

Time-varying Price Flexibility and Inflation Dynamics*

Ivan Petrella[†]
University of Warwick
& CEPR

Emiliano Santoro[‡]
University of Copenhagen

Lasse de la Porte Simonsen[§]
Birkbeck College

November 2018

Abstract

We study how and to what extent inflation dynamics is shaped by time variation in the capacity of nominal demand to stimulate price adjustment. In doing so, we use microdata underlying the UK consumer price index, and estimate a generalized Ss model of lumpy price adjustment, condensing large cross-sectional information on micro price changes into a measure of price flexibility. The latter displays sizable time variation, which maps into a marked non-linearity of inflation dynamics: the half-life of the rate of inflation is twice as large in periods of relatively low flexibility, along with appearing remarkably close to the one observed in a linear setting. State dependence plays a major role for price setting, with the extensive margin of price adjustment becoming prominent when inflation is particularly high and volatile. Neglecting these facts may severely bias our understanding of inflation dynamics.

JEL codes: E30, E31, E37, C22.

Keywords: Inflation, Price flexibility, Ss models.

*We thank for comments Luca Dedola, Federico Di Pace, Huw Dixon, Etienne Gagnon, Ana Galvao, Ricardo Nunes, Giorgio Primiceri, Omar Rachedi, Federico Ravenna, Søren Hove Ravn, Yad Selvakumar, Daniel Villar, and seminar participants at the Bank of Finland, the ESCoE/ONS/BoE Conference on Economic Measurement, the National Institute of Economic and Social Research, the University of Surrey, and the European Central Bank. 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 increased availability of disaggregated data on consumer prices has allowed economists to study in deep detail the role of price setting for the transmission of shocks and the conduct of monetary policy. Micro price data can be usefully employed to compute measures of aggregate *price flexibility*, which is broadly intended as the response of the aggregate price level to monetary shocks. This concept lies at the core of the monetary policy transmission mechanism, ultimately embodying Central Banks’ capacity to affect output and inflation. Despite a large number of empirical contributions being concerned with measuring the response of prices to nominal stimulus, little emphasis has been placed on the sources and the extent of time variation in price flexibility.¹ Accounting for these traits substantially improves our understanding of inflation dynamics. Using micro price data underlying the UK consumer price index (CPI), we show that aggregate inflation is deeply affected by price flexibility, being substantially more persistent and less volatile in periods of relatively low flexibility. Failing to account for this non-linearity in inflation dynamics leads to a substantial bias in predicting future inflation. These results are of crucial importance to both practitioners and central bankers.

Our analysis starts by documenting sizable time variation in the distribution of price changes over the last twenty years. Most notably, while in the first half of the sample the frequency of adjustment has been roughly stable, it has dropped markedly during the last decade. Moreover, after the Great Recession the dispersion of price changes has denoted a sustained increase, and its pairwise correlation with the frequency of adjustment has turned deeply negative.² We interpret time variation in the distribution of price changes through the prism of a stylized menu cost model à la Barro (1972). Within this setting, we show how divergent movements in the dispersion of price changes and the frequency of adjustment may result from changes in the incentives that firms face when setting prices, as prompted by an increase in the fixed cost of adjustment and/or a drop in the cost of deviating from the optimal price. To the extent that the resulting expansion in the *inaction region* (i.e., the area where it is not worth adjusting prices) overcomes the effects of shocks affecting 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.

To test this prediction and discipline our data, we estimate the generalized *Ss* model developed by Caballero and Engel (2007), fitting the distribution of price gaps and the *hazard*

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

²These facts stand in contrast with the behavior of US microdata, where the cross-sectional standard deviation of price changes typically displays positive comovement with the frequency of adjustment (see, e.g., Vavra, 2014).

function (i.e., the probability of a good’s price changing as a function of its price gap) over the price quotes available in each month. Along with encompassing various price-setting protocols, and allowing us to examine different determinants of the distribution of price changes, this empirical framework is particularly well-suited to examine time variation and comovements among various price-setting statistics. The estimation reveals that both the distribution of price gaps and the adjustment hazard vary over time. In line with our prediction, about five times as many firms appear ‘inactive’ by the end of the sample, as compared with the pre-2010 time period, thus reflecting a drop in the frequency of adjustment and a concurrent increase in the dispersion of price changes. This confirms that changes in the width of the inaction region may result not only from shocks to the dispersion of price gaps (Vavra, 2014), but also from changes in firms’ incentives to adjust their prices, as embodied by their adjustment cost structure.

At the macroeconomic level, time variation in the frequency of adjustment and the dispersion of micro price changes are important in that they reflect shifts in the price-setting protocol that have the capacity to affect aggregate price stickiness. To see this, we employ the *Ss* model of Caballero and Engel (2007) to condense large cross-sectional information on micro price changes into a measure of price flexibility. According to our estimates, the response of aggregate inflation to a nominal shock varies substantially, increasing by about 50% between the start of the Great Recession and 2011, thus reverting and attaining its minimum in the first quarter of 2016. As a result, the pass-through of aggregate nominal shocks to inflation has decreased markedly during this period, thus reverting after the Brexit referendum. More generally, changes in price flexibility tend to occur in correspondence with 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: first, inflation has been outside the 1%-3% interval for a total of 22 out of 40 quarters, while this has only occurred for 11 quarters in the previous decade; second, over the same period, inflation has shot above and below the target, reaching both its maximum (+4.8%) and minimum (-0.1%) in the overall sample. In light of this, time variation in price flexibility may help us understand why hitting the inflation target may have proven to be rather arduous over the last decade, with relatively high flexibility exacerbating the impact of inflationary shocks (e.g., movements in the exchange rate and in commodity prices) during and straight after the recession, thus reaching its minimum in correspondence with inflation hitting its historical low in 2015.

Time variation in price flexibility is extremely important to understand inflation dynamics. The half-life of the rate of inflation is twice as large in periods of relatively low flexibility, along with appearing remarkably close to the one observed in a 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 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 22% of the variability in the forecast error, at

a four-quarter horizon. This reflects the fact that forecasters fail to incorporate the faster pass-through of inflationary shocks in periods of relatively high flexibility.

A question that is typically posed in connection with the analysis of micro price data is the following: to what extent does price setting behave in accordance with time-dependent models, whereby the timing of all price changes is predetermined, as opposed to state-dependent protocols, according to which the timing of price changes can itself respond to shocks? We revisit this question through the lens of time variation in the distribution of price changes. To this end, we decompose the time series of price flexibility into predetermined price adjustments—the so-called *intensive margin*—and adjustments triggered or canceled by the shock—the *extensive margin*.³ The latter appears quite important, 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 as if price adjustments only took place over the intensive margin—is particularly large.

To refine our analysis, we examine adjustments along the extensive margin in the occurrence of changes in the value-added tax (VAT). As UK posted prices include the VAT, a key advantage of examining the transmission of large first-moment shock of this type is that they are particularly suitable to understand whether price setting works in line with the predictions of menu cost models (Karadi and Reiff, 2014). In line with Gagnon et al. (2013), massive repricing occurring in the face of a VAT shock does not emerge as a mere shift of the distribution of price gaps, but reflects a major reallocation of probability mass over the price-gap support. Most importantly, many firms seize the opportunity to adjust their prices by more than the VAT change, which implies that inflationary/deflationary pressures from other sources are released in the process. In fact, movements in the hazard function, as compared with those in the distribution of price gaps, have most of the impact on adjustments along the extensive margin and account for the bulk of the adjustment in aggregate inflation. This fact, which has not been highlighted before, implies that price-setting units’ incentives to adjust prices may vary markedly in the face of large first-moment shocks, even if the latter are largely foreseeable, as in the case of a VAT change. Acknowledging this property may be an important avenue to inform the design of structural models.

Related literature Our work relates to a number of studies that have examined the connection between micro price changes and aggregate inflation.⁴ The paper that connects most closely to our analysis is that of Berger and Vavra (2017), who report that price flexibility is time-varying. We build on this, and stress the importance of time variation in price flexibility to improve our understanding of inflation dynamics. In line with what expected on theoretical grounds, we document that inflation is more persistent and less volatile in periods of relatively

³Adjustments occurring over the intensive margin characterize both time- and state-dependent models. The extensive margin, instead, is a defining feature of state-dependent models.

⁴See, among others, Bilal and Klenow (2004), Dotsey and King (2005), Alvarez et al. (2006), Gertler and Leahy (2008), Klenow and Kryvtsov (2008), Nakamura and Steinsson (2008), Gagnon (2009), Costain and Nakov (2011), Midrigan (2011), Nakamura et al. (2011), Alvarez and Lippi (2014), Karadi and Reiff (2014), Berardi et al. (2015), Alvarez et al. (2016), Nakamura et al. (2018).

low price flexibility, and show that neglecting this fact can lead to a large prediction bias. Moreover, we employ the empirical framework of Caballero and Engel (2007) to build counterfactual exercises that stress the prominence of state-dependent price setting, as well as the distinctive role of the adjustment hazard in the occurrence of large first-moment shocks. Using VAT changes as an identification device allows us to highlight the importance of endogenous movements in firms’ incentives to price adjustment, in line with Hobijn et al. (2006). Our work also relates to a number of papers that devise and estimate specific structural models that connect movements in the distribution of price changes to price flexibility (see, e.g., Midrigan, 2011, Alvarez et al., 2016 and Vavra, 2014, among others). As discussed by Berger and Vavra (2017), an empirical limitation of these models is to rely on specific shocks to the price-setting units, while our approach is more agnostic. This represents a strategic advantage, and more so in the analysis of UK microdata, where the 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 changes in the endogenous incentives of firms to adjust their prices. Finally, we connect 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, using VAT changes to devise an event study. Compared with this paper, we employ the generalized Ss model to provide quantitative statements about the importance of both types of adjustment—examining the behavior of the distribution of price gaps in separation from that of the hazard function—and highlighting important asymmetries over the two 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 investigate the key characteristics of the frequency of price setting and the hazard functions implied by the microdata from the Office for National Statistics (ONS), while Dixon et al. (2014) have focused on the impact of the Great Recession on price setting. As compared with these papers, we place particular emphasis on state dependence in inflation dynamics, 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 Kryvtsov, 2018 and references therein). Specifically, we highlight the versatility of the empirical approach of Caballero and Engel (2007), and show how this can be used to map price flexibility into changes in inflation persistence and volatility. Employing UK data, Chu et al. (2018) emphasize that information on the distribution of price changes can be exploited to forecast inflation. Our results are in line with this finding. In fact, we show that accounting for changes in the degree of price flexibility—which condenses valuable information from key micro price statistics—improves our capacity to predict inflation.

Structure The rest of the paper is organized as follows. Section 2 discusses the key characteristics of the ONS microdata on consumer prices. Section 3 reviews the menu cost model that frames our empirical analysis. Section 4 reports the generalized Ss model and takes it to

the data. Section 5 assesses time variation in price flexibility and discusses the relative contribution of adjustments along the intensive and the extensive margin. Section 6 discusses the implications of state dependence in price flexibility for inflation dynamics. Section 7 concludes.

