markov Chain Monte Carlo (MCMC) method developed in Chernozhukov and Hong (2003).

forecast errors as the (absolute value of the) difference between the mean forecast³⁹ and the appropriate forecast target at a given horizon. These are then regressed on a nonlinear function of the flexibility regime indicator, $G(\tilde{\mathcal{F}}_{t-1}; \gamma)$: specifically, we use a quadratic spline function with a knot at 0.5. This function is a rather flexible tool, as it allows us to capture a number of potential shapes characterizing the relationship between the flexibility regime and the forecast errors. Table 6 provides a summary of the results from our regression exercise. The fitted function tends to reach a minimum at about $G(\tilde{\mathcal{F}}_{t-1}; \gamma) = 0.6$, for all forecast horizons. Thus, we report the slope of the function at values of the indicator equal to 0.3 and 0.9 (so as to consider an equal distance from the minimum point), as well as at 0.5. The last two columns of the table also report the p-value associated with the null that no relationship between the forecast error and the flexibility regime exists, as well as the r-squared (adjusted for the number of regressors), so as to get an idea of the strength of the relationship. The results are consistent with the idea that information about the degree of price flexibility is not fully exploited by the Central Bank or by market participants. Consistent with the results in Figure 8, we find that the relationship tends to be stronger at medium-term horizons, while weakening at both short-term and long-term horizons. Specifically, around a four-quarter horizon, price flexibility accounts for roughly 25% of the variability in the absolute forecast error. The relationship is not statistically significant in periods of low flexibility, whereas it is usually significant in periods of high flexibility, with the slope displaying larger values at medium-term forecast horizons. These results are roughly the same, no matter which type of forecasts we examine.

The pronounced time variation in price flexibility after the Great Recession helps us to get a

³⁹Table G.4 in Appendix G reports similar results using squared forecast errors. In both cases, the results are virtually unchanged if we use median in place of mean forecasts.

Table 6: FORECAST ERRORS AND PRICE FLEXIBILITY

(a) BoE MPC RPIX/CPI (Absolute) Forecast Errors								
Horizon	Slope at $G = 0.3$		Slope at $G = 0.5$		Slope at $G = 0.9$		F-stat	\tilde{R}^2
1	0.093	[0.628]	-0.578	[0.232]	0.840	[0.092]	0.229	1.69
2	-0.330	[0.279]	-1.342	[0.060]	2.319	[0.011]	0.045	6.41
3	-0.484	[0.145]	-2.433	[0.027]	4.117	[0.010]	0.003	13.82
4	-0.344	[0.437]	-3.798	[0.007]	6.161	[0.003]	0.000	26.45
5	-0.144	[0.811]	-3.277	[0.024]	5.945	[0.011]	0.000	20.10
6	0.309	[0.603]	-2.542	[0.118]	4.858	[0.032]	0.003	13.70
7	0.634	[0.236]	-2.749	[0.116]	4.402	[0.021]	0.006	12.32
8	0.691	[0.182]	-1.758	[0.302]	3.029	[0.055]	0.063	5.93
(b) Market Participants' (Absolute) Forecast Errors								
Horizon	Slope at $G = 0.3$		Slope at $G = 0.5$		Slope at $G = 0.9$		F-stat	\tilde{R}^2
1	0.265	[0.361]	-1.107	[0.066]	0.826	[0.122]	0.278	1.11
2	-0.383	[0.264]	-1.658	[0.043]	2.448	[0.010]	0.053	6.12
3	-0.561	[0.150]	-2.883	[0.017]	4.293	[0.008]	0.004	13.10
4	-0.382	[0.418]	-4.522	[0.002]	6.398	[0.002]	0.000	25.60
5	-0.103	[0.862]	-4.086	[0.010]	6.042	[0.009]	0.000	18.74
6	0.453	[0.412]	-3.207	[0.101]	4.516	[0.049]	0.013	10.48
7	0.903	[0.052]	-3.380	[0.095]	3.631	[0.052]	0.019	9.47
8	0.883	[0.099]	-2.385	[0.189]	1.935	[0.221]	0.211	2.19

