

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

October 2020

Abstract

Using microdata underlying the UK consumer price index, we examine how time variation in the capacity of nominal demand to stimulate price adjustment—broadly defined as *price flexibility*—shapes inflation dynamics. Price flexibility maps into a distinct non-linearity of the rate of inflation, whose half-life is twice as large in periods of relatively low flexibility. Such non-linearity naturally arises in environments characterized by state-dependent price setting, of which we find ample evidence in the data. Yet, this fact is often neglected in the practice of central banking. Indeed, we illustrate that a sizeable fraction of professional forecasters' inflation prediction error is explained by time variation in price flexibility, especially at medium-term forecast horizons. Overlooking these facts may severely bias our understanding of price setting and inflation dynamics.

JEL codes: E30, E31, E37, C22.

Keywords: Inflation, Price flexibility, Ss models.

*We thank for comments Yoosoon Chang, Andrea Colciago, Luca Dedola, Federico Di Pace, Etienne Gagnon, Ana Galvao, Francesco Lippi, Ricardo Nunes, Giorgio Primiceri, Omar Rachedi, Federico Ravenna, Søren Hove Ravn, Yad Selvakumar, Joseph Vavra, 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, the European Central Bank, De Nederlandsche Bank, the University of Warwick, the University of Padua, the Monetary Policy, Term Premia, and Macro Volatility SNDE session at the 2019 ASSA meetings, the University of Lund, and the University of Verona. 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:* Øster Farimagsgade 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 examine in detail the microeconomics of price setting—often through the lens of alternative protocols of price formation—as well as its macroeconomic implications—through the analysis of inflation dynamics and various measures of aggregate *price flexibility*. The latter—which is to be broadly intended as the response of the aggregate price level to monetary shocks—lies at the core of the monetary policy transmission mechanism, ultimately embodying Central Banks’ capacity to affect output and inflation. We examine the statistical properties of price setting from a micro-data viewpoint, taking a specific stand on time variation in price flexibility and how it shapes aggregate inflation. Despite a large number of empirical contributions measuring the response of prices to nominal stimulus, little emphasis has been placed on the sources and the characteristics of time variation in price flexibility.¹ Most notably, the literature has been silent on how and to which extent time-varying price flexibility maps into inflation dynamics. We contribute to fill this gap in the literature, and show that keeping track of changes in price flexibility is important to produce accurate inflation projections.

Using price microdata underlying the UK consumer price index (CPI) over the 1996-2017 time span, we estimate the generalized *Ss* model developed by Caballero and Engel (2007) over the price quotes available in each month, fitting both the *distribution of price gaps* (i.e., the wedge between actual and optimal reset prices) and the *adjustment hazard* (i.e., the probability of a good’s price changing as a function of its price gap). Along with encompassing various price-setting protocols, the generalized *Ss* model is well suited to examine time variation and comovements among various price-setting statistics. Both functions change substantially over time, displaying some distinct asymmetries that are extremely informative about the underlying microeconomic process of price setting. Notably, changes in the price-adjustment cost structure—as reflected in a marked downward shift and flattening of the adjustment hazard—imply greater inaction in price adjustment and lower price flexibility during the post-recession sample. In turn, this reflects into a marked increase in the dispersion of price changes, as well as a concurrent fall in the frequency of price adjustment.

Our semi-parametric estimates are then used to make inference about the underlying process of price-setting and how it evolves over time. To this end, we back out predetermined price adjustments—the so-called *intensive margin*—and adjustments triggered or canceled by the shock—the *extensive margin*.² State dependence in price setting—as embodied by the extensive

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

²Adjustments occurring over the intensive margin characterize both time- and state-dependent models. The extensive margin, instead, is a defining feature of state-dependent models.

margin—appears quite important, and more so after the Great Recession, in correspondence with a marked increase of inflation volatility. Once again, this is indicative of some persistent movements in the hazard function that, according to conventional menu cost models, may be rationalized by a rise in market power, among other things. This is in line with DeLoecker and Eeckhout (2018) and Bell and Tomlinson (2018), who show that the markup in the UK has displayed only a modest increase in the 1996-2007 period, while rising substantially thereafter.³

From a macroeconomic perspective, the inflation response to nominal shocks varies substantially: price flexibility increases by about 50% between the start of the Great Recession and 2011, thus reverting and attaining its minimum in the first quarter of 2016, to revert again 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, over the last decade of the sample inflation has been outside the 1%-3% interval for a total of 22 quarters, while in the previous decade this has only occurred during 11 quarters. Moreover, after the Great Recession 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 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 right after the recession, thus reaching its minimum in correspondence with inflation at the historical low in 2015.

The key message of this paper is that 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, 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 take 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.

Our evidence highlights the practical relevance of state-dependent pricing for understanding inflation dynamics. As such, it has crucial implications for monetary authorities seeking to stabilize inflation. In a relatively low price-flexibility environment, nominal shocks are likely to dissipate slowly, so that the Central Bank might struggle to keep inflation at its target. By contrast, under relatively high price flexibility the same shock would transmit proportionally more to quantities than prices. It is plausible to expect that such state dependence would map into the trade off that central banks typically face when stabilizing output and inflation. Therefore, it is advisable to keep track of the real-time evolution of price flexibility, to have a reliable appreciation of inflation persistence.

³In ongoing work we examine the implications of secular changes in the markup of both the UK and other advanced economies for price-setting and the slope of the Phillips curve (see Petrella et al., 2020).

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. Relative to this paper, our novel contribution is to show how accounting for time variation in price flexibility improves our understanding of inflation dynamics: in this respect, we document that inflation is more persistent and less volatile in periods of relatively low price flexibility, and show that neglecting this fact can lead to a large prediction bias. We also relate to Luo and Villar (2017b), in that we highlight the importance of time variation in the hazard function for quantifying the pass-through of monetary shocks on prices. In fact, movements in firms' incentives to price adjustment are shown to be prominent (in this respect, see also Hobijn et al., 2006).

Our work also 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). An empirical limitation of these models is to rely on specific shocks to the price-setting units, while the approach we follow 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.

We also relate to some recent empirical contributions employing individual consumer prices from the UK. In this respect, Bunn and Ellis (2012) have been among the first to investigate the key characteristics of the frequency of price setting and the hazard functions implied by the microdata from the Office for National Statistics (ONS), while Dixon et al. (2014) have focused on the impact of the Great Recession on price setting. The latter, in particular, attributes little importance to endogenous macroeconomics effects on pricing, while our evidence points to a certain prominence of state dependence in price setting, and more so during the sample that is not accounted for in their analysis (i.e., post-2013), during which the extensive margin of price adjustment overcomes the intensive one (i.e., the frequency of adjustment) in the contribution to price flexibility. In fact, the novelty of the approach rests on tracking time changes in both margins of adjustment, rather than focusing on their average importance. Our application underlines the importance of the selection effect for aggregate inflation (see, on this, Carvalho and Kryvtsov, 2018 and references therein), along with stressing its time variation. Lastly, Chu et al. (2018) emphasize that information on the distribution of price changes can be exploited to forecast inflation. We dig deeper in this, showing that price flexibility—which condenses valuable information from key micro price statistics—also contains valuable information for predicting inflation persistence.

⁴See, among others, Bils and Klenow (2004), Dotsey and King (2005), Alvarez et al. (2006), Gertler and Leahy (2008), Klenow and Kryvtsov (2008), Nakamura and Steinsson (2008), Gagnon (2009), Costain and Nakov (2011), Midrigan (2011), Nakamura et al. (2011), Alvarez and Lippi (2014), Karadi and Reiff (2019), Berardi et al. (2015), Alvarez et al. (2016), Nakamura et al. (2018).

Structure The rest of the paper is organized as follows. Section 2 discusses the key characteristics of the ONS microdata on consumer prices. Section 3 reviews the generalized S_s model and takes it to the data. Section 4 examines time variation in the distribution of price gaps and the adjustment hazard, as well as the relative importance of adjustments along the intensive and the extensive margin at different points in time. Section 5 discusses the implications of state dependence in price flexibility for inflation dynamics. Section 6 concludes.

2 Microdata on consumer prices

We use ONS microdata underpinning the UK CPI. Prices are collected on a monthly basis, for more than 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 might allow 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

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 CO*nsumption 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.

not be uniquely identified. This is typically the case when the ONS samples the same item, in outlets that are part of a chain, but for multiple locations within the same region. As an example, in March 2013 we pick an item with the following characteristics: ‘Women’s Long Sleeves Top’ (*id*: 510223) sold in multiple outlets (*shop type*: 1) within the region of London (*region*: 2). With these coordinates at hand we retrieve two different price quotes: one location sells the item for £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 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

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

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.

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 results are robust to this approach.

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).

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 Stylized facts

This section presents some key facts about the behavior of the ONS microdata, and their implications for inflation dynamics. We start by re-writing inflation as the product of the frequency of adjustment (fr_t)—defined as the share of prices being adjusted in every month—and the average price change in every month (Δp_t):

$$\pi_t = fr_t \times \Delta p_t. \quad (1)$$

The frequency is computed as $\sum_i \omega_{i,t} \mathbb{1}_{\{\Delta p_{i,t} \neq 0\}}$, with $\omega_{i,t}$ denoting the CPI weight associated with good i at time t , and $\mathbb{1}_{\{\Delta p_{i,t} \neq 0\}} = 1$ if $\Delta p_{i,t} \neq 0$, and zero otherwise. The average price, instead, is computed as $fr_t^{-1} \sum_i \omega_{i,t} \mathbb{1}_{\{\Delta p_{i,t} \neq 0\}} \Delta p_{i,t}$. All the statistics derived from microdata display a pronounced seasonality (see, e.g., Alvarez et al., 2006), which we remove by computing the 12-month moving average.

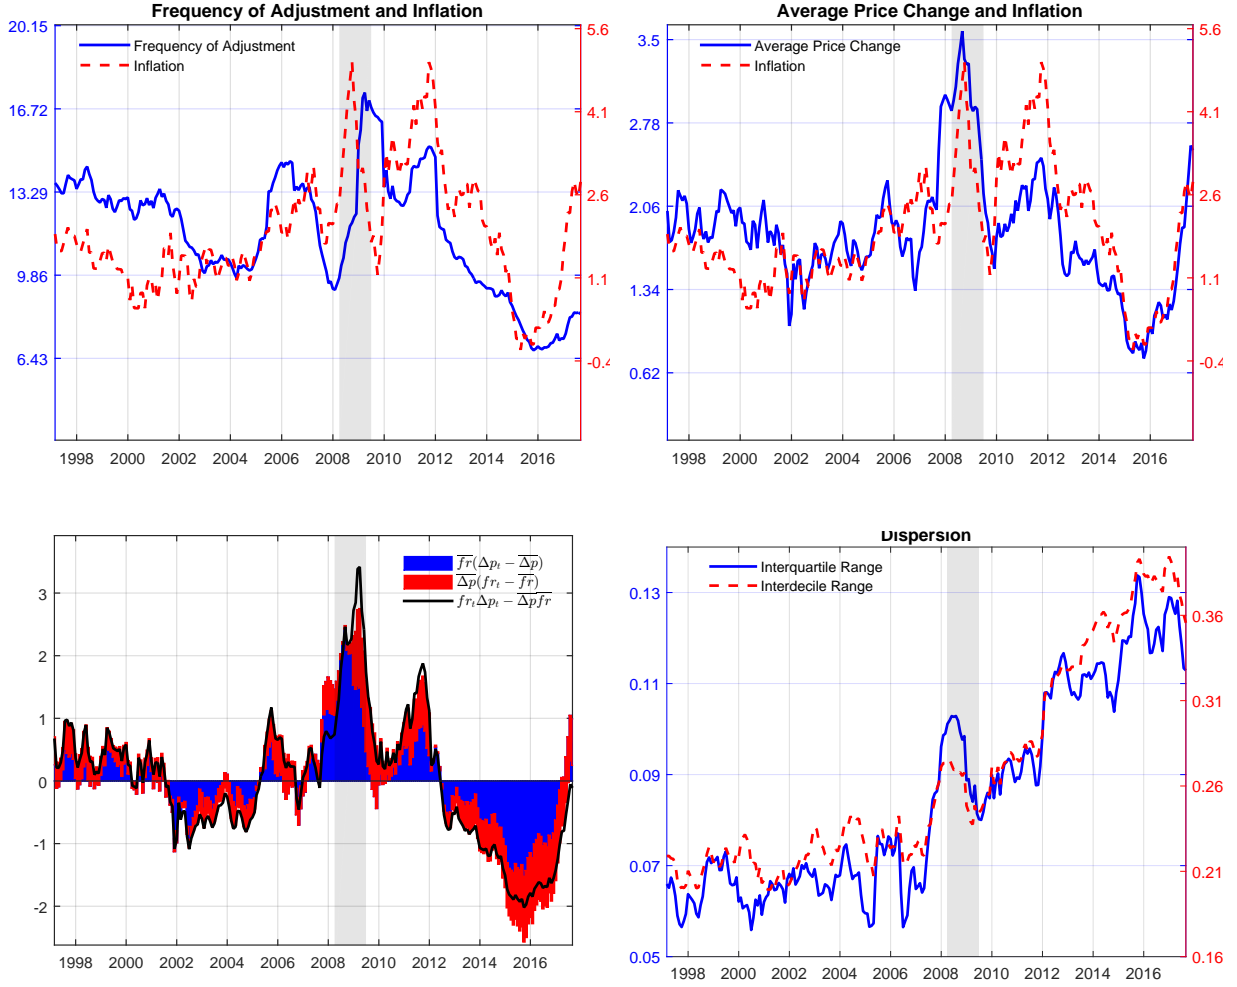
The top panels of Figure 1 report fr_t and Δp_t , respectively. As expected, the average price change displays a strong degree of positive comovement with CPI inflation, at least until the end of the Great Recession. Thus, in the last part of 2015 the two series are back moving in tandem. As for the frequency of adjustment, this tracks very closely the contraction in CPI 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 2015.¹⁶