2 Microdata on consumer prices

We use ONS microdata underpinning 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 the identification of some price setters, or because of the need to frequently adjust for quality changes.⁵ Price quotes are recorded on or around the second or third Tuesday of the month, with the exact date being kept secret to avoid abnormal prices that, among other things, may be due to the collection of prices during bank-holiday weeks, or to price manipulations by service providers and retailers. Furthermore, to make sure the collected price quotes are valid prices, the ONS has set various checks in place, both at the collection point and at later stages in the process. As a preliminary step in handling the dataset, we only employ price quotes that have been marked as being validated by the system or accepted by the ONS. Thus, any price quote that has been marked as missing, non-comparable, or temporarily out of stock is excluded from the sample. We refer to the remaining subset of prices—which make for approximately 60% of those included in the CPI—as *Classification Of Individual CO*nsumption by Purpose (COICOP) price quotes.

Each price quote is classified by region, location, outlet and item. The region refers to the geographical entity within the UK from which a given price quote is recorded. The location is intended as a shopping district within a given region: on price-collection days, 146 different locations are visited.⁶ For a given location, the shop code is a unique but anonymized *id* associated with the outlet from which the quote is recorded. In turn, each shop is classified according to whether it is independent (i.e., part of a group comprising less than 10 outlets at the national level) or part of a chain (i.e., more than 10 outlets). Due to a confidentiality agreement between the ONS and the individual shops, for each price quote only the region, outlet and item classifications are published. In light of this, some of the price quotes may not be uniquely identified. This is typically the case when the ONS samples the same item, in 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’

⁵This is typically the case for personal computers, whose frequent model upgrades impose the use of hedonic regressions, so as to enhance comparisons across time.

⁶Until 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				
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 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 categories, 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 are obtained by accounting for CPI, item-specific, stratum and shop (i.e., elementary aggregate) weights. Sample period: 1996:M2-2017:M8.

(*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 type of good were recorded at £25 and £26, respectively. The price quotes are so close that telling the two price trajectories apart may be challenging. To make sure that price trajectories can be uniquely identified, we look at ‘base prices’, which are intended as the January’s price for each of the items under scrutiny.⁷ Given this information, the price trajectories can be identified. Even after conditioning on base prices, though, a small portion of price trajectories are still not uniquely identified (about 0.1%, on average): we opt for discarding them. In Table 1 the column labeled ‘History’ refers to the price quotes with an identifiable history that spans at least two consecutive periods. Following the criteria outlined above, we drop about 12,000 quotes per month.^{8,9}

⁷The base price is typically relied upon to normalize price quotes and calculate price indices, or to adjust for changes in the quality and/or quantity of a given good.

⁸Due 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 to be due to measurement errors. These take place rarely (< 0.01%). Appendix A provides additional details on the construction of the dataset.

⁹The total number of available price quotes denotes a weak downward trend. However, it is important to stress that the composition in terms of categories accounted for by Table 1 is roughly stable over time. This implies the presence of no particular trends in the behavior of product substitutions and sales.

To aggregate the individual price quotes into a single price, we also make use of the following weights produced by the ONS:¹⁰ the *shop* weights, which are employed to account for the fact that a single item’s price is the same in different shops of the same chain (e.g., a pint of milk at a Tesco store);¹¹ the *stratification* weights, which reflect the fact that purchasing patterns may differ markedly by region or type of outlet;¹² finally, the *item* and *COICOP* weights reflect consumers’ expenditure shares in the national accounts.

2.1 Variable definition

After deriving our price quotes in line with the criteria set out above, it is important to make a distinction between regular and temporary price changes such as sales, which tend to behave significantly differently from that of regular prices (see Eichenbaum et al., 2011 and Kehoe and Midrigan, 2015). To this end, we first exclude all the price quotes to which the ONS attaches a sales indicator.¹³ As a second step, we implement a symmetric V-shaped filter, as defined by Nakamura and Steinsson (2010b), for the remaining price quotes. According to the filter, the sale price of item i at time t , $P_{i,t}^s$, 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.¹⁴

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

¹⁰See Chapter 7 of the ONS CPI Manual (ONS, 2014).

¹¹In this case the ONS enters a single price for a pint of milk, but the weight attached to this is large, so as to reflect that all Tesco stores within the region have posted the same price.

¹²In 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).

¹³For a price to be marked as being associated with a sale, the ONS requires the latter to be available to all potential costumers—so as to exclude quantity discounts and membership deals—and that it only entails a temporary or an end-of-season price reduction. This definition excludes clearance sales of products that have reached the end of their life cycle.

¹⁴See also Nakamura and Steinsson (2008) and Vavra (2014). As an alternative approach, in place of the price associated with a sale, Klenow and Kryvtsov (2008) report the last regular price, until a new regular price is observed. Our analysis is robust to pursuing this approach.

Table 1 shows that, after these preliminary 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 $\Delta p_{i,t} = \log(P_{i,t}/P_{i,t-1})$.¹⁵

2.2 Data facts

This section unveils key facts about the behavior of the ONS microdata.¹⁶ The top panels of Figure 1 report the frequency of adjustment and the average magnitude of price changes: decomposing inflation as the product of these statistics carries important information about 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 high degree of positive co-movement 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, this tracks very closely the contraction in the rate of inflation that starts in 2012—moving well below its sample average up to that point—while only displaying a weak reversion towards the end of 2015.¹⁷

In the bottom panels of the figure, both statistics are split between positive and negative price changes. Throughout the entire sample, the frequency of positive price changes is greater than that associated with negative adjustments, while the opposite broadly holds true when comparing 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 depicted 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.¹⁸ Both facts point to a certain degree of asymmetry in price adjustment.

Figure 2 plots measures of dispersion of the distribution of (non-zero) price changes. Both the cross-sectional standard deviation and the interdecile range display a very large increase in the aftermath of the Great Recession. In fact, as displayed by the right panel of the figure, dispersion increases on either side of the median, though negative price changes become more dispersed than positive price changes. In light of this, it should be stressed that the fall in CPI inflation occurred in the post-2010 sample is to a large extent a manifestation of the trend in the

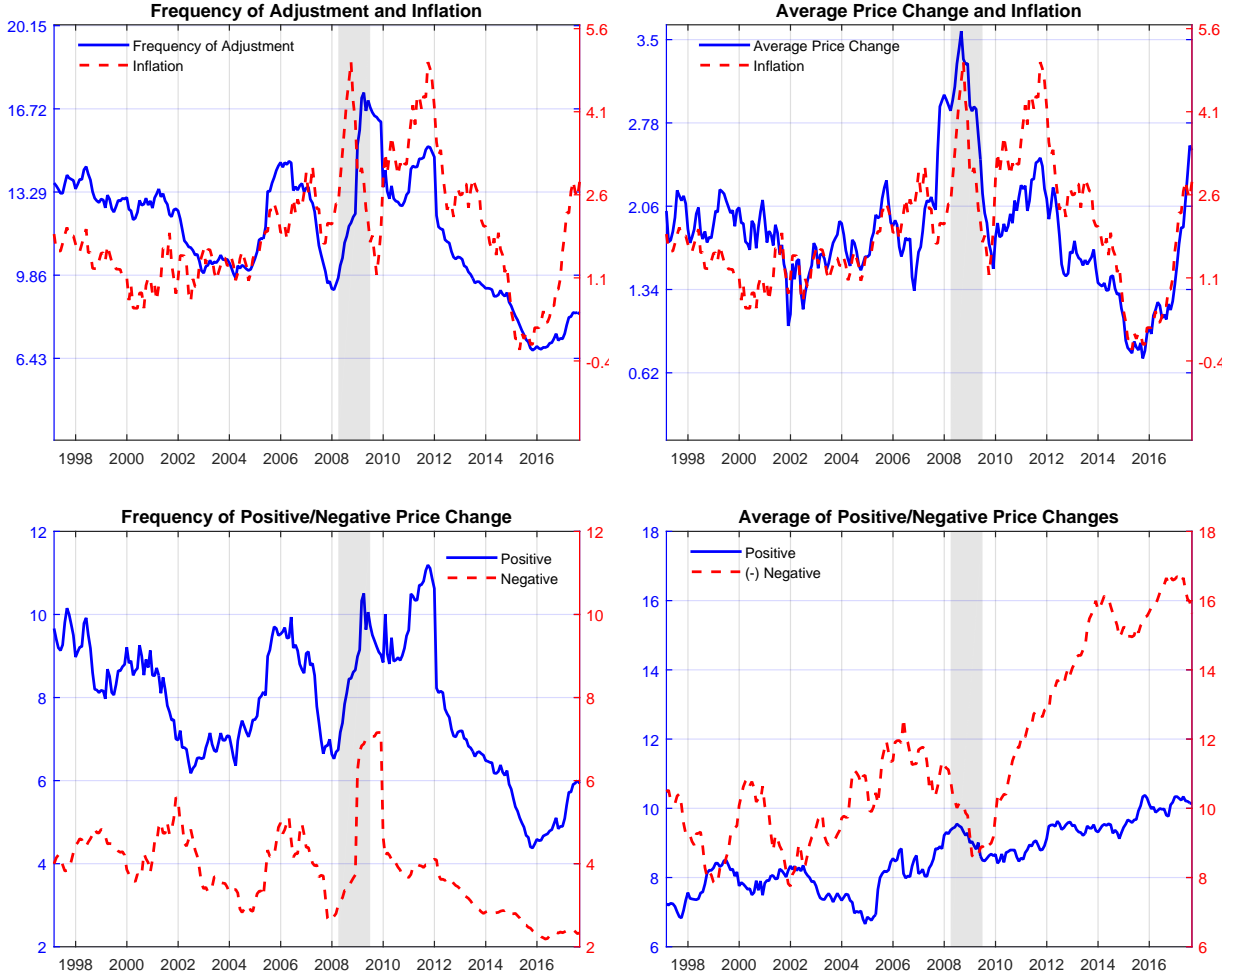
¹⁵We also compute price changes as $\Delta p_{i,t} = 2 \frac{P_{i,t} - P_{i,t-1}}{P_{i,t} + P_{i,t-1}}$. This definition has the advantage of being bounded and less sensitive to outliers. The results—virtually unchanged with respect to the ones we report—are available from the authors, upon request.

¹⁶Throughout the paper all the statistics derived from microdata on prices are reported as a 12-month moving average, so as to get rid of seasonality.

¹⁷The average frequency of price adjustment prior to the fall is broadly in line with the figures reported by previous studies on UK micro price data. To see this, one has to account for the fact that we exclude both utility prices (COICOP 4) and sales. Bunn and Ellis (2012), instead, consider both categories, while Dixon and LeBihan (2012) and Dixon and Tian (2017) include sales, but exclude utility prices.

¹⁸Figure B.1 in Appendix B shows that composition effects have no role in generating the facts presented in this subsection. To this end, we compare the moments of the distribution of price changes with their counterparts obtained by averaging the corresponding moments of the price quotes, for each of the 25 COICOP group categories.