Notes: The table reports the results of a quadratic spline regression of the absolute forecast errors $e_{T+h|T}$ (for different forecast horizons, h) on a quarterly average of an indicator of the normalized price flexibility index, $G_t = G(\tilde{\mathcal{F}}_t; \gamma) = (1 + e^{-\gamma \tilde{\mathcal{F}}_t})^{-1}$, where $\tilde{\mathcal{F}}$ denotes the normalized flexibility index. The regression takes the form: $|e_{T+h|T}| = a_0 + a_1 G_t + a_2 G_t^2 + a_3 G_t^2 I_{\{G_t > 0.5\}}$. The upper panel refers to the Bank of England MPC's RPIX/CPI forecast errors, while the bottom panel considers market participants' forecast errors. In each panel, the first three pairs of columns report the slope of the relationship evaluated at different levels of the indicator, together the p-value associated with the null hypothesis that the slope is equal to 0 (this is calculated using Newey-West standard errors). 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. The last column reports the adjusted r-squared, denoted by \tilde{R}^2 .

better understanding of the concurrent dynamics of the inflation rate. Over this time window inflation peaks twice between 2008 and 2011, while reaching its sample minimum in 2016, partially reflecting sharp movements in the value of the GBP and commodity prices.⁴⁰ The Bank of England has generally underestimated the speed and impact of shocks to inflation in the 2008-2011 period. In light of our results, this points to a potential failure in appreciating that price flexibility was itself at its historical peak, possibly as a reflection of the three VAT adjustments taking place in a rather short time window. Conversely, the low-flexibility regime can explain the protracted period of low inflation towards the end of the sample, during which the Bank of England has displayed greater predictive accuracy. This regime of low price flexibility has then reversed around the summer of 2016, in coincidence with the sharp movements of the GBP in

⁴⁰Two main facts are worth noticing with respect to the UK experience in the post-recession sample: i) inflation has been outside the 1%-3% interval for a total of 22 out of 40 quarters, whereas the same happened only during 11 quarters in the previous decade; ii) over the same period, inflation has also shot above and below the target, reaching both its maximum (+4.8%) and minimum value (-0.1%) in the overall sample.

the aftermath of the Brexit referendum.

7 Concluding remarks

We document some distinctive patterns in the evolution of the distribution of micro price changes in the UK, and discuss their implications for the transmission of nominal stimulus to output and inflation. By estimating the generalized S_s model of Caballero and Engel (2007), we are able to report that price flexibility displays pronounced time variation. In fact, over the last decade the capacity of nominal stimulus to generate inflation has decreased substantially. Despite the marked non-linearity in the price response to inflationary shocks—which is crucially dictated by the degree of price flexibility—neither the Bank of England nor professional forecasters appear to account for this type of state dependence when forecasting CPI inflation. In fact, both of them tend to overestimate the impact of inflationary shocks in periods of relatively high price flexibility, especially at medium-term forecast horizons. In light of this, we point to price flexibility as a state variable that both practitioners and policy makers should carefully account for in their forecasting routine. In this respect, we also suggest that time variation in price flexibility should be considered as a key dimension of monetary-policy making. To this end, we observe that changes in price flexibility correlate with departures of CPI inflation from the target, potentially providing the policy-maker with a basis for assessing her state-contingent capability to influence price and output growth.

A final note on the implications of our results for modeling price setting: by imposing a Calvo price-setting protocol to match the frequency of adjustment one could understate time-variation in price flexibility, which is heavily influenced by the extensive margin of price setting, especially during periods of high volatility in inflation dynamics. In this respect, our work does not just emphasize the importance of time variation in higher moments of the distribution of price changes and their connection with price flexibility—one of the main conclusions of works in this area of research—but also assigns a prominent role to state-dependent price setting in order to understand inflation dynamics, which is what Central Banks and practitioners are ultimately concerned with.