The bottom-left panel of the figure looks deeper into the cyclical behavior of the rate of inflation, decomposing inflation deviations from its sample mean, $\pi_t - \overline{fr \Delta p}$, into three components: *i*) one term, $\overline{fr}(\Delta p_t - \overline{\Delta p})$, capturing the variation in aggregate inflation associated with changes in the average change of those prices that change from one period to the next one; *ii*) one term, $\overline{\Delta p}(fr_t - \overline{fr})$, accounting for variation in the frequency of adjustment; *iii*) the last term, $(\Delta p_t - \overline{\Delta p})(fr_t - \overline{fr})$, that accounts for the joint variation of the two moments around their respective sample means. Notably, only about half of the variability of inflation is explained by the variability of the average price change, whereas the remaining part is due to changes in the frequency (either directly or indirectly, through its positive comovement with

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.

¹⁶The average frequency of price adjustment prior to its drop is slightly below the estimates reported by previous studies on UK micro price data conducted over roughly the same time span. This reflects the fact that we exclude from our sample both sales and utility prices (COICOP 4), with the latter being a particularly volatile component of the CPI index. By contrast, Bunn and Ellis (2012) include utility prices and sales, while Dixon and LeBihan (2012), Dixon et al. (2014) and Dixon and Tian (2017) only include sales.

Figure 1: FREQUENCY, AVERAGE PRICE CHANGES, AND DISPERSION



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

the average price change). The contribution of the frequency is particularly evident in the post-recession period. In fact, if one considers the 2011 inflation spike, or the period of weak inflation in the last part of the sample, about half of inflation's deviation from its mean is explained by a relatively high or a relatively low frequency of price adjustment, respectively.

Finally, the bottom-right panel of the figure plots different measures of dispersion of the distribution of (non-zero) price changes. Both the interquartile and the interdecile range display a very large increase in the aftermath of the Great Recession.¹⁷ In this respect, a key

¹⁷Also the standard deviation displays a similar increase. However, this measure is often influenced by outliers. This type of problem does not plague the interquartile and the interdecile range.

Table 2: PRICING MOMENTS AND MACROECONOMIC VARIABLES

	Rotemberg Filter			Year-over-Year Filter		
	fr_t	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	fr_t	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$
$\mathbb{I}(\text{Rec.})$	1.643***	-0.448***	-0.955***	1.570	-1.162**	-0.882
y_t	-0.569***	0.334**	0.422***	-0.489***	0.207	0.219
π_t	0.169	-0.016	-0.147	0.128	0.113	0.115
fr_t	–	-0.510***	-0.737***	–	-0.461***	-0.687***

	Quadratic Detrending			Hamilton Filter		
	fr_t	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	fr_t	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$
$\mathbb{I}(\text{Rec.})$	3.383***	-0.316***	-1.503***	5.747***	-1.052**	-1.850***
y_t	-0.486***	0.143	0.330**	-0.428***	-0.085	-0.064
π_t	0.497***	0.055	-0.265*	0.125	0.239	0.199
fr_t	–	-0.186	-0.575***	–	-0.413***	-0.652***

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

fact emerging from the graphical inspection is that the dispersion of price changes and the frequency of adjustment tend to comove negatively, and more so after the Great Recession.¹⁸ As stressed by Vavra (2014), the joint dynamics of these statistics and their cyclical properties are key to understand the endogenous and exogenous determinants of price adjustment.

While the frequency of adjustment and the dispersion of price changes display diverging trends over the last decade—with the former falling and the latter rising—Table 2 looks at the correlation between their detrended counterparts, as well as at the comovement of each of these with a business cycle indicator.¹⁹ We set aside potential spurious correlation emanating from the low-frequency behavior of these series, and detrend all of them—using different filters—apart from the inflation rate. The frequency of price adjustment moves countercyclically, while the dispersion of price changes shows some mild procyclicality, so that prices tend to be adjusted more frequently during recessions and, concurrently, smaller price changes are observed. Most

¹⁸Figure B.1 in Appendix B shows that composition effects have no role in generating the facts presented in this section. 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.

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

notably, the two statistics consistently display negative correlation.²⁰ As we explain in the next section, diverging movements in these moments imply some distinctive dynamics in the underlying determinants of price adjustment, being key to track time variation in the pass-through of nominal shocks to inflation.

3 Framing the analysis

To dig deeper into the determinants of price adjustment and their aggregate implications we use the generalized *Ss* model developed by Caballero and Engel (2007). This 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.²¹ 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 allow for time variation in the determinants of price adjustment, we can estimate the model over each cross section of price microdata, matching different price-setting statistics.

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

Caballero and Engel (2007) assume *iid* idiosyncratic shocks to the adjustment cost. Thus, by integrating over their possible realizations, we obtain 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

²⁰We also detect a certain tendency for the skewness to behave countercyclically, while kurtosis appears acyclical. Most importantly, the skewness displays no correlation with the rate of inflation. This is somewhat puzzling, given that a wide range of structural models tend to produce non-zero inflation-skewness comovement (see, e.g., Luo and Villar, 2017a). However, we should stress that the correlation turns positive and significant in the post-recession period, thus emphasizing the role of the extensive margin of price adjustment and, therefore, state dependence.

²¹To focus on two somewhat extreme examples, the generalized *Ss* model can account for both price setting à la Calvo (1983)—where firms are selected to adjust prices at random and price flexibility is fully determined by the frequency of adjustment—as well as for schemes à la Caplin and Spulber (1987)—where adjusting firms change prices by such large amounts that the aggregate price is fully flexible, regardless of the frequency of adjustment.

price gaps immediately before an adjustment takes place at time t , aggregate inflation can be recovered as

$$\pi_t = - \int x \Lambda_t(x) f_t(x) dx. \quad (2)$$

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

3.1 Taking the model to the data

To take the model to the data we need to specify generic functional forms for the distribution of price gaps and the hazard function.²² 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

$$f_t(x) = \begin{cases} \frac{\delta(\varrho_t, \nu_t)^{1/\nu_t}}{\psi_t \Gamma(1+1/\nu_t)} \exp \left[-\frac{\delta(\varrho_t, \nu_t)}{\varrho_t^{\nu_t}} \left| \frac{x-\theta_t}{\psi_t} \right|^{\nu_t} \right] & \text{if } x \leq \theta_t \\ \frac{\delta(\varrho_t, \nu_t)^{1/\nu_t}}{\psi_t \Gamma(1+1/\nu_t)} \exp \left[-\frac{\delta(\varrho_t, \nu_t)}{(1-\varrho_t)^{\nu_t}} \left| \frac{x-\theta_t}{\psi_t} \right|^{\nu_t} \right] & \text{if } x > \theta_t \end{cases}, \quad (3)$$

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:²³

$$\Lambda_t(x) = \min \{ a_t + b_t x^2 \mathbb{1}_{\{x>0\}} + c_t x^2 \mathbb{1}_{\{x<0\}}, 1 \}, \quad (4)$$

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, which has recently been supported by

²²Alvarez et al. (2020) highlight that the moments of the price gap distribution, together with the frequency of price changes, provide enough information to identify the distribution of price gaps and the hazard function. In fact, they show—for the case of symmetric functional forms—that $\Lambda_t(x)$ and $f_t(x)$ are fully encoded in distribution of price changes and fr_t . While our estimates do not explicitly take into account this mapping, in Section 4 we show that our estimates return measures of the cumulative response of prices to a monetary stimulus which are consistent with alternative measures of money non-neutrality, such as the one proposed by Alvarez et al. (2016).

²³Unlike Berger and Vavra (2017), we allow for asymmetry in the hazard function, along with the distribution of price gaps. In fact, the distribution of price changes is characterized by some sizeable asymmetry, as well as a marked tendency of skewness to vary over time. Moreover, there is a sizable difference between the frequency of price adjustment associated with positive and negative price changes (both statistics are reported in Figure I.1). By allowing both the distribution of price gaps and the hazard function to be asymmetric, we avoid assuming that the underlying asymmetries in the data are entirely driven by either of the two functions. In support of this choice, the assumption of symmetric branches in the hazard function is generally rejected by the data.

Luo and Villar (2017b).²⁴

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.²⁵ We also match the frequency and the average size of prices movements, conditioning on positive and negative price changes. Last, we match the observed rate of inflation. The estimates are obtained by simulated minimum distance, using the identity matrix to weight different moments.²⁶ All the estimated parameters and the derived statistics inherit some pronounced seasonal variation from the raw data. Thus, we report their 12-month moving-average counterparts.

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

3.2 Estimates

Figure 2 graphs the estimated price gap distributions and the hazard functions.²⁷ Notably, $f_t(x)$ spreads considerably after 2010. As for $\Lambda_t(x)$, its shape changes significantly over the entire sample, with periods of high and low probability of adjusting prices that alternate over

²⁴We have also checked that our results are robust to plausible variations to the specification of these functional forms. Using a Pearson Type 7 distribution, a mixture of two Normal distributions, or a mixture of a Laplace and a Normal distribution for the price gap, as well as an asymmetric inverted normal function for the hazard function, delivers results that are qualitatively similar to those reported in the next section.

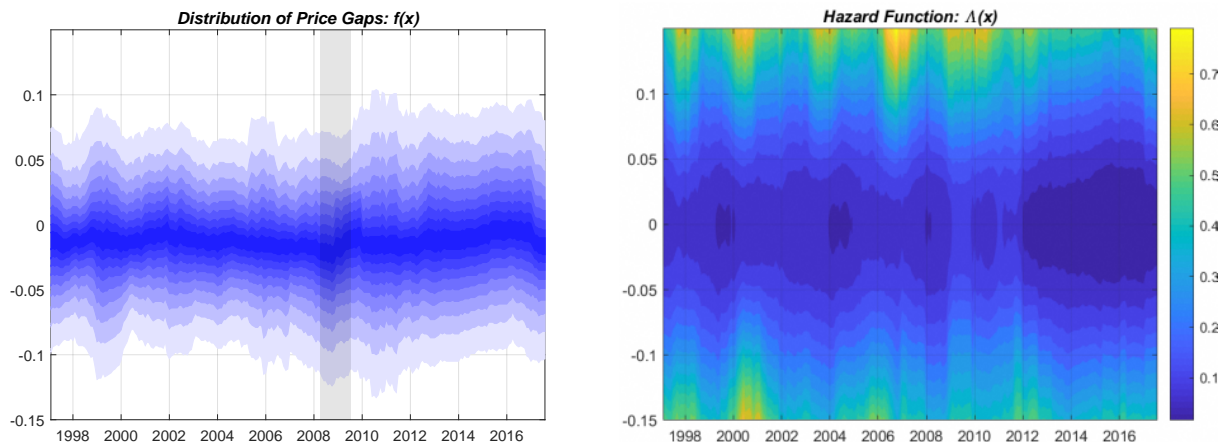
²⁵We match quantilic moments, as the 3rd and 4th moments of the cross-sectional distribution are quite sensitive to outliers.