Figure 1: FREQUENCY OF ADJUSTMENT AND AVERAGE PRICE CHANGES



Notes: The frequency of adjustment, fr_t , is computed as $\sum_i \omega_{i,t} \mathbb{1}_{\{\Delta p_{i,t} \neq 0\}}$, where $\omega_{i,t}$ denotes the CPI weight associated to good i at time t , and $\mathbb{1}_{\{\Delta p_{i,t} \neq 0\}} = 1$ if $\Delta p_{i,t} \neq 0$ and zero otherwise. The average price change, instead, is computed as $fr_t^{-1} \sum_i \omega_{i,t} \mathbb{1}_{\{\Delta p_{i,t} \neq 0\}} \Delta p_{i,t}$. The positive and negative counterparts of these statistics are obtained by conditioning them on positive and negative price changes, respectively. All series are in percentage terms. In the bottom-right panel we report the mirror image of the average of negative price changes. The inflation rate graphed in the upper panel of the figure is the official CPI inflation rate published by the ONS. The shaded vertical band indicates the duration of the Great Recession.


dispersion of negative price changes—relative to that of positive ones—rather than reflecting a mere shift in the mode of the density.

Table 2 reports the correlation between moments of the distribution of price changes, CPI inflation and a business cycle indicator.¹⁹ To set aside potential spurious correlation emanating from the low-frequency behavior of the series under examination, we detrend all of them, aside of the inflation rate.²⁰ Turning our attention to the frequency of adjustment and the

¹⁹Appendix C contains more details on the derivation of the monthly coincident indicator of economic activity.

²⁰Moreover, 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 macroeconomic turmoil in that period. In light of this, it is worth stressing that, when interpreting the cyclical properties of the data in the two subsamples, the correlations are likely to be captured by the behavior of the series in periods of relatively stronger/weaker expansion, rather than by different cyclical phases.

Figure 2: DISPERSION 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 shaded vertical band indicates the duration of the Great Recession.

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.

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. Analogous properties hold in the post-recession time span, while in the previous decade both statistics behave 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 points to an invariantly negative correlation.

Skewness displays a marked countercyclical behavior. Interestingly, its correlation with the rate of inflation switches sign over the two subsamples. In this respect, Luo and Villar (2017a) show how different models of price setting may account for different signs of the correlation between inflation and the skewness of price changes. According to their insights, menu cost models—which feature the price change distribution becoming less skewed as inflation rises—could well rationalize our data in the second part of the sample. On the other hand, the Calvo model—which features a positive correlation—could better account for the first subsample. In the remainder of the analysis we will show how this type of characterization is also supported by the behavior of the extensive margin of price adjustment—a hallmark of menu cost models—which assumes a prominent role in the aftermath of the Great Recession.

To summarize, we observe a lot of time variation in the distribution of price changes. Most notably, while dispersion has been procyclical throughout the entire sample, the frequency of adjustment has gone from being procyclical before the Great Recession, to behaving countercyclically during the last decade. As a result, the pairwise correlation between these statistics has turned deeply negative in the post-recession sample. This picture stands in contrast with the analysis on US microdata by Vavra (2014), who reports that the cross-sectional standard deviation of price changes is strongly countercyclical and comoves positively 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 a key role in accounting for positive comovement between the dispersion of price changes and the frequency of adjustment. In the next section, we use the same framework to show how changes in the incentives firms face when deciding to change prices can provide us with a rationale for the emergence of negative comovement.

3 Analytical framework

To frame the 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). For the sake of our analysis, we will use this model as a prism through which interpreting distinctive time-varying phenomena behind price

setting.

Firms face a dynamic control problem where x —the deviation of the current price from the optimal price—is a state variable. A wedge between the state variable and zero entails an out-of-equilibrium cost αx^2 , where α can be inversely related to market power. When not adjusting, x follows a Brownian motion $dx = \phi dW$, where W is the increment to the Wiener process. It is possible to change the value of x by applying an instantly effective control at a lump-sum cost λ . From this environment, a simple *Ss* rule emerges, according to which the optimal policy is ‘do not adjust’ when $|x| < \sigma$ and ‘adjust to zero’ when $|x| \geq \sigma$, where $\sigma = (6\lambda\phi^2/\alpha)^{1/4}$ denotes the standard deviation of price changes. Moreover, $fr = (\alpha/6\lambda)^{1/4} \phi$ is the frequency of adjustment.²¹

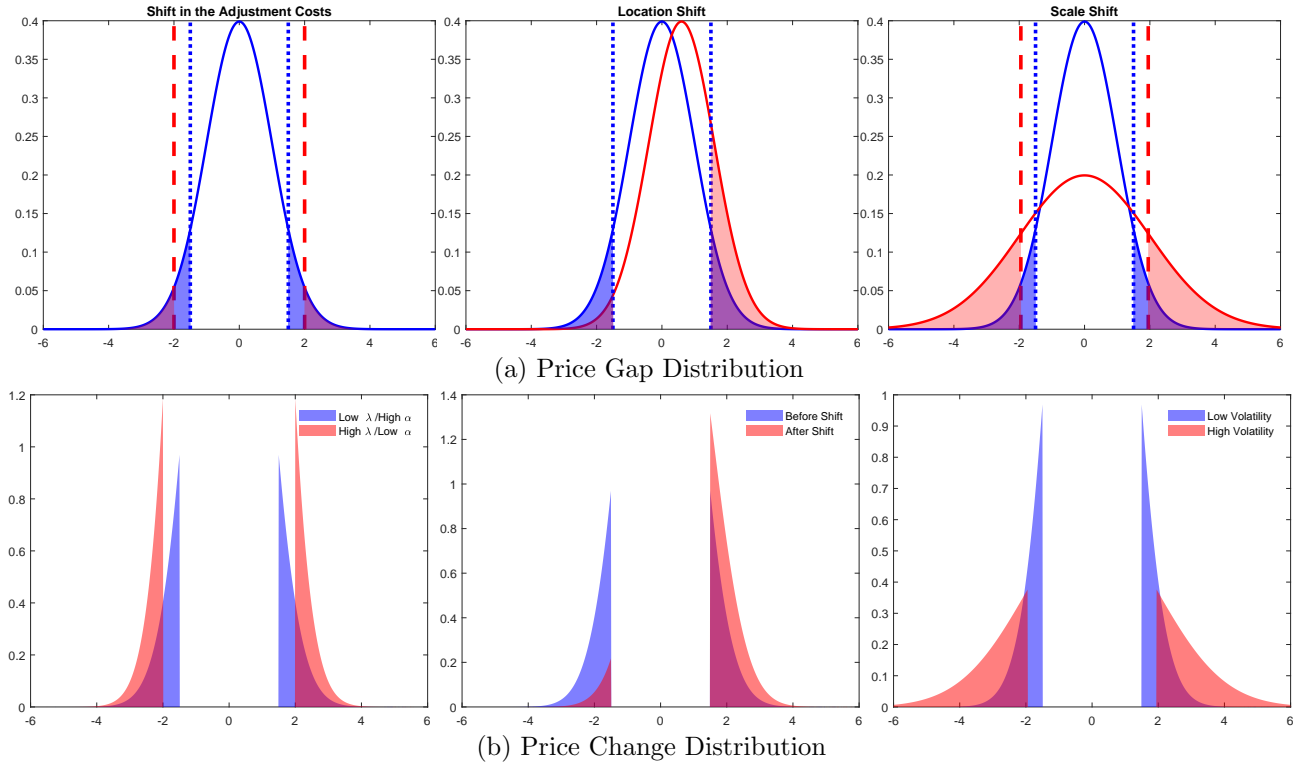
To provide an overview of different determinants of the distribution of price gaps and the associated distribution of price changes, Figure 3 considers three possible scenarios: i) a positive shift in the cost of adjustment λ (or, equivalently, a negative shift in α) that affects the inaction region, 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 widens the inaction region, translating into a reduction 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 (iii), induces increased dispersion in the price gap distribution and an expansion in the inaction region. As a result, both fr and σ increase.

Vavra (2014) points to second-moment shocks as potential drivers of the positive comovement between the frequency of adjustment and the price-change dispersion in U.S. CPI data. However, in the 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 divergent behaviors. In light of this, second-moment shocks might provide a good account of what has happened only up to the Great Recession. Moreover, 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.

When focusing on the post-recession experience, only an increase in the fixed cost of adjustment and/or a drop in the cost of deviating from the optimal price may account for negative

²¹For 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.-

comovement between the frequency of adjustment and the dispersion of price changes, conditional on the resulting expansion of the inaction region dominating the effects of positive shifts in the dispersion of price gaps. To dig deeper into this specific aspect of the data and its repercussions for aggregate price stickiness, the next section introduces an accounting framework that is particularly useful at quantifying the link between changes in the timing of individual price adjustments, price flexibility and inflation dynamics.

4 The generalized S_s model

The generalized S_s model developed by Caballero and Engel (2007) has two clear advantages that make it particularly indicated to discipline our data. First, it is consistent with lumpy and infrequent price adjustments—which are typically seen as distinctive traits of price setting—along with encompassing several pricing protocols.²² In this respect, Berger and Vavra (2017) show that this empirical setting provides a good fit to the data generated by different structural models (e.g., Golosov and Lucas, 2007 and Nakamura and Steinsson, 2010a). Second, as we

²²To mention two extreme examples, the generalized S_s model can account for both price setting à la Calvo (1983)—where firms are selected to adjust prices at random and price flexibility is fully determined by the frequency of adjustment—as well as for schemes à la Caplin and Spulber (1987)—where adjusting firms change prices by such large amounts that the aggregate price is fully flexible, regardless of the frequency of adjustment.

allow for time variation in the determinants of price adjustment, we can estimate the model over each cross section of price microdata, matching different price-setting statistics. 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, the log of firm i 's actual price may deviate from the log of the target or reset price, which is denoted by p_{it}^* . Thus, we define the price gap as $x_{it} \equiv p_{it-1} - p_{it}^*$, implying that a positive (negative) price gap is associated with a falling (increasing) price when the adjustment is actually made. In a simple Ss model as the one detailed in the previous section, a price is adjusted when the associated price gap is large enough, and $p_{it} = p_{it}^*$ after the adjustment has taken place. Assuming l_{it} periods since the last price change, the adjustment reflects the cumulated shocks: $\Delta p_{it} = \sum_{j=0}^{l_{it}} \Delta p_{it-j}^*$, with $\Delta p_{it}^* = \mu_t + v_{it}$, where μ_t is a shock to nominal demand and v_{it} is an idiosyncratic shock.

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

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

$$f_t(x) = \begin{cases} \frac{\delta(\varrho_t, \nu_t)^{1/\nu_t}}{\psi_t \Gamma(1+1/\nu_t)} \exp \left[- \frac{\delta(\varrho_t, \nu_t)}{\varrho_t^{\nu_t}} \left| \frac{x-\theta_t}{\psi_t} \right|^{\nu_t} \right] & \text{if } x \leq \theta_t \\ \frac{\delta(\varrho_t, \nu_t)^{1/\nu_t}}{\psi_t \Gamma(1+1/\nu_t)} \exp \left[- \frac{\delta(\varrho_t, \nu_t)}{(1-\varrho_t)^{\nu_t}} \left| \frac{x-\theta_t}{\psi_t} \right|^{\nu_t} \right] & \text{if } x > \theta_t \end{cases}, \quad (2)$$

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

($\nu_t \rightarrow \infty$). 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:

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

where $\mathbb{1}_{\{z\}}$ is an indicator function taking value 1 when condition z is verified, and zero otherwise. This parsimonious specification nests the Calvo pricing protocol for $b_t = c_t = 0$, while potentially allowing for asymmetric costs of adjustment.²³

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.²⁴ 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.²⁵ Appendix F reports a number of exercises that highlight how our indirect inference approach is able to identify the shape of the price gap distribution and the hazard function. To this end, we simulate price-change data from the model, under different parameterizations, and then contrast the true price gap distribution and the hazard function to their estimated counterparts. Overall, discrepancy is minimal.²⁶

4.2 Making sense of changing comovement between the frequency and dispersion of price changes

The first two panels of Figure 4 report the estimated scale parameter of $f_t(x)$ and the inaction region associated with two hazard probabilities (namely, 5% and 7%). Both statistics are time-varying, and increase markedly in the second decade of the sample. According to our comparative statics analysis in Section 3, a prolonged decline in the frequency of adjustment, coupled with a surge in its dispersion, may be rationalized by an expansion in the inaction


²³We 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.