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. L. BIHAN, 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, L. J., E. DHYNE, M. HOEBERICHTS, C. KWAPIL, H. L. BIHAN, P. LNNEMANN, 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.
- 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.
- 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.
- 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 (1991): “Dynamic (S, s) Economies,” *Econometrica*, 59, 1659–1686.
- (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 (2017): “Price Selection,” Manuscript.
- CHERNOZHUKOV, V. AND H. HONG (2003): “An MCMC approach to classical estimation,” *Journal of Econometrics*, 115, 293–346.
- COSTAIN, J. AND A. NAKOV (2011): “Distributional dynamics under smoothly state-dependent pricing,” *Journal of Monetary Economics*, 58, 646–665.
- D’AGOSTINO, A., D. GIANNONE, M. LENZA, AND M. MODUGNO (2016): “Nowcasting Business Cycles: A Bayesian Approach to Dynamic Heterogeneous Factor Models,” *Advances in Econometrics*, in: *Dynamic Factor Models*, 35, 569–594.
- DEGOOIJER, J. G. (2017): *Elements of Nonlinear Time Series Analysis and Forecasting*, Springer.
- DIXIT, A. (1991): “Analytical Approximations in Models of Hysteresis,” *Review of Economic Studies*, 58, 141–151.
- DIXON, H. D., K. B. LUINTEL, AND K. TIAN (2014): “The impact of the 2008 crisis on UK prices: what we can learn from the CPI microdata,” Cardiff Economics Working Papers E2014/7, Cardiff University, Cardiff Business School, Economics Section.
- 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.
- GAGNON, E. (2009): “Price Setting during Low and High Inflation: Evidence from Mexico,” *The Quarterly Journal of Economics*, 124, 1221–1263.
- GAGNON, E., D. LOPEZ-SALIDO, AND N. VINCENT (2013): “Individual Price Adjustment along the Extensive Margin,” *NBER Macroeconomics Annual*, 27, 235–281.

- 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.
- GILCHRIST, S., R. SCHOENLE, J. SIM, AND E. ZAKRAJEK (2017): “Inflation Dynamics during the Financial Crisis,” *American Economic Review*, 107, 785–823.
- 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.
- HARVEY, A. C. (1990): *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press.
- 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 (2017): “The Cyclicalities of Sales and Aggregate Price Flexibility,” Manuscript.
- MARIANO, R. S. AND Y. MURASAWA (2003): “A new coincident index of business cycles based on monthly and quarterly series,” *Journal of Applied Econometrics*, 18, 427–443.
- MIDRIGAN, V. (2011): “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations,” *Econometrica*, 79, 1139–1180.
- 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.
- (2010): “More Facts About Prices – Supplement to “Five Facts About Prices: A Reevaluation of Menu Cost Models,”” Online Supplement.