²⁶Altonji and Segal (1996) highlight that matching the unweighted distance between moments often performs better in small samples, as compared with using optimal weights. The moments of the simulated distribution are estimated by drawing 100,000 price quotes. We use the Genetic Algorithm to minimize the quadratic distance between data moments and simulated moments, so as avoid ending up in local minima (see, e.g., Dorsey and Mayer, 1995).

²⁷Figures E.1 and E.2 report the estimated parameters, while Figure E.3 reports the fit of selected data moments, and shows that the empirical model is able summarize the main stylized facts in the data.

time. While the 2009-2012 time span is characterized by a somewhat high and steep hazard function—most likely as the result of three VAT changes taking place over that short time window—²⁸ this shifts downwards and flattens quite markedly and persistently thereafter.

Figure 2: ESTIMATED PRICE GAP DISTRIBUTIONS AND HAZARD FUNCTIONS



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

The joint movements in $f_t(x)$ and $\Lambda_t(x)$ are key to explain the emergence of diverging trends in the frequency of adjustment and the dispersion of price changes that occur in the post-recession period. Figure 3 digs deeper on this: in the left panel we report the adjustment thresholds associated with two hazard probabilities, namely 5% and 7%, while the right panel reports the share of prices within the same thresholds:²⁹ notably, in the last part of the sample an increasing number of price quotes falls in this range. To make an example, whereas in 2009 only price cuts of greater than 2% had a sizeable probability of being enacted, in 2016 it was price cuts larger than 7% having the same chance of taking place. Therefore, the resulting distribution of price changes in the last part of the sample is characterized by few and large price changes, along with being skewed towards price cuts.

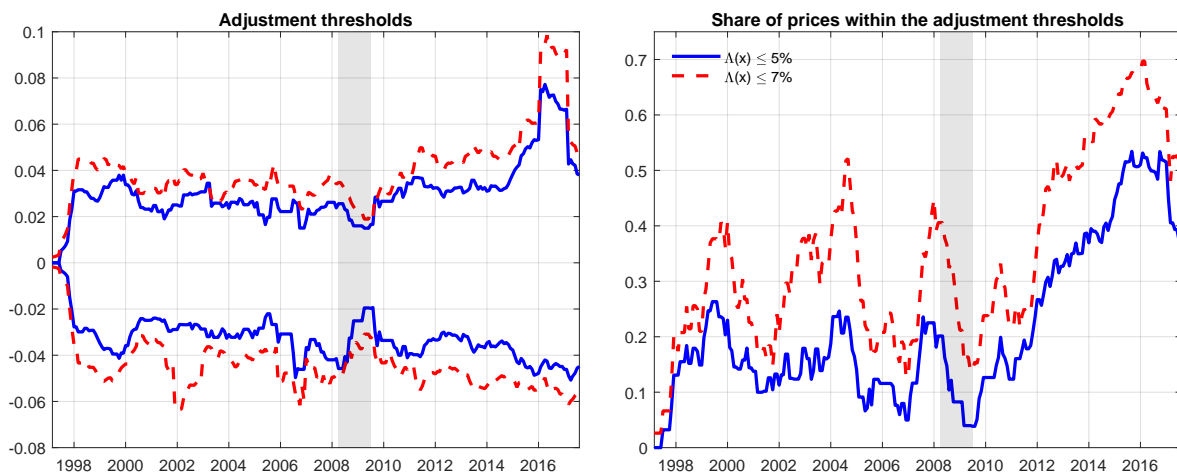
According to conventional menu cost models, a decline in the frequency of adjustment, coupled with a surge in its dispersion, may be rationalized by an expansion in the inaction region that overcomes the effects of a positive shift in the dispersion of price gaps.³⁰ By the end of the sample about five times as many firms are likely not to adjust, as compared with the pre-2010 time window. This stands as indirect evidence that increasing price rigidity, as captured by the downward shift in the hazard function, dominates the increase in the dispersion

²⁸Namely, a reduction, from 17.5% to 15%, on December 1, 2008, followed by two hikes: one, up to 17.5%, on January 1, 2010, and one, further up to 20%, on January 4, 2011. In particular, the first two VAT changes are associated with a marked shift in the estimated parameters of the hazard function, as it is visible from Figure E.2.

²⁹In the present setting the hazard rate is never zero, irrespective of the price gap. Thus, there is not inaction region as such.

³⁰Appendix D reports a stylized menu cost model that stresses how changes in the incentives firms face when deciding to change prices can provide us with a rationale for diverging movements in the dispersion of price changes and the frequency of adjustment.

Figure 3: INACTION IN PRICE ADJUSTMENT



Note: Left panel: adjustment thresholds associated with 5% and 7% hazard probabilities; right-panel: share of prices within the adjustment thresholds. The shaded vertical band indicates the duration of the Great Recession.

of $f_t(x)$. Note also that greater nominal rigidity 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 a flattening of the hazard function, imply that non-predetermined price adjustments—which are more likely to occur for large price gaps—have played an increasingly important role in the recent past, as the analysis in the next section will confirm.


4 Inspecting price setting in a time-varying environment

The estimates of the generalized Ss model emphasize the importance of tracking changes in the distribution of price gaps and the hazard function. Caballero and Engel (2007) show that, within their accounting framework, one can derive a measure of aggregate price flexibility that captures the impact response of realized inflation to a one-off aggregate nominal shock:

$$\mathcal{F}_t = \lim_{\mu_t \rightarrow 0} \frac{\partial \pi_t}{\partial \mu_t} = \underbrace{\int \Lambda_t(x) f_t(x) dx}_{\text{Intensive Margin}} + \underbrace{\int x \Lambda_t'(x) f_t(x) dx}_{\text{Extensive Margin}}. \quad (5)$$

In turn, 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

Figure 4: PRICE FLEXIBILITY AND DIFFERENT MARGINS OF PRICE ADJUSTMENT



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

adjusted their price but choose not to do it.³¹ It is also important to stress that, since \mathcal{F}_t is simply derived from the accounting identity (2), its validity as a measure of aggregate flexibility does not require that we take a stand on a specific model of price setting.

Figure 4 reports the estimated index of price flexibility, as well as its decomposition into the intensive and the extensive margin of price adjustment. \mathcal{F}_t displays sizable variation 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. 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 set of price gaps with an extremely low likelihood of being reset implies that fewer firms are pushed near the adjustment boundaries. In fact, 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 (5):

$$\mathcal{F}_t = \int \Lambda_t(x) f_t(x) [1 + \eta_t(x)] dx \quad (6)$$

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

³¹In 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.

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.2 in Appendix E). Alternatively, the same point can be made by approximating the flexibility index as $F_t \cong Int_t + 2[Int_t - \Lambda_t(0)]$:³² 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*. It is important to recognize that such a shift in the hazard function is in line with an increase in market power that determines a drop in the cost of being away from the optimal price. In fact, this view is consistent with the sizable increase in the markup that has been observed during the post-recession period, as recently documented by DeLoecker and Eeckhout (2018) and Bell and Tomlinson (2018) for the UK economy.³³

From a cyclical perspective, movements in price flexibility do not seem to occur at random: in fact, \mathcal{F}_t goes from being positively correlated with output growth in the first part of the sample (0.456), to comoving negatively during the last decade (-0.577). Therefore, in the second half of the sample monetary policy has become more effective at stimulating demand, though at the cost of losing traction on inflation control. This can help explain the marked difficulty in keeping inflation on target over this period of time, along with other factors that will be mentioned in the next section.

As for the correlation with the rate of inflation, this is generally positive, and more so in the post-recession sample (0.380), 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 4).


Price flexibility and money non-neutrality Our analysis highlights a great deal of variation in aggregate price stickiness, as measured by the Caballero and Engel (2007)’s flexibility index. However, we acknowledge this is not the only measure of price flexibility available. In particular, Alvarez et al. (2016) put forward a sufficient statistic for money non-neutrality, intended as the cumulative output response to a nominal shock, and prove that this proportional to the steady-state ratio of kurtosis to the frequency of price changes, in a variety of sticky-price models. In fact, Alvarez et al. (2020) note that Caballero and Engel (2007)’s flexibility index may correspond to different values of the cumulative impulse response of output to a monetary shock and that, as such, it does not represent a sufficient statistic of aggregate price flexibility. Our estimates of $\Lambda_t(x)$ and $f_t(x)$ imply a dynamics of price flexibility in line with that of Alvarez et al. (2016)’s sufficient statistic for money non-neutrality.³⁴ The left panel of Figure 5 provides a direct comparison between \mathcal{F}_t and Alvarez et al. (2016)’s statistic. A clear (convex)

³²For a formal proof, please refer to Caballero and Engel (2007).

³³Both papers show that the markup has displayed only a modest increase in the 1996-2007 period, while increasing substantially afterwards.

³⁴The latter is constructed from the frequency of price changes and a quantilic measure of kurtosis, $\frac{q_{90,t} - q_{62.5,t} + q_{37.5,t} - q_{10,t}}{q_{75,t} - q_{25,t}}$ (see, e.g., Groeneveld, 1998), both computed for each month of the sample.

Figure 5: COMPARISON WITH ALVAREZ ET AL. (2016)



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


negative relationship emerges, despite the fact that in theory important differences between the two can often arise, as underlined by Alvarez et al. (2020). However, the two statistics are not directly comparable, as one is measuring the *instantaneous* pass-through of nominal shocks to prices, whereas the other focuses on the *cumulative* impact of nominal shocks to output. In fact, it can well be the case that a shock has low impact on prices, taking long time to be fully absorbed, and leading to large cumulative output response. To account for this we compute the cumulated response of inflation over the 18 months following a one-off 1% nominal shock. The right panel of Figure 5 shows a striking (negative) correlation of this cumulative measure of stickiness in price adjustment with the non-neutrality statistic of Alvarez et al. (2016).³⁵ This finding reinforces our confidence in the empirical framework we use to identify shifts in the price gap distribution and the hazard function. The next section extensively examines the connection between time variation in price flexibility and inflation dynamics.

5 State dependence in inflation dynamics

Having established that price flexibility displays some large swings over the sample under examination, a natural question is whether these movements matter for our understanding of inflation dynamics. A simple exercise may serve at contextualizing the analysis on the connection between price flexibility and inflation dynamics. To this end, we employ the estimates from the Ss model to back out the response of inflation to an aggregate nominal shock in two different periods, one characterized by a relatively strong and one by a relatively weak pass-through

³⁵The time-series profile of the two measures of money non-neutrality can be appreciated in Figure I.2.

Figure 6: 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 indicates the half-life of the shock.

of nominal shocks to inflation, respectively.³⁶ Figure 6 shows that inflation is more responsive and less persistent in periods of relatively high price flexibility. In light of this, one would expect price flexibility to contain valuable information for the analysis of inflation dynamics. This message arises naturally in environments characterized by state-dependent pricing. The remainder of this section investigates whether aggregate inflation displays non-linearities that are consistent with these properties, and discusses some implications for the practice of central banking.

5.1 Price flexibility and inflation dynamics

We seek to examine how inflation generally behaves in periods of relatively high and low flexibility. To this end, we employ a regime-switching autoregressive moving average model, where the transition across regimes is a smooth function of the degree of price flexibility. The STARMA(p,q) model is a generalization of the smooth transition autoregression model proposed by Granger and Terasvirta (1993).³⁷ 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

³⁶As 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 G.

³⁷In this respect, the STARMA(p,q) model also generalizes the threshold ARMA(p,q) model (DeGooijer, 2017).

are based on a larger set of observations (Auerbach and Gorodnichenko, 2012).³⁸

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

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

with $\varepsilon_t^i \sim N(0, \sigma_i^2)$ for $i = \{L, H\}$. Moreover, we set $G\left(\tilde{\mathcal{F}}, \gamma\right) = (1 + e^{-\gamma\tilde{\mathcal{F}}})^{-1}$, where $\tilde{\mathcal{F}}$ denotes the normalized flexibility index and γ is the speed of transition across regimes.³⁹ 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.⁴⁰

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 non-linear 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 $G\left(\tilde{\mathcal{F}}_{t-1}; \gamma\right) > 0.8$ ($G\left(\tilde{\mathcal{F}}_{t-1}; \gamma\right) < 0.2$).⁴¹ Thus, based on the Akaike criterion, we choose $p = 1$ and $q = 7$.⁴²

Figure 7 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. Consistent with the theoretical impulse-response functions reported in Figure 6, inflation is way 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

³⁸Estimating 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.