²⁴We match quantilic moments, as the 3rd and 4th moments of the cross-sectional distribution are quite sensitive to outliers. Figure E.1 graphs the dynamics of both $f_t(x)$ and $\Lambda_t(x)$, while Figures E.2 and E.3 report the estimated parameters. Finally, Figure E.4 reports the fit of selected data moments, and shows that the empirical model is able to summarize the main stylized facts in the data.

²⁵Altonji 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 to avoid ending up in local minima (see, e.g., Dorsey and Mayer, 1995).

²⁶In a similar context, Berger and Vavra (2017) report additional exercises in support of the indirect inference approach.

Figure 4: 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.

region—as triggered 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.²⁷ To verify this is indeed the case, the last panel of Figure 4 reports the share of prices in the inaction region, defined as the proportion of prices whose $\Lambda_t(x)$ 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 $f_t(x)$. Note also that greater inactivity appears more evident in correspondence with positive price gaps, as compared with the negative ones, thus implying an increased degree of downward price stickiness.


On a more general note, 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. The next section digs into this, through the lens of the accounting framework of Caballero and Engel (2007).

5 On the importance of state dependence in price adjustment

The estimation of the generalized S_s model highlights the importance of tracking changes in the distribution of price gaps and the hazard function. To dig deeper into the connection between individual price adjustment and the response of aggregate inflation to nominal demand, Caballero and Engel (2007) show that, within their accounting framework, one can derive a measure of aggregate price flexibility that captures the impact response of realized inflation to

²⁷Appendix D frames this situation within the menu cost model of Section 3.

Figure 5: PRICE FLEXIBILITY AND DIFFERENT MARGINS OF PRICE ADJUSTMENT



Notes: The left panel reports the estimated index of price flexibility, which is decomposed in the right panel between the intensive and the extensive margin of price adjustment. The shaded vertical band indicates the duration of the Great Recession.

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.²⁸

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

Figure 5 reports the estimated index of price flexibility and its decomposition into the intensive and the extensive margin of price adjustment. Price flexibility displays sizable variation

²⁸In this respect, Alvarez et al. (2016) show that the steady-state ratio of kurtosis to the frequency of adjustment is a sufficient statistic of 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.

over time, and 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 contract, though the fall in the former is much more abrupt, in line with the sustained drop 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. It should be stressed that, over most of the decline, the extensive margin 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 and 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 result, 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 E.3 in Appendix E). Alternatively, the same point can be made by approximating the flexibility index as $F_t \cong Int_t + 2[Int_t - \Lambda_t(0)]$:²⁹ from this expression it is clear how a downward shift in a_t —which is equivalent to lowering $\Lambda_t(0)$ —translates into an increase in the importance of the extensive margin relative to the intensive one, *ceteris paribus*.

From a cyclical 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 inverting its cyclicity during the last decade (see Table K.1 in Appendix K). 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. On a more general note, it is worth emphasizing how changes in the correlation structure over the two subsamples are consistent with a shift from an environment where the intensive margin dominates the extensive one, to one where the extensive margin assumes a prominent role and inflation volatility is particularly marked (see Figure 5).

Having established that the extensive margin of price adjustment may be quantitatively important to price flexibility, we turn to examining its role for inflation dynamics. Figure 6 reports both the 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 by assuming that the probability of price adjustment, while exogenous to the firm, can vary with the state of the economy (i.e., $\pi_t^{Calvo} = -fr_t^{Calvo} \int x f_t(x) dx$, where $fr_t^{Calvo} = \int \Lambda_t(x) f_t(x) dx$).

²⁹For a formal proof, please refer to Caballero and Engel (2007).

Figure 6: A CALVO COUNTERFACTUAL



Notes: The rate of inflation is obtained from our sample of ONS price quotes (blue-solid line), while its counterfactual (red-dashed line) is 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.

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, when the difference between the latter and its ‘Calvo counterfactual’ is sizable.

5.1 On the impact of large first-moment shocks: a VAT event study

Examining the relative contribution of price adjustment along the extensive margin is of key 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 this end, we can usefully exploit episodes of major repricing activity triggered by changes in the VAT. These are typically useful for two reasons. First, they are relatively simple to study, because their timing and size are directly observable. Second, changes in the VAT are particularly suitable to understand whether price setting works in line with the predictions of menu cost models (Karadi and Reiff, 2014): as in the UK posted prices include the VAT, price-setting units need to post new prices—and, thus, bear a menu cost—if they choose

to incorporate the tax change into their prices. Furthermore, large first-moment shock are particularly well-suited to disentangle movements in the price gap distribution from those in the adjustment hazard. Otherwise, as shown in Section 3, both $f_t(x)$ and $\Lambda_t(x)$ would vary—albeit to different extents—in the face of second-moment shocks (see also Appendix G).

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.³⁰

Looking at the inflation rate in the month corresponding to a VAT change, we note 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 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. Nonetheless, Table 3 shows that the intensive margin is much higher in correspondence with a VAT hike, as compared with a negative VAT change. In this respect, our evidence is consistent with Karadi and Reiff (2014).

To dig deeper into the role of state-dependent pricing, Table 3 also reports some statistics in coincidence with the three VAT changes, as well as two counterfactual scenarios.³¹ In the *no $\Lambda(x)$ change* scenario, we keep the hazard function as that computed in the counterfactual exercise of Figure 7, but let the price gap distribution vary as a result of the VAT change. Thus, we abstract from any amplification that could be induced by state-dependent pricing through upward shifts of the hazard function. The *no VAT change* scenario, instead, considers a situation in which neither the price gap distribution nor the hazard function are affected by the VAT change.³²

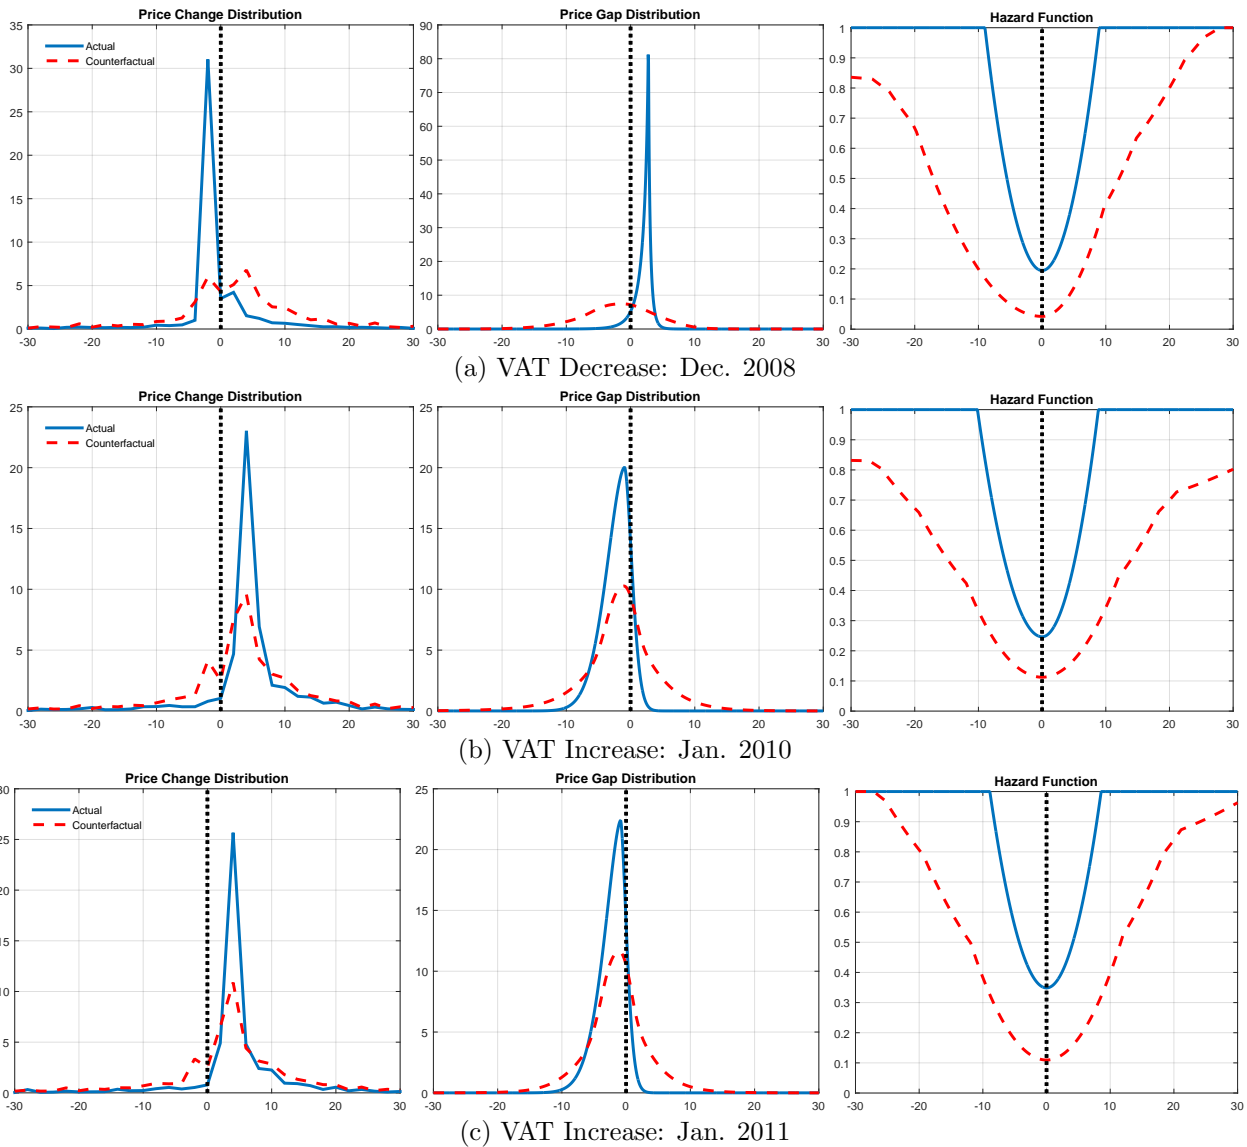
From the comparison between inflation in the occurrence of a VAT change and its counterfactuals, two observations are worth emphasizing. First, state-dependent pricing accounts for most of the change in the rate of inflation in the presence of a VAT change. Otherwise, inflation would have been not very different from its counterfactual in the *no VAT change* scenario. This is particularly evident when the VAT is raised. Second, movements in the hazard function, as compared with those in the distribution of price gaps, have most of the impact on adjustments along the extensive margin. This implies that price-setting units' incentives

³⁰January 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.