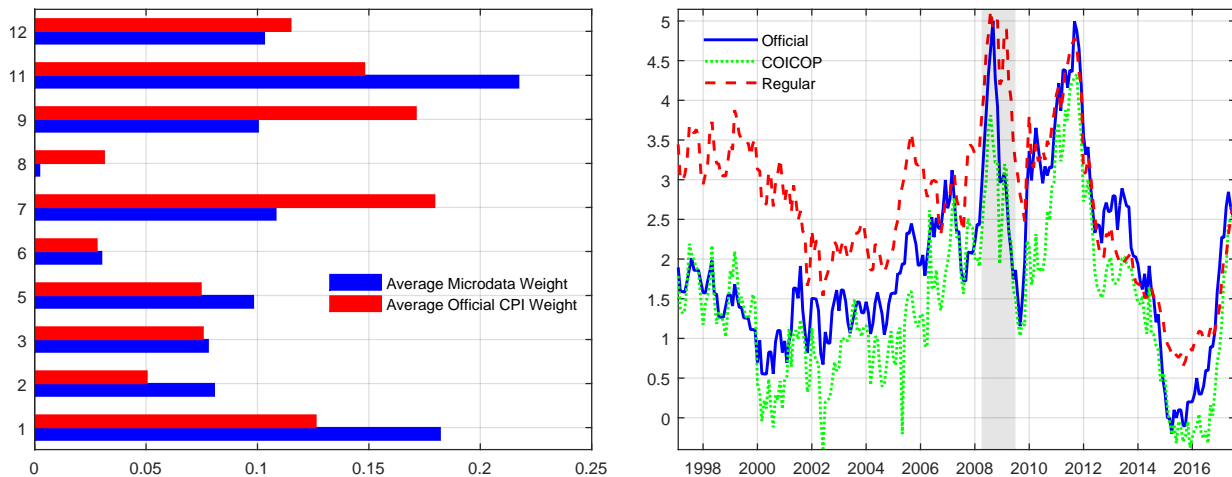
- NAKAMURA, E., J. STEINSSON, P. SUN, AND D. VILLAR (2018): “The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation,” *Quarterly Journal of Economics*, Forthcoming.
- ONS (2014): *Consumer Price Indices: Technical Manual*, Office for National Statistics, 2014 ed.
- ROTEMBERG, J. J. (1999): “A Heuristic Method for Extracting Smooth Trends from Economic Time Series,” NBER Working Papers 7439, National Bureau of Economic Research, Inc.
- VAVRA, J. (2014): “Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation,” *The Quarterly Journal of Economics*, 129, 215–258.
- VILLAR, D. V. AND S. LUO (2017): “The Skewness of the Price Change Distribution: A New Touchstone for Sticky Price Models,” Finance and Economics Discussion Series 2017-028, Board of Governors of the Federal Reserve System (U.S.).
- YELLEN, J. L. (2016): “Macroeconomic Research After the Crisis: a speech at the 60th annual economic conference sponsored by the Federal Reserve Bank of Boston, Boston, Massachusetts, October 14, 2016,” Speech, Board of Governors of the Federal Reserve System.

A On the representativeness of the data

This section provides additional details on the construction of the dataset used in the empirical analysis. The ONS data have a good coverage of all COICOP sectors with the exception of Housing, Water, Electricity, Gas And Other Fuels (COICOP 4), Communication (COICOP 8) and Education (COICOP 10), whose coverage are less than 15%, 4%, and 3%, respectively. Given the extremely low coverage, we exclude COICOP 4 and 10, whereas we keep COICOP 8 as the available price quotes are clustered in a small subset of item goods for the sector, such as Flower Delivery, Telephone for home use and Phone Accessories.⁴¹

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 (labelled *COICOP*) and the inflation obtained once all filters described in Section 2 are applied (labelled *Regular*). Unfiltered data track quite closely the official numbers, whereas the regular series displays a robust correlation with the official data (roughly 0.7), 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 (labelled *COICOP*) and the inflation obtained once all filters described in Section 2 are applied (labelled *Regular*). The COICOP codes are (1) Food And Non-Alcoholic Beverages, (2) Alcoholic Beverages, Tobacco And Narcotics, Clothing And Footwear (3), Furnishings, Household Equipment And Routine Household Maintenance (5), Health (6), Transport (7), Communication (8), Recreation And Culture (9), Hotels, Cafes And Restaurants (11), Miscellaneous Goods And Services (12).

⁴¹Due to the small number of price quotes in this sector, the results would be little affected by its exclusion from the analysis.

B On the role of 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 homologues 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 the Great Recession.

C A monthly coincident indicator of economic activity

In Tables 2, 3 and 4 we report the correlation of different variables with respect to a business cycle indicator. The latter is computed as a monthly coincident indicator of GDP growth, where we use monthly information on a number of monthly macroeconomic indicators of economic activity to infer the underlying movements of GDP at the monthly frequency. Following Mariano and Murasawa (2003), we approximate the (normalized) quarterly growth of real GDP, Δy_t^q , as a moving average of an unobserved month-on-month GDP growth rate, Δy_t^* :

$$\Delta y_t^q = \frac{1}{3}\Delta y_t^* + \frac{2}{3}\Delta y_{t-1}^* + \Delta y_{t-2}^* + \frac{2}{3}\Delta y_{t-3}^* + \frac{1}{3}\Delta y_{t-4}^*.$$