³⁹We 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.

⁴⁰See Appendix H for further details.

⁴¹Figure I.3 in Appendix I reports the dynamics of $G\left(\tilde{\mathcal{F}}_{t-1}; \gamma\right)$. 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 .

⁴²Note that the modified AIC information criterion indicates a STARMA(1,3). Figures I.4 and I.5 in Appendix I report the results for this alternative setting. Our key insights are not affected by the exact specification of the STARMA(p,q) model.

Figure 7: 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 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 indicates the half-life of the shock.

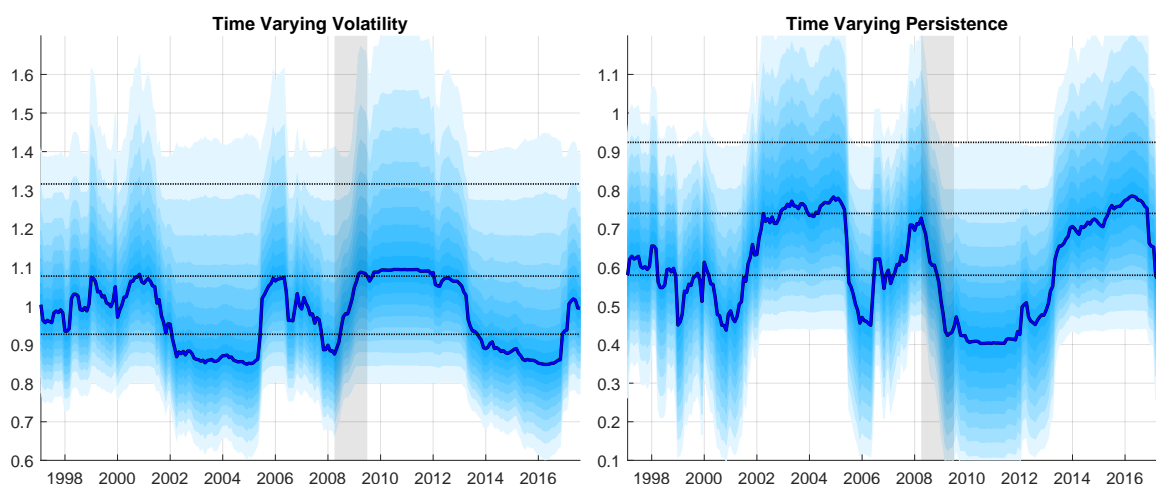
regime. Moreover, the impulse-response function for the low flexibility regime is remarkably close to the one we would estimate by imposing a constant inflation-dynamics regime over the entire sample. This evidence has a main takeaway: neglecting this type of non-linearity implies that one might well end up overestimating inflation persistence.

While the overall level of inflation will ultimately depend on the magnitude of the shocks hitting the economy and their severity, the degree of price flexibility may play an important role in affecting their propagation. Given our impulse-response evidence, it is important to investigate whether variations in price flexibility can account for meaningful changes in inflation dynamics. In this respect, using the popular statistical model of Stock and Watson (2007), Forbes et al. (2018) document a number of stylized facts about inflation dynamics in the UK. In particular, they highlight that inflation in the UK has shown a relatively high degree of volatility and low persistence in the 2008-2012 period and, to a lesser extent, around the early 2000s.⁴³ These periods have been associated with large departures of inflation from the target. Notably, our estimates of price flexibility attains relatively high values in both periods.

Figure 8 reports the estimates of inflation volatility and persistence—with the latter being measured as the one-year-ahead response of inflation to a 1% shock—as resulting from the STARMA(1,7) model. Inflation volatility is relatively high during periods associated with a fast mean reversion of the shocks. Consistent with the evidence of Forbes et al. (2018), these periods are also associated with high-volatility and a somewhat reliable predictability of inflation. By contrast, inflation volatility is particularly low, and so persistence, between 2014 and 2016. This is a prolonged period in which year-on-year inflation has been substantially

⁴³Volatility is measured by standard deviation of the mean reverting component of their model of inflation.

Figure 8: TIME VARYING VOLATILITY AND PERSISTENCE



Note: This figure reports the the implied time-varying volatility and the persistence of inflation retrieved from the STARMA(1,7) model. Time-varying persistence is measured as the 1 year-ahead response of inflation to a 1% shock. In both cases, we also report the response from a (linear) ARMA(1,7) model (in black), together with the associated 68% confidence intervals. These are built based on the Markov Chain Monte Carlo (MCMC) method developed in Chernozhukov and Hong (2003).

low, diving below zero for the first time in the post-WWII sample. Various analyses conducted by the Bank of England clearly highlight that weak inflation over this period reflects the fall in the oil price and the weakness of the pound.⁴⁴ Our analysis suggests that relatively low price flexibility might have contributed to prolonging this period even further. The next subsection hinges on this to explain why the Bank of England and professional forecasters at large have been displaying a certain belief that the impact of these inflationary shocks would have lasted substantially less than it eventually did.

5.2 State dependence and inflation projections

The impulse-response function from the linear model is remarkably close to that in the low-flexibility regime (see, again, Figure 7). As a direct implication, neglecting that shocks are propagated at different speeds—depending on the overall degree of price flexibility—may 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 non-linear 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? 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)

⁴⁴See, e.g., the Inflation Report published on February 12, 2015.

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.⁴⁵ These are then regressed on a non-linear 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:

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

where $\mathbb{1}_{\{G_{t-1} > 0.5\}}$ is an indicator function taking value 1 when $G_{t-1} > 0.5$ and zero otherwise. This function represents 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 3 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 7, 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.

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. 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 that time frame. 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 aftermath of the Brexit

⁴⁵Table I.1 in Appendix I 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.

Table 3: 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: $e_{t+h|t} = a_0 + a_1 G_{t-1} + a_2 G_{t-1}^2 + a_3 \mathbb{1}_{\{G_{t-1} > 0.5\}} G_{t-1}^2$, where $\mathbb{1}_{\{G_{t-1} > 0.5\}}$ is an indicator function taking value 1 when $G_{t-1} > 0.5$ and zero otherwise. 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 .

referendum.

6 Concluding remarks

Looking at UK micro price data, we document distinctive patterns of time variation in some key moments of the underlying process of price adjustment. A key implication of our analysis is that imposing a Calvo price-setting protocol to match the frequency of adjustment would 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. Our work 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. In doing so, we point to the importance of allowing for asymmetry and time variation in the hazard function, a

long with the distribution of price gaps. In this respect, more research should be devoted to understanding the sources of such time variation, and to what extent this is connected with firm dynamics, the degree of market concentration, and other relevant microeconomic and macroeconomic features.

Importantly, we highlight a marked non-linearity in the price response to inflationary shocks, and show how this is crucially dictated by the degree of price flexibility. Neither the Bank of England nor professional forecasters appear to account for state dependence when forecasting CPI inflation. In fact, both forecasters 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. Reliable proxies of aggregate price flexibility can easily be constructed from timely available micro prices, and may successfully be employed to improve inflation projections.

On a more general note, this study bears important implications for monetary authorities seeking to stabilize inflation. In a relatively low price-flexibility environment, nominal shocks are likely to dissipate slowly, so that the Central Bank might struggle to keep inflation at its target. By contrast, under relatively high price flexibility the same shock would transmit proportionally more to quantities than prices. It is plausible to expect that such state dependence would map into the trade-off that central banks typically face when stabilizing output and inflation. Although this message arises naturally from state dependent models of price setting, we do not see much of this being picked-up in the practice and narrative of central banks. In this respect, the present work highlights the practical importance of keeping track of the real-time evolution of money non-neutrality, so as to have a reliable appreciation of the evolving inflation persistence and the pass-through of nominal shocks, when setting monetary policy.

References

- ALTONJI, J. G. AND L. M. SEGAL (1996): “Small-Sample Bias in GMM Estimation of Covariance Structures,” *Journal of Business & Economic Statistics*, 14, 353–366.
- ALVAREZ, F., H. LEBIHAN, AND F. LIPPI (2016): “The Real Effect of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach,” *American Economic Review*, 106, 2817–2851.
- ALVAREZ, F. AND F. LIPPI (2014): “Price Setting With Menu Cost for Multiproduct Firms,” *Econometrica*, 82, 89–135.
- ALVAREZ, F. E., F. LIPPI, AND A. OSKOLKOV (2020): “The Macroeconomics of Sticky Prices with Generalized Hazard Functions,” Working Paper 27434, National Bureau of Economic Research.
- ALVAREZ, L. J., E. DHYNE, M. HOEBERICHTS, C. KWAPIL, H. L. BIHAN, P. LÜNNEMANN, F. MARTINS, R. SABBATINI, H. STAHL, P. VERMEULEN, AND J. VILMUNEN (2006): “Sticky Prices in the Euro Area: A Summary of New Micro-Evidence,” *Journal of the European Economic Association*, 4, 575–584.
- AUERBACH, A. J. AND Y. GORODNICHENKO (2012): “Measuring the Output Responses to Fiscal Policy,” *American Economic Journal: Economic Policy*, 4, 1–27.
- BARRO, R. J. (1972): “A Theory of Monopolistic Price Adjustment,” *Review of Economic Studies*, 39, 17–26.
- BELL, T. AND D. TOMLINSON (2018): “Is everybody concentrating? Recent trends in product and labour market concentration in the UK,” Briefing, Resolution Foundation.
- BERARDI, N., E. GAUTIER, AND H. L. BIHAN (2015): “More Facts about Prices: France Before and During the Great Recession,” *Journal of Money, Credit and Banking*, 47, 1465–1502.
- BERGER, D. AND J. VAVRA (2017): “Dynamics of the U.S. Price Distribution,” *European Economic Review*, Forthcoming.
- 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 (1993a): “Microeconomic Adjustment Hazards and Aggregate Dynamics,” *The Quarterly Journal of Economics*, 108, 359–383.
- (1993b): “Microeconomic rigidities and aggregate price dynamics,” *European Economic Review*, 37, 697–711.
- (2007): “Price stickiness in Ss models: New interpretations of old results,” *Journal of Monetary Economics*, 54, 100–121.
- CALVO, G. A. (1983): “Staggered prices in a utility-maximizing framework,” *Journal of Monetary Economics*, 12, 383–398.
- CAPLIN, A. S. AND D. F. SPULBER (1987): “Menu Costs and the Neutrality of Money,” *The Quarterly Journal of Economics*, 102, 703–725.
- CARVALHO, C. AND O. KRYVTSOV (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.
- DELOECKER, J. AND J. EECKHOUT (2018): “Global Market Power,” NBER Working Papers 24768, National Bureau of Economic Research.
- DIXIT, A. (1991): “Analytical Approximations in Models of Hysteresis,” *Review of Economic Studies*, 58, 141–151.
- DIXON, H. D. AND H. LEBIHAN (2012): “Generalised Taylor and Generalised Calvo Price and Wage Setting: Micro-evidence with Macro Implications,” *Economic Journal*, 122, 532–554.
- DIXON, H. D., K. B. LUINTEL, AND K. TIAN (2014): “The impact of the 2008 crisis on UK prices: what we can learn from the CPI microdata,” Cardiff Economics Working Papers E2014/7, Cardiff University, Cardiff Business School, Economics Section.