³¹More details on the computation of two alternative scenarios are provided in Appendix H.

³²This amounts to keeping both the price gap distribution and the hazard function to their counterfactuals in Figure 7.

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, one up to 17.5% on January 1, 2010, and further up to 20% on January 4, 2011.

to adjust prices—as embodied by their adjustment costs structure—may display substantial variation in the face of large first-moment shocks, even if the latter are largely foreseeable, as in the case of VAT changes. In this respect, the price-setting behavior we portray bears close resemblance to price adjustment as described by [Hobijn et al. \(2006\)](#) in the occurrence of the Euro changeover. As in this case, the VAT-adjustment decision could result from the interplay between a *churning effect*—whereby price-setting units concentrate otherwise staggered price increases around the VAT change—and a *horizon effect*, which depends on the fact that prices adjusted before the VAT change do not reflect the marginal cost increases expected to occur

Table 3: 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
No $\Lambda(x)$ change	-1.604	0.101	0.060	0.041	0.055	0.005	0.040	0.001
No VAT change	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
No $\Lambda(x)$ change	4.580	0.181	0.135	0.045	0.008	0.127	0.001	0.045
No VAT change	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
No $\Lambda(x)$ change	4.708	0.190	0.136	0.053	0.006	0.130	0.001	0.053
No VAT change	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), all 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). The the extensive margin associated with positive and negative price gaps are computed by decomposing the extensive margin as $Ext_t = \int_{-\infty}^0 x \Lambda'_t(x) f_t(x) dx + \int_0^{\infty} x \Lambda'_t(x) f_t(x) dx$, where Ext_t^- (Ext_t^+) is the first (second) term on the right side of the equality. For every VAT change episode, we contrast the actual numbers with two alternative scenarios. In the *no $\Lambda(x)$ change* scenario, the VAT change only impacts on the distribution of price gaps, while the hazard function is kept at the counterfactual (see Figure 8). The *no VAT change* scenario, instead, considers an alternative case in which neither the hazard function nor the price gap distribution change.

afterwards.

When comparing the two margins of adjustment, the intensive one is typically much larger than its counterparts in the alternative scenarios—indicating that upward shifts in $\Lambda_t(0)$ are the most prominent feature in the occurrence of a VAT change—while movements along the extensive margin appear less dramatic. However, such a conclusion is not warranted after conditioning both margins to positive and negative price changes. In this case, substantial variation also takes place along the extensive margin coherent with the sign of the underlying price change. For instance, in the occurrence of the VAT drop, Ext^+ is more than twice as big as its counterfactuals. The same order of magnitude can be appreciated when making the same comparison for the 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 three episodes 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 VAT change. On one

hand, this necessarily implies that the extensive margin associated with price gaps coherent with the sign of the adjustment is large. On the other hand, the extensive margin associated with price gaps of the opposite sign is very low, in light of the hazard function being weighed by a very small probability mass, after a shift in the distribution of price gaps has occurred.

6 Price flexibility and inflation dynamics

The estimation of Caballero and Engel’s model of lumpy price adjustment shows that the pass-through of nominal shocks to inflation is highly variable. We also report that—while not hinging on a specific margin of adjustment—flexibility is higher in connection with positive price changes. These properties bear major implications for evaluating the transmission of shocks to nominal demand. Using the estimated Ss model, we are able to examine the response of inflation to an aggregate nominal shock in two different periods, characterized by relatively high and low price flexibility, respectively.³³ As expected, Figure 8 shows that inflation is more responsive and less persistent in a period of relatively high price flexibility. In light of this simple exercise, one would expect price flexibility to contain valuable information for the analysis of inflation dynamics. This section aims at substantiating this claim.

We seek to examine how inflation behaves in periods of relatively high and low flexibility, and to contrast its inherent non-linear dynamics to its behavior in a linear setting. 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).³⁴ Estimating a traditional ARMA(p,q) for each regime separately entails a certain disadvantage in that we may end up with relatively few observations in a given regime, which typically renders the estimates unstable and imprecise. By contrast, we can effectively rely upon more information by exploiting variation in the probability of being in a particular regime, so that estimation and inference for each regime are based on a larger set of observations (Auerbach and Gorodnichenko, 2012).³⁵

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


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

³³As we only identify the price gap distribution at each point in time, we are not able to disentangle the contribution of the aggregate shock from that of idiosyncratic shocks. Therefore, for purely illustrative purposes, we choose an autoregressive specification for the first-moment shock. More details are available in Appendix I.

³⁴In this respect, the STARMA(p,q) model also generalizes the threshold ARMA(p,q) model (DeGooijer, 2017).

³⁵Estimating the properties of a given regime by relying on the dynamics of inflation in a different regime would bias our results towards not finding any evidence of non-linearity. In light of this, the asymmetries we will be reporting in the remainder of this section acquire even more statistical relevance.

Figure 8: IMPULSE RESPONSES FROM THE S_s MODEL



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

with $\varepsilon_t^i \sim N(0, \sigma_i^2)$ for $i = \{L, H\}$. Moreover, we set $G(\tilde{\mathcal{F}}, \gamma) = (1 + e^{-\gamma\tilde{\mathcal{F}}})^{-1}$, where $\tilde{\mathcal{F}}$ denotes the normalized flexibility index and γ is the speed of transition across regimes.³⁶ We allow for different degrees of inflation persistence across the two regimes, as captured by the regime-specific autoregressive and moving average coefficients, as well as for different volatilities of the innovations in either regime. The likelihood of the model can be easily computed by recasting the system in state space (see, e.g., Harvey, 1990). We use Monte Carlo Markov-chain methods developed in Chernozhukov and Hong (2003) for estimation and inference. The parameter estimates, as well as their standard errors, are directly computed from the generated chains.³⁷

As we focus on the post-1996 sample, we estimate the model by imposing that, in both regimes, the long-run inflation forecast is 2%, consistent with the 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% of the observations are classified to be in the high-flexibility (low-flexibility) regime, where this is defined by $G(\tilde{\mathcal{F}}_{t-1}; \gamma) > 0.8$ ($G(\tilde{\mathcal{F}}_{t-1}; \gamma) < 0.2$).³⁸ Thus, based on the Akaike criterion, we choose $p = 1$ and $q = 7$.³⁹


³⁶We 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.

³⁷See Appendix J for further details.

³⁸Figure K.1 in Appendix K 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 .

³⁹Note that the modified AIC information criterion indicates a STARMA(1,3). Figures K.2 and K.3 in Appendix K report the results for this alternative setting. Our key insights are not affected by the exact

Figure 9: 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). In each of the two charts the vertical line delineates the half-life of the shock.

Figure 9 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 price 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 and/or 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)

specification of the STARMA(p,q) model.

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 December 2003. Thus, we construct quarterly forecast errors as the difference between realized inflation and the appropriate (mean) forecast at a given horizon.⁴⁰ 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 4 provides a summary of the results from our regression exercise. The first four columns report the slope coefficients and the associated p-values at relatively low and high levels of flexibility (i.e., $G = 0.2$ vs. $G = 0.8$). We recover an inclined L-shaped relationship between the forecast errors and price flexibility, which confirms that inflation tends to be overpredicted when prices are relatively flexible. 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. In line with Figure 9, 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 22% of the variability in the absolute forecast error. The relationship is not statistically significant in periods of relatively low flexibility, whereas it is typically positive and statistically significant when flexibility is relatively high, with the slope displaying larger values at medium-term forecast horizons. The results are roughly the same, no matter which source of forecasts we consider.

The pronounced time variation in price flexibility after the Great Recession helps us to get a better understanding of the concurrent dynamics of the inflation rate. 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.⁴¹ The Bank of England has generally underestimated the speed and impact of shocks to inflation in the 2008-2011 period. In light of our evidence, this points to a potential failure in appreciating that price flexibility was itself at the historical peak, possibly as a reflection of the three VAT adjustments taking place over 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 in the summer of 2016, in coincidence with the sharp movements of the GBP in the

⁴⁰Table K.4 in Appendix K returns similar evidence when we use absolute and squared forecast errors. The results are also virtually unchanged if we use median in place of mean forecasts.

⁴¹Two 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, while this has only occurred for 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 (-0.1%) in the overall sample.

Table 4: INFLATION FORECAST ERRORS AND PRICE FLEXIBILITY

(a) BoE MPC RPIX/CPI Forecast Errors						
Horizon	Slope at $G = 0.2$		Slope at $G = 0.8$		F-stat	\tilde{R}^2
1	-0.195	[0.695]	0.797	[0.172]	0.168	2.61
2	-0.920	[0.261]	2.059	[0.031]	0.004	12.88
3	-1.341	[0.241]	2.927	[0.041]	0.000	18.33
4	-0.925	[0.563]	3.919	[0.025]	0.000	21.98
5	-0.493	[0.796]	4.067	[0.016]	0.000	22.86
6	-0.249	[0.901]	3.596	[0.033]	0.000	21.59
7	-0.275	[0.895]	3.555	[0.016]	0.000	19.96
8	-0.903	[0.621]	3.543	[0.003]	0.001	16.33
(b) Market Participants' Forecast Errors						
Horizon	Slope at $G = 0.2$		Slope at $G = 0.8$		F-stat	\tilde{R}^2
1	0.317	[0.706]	0.636	[0.305]	0.468	-0.60
2	-1.117	[0.213]	2.097	[0.030]	0.003	13.50
3	-1.567	[0.224]	2.950	[0.041]	0.000	18.69
4	-1.045	[0.569]	3.860	[0.028]	0.000	21.03
5	-0.504	[0.815]	3.866	[0.022]	0.000	21.36
6	-0.085	[0.970]	3.161	[0.055]	0.000	19.45
7	-0.005	[0.998]	2.808	[0.045]	0.002	15.74
8	-0.665	[0.745]	2.431	[0.030]	0.022	9.27

Notes: The table reports the results of a quadratic spline regression of the forecast errors $e_{t+h|t}$ (for different forecast horizons, h , measured in quarters) on a quarterly average of an indicator of the normalized price flexibility index, $G_{t-1} = G(\tilde{\mathcal{F}}_{t-1}; \gamma) = (1 + e^{-\gamma\tilde{\mathcal{F}}_{t-1}})^{-1}$, where $\tilde{\mathcal{F}}$ denotes the normalized flexibility index. The regression takes the form: $e_{t+h|t} = a_0 + a_1G_{t-1} + a_2G_{t-1}^2 + a_3\mathbb{1}_{\{G_{t-1} > 0.5\}}G_{t-1}^2$, where $\mathbb{1}_{\{G_{t-1} > 0.5\}}$ is an indicator function taking value 1 when $G_{t-1} > 0.5$ and zero otherwise. The upper panel refers to the Bank of England MPC's RPIX/CPI forecast errors, while the bottom panel considers market participants' forecast errors. In each panel, the first two 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 (i.e., $H_0 : a_1 = a_2 = a_3 = 0$). The last column reports the adjusted R-squared, denoted by \tilde{R}^2 .

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 Ss model of Caballero and Engel (2007), we are able to report that price flexibility displays pronounced time variation, especially during the last decade. 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.

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, but also assigns a prominent role to state-dependent price setting for the study of 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. LEBIHAN, AND F. LIPPI (2016): “The Real Effect of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach,” *American Economic Review*, 106, 2817–2851.
- ALVAREZ, F. AND F. LIPPI (2014): “Price Setting With Menu Cost for Multiproduct Firms,” *Econometrica*, 82, 89–135.
- ALVAREZ, 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.