We then assume that Δy_t^* can be decomposed into an aggregate component, α_t , which is common across a number of other macroeconomic indicators, and an idiosyncratic component, ε_t :

$$\Delta y_t^* = \alpha_t + \varepsilon_t.$$

We assume that the idiosyncratic component follows an autoregressive process of order one:

$$\varepsilon_t = \psi\varepsilon_{t-1} + \eta_t.$$

The other macroeconomic indicators are available at a monthly frequency. We specify (the standardized value of) each of them as the sum of two mutually orthogonal components, a common and an idiosyncratic one. The former is captured by the current and lagged values of the aggregate common factor (see, e.g., D'Agostino et al., 2016). Specifically, denoting with Δx_{it} the generic i -th macroeconomic indicator, we have that

$$\Delta x_{it} = \sum_{j=1}^l \lambda_{ij}\alpha_{t-j} + e_{it},$$

where e_{it} follows an autoregressive process of order one:

$$e_{it} = \rho_i e_{it-1} + v_{it},$$

where the innovations to the idiosyncratic process are *iid* and uncorrelated across the indicators (i.e., $E(v_{it}v_{jt}) = 0, \forall i \neq j$, and $E(v_{it}\eta_t) = 0, \forall i$).

We let the aggregate factor follow an autoregressive process of order two:

$$\alpha_t = \phi_1\alpha_{t-1} + \phi_2\alpha_{t-2} + u_t.$$

In our specific application, we set $l = 3$ and all autoregressive processes are restricted to be stationary. The model can be cast in state space. Therefore, the likelihood can be easily computed through the Kalman filter and the factor is retrieved by using the Kalman smoother (see Harvey, 1990).

Together with the GDP data, we use following short term (monthly) macroeconomic indicators: (1) the index of manufacturing, (2) the index of services, (3) retail sales (excl. Auto Fuel), (4) Employment and (5) unemployment (claimants count). We use data starting on January 1990: we rely on a sample that is longer than the one employed in our analysis, so as to include two recessionary episodes. The dataset is unbalanced, as some of the indicators are not available from the starting date (and GDP is observed only once in the quarter). This is not an issue, as the Kalman filter can easily deal with an arbitrary pattern of missing observations in the sample.

Table C.1 reports the fit of the aggregate components for the quarter-on-quarter growth rates of each of the variables being employed. Clearly, the single-factor specification is able to capture a large fraction of the variation in the set of indicators considered here. Figure C.1 reports quarter-on-quarter variations in the aggregate factor ($\alpha_t^q = \frac{1}{3}\alpha_t + \frac{2}{3}\alpha_{t-1} + \alpha_{t-2} + \frac{2}{3}\alpha_{t-3} + \frac{1}{3}\alpha_{t-4}$), together with the GDP growth. The level of the business cycle indicator is then computed by cumulating the common factor over time, and

assuming that trend growth equals the mean of GDP growth over the sample (this is denoted by μ):

$$z_t = \sum_{\tau=1}^t (\hat{\mu} + \hat{\alpha}_{\tau}),$$


where $\hat{\alpha}_{\tau}$ is retrieved by using the Kalman smoother. The business cycle indicator is then computed by applying a simple filter to z_t . For the baseline results in the paper we use Rotemberg (1999) version of the HP filter, which chooses the smoothing coefficient of the HP filter so as to minimize the correlation between the cycle and the first difference of the trend estimate.

Table C.1: COINCIDENT INDICATOR - MODEL FIT

	$R^2(\%)$
GDP	87.9
Index of Manufacturing	39.6
Index of Services	82.4
Retail Sales	14.7
Employment	23.3
Unemployment	22.4

Notes: The table reports the fit of the coincident business cycle indicator on the quarter-on-quarter growth rate of the indicators used to retrieve it.

Figure C.1: MONTHLY GDP AND DETRENDED COINCIDENT INDICATOR



Note: The left panel shows the fit of the (monthly) coincident indicator on the (annualized) quarter-on-quarter growth of real GDP. The right panel reports the detrended GDP using 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. The vertical green lines denote the end and the beginning of the subsamples used to exclude the Great Recession from the analysis.