- DIXON, H. D. AND K. TIAN (2017): “What We can Learn About the Behaviour of Firms from the Average Monthly Frequency of Price-Changes: An Application to the UK CPI Data,” *Oxford Bulletin of Economics and Statistics*, 79, 907–932.
- DORSEY, R. E. AND W. J. MAYER (1995): “Genetic Algorithms for Estimation Problems with Multiple Optima, Nondifferentiability, and Other Irregular Features,” *Journal of Business & Economic Statistics*, 13, 53–66.
- DOTSEY, M. AND R. G. KING (2005): “Implications of state-dependent pricing for dynamic macroeconomic models,” *Journal of Monetary Economics*, 52, 213–242.
- EICHENBAUM, M., N. JAIMOVICH, AND S. REBELO (2011): “Reference Prices, Costs, and Nominal Rigidities,” *American Economic Review*, 101, 234–262.
- FORBES, K., L. KIRKHAM, AND K. THEODORIDIS (2018): “A Trendy Approach to UK Inflation Dynamics,” CEPR Discussion Papers 12652, C.E.P.R. Discussion Papers.
- GAGNON, E. (2009): “Price Setting during Low and High Inflation: Evidence from Mexico,” *The Quarterly Journal of Economics*, 124, 1221–1263.
- GELMAN, A., D. DUNSON, D. RUBIN, H. S. STERN, AND J. B. CARLIN (2004): *Bayesian Data Analysis*, Chapman & Hall/CRC Texts in Statistical Science.
- GERTLER, M. AND J. LEAHY (2008): “A Phillips Curve with an Ss Foundation,” *Journal of Political Economy*, 116, 533–572.
- GOLOSOV, M. AND R. E. LUCAS (2007): “Menu Costs and Phillips Curves,” *Journal of Political Economy*, 115, 171–199.
- GRANGER, C. W. J. AND T. TERASVIRTA (1993): *Modelling Non-Linear Economic Relationships*, Oxford University Press.
- GROENEVELD, R. A. (1998): “A Class of Quantile Measures for Kurtosis,” *The American Statistician*, 52, 325–329.
- HAMILTON, J. D. (2018): “Why You Should Never Use the Hodrick-Prescott Filter,” *The Review of Economics and Statistics*, 100, 831–843.
- HARVEY, A. C. (1990): *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press.
- HOBijn, B., F. RAVENNA, AND A. TAMBALOTTI (2006): “Menu Costs at Work: Restaurant Prices and the Introduction of the Euro,” *The Quarterly Journal of Economics*, 121, 1103–1131.
- KARADI, P. AND A. REIFF (2019): “Menu Costs, Aggregate Fluctuations, and Large Shocks,” *American Economic Journal: Macroeconomics*, 11, 111–146.

- KEHOE, P. AND V. MIDRIGAN (2015): “Prices are sticky after all,” *Journal of Monetary Economics*, 75, 35–53.
- KLENOW, P. J. AND O. KRYVTSOV (2008): “State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?” *The Quarterly Journal of Economics*, 123, 863–904.
- KOMUNJER, I. (2007): “Asymmetric power distribution: Theory and applications to risk measurement,” *Journal of Applied Econometrics*, 22, 891–921.
- KRYVTSOV, O. AND N. VINCENT (2017): “The Cyclicalities 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.
- 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.
- PETRELLA, I., L. DE LA PORTE SIMONSEN, AND E. SANTORO (2020): “Trends in markups and the slope of the Phillips curve,” Tech. rep.
- ROTEMBERG, J. J. (1999): “A Heuristic Method for Extracting Smooth Trends from Economic Time Series,” NBER Working Papers 7439, National Bureau of Economic Research, Inc.

STOCK, J. H. AND M. W. WATSON (2007): “Why Has U.S. Inflation Become Harder to Forecast?” *Journal of Money, Credit and Banking*, 39, 3–33.

——— (2019): “Slack and Cyclically Sensitive Inflation,” NBER Working Papers 25987.

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

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.

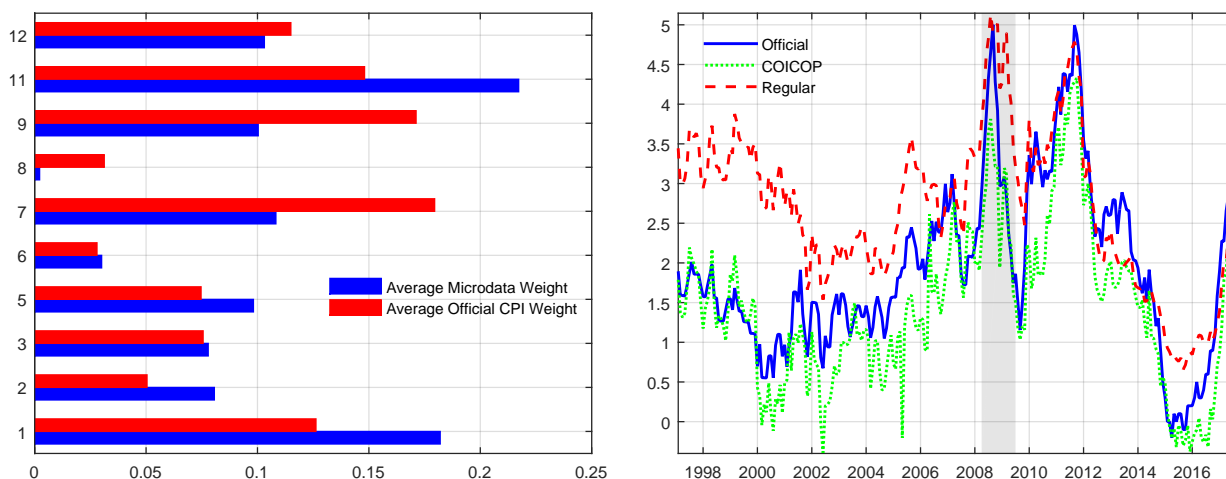
Email: l.simonsen@mail.bbk.ac.uk

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

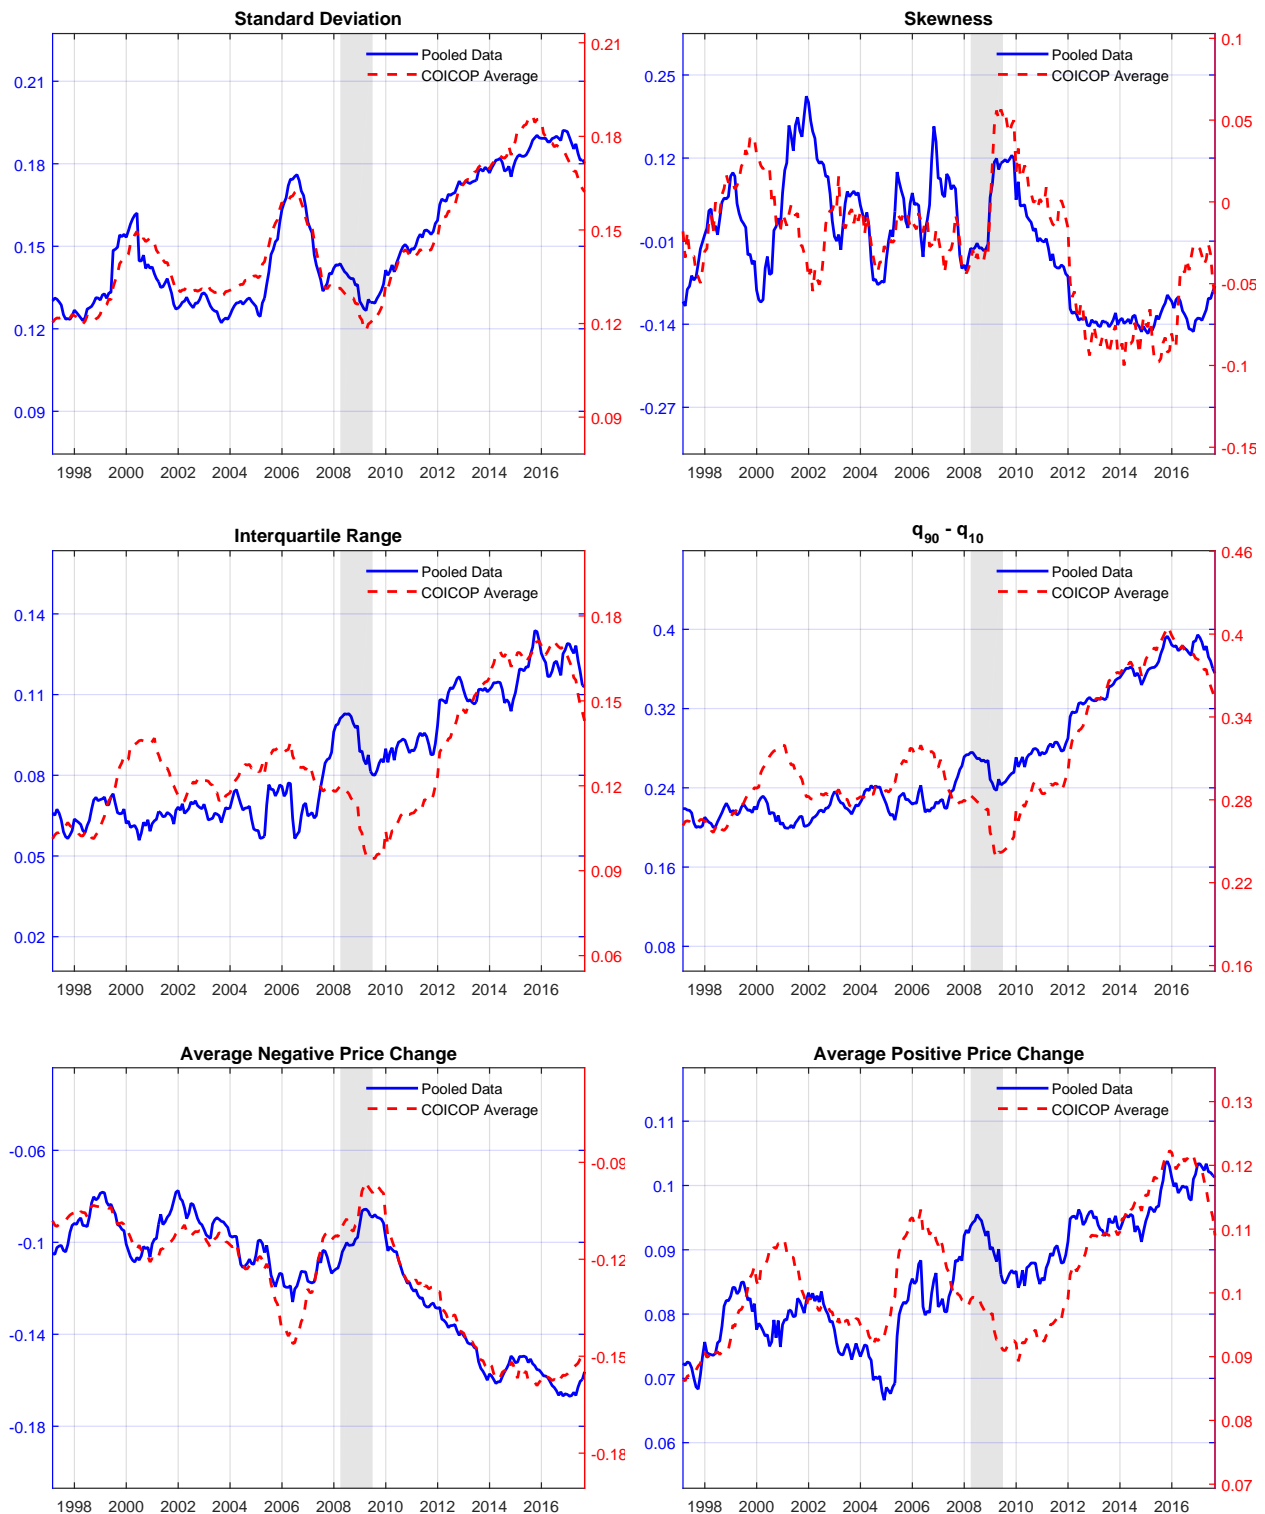


Notes: The left panel contrasts the weights assigned to each of the COICOP sectors to those assigned to build the CPI (re-normalized to exclude COICOP 4 and 10). The right panel reports the official CPI inflation, together with the inflation series retrieved from all the available price quotes (labeled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