- (2006): “Price stickiness in Ss models: Basic properties,” Unpublished manuscript.
- (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 (2018): “Price Selection,” Staff Working Papers 18-44, Bank of Canada.
- CHERNOZHUKOV, V. AND H. HONG (2003): “An MCMC approach to classical estimation,” *Journal of Econometrics*, 115, 293–346.
- CHU, B. M., K. P. HUYNH, D. T. JACHO-CHÁEZ, AND O. KRYVTSOV (2018): “On the Evolution of the United Kingdom Price Distributions,” *Annals of Applied Statistics*.
- COSTAIN, J. AND A. NAKOV (2011): “Distributional dynamics under smoothly state-dependent pricing,” *Journal of Monetary Economics*, 58, 646–665.
- 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. AND H. LEBIHAN (2012): “Generalised Taylor and Generalised Calvo Price and Wage Setting: Microevidence with Macro Implications,” *Economic Journal*, 122, 532–554.
- DIXON, H. D., K. B. LUNTEL, 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.
- DIXON, H. D. AND K. TIAN (2017): “What We can Learn About the Behaviour of Firms from the Average Monthly Frequency of Price-Changes: An Application to the UK CPI Data,” *Oxford Bulletin of Economics and Statistics*, 79, 907–932.
- DORSEY, R. E. AND W. J. MAYER (1995): “Genetic Algorithms for Estimation Problems with Multiple Optima, Nondifferentiability, and Other Irregular Features,” *Journal of Business & Economic Statistics*, 13, 53–66.

- DOTSEY, M. AND R. G. KING (2005): “Implications of state-dependent pricing for dynamic macroeconomic models,” *Journal of Monetary Economics*, 52, 213–242.
- EICHENBAUM, M., N. JAIMOVICH, AND S. REBELO (2011): “Reference Prices, Costs, and Nominal Rigidities,” *American Economic Review*, 101, 234–262.
- 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.
- 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.
- HOBIIJN, B., F. RAVENNA, AND A. TAMBALOTTI (2006): “Menu Costs at Work: Restaurant Prices and the Introduction of the Euro,” *The Quarterly Journal of Economics*, 121, 1103–1131.
- KARADI, P. AND A. REIFF (2014): “Menu Costs, Aggregate Fluctuations, and Large Shocks,” CEPR Discussion Papers 10138.
- 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 Cyclicity of Sales and Aggregate Price Flexibility,” Manuscript.

- LUO, S. AND D. V. VILLAR (2017a): “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.).
- (2017b): “The State-dependent Price Adjustment Hazard Function: Evidence from High Inflation Environments,” Unpublished 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.
- (2010a): “Monetary Non-neutrality in a Multisector Menu Cost Model,” *The Quarterly Journal of Economics*, 125, 961–1013.
- (2010b): “More Facts About Prices – Supplement to “Five Facts About Prices: A Reevaluation of Menu Cost Models”,” Online Supplement.
- NAKAMURA, E., J. STEINSSON, P. SUN, AND D. VILLAR (2018): “The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation,” *Quarterly Journal of Economics*, Forthcoming.
- ONS (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.
- 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.

Time-varying Price Flexibility and Inflation Dynamics

Supplementary Material

Ivan Petrella

Warwick Business School, University of Warwick and CEPR.

Email: ivan.petrella@wbs.ac.uk

Emiliano Santoro

Department of Economics, University of Copenhagen.

Email: emiliano.santoro@econ.ku.dk

Lasse de la Porte Simonsen

Department of Economics, Mathematics and Statistics, Birkbeck College, London.


Email: l.simonsen@mail.bbk.ac.uk

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. We keep COICOP 8, as the available price quotes are clustered in a small subset of items, such as Flower Delivery, Telephone for home use and Phone Accessories.¹

The left panel of Figure A.1 contrasts the weights assigned to each of the COICOP sectors to those employed to build the CPI (re-normalized to exclude COICOP 4 and 10). Overall, we observe that using the available price quotes results into relatively larger weights for COICOP 1 and 11, whereas sectors 7 and 9 are underweighed. The right panel of Figure A.1 reports the official CPI inflation together with the inflation series retrieved from all the available price quotes (labeled *COICOP*) and the inflation obtained once all filters described in Section 2 are applied (labeled *Regular*). Unfiltered data track quite closely the official numbers, whereas the regular series displays a robust correlation with the official data (roughly 0.7), and shows a positive bias. The latter mainly emerges from the exclusion of sales from the sample.

Figure A.1: REPRESENTATIVENESS

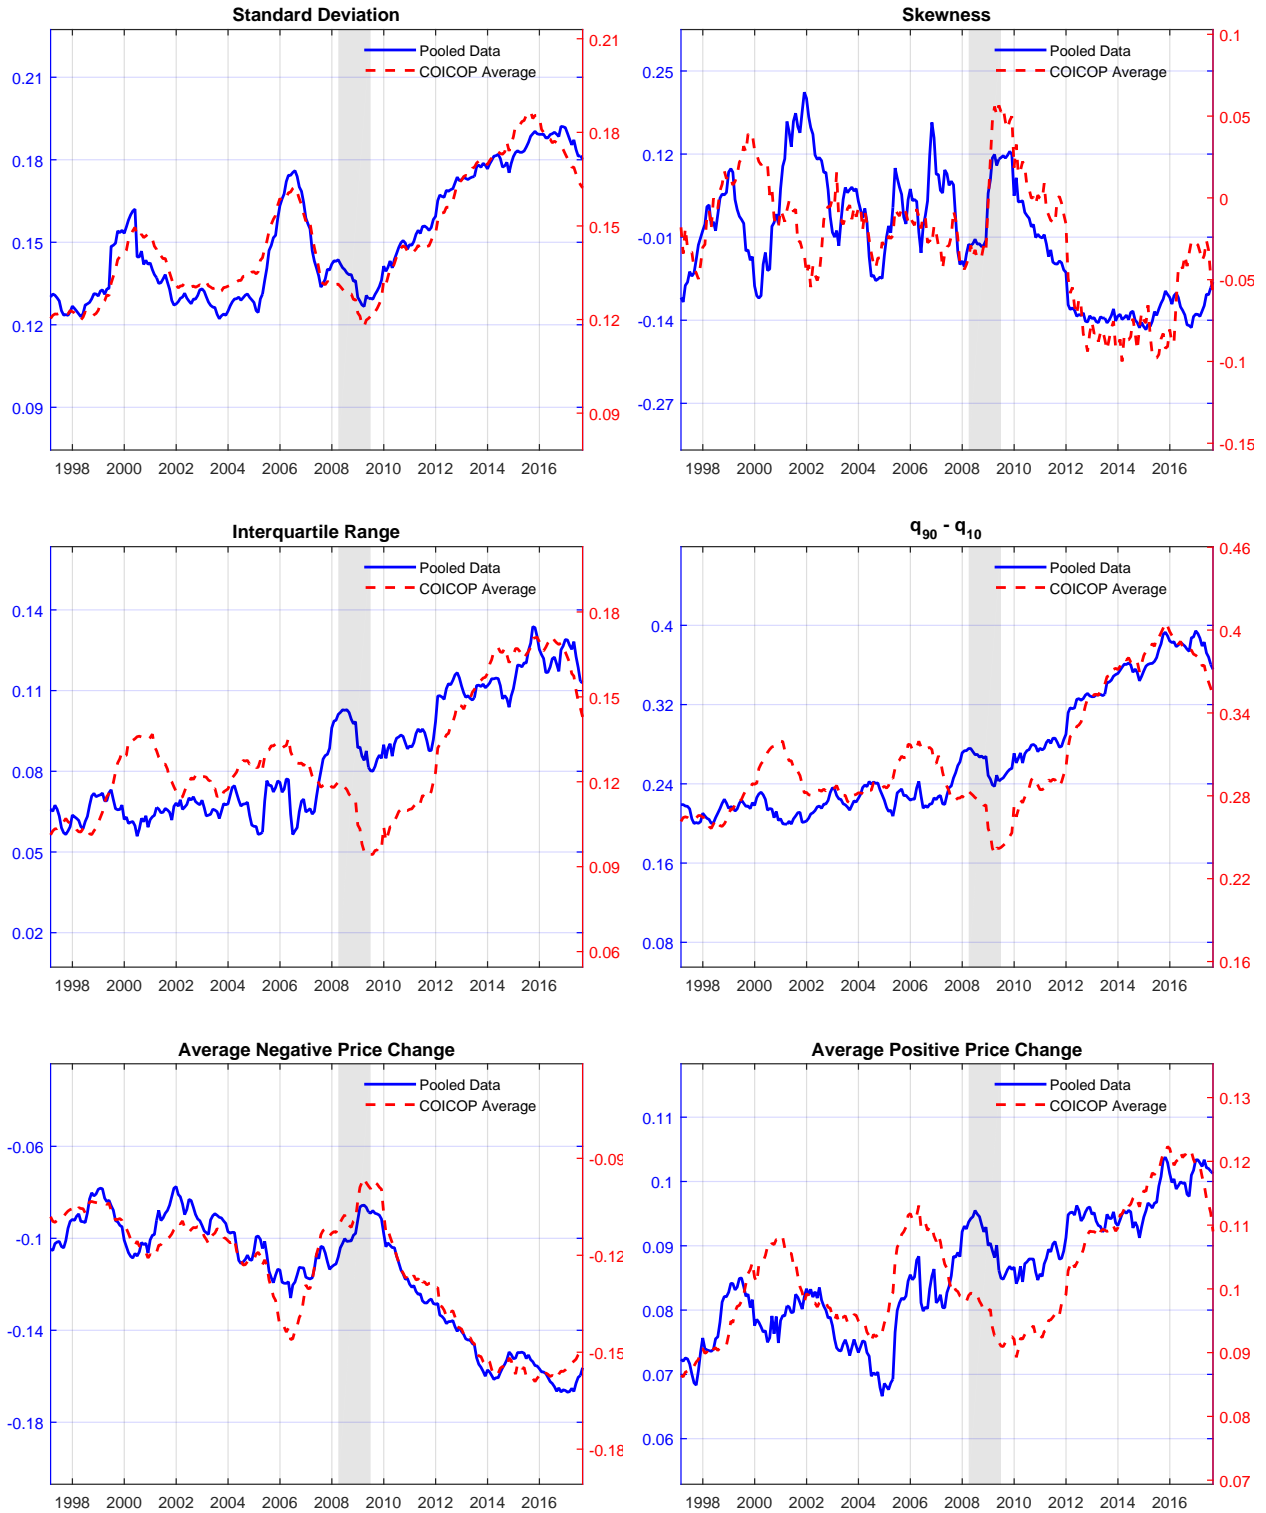


Notes: The left panel contrasts the weights assigned to each of the COICOP sectors to those assigned to build the CPI (re-normalized to exclude COICOP 4 and 10). The right panel reports the official CPI inflation, together with the inflation series retrieved from all the available price quotes (labeled *COICOP*) and the inflation obtained once all filters described in Section 2 are applied (labeled *Regular*). The COICOP codes are (1) Food And Non-Alcoholic Beverages, (2) Alcoholic Beverages, Tobacco And Narcotics, 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).