D Alternative scenarios in the occurrence of a VAT change

Recall that inflation in the occurrence of a VAT change is computed as

$$\pi_t^{VAT} = - \int x \Lambda_t^{VAT}(x) f_t^{VAT}(x) dx,$$

implying that the observed inflation results from both changes in the distribution of price gaps, as well as from shifts in the hazard function. Based on this benchmark, one can envisage two relevant scenarios:

- *Scenario 1*: What rate of inflation would have been observed, had the VAT change only been associated with a change in the price gap distribution, while keeping the incentives of changing prices fixed? To address this question, we compute the following counterfactual rate of inflation

$$\pi_t^{VAT,1} = - \int x \Lambda_t^{No-VAT}(x) f_t^{VAT}(x) dx$$

- *Scenario 2*: What inflation would have been observed in absence of changes in the price gap distribution and the hazard function? This can be retrieved as

$$\pi_t^{VAT,2} = - \int x \Lambda_t^{No-VAT}(x) f_t^{No-VAT}(x) dx$$

The No-VAT counterfactual is computed by averaging the same function, for the same month of the year in the 6 years before the VAT change.

Comparing $\pi_t^{VAT,2}$ with the actual rate of inflation highlights the overall effects of the VAT, whereas the comparison between $\pi_t^{VAT,1}$ and observed inflation quantifies the relevance of the state dependence in price setting (i.e., the fact that incentives to change prices are themselves a function of the underlying environment).

E Estimation of the STARMA (p,q) model

Recall the smooth transition ARMA model, STARMA(p,q), in Section 6:

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

This can be easily cast 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

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


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 $\bar{\theta} = \frac{1}{N} \sum_{i=1}^N \theta^{(i)}$ 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.

F Model estimates

Figure F.1: ESTIMATED PRICE GAP DISTRIBUTIONS AND HAZARD FUNCTIONS




(a) Price Gap Distributions



(b) Hazard Functions


Note: The red lines denote the three VAT changes in the sample. The shaded vertical band indicates the duration of the Great Recession.

Figure F.2: PARAMETERS OF THE GAP DISTRIBUTION



Note: The red lines denote the three VAT changes in the sample. The shaded vertical band indicates the duration of the Great Recession.

Figure F.3: PARAMETERS OF THE HAZARD FUNCTION



Note: The red lines denote the three VAT changes in the sample. The shaded vertical band indicates the duration of the Great Recession.

G Additional figures and tables

Table G.1: CORRELATIONS OF PRICING MOMENTS WITH MACROECONOMIC VARIABLES (QUADRATIC TRENDS)

Full Sample						
	fr_t	σ_t^2	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	$Skew_t$	$Kurt_t$
y_t	-0.486***	0.452***	0.143**	0.330***	-0.318***	-0.015
π_t	0.497***	-0.182***	0.055	-0.265***	-0.065	-0.381***
fr_t	—	-0.098	-0.186***	-0.575***	0.267***	0.004
Pre-Recession						
	fr_t	σ_t^2	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	$Skew_t$	$Kurt_t$
y_t	0.381***	0.576***	-0.492***	-0.368***	-0.141	0.406***
π_t	0.393***	0.206**	-0.420***	-0.539***	0.111	0.169*
fr_t	—	0.402***	0.067	-0.484***	0.122	-0.168*
Post-Recession						
	fr_t	σ_t^2	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	$Skew_t$	$Kurt_t$
y_t	-0.733***	0.652***	0.197*	0.578***	-0.172	0.634***
π_t	0.918***	-0.449***	-0.141	-0.372***	-0.220**	-0.704***
fr_t	—	-0.587***	-0.272**	-0.511***	-0.074	-0.619***

Notes: fr_t denotes the frequency of adjustment; σ_t^2 stands for the volatility of the distribution of price changes; $q_{n,t}$ measures the n -th quantile of the distribution of price changes; $Skew_t$ denotes the skewness of the distribution of price changes and is measured as $\frac{q_{90,t} + q_{10,t} - 2q_{50,t}}{q_{90,t} - q_{10,t}}$; $Kurt_t$ denotes the kurtosis of the distribution of price changes and is measured as $\frac{q_{90,t} - q_{62.5,t} + q_{37.5,t} - q_{10,t}}{q_{75,t} - q_{25,t}}$; y_t is a business cycle indicator; π_t indicates aggregate CPI inflation. Aside of the inflation rate, all series are detrended with a linear and a quadratic trend. ***/**/* indicates statistical significance at the 1/5/10% level, respectively.