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

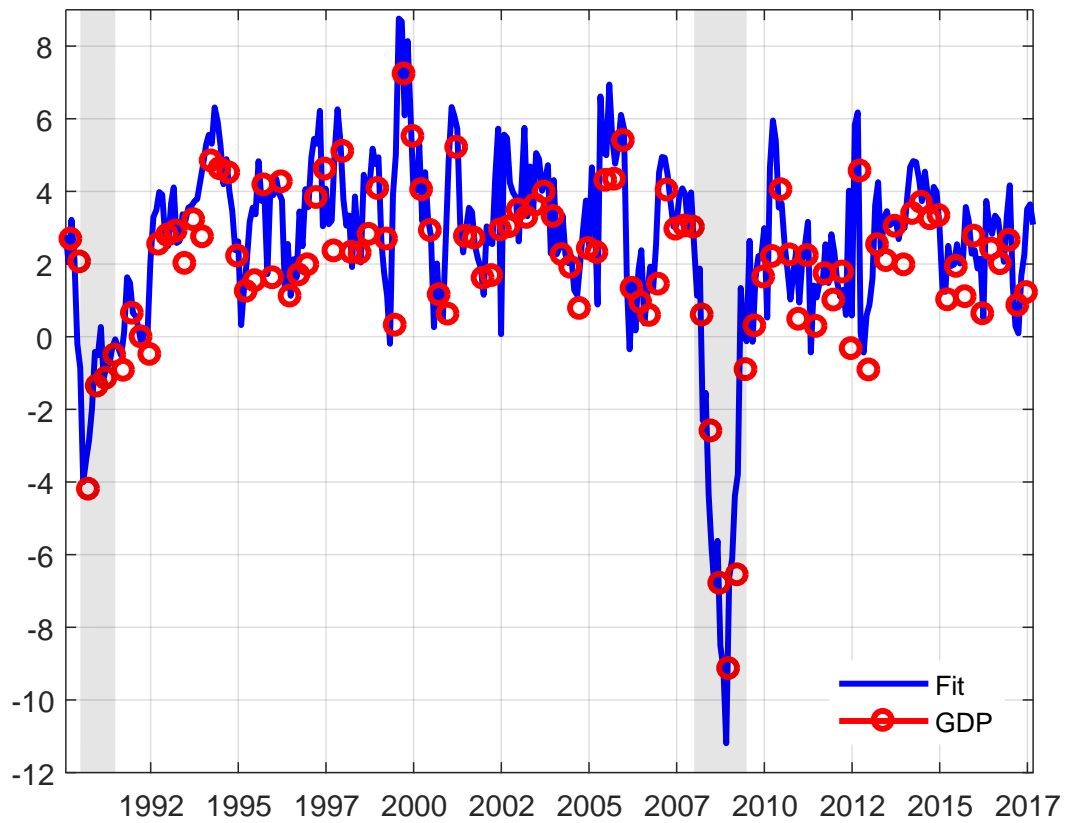
where $\hat{\alpha}_\tau$ is retrieved by using the Kalman smoother. The business cycle indicator is then computed by applying a simple filter to z_t . For the baseline results in the paper we use the Rotemberg (1999) version of the HP filter, which chooses the smoothing coefficient of the HP filter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



Note: The figure shows the fit of the (monthly) coincident indicator on the (annualized) quarter-on-quarter growth of real GDP.

D A simple analytical setting to frame the stylized facts

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

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

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

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

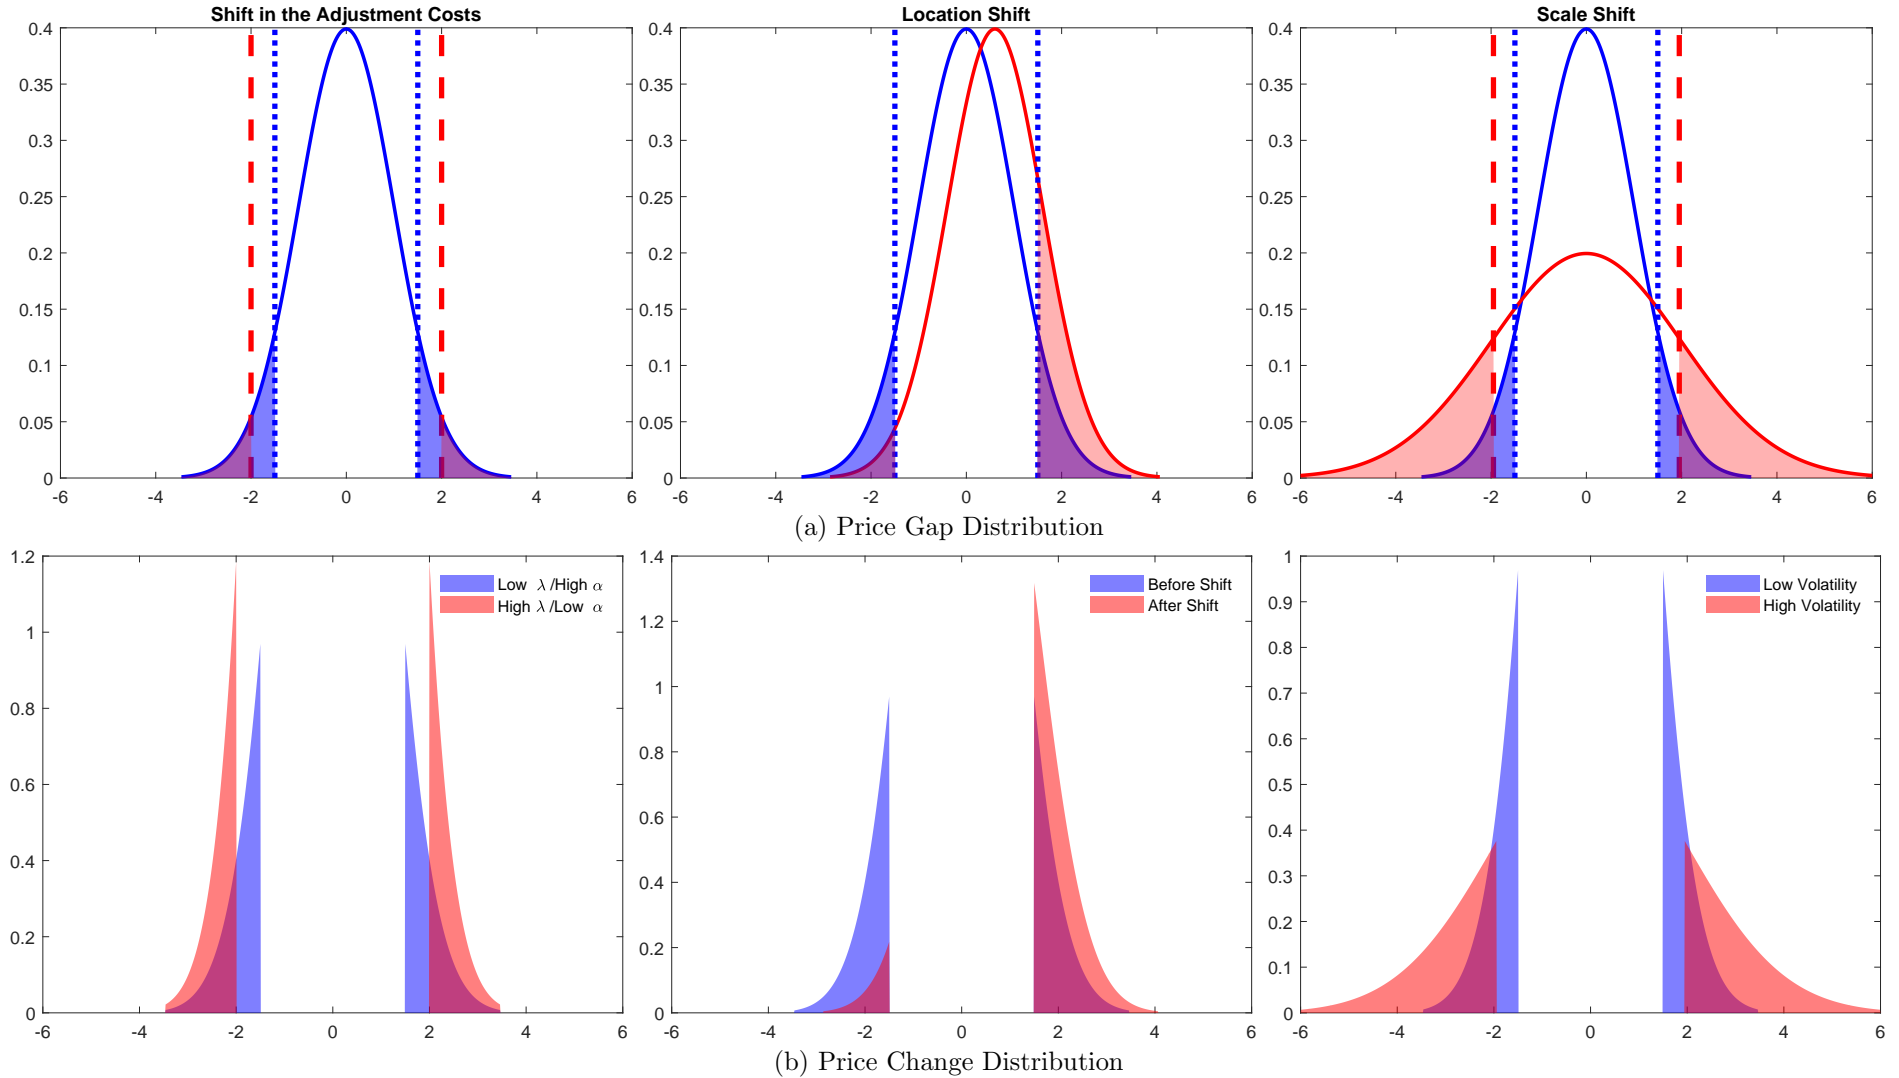
Vavra (2014) points to second-moment shocks as potential drivers of the positive comovement between the frequency of adjustment and price-change dispersion in U.S. CPI data. It is clear how this type of shock would not be suitable to rationalize negative comovement. In fact, only an increase in the fixed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

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

³In 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).

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

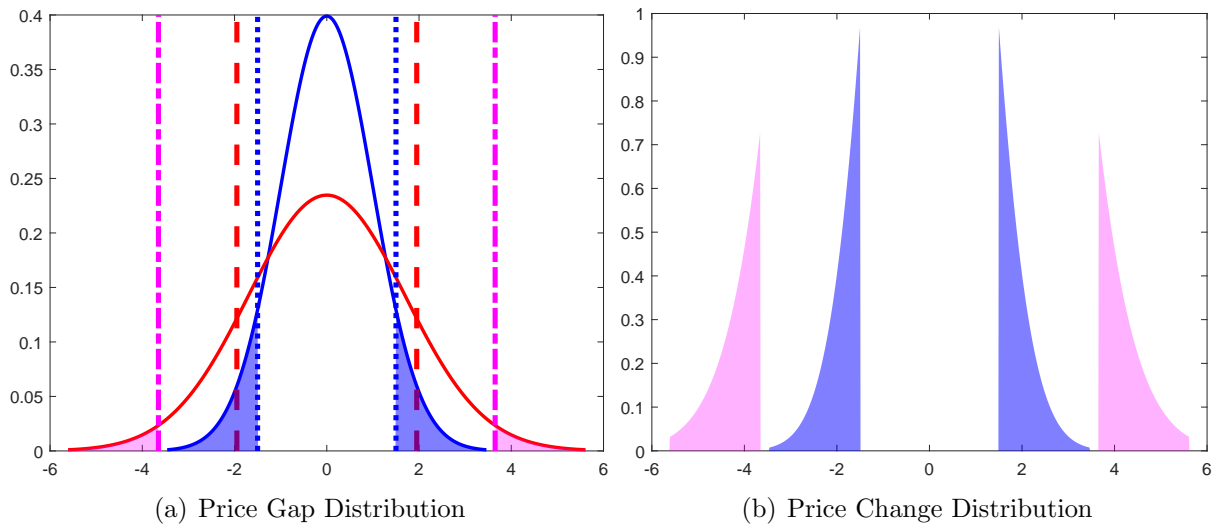
Figure D.1: ANALYTICAL FRAMEWORK



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.