¹Due 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 counterparts obtained by averaging the corresponding moments of the price quotes obtained for each of the 25 COICOP group categories. The shaded vertical band indicates the duration of the Great Recession.

C A monthly coincident indicator of economic activity

We use monthly information on a number of 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 the Rotemberg (1999) version of the HP filter, which chooses the smoothing coefficient of the HP filter 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 underlying variables.

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 the Rotemberg (1999) version of the HP filter, which sets the smoothing coefficient 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 The combined effect of second-moment shocks and changes the incentives to adjust prices

Figure D.1 considers a situation in which both ϕ and λ increase.² The rise in the dispersion of price changes determines an expansion in the inaction region, thus increasing the density outside the adjustment bands and, in turn, the frequency of adjustment. This effect is counteracted by the rise in λ , which widens the inaction region further and restricts the density outside the adjustment bands beyond the initial situation. If the expansion in the inaction region is large enough to overcome the increase in dispersion, we observe opposite movements in the cross-sectional dispersion of prices and the frequency of adjustment. This is in line with what we observe in the post-recession period.

Figure D.1: 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 the dispersion of the distribution of price gaps, ϕ . The left panel reports the transformations occurring to the distribution of price gaps and the corresponding bands delimiting the inaction region: the dotted (blue) line refers to the ex-ante situation, the dashed (red) line denotes the effects of the volatility shift, while the dashed-dotted (magenta) line refers to the effects produced by the joint increase in ϕ and λ . The right panel reports the distributions of price changes, both in the ex-ante situation and in the case of a combined increase in ϕ and λ .


²Once again, a drop in α would lead to qualitatively similar results.

E Model estimates

Figure E.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 E.2: PARAMETERS OF THE PRICE GAP DISTRIBUTION



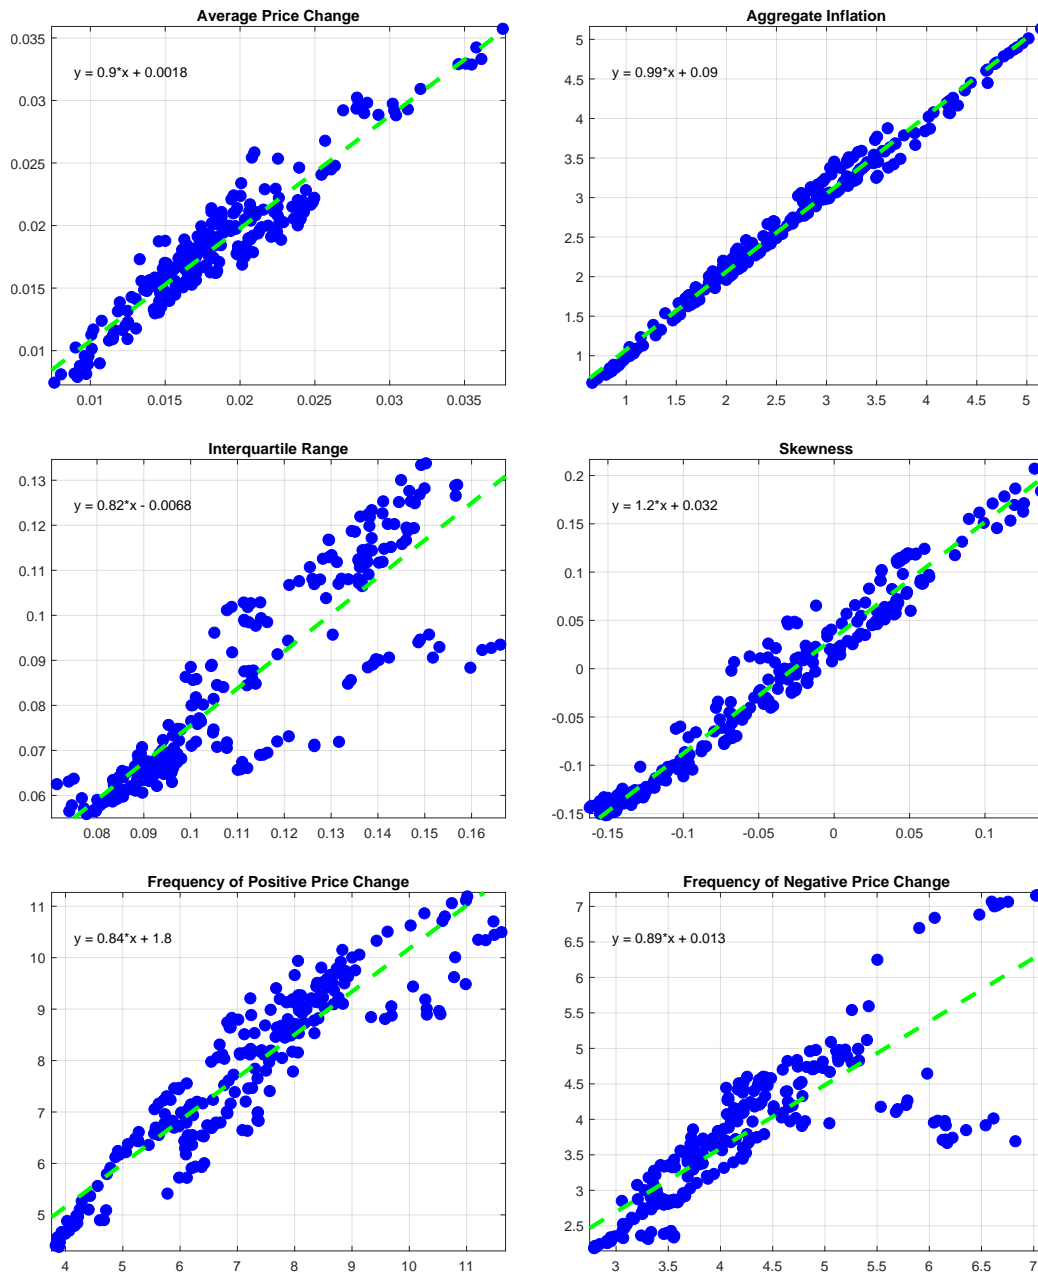
Note: The red lines denote the three VAT changes in the sample. The shaded vertical band indicates the duration of the Great Recession.

Figure E.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.

Figure E.4: FIT OF THE S_s MODEL (SELECTED MOMENTS)



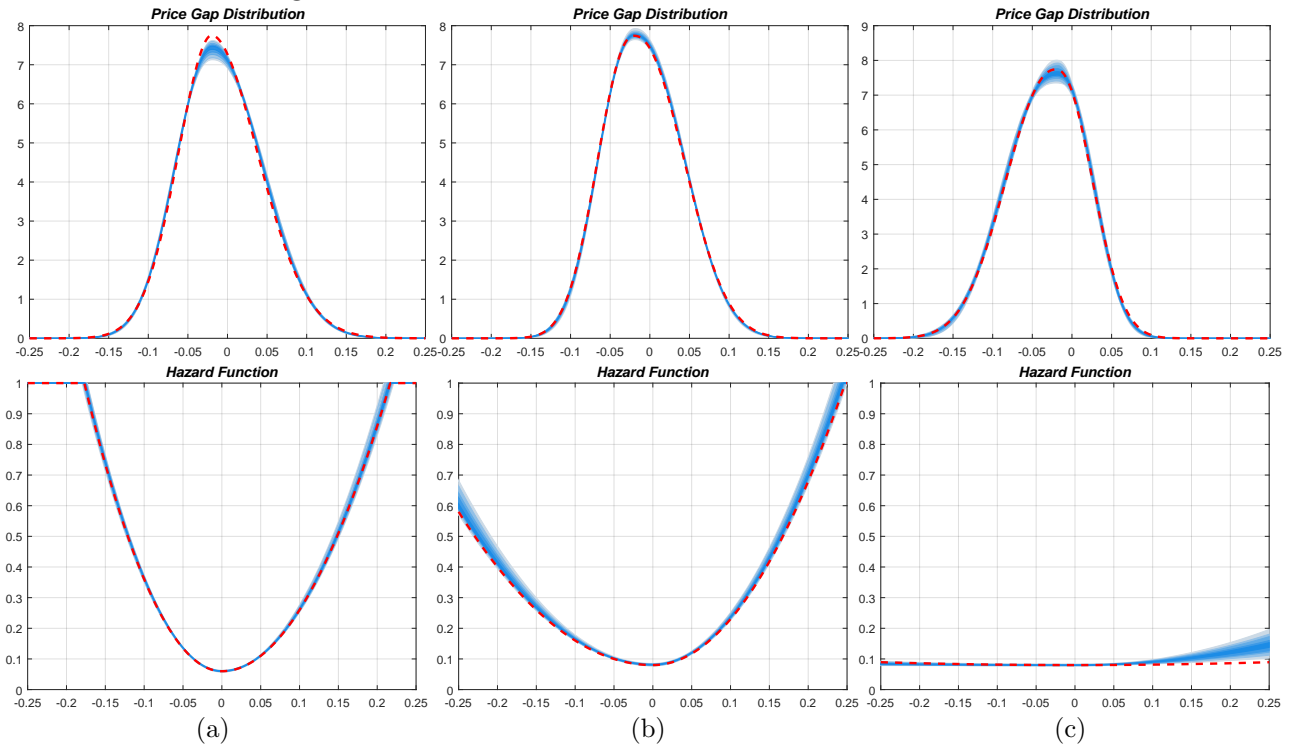
Notes: The figure compares the estimated moments from the S_s model in Section 4 (x-axis) to the moments estimated from the raw data (y-axis). Each chart also reports the linear fit (green/broken) line.

F Model identification

In this appendix we check whether the SMM estimation strategy we employ for the estimation of the generalized Ss model is able to separately identify the shape of the price gap distribution and the hazard function. To this end, we simulate samples of 100,000 price changes from the model, and then fit the model on each of these synthetic samples by SMM, matching the same moments we use in the baseline estimation (see Section 4.1). Figure F.1 contrasts the true price gap distribution (upper panel) and hazard function (lower panel) to the estimated counterparts. We look at three possible different parameterizations of the model, and report the ‘fan charts’ of the estimated functions. The specific parameterizations are merely meant to serve for illustrative purposes: we would obtain very similar evidence by imposing alternative specifications. Finally, for each set of calibrations, we simulate and estimate the model over 200 different samples.

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

Figure F.1: MODEL SIMULATIONS AND EMPIRICAL FIT



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


G Time variation in the adjustment hazard

In this appendix we are set to understand the importance of time variation in the hazard function. We are motivated by the fact that a number of earlier studies have considered empirical specifications of the Caballero and Engel (2007) framework with a fixed hazard function (see, e.g., Caballero and Engel, 2006, and Luo and Villar, 2017b and references therein). The hazard function typically moves in the face of second-moment shocks, as well as in reaction to changes in the cost firms face when deciding to change their prices. Both channels are abstracted away with a fixed adjustment hazard. In the exercise we devise in this appendix we fix the parameters of Λ_t to their sample average, and derive the resulting counterfactuals for price flexibility and inflation.³

Figure G.1 depicts the counterfactual flexibility index and the rate of inflation against their actual counterparts, respectively. Neglecting time variation in the adjustment hazard would generally lead to higher estimates of price flexibility: in this case, most of the adjustment would be triggered by second-moment shocks, as implied by the fact that the counterfactual index of price flexibility is highly correlated with the estimated scale parameter of the price gap distribution (see Figure G.2). This is pretty much in line with the analysis of Vavra (2014), who shows how greater price flexibility may emerge from greater volatility.