Table G.2: CORRELATIONS OF PRICING MOMENTS WITH MACROECONOMIC VARIABLES: THE ROLE OF ASYMMETRY (QUADRATIC TRENDS)

Full Sample								
	fr_t^+	fr_t^-	dp_t^+	$-dp_t^-$	$q_{75,t} - q_{50,t}$	$q_{50,t} - q_{25,t}$	$q_{90,t} - q_{50,t}$	$q_{50,t} - q_{10,t}$
y_t	-0.346***	-0.541***	0.120*	0.572***	-0.078	0.233***	0.058	0.333***
π_t	0.717***	0.149**	0.135**	-0.105	0.148**	-0.034	-0.411***	-0.127**
Pre-Recession								
	fr_t^+	fr_t^-	dp_t^+	$-dp_t^-$	$q_{75,t} - q_{50,t}$	$q_{50,t} - q_{25,t}$	$q_{90,t} - q_{50,t}$	$q_{50,t} - q_{10,t}$
y_t	0.480***	0.049	-0.141	0.705***	-0.547***	-0.065	-0.399***	-0.120
π_t	0.733***	0.447***	0.110	-0.039	-0.121	0.242***	-0.295***	-0.179**
Post-Recession								
	fr_t^+	fr_t^-	dp_t^+	$-dp_t^-$	$q_{75,t} - q_{50,t}$	$q_{50,t} - q_{25,t}$	$q_{90,t} - q_{50,t}$	$q_{50,t} - q_{10,t}$
y_t	-0.721***	-0.684***	-0.098	0.376***	-0.527***	0.412***	0.525***	0.449***
π_t	0.891***	0.767***	-0.062	-0.083	0.352***	-0.392***	-0.703***	-0.269**

Notes: fr_t^+/fr_t^- stands for the frequency of positive/negative price changes; dp_t^+/dp_t^- indicates the average size of positive/negative price changes; $q_{n,t}$ measures the n -th quantile of the distribution of price changes; y_t is a business cycle indicator; π_t indicates aggregate CPI inflation. Aside of the inflation rate, all series are detrended with a linear and a quadratic trend. ***/**/* indicates statistical significance at the 1/5/10% level, respectively.

Table G.3: FLEXIBILITY IN PRICE ADJUSTMENT: CORRELATION WITH REAL ACTIVITY AND INFLATION (QUADRATIC TRENDS)

Full Sample							
	\mathcal{F}_t	Int_t	Ext_t	Int_t^+	Int_t^-	Ext_t^+	Ext_t^-
y_t	-0.492***	-0.502***	-0.398***	-0.615***	-0.396***	-0.462***	-0.297***
π_t	0.584***	0.620***	0.443***	0.293***	0.725***	0.124*	0.578***
Pre-Recession							
	\mathcal{F}_t	Int_t	Ext_t	Int_t^+	Int_t^-	Ext_t^+	Ext_t^-
y_t	0.093	0.202**	-0.035	-0.041	0.296***	-0.166*	0.043
π_t	0.495***	0.705***	0.181**	0.393***	0.772***	-0.033	0.272***
Post-Recession							
	\mathcal{F}_t	Int_t	Ext_t	Int_t^+	Int_t^-	Ext_t^+	Ext_t^-
y_t	-0.800***	-0.787***	-0.713***	-0.648***	-0.799***	-0.453***	-0.791***
π_t	0.769***	0.788***	0.645***	0.562***	0.832***	0.304***	0.775***


Notes: The table reports pairwise correlations of output and inflation with the flexibility index, as well as the intensive margin and the extensive margin of price adjustment (together with their counterparts corresponding to positive and negative price gaps). Aside of the inflation rate, all series are detrended with a linear and a quadratic trend. ***/**/* indicates statistical significance at the 1/5/10% level, respectively.