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

Figure D.2: A COMBINED INCREASE IN ϕ AND λ

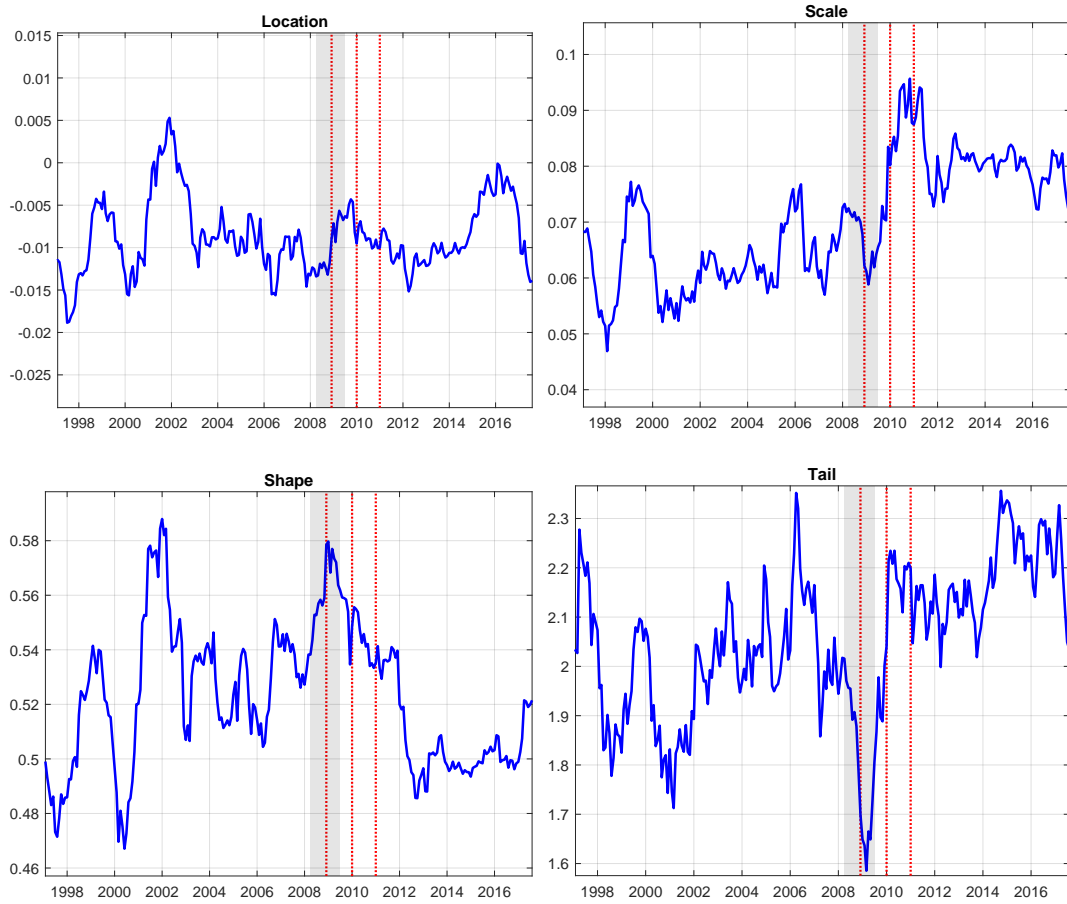


Note: We consider a positive shift in λ that affects the inaction region (while leaving the distribution of price gaps unaffected), combined with an increase in the dispersion of the distribution of price gaps, ϕ . The left panel reports the transformations occurring to the distribution of price gaps 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 .

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


E Model estimates

Figure E.1: PARAMETERS OF THE PRICE GAP DISTRIBUTION



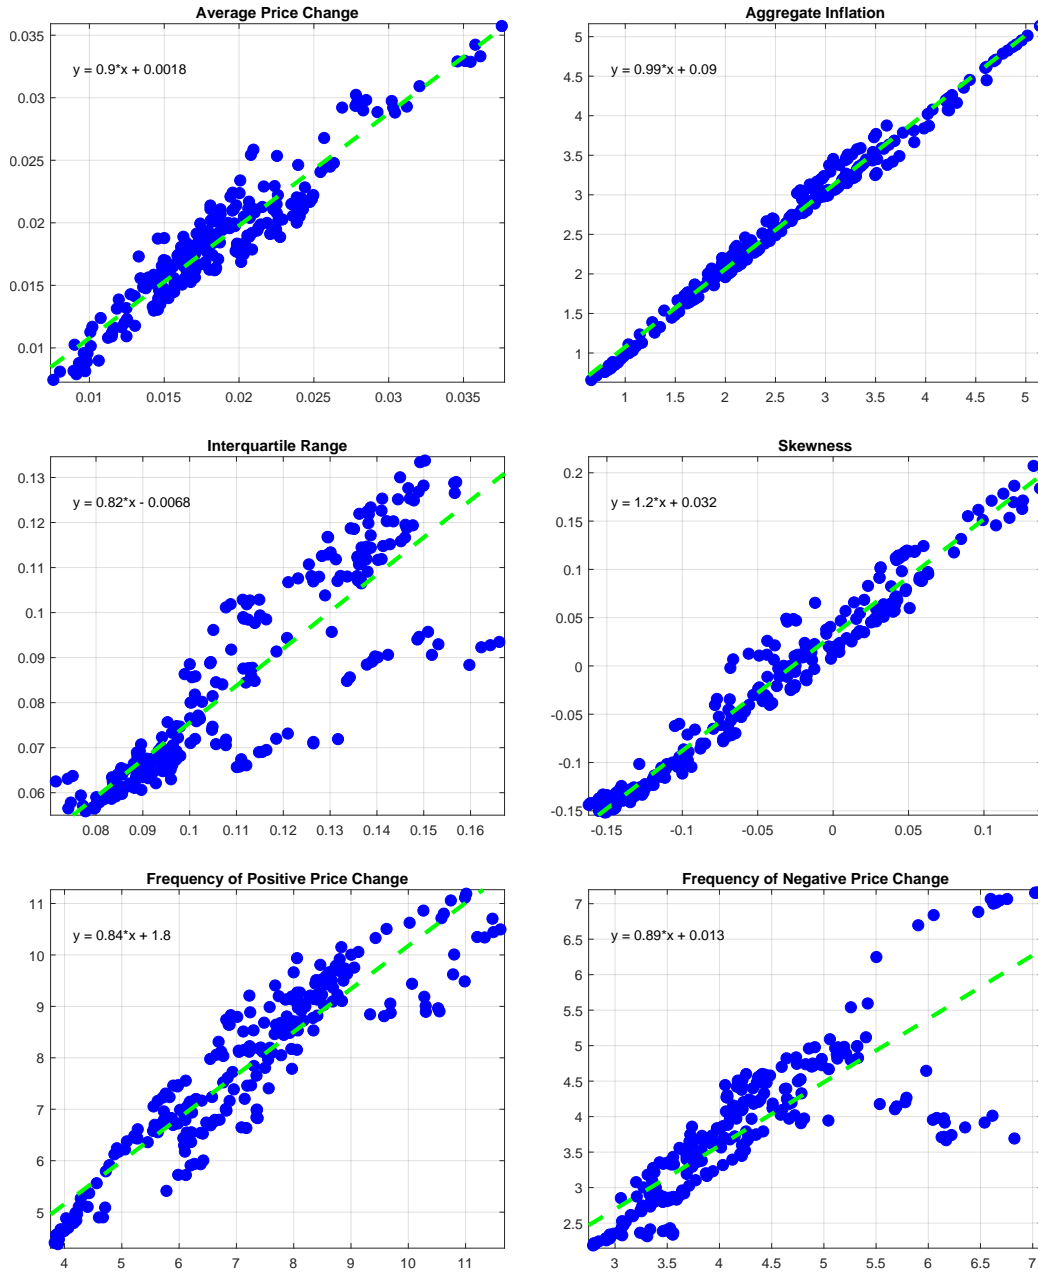
Note: The red lines denote the three VAT changes in the sample. The shaded vertical band indicates the duration of the Great Recession.

Figure E.2: 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.3: FIT OF THE S_s MODEL (SELECTED MOMENTS)



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

F Model identification

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

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

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

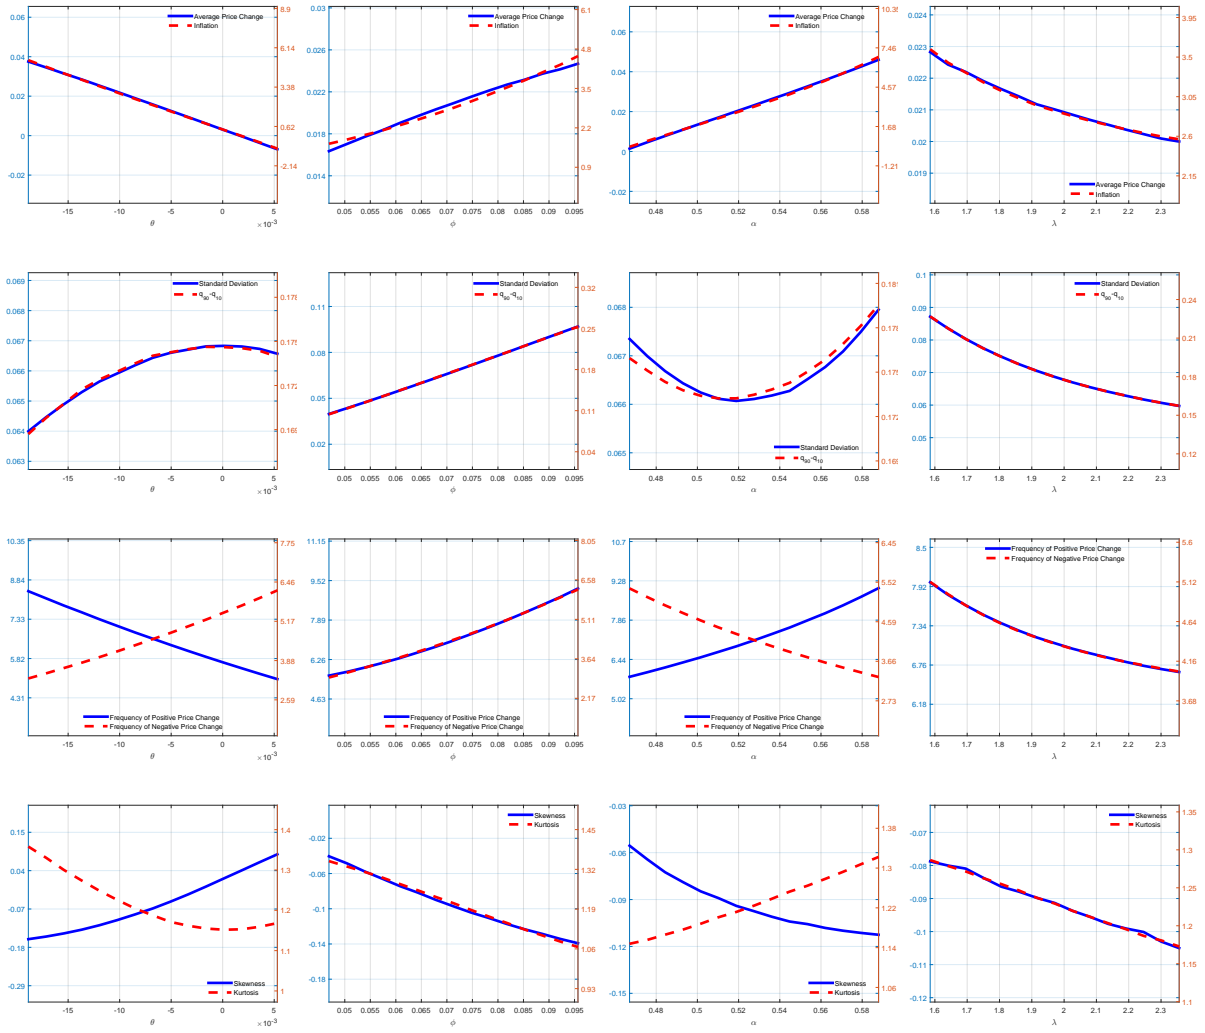
All parameters have an impact on a number of moments, and in the expected direction. For instance, increasing the scale (tail) parameter of the price gap distribution increases (decreases) monotonically the implied dispersion of the distribution of (non-zero) price changes, and in both cases decreases the skewness and the kurtosis. Instead, changing the location or the shape parameter has an opposite impact on skewness and kurtosis, and affects non-monotonically the dispersion (with higher dispersion obtained for a more skewed distribution, regardless of the sign of the skewness). As for the parameters of the hazard function, changing the constant term affects equally the frequency of price adjustment, whereas changes in the slope for positive (negative) price gaps impacts the frequency of negative (positive) price changes and the average negative (positive) price changes, leaving invariate the positive (negative) side. These results confirm the observation of Berger and Vavra (2017) for the specific functional forms of the price gap distribution and the hazard function we employ.