When examining the counterfactual rate of inflation, a notable fact is that changes in the hazard function appear to counteract changes in the price gap distribution, as implied by inflation being substantially more volatile in the absence of time variation in the hazard function.⁴ In this respect, it is important to recall that the hazard function moves both as a result of changes in the volatility of price gaps—whose movements also affect the price gap distribution—as well as a result of changes in the adjustment cost structure (recall the analysis of Section 3). When imposing an adjustment hazard with constant parameters, both effects are switched off, so that only the price gap distribution varies, as the result of first- or second-moment shocks. This is indirectly confirmed by the behavior of the actual and

Figure G.1: COUNTERFACTUALS WITH A FIXED HAZARD FUNCTION




Note: The counterfactuals are computed by setting the parameters of the adjustment hazard to the time average of their estimates. The shaded vertical band indicates the duration of the Great Recession.

³The counterfactual rate of inflation is produced by imposing a restriction to the adjustment hazard in each period, based on our estimates. For this reason, we only focus on period-by-period adjustments in the counterfactual scenarios, as modifying the hazard function in a given period should imply a different estimate of the price gap distribution in the next periods. We necessarily abstract from this dynamic implication, while focusing on the period wedge between the actual and the counterfactual inflation rate.

⁴This behavior is consistent with a regime of relatively large price flexibility, as it is effectively the case in this counterfactual scenario (see Section 6).

the counterfactual rates of inflation in the 2011-2015 time window: during this time span, the two series tend to display only minor differences, as implied by the fact that changes in the dispersion of the price gap distribution are dominated by the expansion in the inaction region.

Figure G.2: COUNTERFACTUAL FLEXIBILITY AND THE SCALE OF THE PRICE GAP DISTRIBUTION



Notes: Counterfactual flexibility index against the estimated scale of the price gap distribution.

H Alternative scenarios in the occurrence of a VAT change

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

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

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:

- *No $\Lambda(x)$ change*: 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

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

- *No VAT change*: What inflation would have been observed in absence of changes in the price gap distribution and the hazard function? This can be retrieved as

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

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

Comparing $\pi_t^{No\ VAT\ change}$ with the actual rate of inflation highlights the overall effects of the VAT, whereas the comparison between $\pi_t^{No\ \Lambda(x)\ change}$ 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).

I Details on the computation of the impulse response function from the Ss model

This appendix gives a brief account of how we compute the impulse response functions from the generalized Ss model presented in Section 4. We start by specifying a process for the exogenous (first-moment) shock.⁵ Specifically, we assume that:

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

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

The impulse responses are then calculated as:

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

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

J Estimation of the STARMA (p,q) model

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

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

This can be easily casted in state space. Therefore the likelihood can be calculated recursively using the Kalman filter (see Harvey, 1990). Since the model is highly non-linear in the parameters, it is possible to have several local optima and one must try different starting values of the parameters. Furthermore, given the non-linearity of the problem, it may be difficult to construct confidence intervals for parameter estimates, as well as impulse responses. To address these issues, we use a Markov Chain Monte Carlo (MCMC) method developed in Chernozhukov and Hong (2003; henceforth CH). This method delivers not only a global optimum but also distributions of parameter estimates.

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

- *Step 1:* Draw $\vartheta^{(n+1)}$, a candidate vector of parameter values for the chain's $n + 1$ state, as $\vartheta^{(n+1)} = \theta^{(n)} + \mathbf{u}_n$ where \mathbf{u}_n is a vector of *iid* shocks taken from a student-t distribution with zero mean, $\nu = 5$ degrees of freedom and variance Ω .
- *Step 2:* Take the $n + 1$ state of the chain as

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

⁵Since we assume that the shock has the same impact on all price quotes, the shock acts as a location shifter of the price gap distribution.

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

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

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

K Additional figures and tables

Table K.1: 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 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.

Table K.2: 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 K.3: 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.

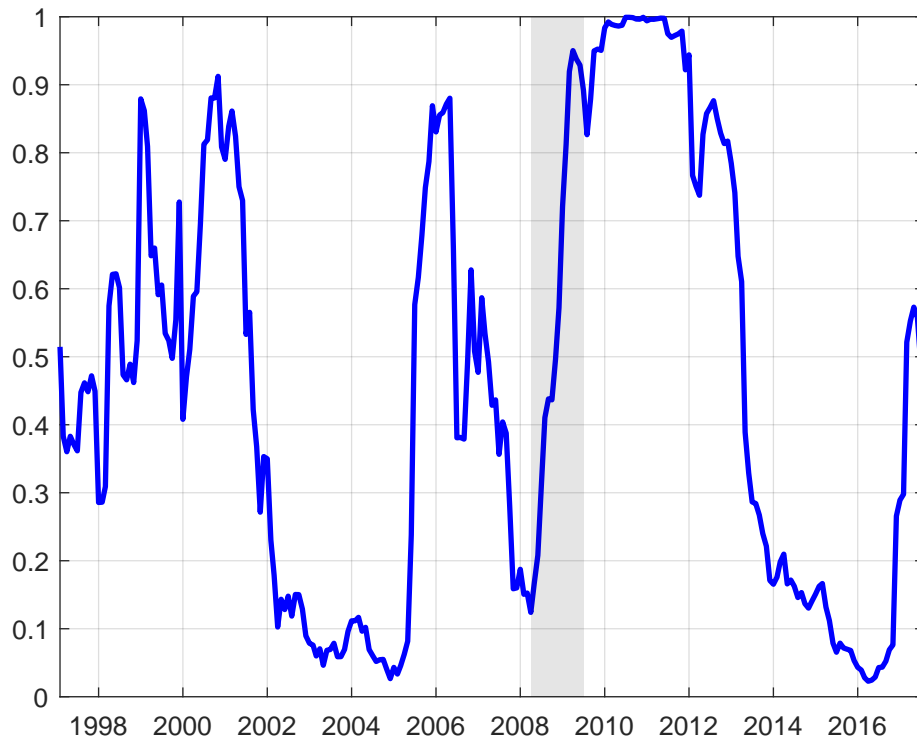
Table K.4: FORECAST ERRORS AND PRICE FLEXIBILITY: ROBUSTNESS (ABSOLUTE AND SQUARED FORECAST ERRORS)

(a) BoE MPC RPIX/CPI (Absolute) Forecast Errors						(b) BoE MPC RPIX/CPI (Squared) Forecast Errors							
Horizon	Slope at $G = 0.3$		Slope at $G = 0.9$		F-stat	\tilde{R}^2	Horizon	Slope at $G = 0.3$		Slope at $G = 0.9$		F-stat	\tilde{R}^2
1	0.093	[0.628]	0.840	[0.092]	0.229	1.69	1	0.078	[0.679]	0.606	[0.183]	0.507	-0.87
2	-0.330	[0.279]	2.319	[0.011]	0.045	6.41	2	-0.317	[0.490]	3.242	[0.008]	0.124	3.55
3	-0.484	[0.145]	4.117	[0.010]	0.003	13.82	3	-0.588	[0.303]	8.723	[0.011]	0.003	13.16
4	-0.344	[0.437]	6.161	[0.003]	0.000	26.45	4	-0.485	[0.584]	15.984	[0.014]	0.000	26.28
5	-0.144	[0.811]	5.945	[0.011]	0.000	20.10	5	-0.010	[0.994]	17.957	[0.022]	0.000	23.22
6	0.309	[0.603]	4.858	[0.032]	0.003	13.70	6	0.800	[0.554]	15.398	[0.050]	0.001	16.92
7	0.634	[0.236]	4.402	[0.021]	0.006	12.32	7	1.551	[0.225]	12.104	[0.078]	0.006	12.18
8	0.691	[0.182]	3.029	[0.055]	0.063	5.93	8	2.123	[0.143]	7.055	[0.244]	0.094	4.71

(c) Market Participants' (Absolute) Forecast Errors						(d) Market Participants' (Squared) Forecast Errors							
Horizon	Slope at $G = 0.3$		Slope at $G = 0.9$		F-stat	\tilde{R}^2	Horizon	Slope at $G = 0.3$		Slope at $G = 0.9$		F-stat	\tilde{R}^2
1	0.265	[0.361]	0.826	[0.122]	0.278	1.11	1	0.713	[0.291]	0.426	[0.497]	0.363	0.25
2	-0.383	[0.264]	2.448	[0.010]	0.053	6.12	2	-0.396	[0.464]	3.491	[0.007]	0.123	3.65
3	-0.561	[0.150]	4.293	[0.008]	0.004	13.10	3	-0.763	[0.287]	9.235	[0.008]	0.007	11.63
4	-0.382	[0.418]	6.398	[0.002]	0.000	25.60	4	-0.608	[0.517]	16.589	[0.010]	0.000	24.46
5	-0.103	[0.862]	6.042	[0.009]	0.000	18.74	5	-0.063	[0.960]	18.043	[0.016]	0.000	20.81
6	0.453	[0.412]	4.516	[0.049]	0.013	10.48	6	0.923	[0.465]	14.287	[0.045]	0.005	13.17
7	0.903	[0.052]	3.631	[0.052]	0.019	9.47	7	1.789	[0.129]	9.562	[0.099]	0.043	7.16
8	0.883	[0.099]	1.935	[0.221]	0.211	2.19	8	2.315	[0.091]	3.916	[0.431]	0.390	0.02


Notes: The table reports the results of a quadratic spline regression of the absolute (LHS) and squared (RHS) forecast errors (for different forecast horizons, h , measured in quarters) on a quarterly average of an indicator of the normalized price flexibility index, $G_{t-1} = G(\tilde{\mathcal{F}}_{t-1}; \gamma) = (1 + e^{-\gamma \tilde{\mathcal{F}}_{t-1}})^{-1}$, where $\tilde{\mathcal{F}}$ denotes the normalized flexibility index. The regression takes the form: $z_t = a_0 + a_1 G_{t-1} + a_2 G_{t-1}^2 + a_3 \mathbb{1}_{\{G_{t-1} > 0.5\}} G_{t-1}^2$, where $\mathbb{1}_{\{G_{t-1} > 0.5\}}$ is an indicator function taking value 1 when $G_{t-1} > 0.5$ and zero otherwise, $z_t = |e_{t+h|t}|$ (tables (a) and (c)) and $z_t = e_{t+h|t}^2$ (tables (b) and (d)). The upper panels refer to the Bank of England MPC's RPIX/CPI forecast errors, while the bottom panels consider market participants' forecast errors. In each panel, the first two 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). Since the fitted function tends to reach a minimum at about $G = 0.6$, for most forecast horizons, 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). The penultimate column (F-stat) reports the p-value of the null hypothesis that all the coefficients associated to the flexibility regime are equal to 0 (i.e., $H_0 : a_1 = a_2 = a_3 = 0$). The last column reports the adjusted R-squared, denoted by \tilde{R}^2 .

Figure K.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 K.2: PRICE FLEXIBILITY AND INFLATION PERSISTENCE



Note: Figure K.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). In each of the two charts the vertical line delineates the half-life of the shock.

Figure K.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).