Table G.4: FORECAST ERRORS AND PRICE FLEXIBILITY: ROBUSTNESS (MSE)

(a) BoE MPC RPIX/CPI (Squared) Forecast Errors								
Horizon	Slope at $G = 0.3$		Slope at $G = 0.5$		Slope at $G = 0.9$		F-stat	\tilde{R}^2
1	0.082	[0.668]	-0.584	[0.221]	0.544	[0.222]	0.559	-1.21
2	-0.304	[0.512]	-2.455	[0.018]	3.073	[0.009]	0.150	2.98
3	-0.553	[0.333]	-5.411	[0.011]	8.413	[0.011]	0.005	12.45
4	-0.440	[0.617]	-9.238	[0.008]	15.556	[0.014]	0.000	25.80
5	0.018	[0.989]	-8.985	[0.014]	17.463	[0.023]	0.000	22.44
6	0.818	[0.540]	-5.907	[0.154]	14.810	[0.054]	0.001	16.12
7	1.564	[0.212]	-4.600	[0.353]	11.514	[0.091]	0.009	11.29
8	2.145	[0.135]	-1.640	[0.789]	6.578	[0.285]	0.117	4.02
(b) Market Participants' (Squared) Forecast Errors								
Horizon	Slope at $G = 0.3$		Slope at $G = 0.5$		Slope at $G = 0.9$		F-stat	\tilde{R}^2
1	0.713	[0.291]	-1.977	[0.110]	0.426	[0.497]	0.363	0.25
2	-0.396	[0.464]	-3.170	[0.011]	3.491	[0.007]	0.123	3.65
3	-0.763	[0.287]	-6.768	[0.005]	9.235	[0.008]	0.007	11.63
4	-0.608	[0.517]	-11.327	[0.003]	16.589	[0.010]	0.000	24.46
5	-0.063	[0.960]	-11.353	[0.007]	18.043	[0.016]	0.000	20.81
6	0.923	[0.465]	-8.287	[0.103]	14.287	[0.045]	0.005	13.17
7	1.789	[0.129]	-6.908	[0.209]	9.562	[0.099]	0.043	7.16
8	2.315	[0.091]	-4.113	[0.467]	3.916	[0.431]	0.390	0.02


Notes: The table reports the results of a quadratic spline regression of the squared forecast errors $e_{T+h|T}$ (for different forecast horizons, h) on a quarterly average of an indicator of the normalized price flexibility index, $G_t = G(\tilde{\mathcal{F}}_t; \gamma) = (1 + e^{-\gamma \tilde{\mathcal{F}}_t})^{-1}$, where $\tilde{\mathcal{F}}$ denotes the normalized flexibility index. The regression takes the form: $e_{T+h|T}^2 = a_0 + a_1 G_t + a_2 G_t^2 + a_3 G_t^2 I_{\{G_t > 0.5\}}$. The upper panel refers to the Bank of England MPC's RPIX/CPI forecast errors, while the bottom panel considers market participants' forecast errors. In each panel, the first three pairs of columns report the slope of the relationship evaluated at different levels of the indicator, together the p-value associated with the null hypothesis that the slope is equal to 0 (this is calculated using Newey-West standard errors). 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. The last column reports the adjusted r-squared, denoted by \tilde{R}^2 .

Figure G.1: PROBABILITY OF A HIGH-FLEXIBILITY REGIME




Note: The figure reports the probability of ending up in a high-flexibility regime, obtained in accordance with the STARMA(1,7) model presented in Section 6. The shaded vertical band indicates the duration of the Great Recession.

Figure G.2: PRICE FLEXIBILITY AND INFLATION PERSISTENCE



Note: Figure G.2 reports the responses of inflation to a 1% shock in the STARMA(1,3) model. The 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(1,3) model. 68% confidence intervals are built based on the Markov Chain Monte Carlo (MCMC) method developed in Chernozhukov and Hong (2003).

Figure G.3: PRICE FLEXIBILITY AND INFLATION VOLATILITY



Notes: Each panel reports the distribution of the estimated inflation volatility in the two regimes. The left panel refers to the STARMA(1,7), while the right panel refers to the STARMA(1,3).