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

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

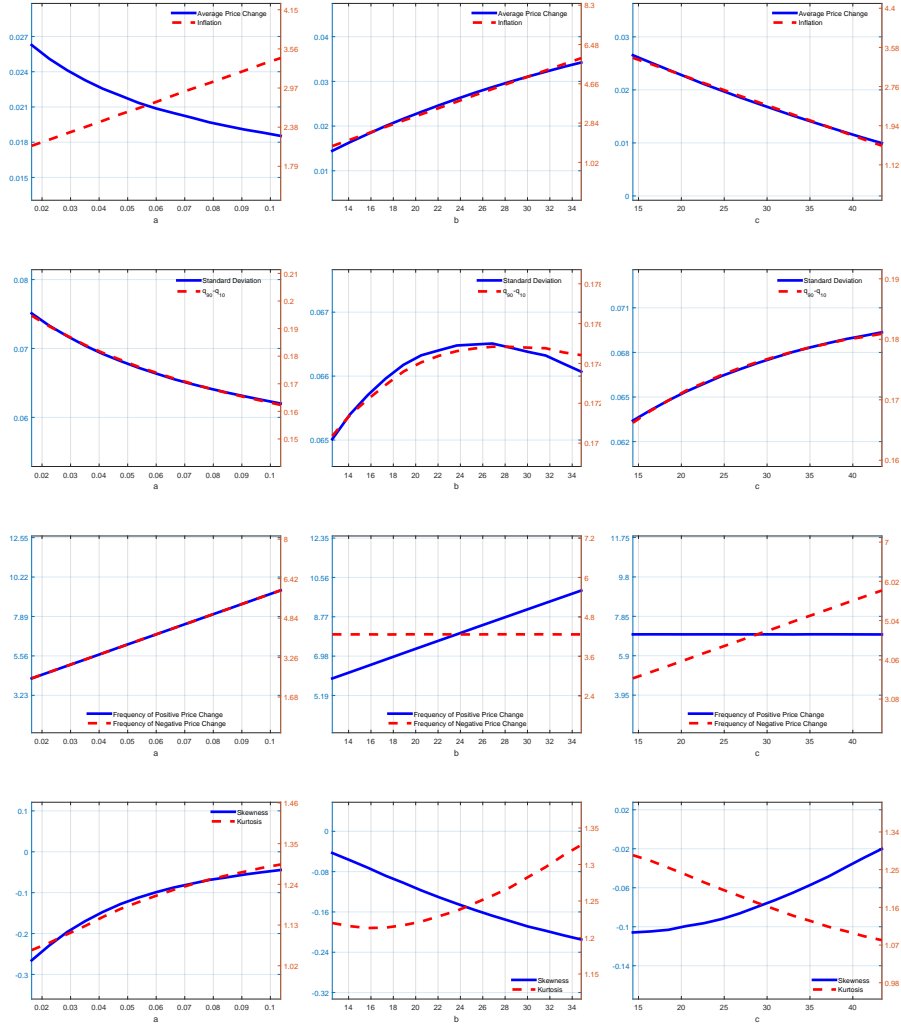
It is also important to stress that Berger and Vavra (2017) produce a battery of exercises in support of our approach. Most importantly, they address how well the resulting measure of price flexibility—which only captures the impact response of prices to a nominal shock—reflects overall non-neutrality. To this end, they estimate simulated data from the CalvoPlus model of Nakamura and Steinsson (2008), and report close comovement between the impact response from the structural model and the estimated index of price flexibility from the accounting framework. Notably, this exercise also addresses the criticism towards estimating the generalized Ss model in every period, as if observations were independent across time. In this respect, we should stress that standard structural frameworks tend to impose a rather tight relationship between distributions at a given point in time and how they evolve. In line with our predecessors, we claim that imposing flexible functional forms within a period—in a way that represents an intermediate step between a fully structural approach and a non parametric one—allows us to exploit valuable information, in the perspective of studying time variation in aggregate price flexibility.

Figure F.1: IDENTIFICATION AND THE PARAMETERS OF $f_t(x)$



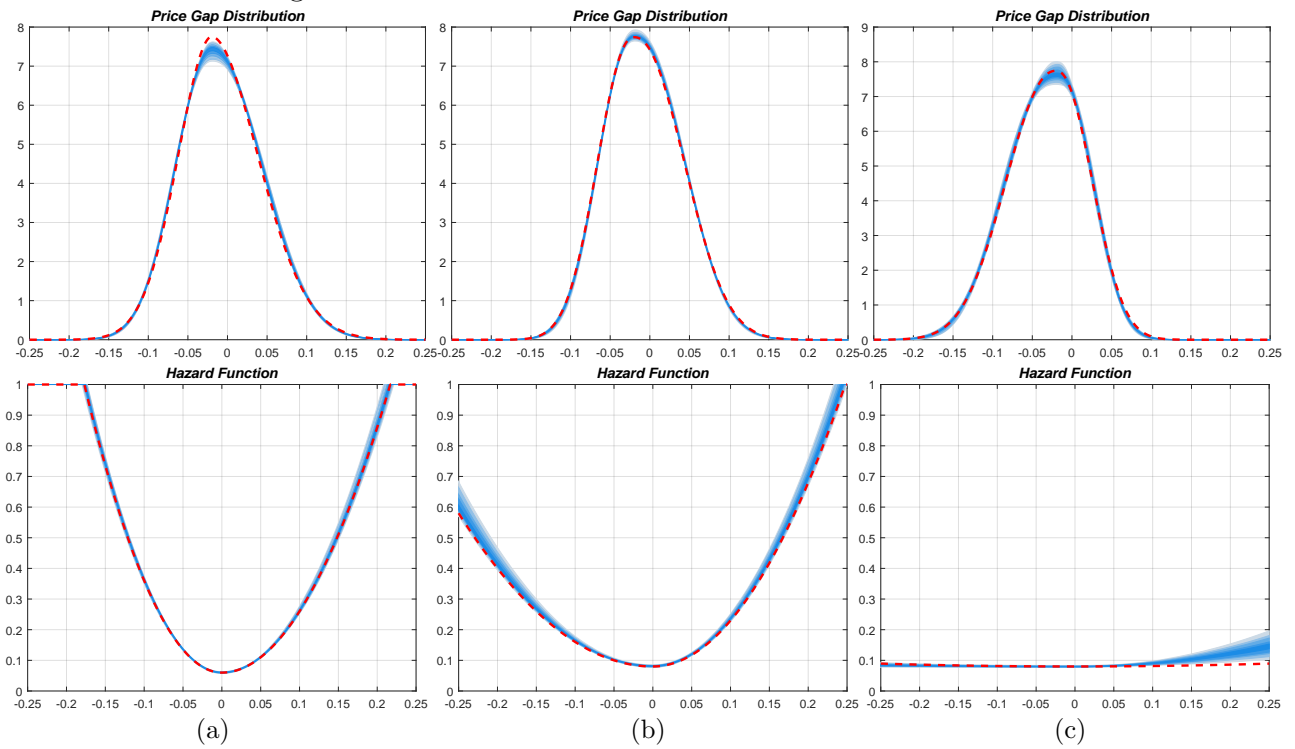
Notes: In each panel, we vary one of the parameters of $f_t(x)$ at the time—while keeping the other coefficients at their baseline estimate—and report its effect on key moments of the price change distribution, as well as the resulting rate of inflation.

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



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

Figure F.3: MODEL SIMULATIONS AND EMPIRICAL FIT



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

G Details on the computation of the impulse response function from the S_s model

This appendix gives a brief account of how we compute the impulse response functions from the generalized S_s model presented in Section 3. We start by specifying a process for the exogenous (first-moment) shock.⁶ Specifically, we assume that:

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

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

The impulse responses are then calculated as:

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

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

H Estimation of the STARMA (p,q) model

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

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

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

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

⁶Since 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.

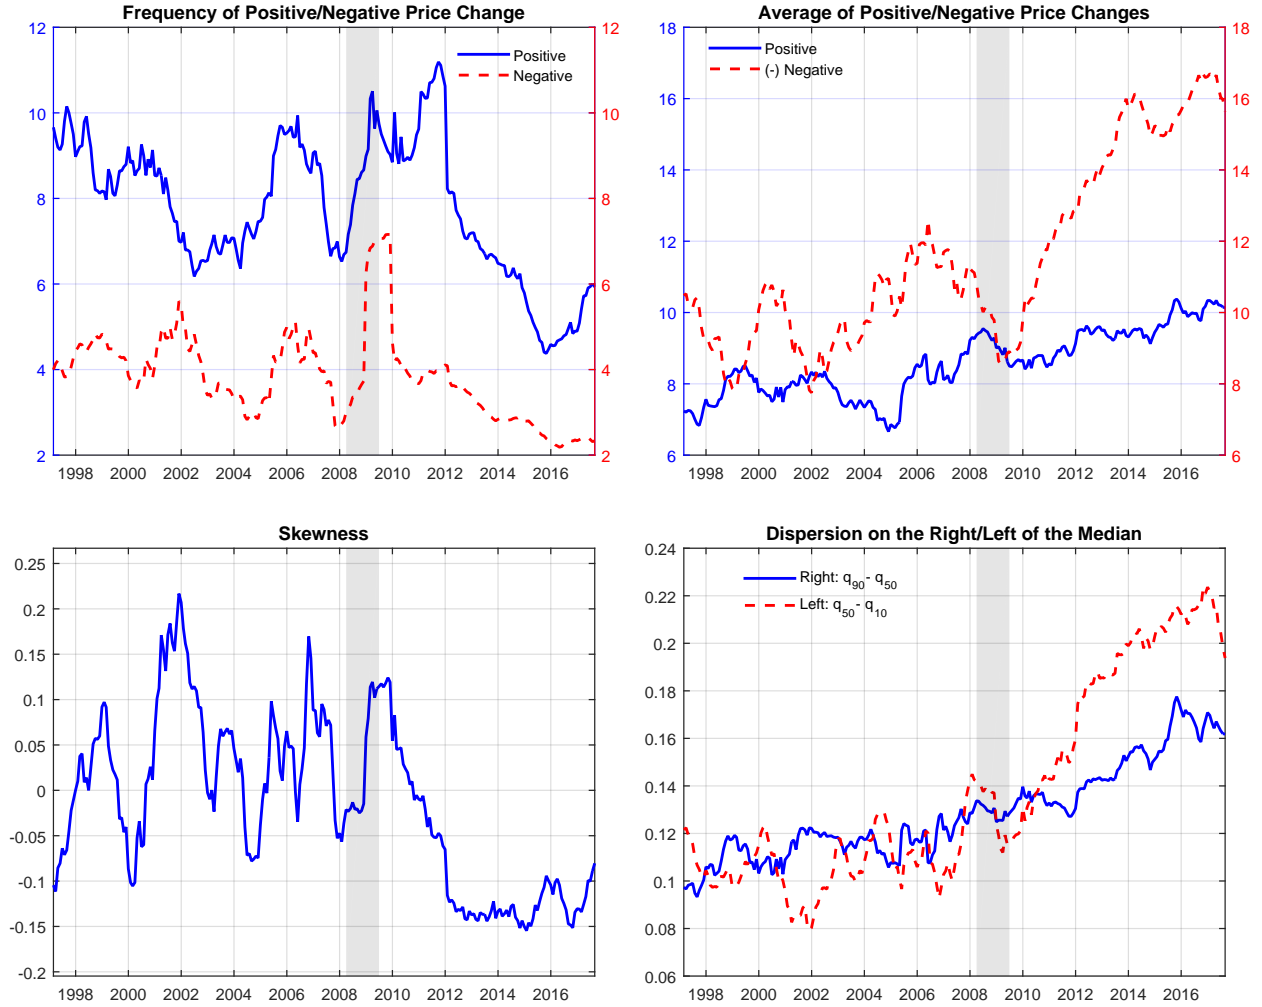
where $L(\theta)$ denotes the value of the likelihood of the model evaluated at the parameters values θ .

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.


I Additional figures and tables

Figure I.1: ADDITIONAL STATISTICS FROM MICRO PRICE DATA




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

Figure I.2: COMPARISON WITH ALVAREZ ET AL. (2016)




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

Figure I.3: PROBABILITY OF A HIGH-FLEXIBILITY REGIME




Note: The figure reports the probability of ending up in a high-flexibility regime used in the STARMA model of Section 5.1. The shaded vertical band indicates the duration of the Great Recession.

Figure I.4: PRICE FLEXIBILITY AND INFLATION PERSISTENCE



Note: Figure I.4 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 indicates the half-life of the shock.

Figure I.5: 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).

Table I.1: 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 .