Reassessing Sticky Price Models Through the Lens of Scraped Price Data

Tomás Alejandro Carrera de Souza

Primer Premio / Categoría Jóvenes Profesionales

13° Premio de Investigación Económica

"Dr. Raúl Prebisch" 2021

BANCO CENTRAL
DE LA REPÚBLICA ARGENTINA

Reassessing Sticky Price Models Through the Lens of Scraped Price Data

What micro facts of price changes should be taken into account in the incorporation of price rigidities into macro models? To answer this, I exploit a novel micro data set obtained with web scraping techniques, containing daily prices of eight retailers from six countries with heterogeneous macroeconomic conditions. I find that: (1) There is a relation between the main statistics (related to the size and frequency of price adjustment) and the inflation rate of a country; (2) The distribution of the size of price changes has a relatively small, yet nontrivial mass around zero; (3) Familiar products from the same manufacturer have greater similarity in the timing and magnitude of price adjustment than heterogeneous products. I show that incorporating a three-dimensional cost –composed by a general cost, a product-specific cost, and a cost curtail for price changes in familiar products– makes an otherwise standard menu cost model reproduce these facts.

1. Introduction

Firms' price setting behavior plays a central role in modern macroeconomics. Without nominal rigidities limiting firms' capacity to adjust prices, macro models would not display monetary non-neutrality. Moreover, strikingly different aggregate implications are predicted by macro models depending on the degree of price stickiness: more flexible prices enlarge the immediate response of the price level to money supply innovations and therefore weaken the real effect of monetary policy.

A majority of the macroeconomic literature introduces Calvo (1983) time dependent pricing as the cause of nominal price rigidity, which enables tractability and simplifies the characterization of models that aim to answer questions related to components of the general equilibrium that are not directly associated to the firms' price-setting behavior. However, the inconsistency of Calvo pricing predictions with the micro data of price changes motivated several recent papers to give close attention to the source of nominal price rigidity. Most of these papers found that state dependent pricing –where the timing of firms' price adjusting decision is endogenous of their profit maximization problem, such as menu cost models– performs better in matching the micro price facts, emerging a still active debate about the correct modelling of price stickiness and its aggregate implications for macro models [e.g. Golosov and Lucas (2007); Klenow and Kryvtsov (2008); Nakamura and Steinsson (2008, 2010); Midrigan (2011); Alvarez and Lippi (2014, 2020);

Kehoe and Midrigan (2015); Alvarez, Le Bihan, and Lippi (2016); and Alvarez, Lippi, and Paciello (2018)].

In this paper, I follow up on this discussion by providing new evidence about firms' price-setting behavior and formalizing my findings in a menu cost model. With this motivation, I exploit a new micro daily data set, which I collected with web scraping techniques from eight large multichannel consumer-goods retailers, operating in six different countries with heterogeneous macroeconomic conditions.

In the empirical part of this study, I present the salient facts observed in the data, compare them with analogous results documented in earlier papers that use traditional data sources, and discuss the implications of my findings for the characterization of nominal rigidities. The three main novel findings of this part are the following: (1) There is a relation between the main statistics (related to the size and frequency of price adjustment) and the inflation rate of a country; (2) The distribution of the size of price changes has a relatively small, yet nontrivial mass around zero; (3) Familiar products from the same manufacturer have greater similarity in the timing and magnitude of price adjustment than utterly different products.

I begin focusing on the main statistics that drew interest in previous studies of price stickiness, such as the distribution and timing of price changes, and the degree of synchronization of adjustments across products of the same firm. I show that some statistics found in my data differ from previously documented values, which is partly caused by the presence of measurement errors in traditional data sets (namely CPI and scanner data) as suggested in Cavallo (2018). I also investigate the relationship between the main facts of price changes and inflation rate, which is an advantage of the use of scraped data, since it allows me to compare data from different countries collected for similar retailers and with the same process. I show that a higher inflation rate is associated with a larger average size of price adjustments, a larger share of price changes that are increases, and with a lower duration of price spells.

Following, I study the behavior of temporary price changes (sales). Temporary changes represent, on average across retailers, 73% of all price changes in my data set, which is compatible with other studies that also report a very large share of sales [Midrigan (2011), Kehoe and Midrigan (2015)]. The large number of temporary changes is the reason why the treatment of sales is a crucial matter in the modelling of price stickiness: should we consider sales in the same way as regular price changes; differentiate regular and temporary adjustments as two different types of price changes; or exclude temporary changes and instead focus only on regular? Or, in words of Nakamura and Steinsson (2008), "is a price change just a price change?".

I provide new evidence against the relevance of sales for the degree of price stickiness in an economy. I show that most sales return to their previous value after a short period and that there is a high concentration of the duration and magnitudes of sales in a very few values of these statistics. The fact that temporary price changes automatically revert to previous value after a given period hints that retailers do not use them as a response to new macroeconomic information as they do with regular changes. This is also aligned with the idea that temporary changes are typically pre-established with substantial anticipation and defined by each retailer's marketing strategy.

Moreover, I show that many of the distributions of the size of price changes have a low density around zero, which causes a simultaneous near bimodality and normality. Also using scraped data, Cavallo (2018) reports a similar finding, which contrasts with those obtained from traditional data sources. Importantly, the shape of the distribution also differs from the predictions of the two standard pure menu costs types of models. These models predict either a clearly bimodal shape with no changes around the threshold given by the cost that has to be paid for *each* price that is changed (Golosov and Lucas, 2007); or a bell-shaped distribution with a large mass around zero, thanks to the presence of economies of scope in price adjustment (the firm pays only one fixed menu cost to change *any* number of prices) that makes it optimal to change every price that differs from its optimum no matter how small is that gap (Midrigan, 2011). An immediate relevant implication of this finding is that a model that matches the micro data has to display a selection effect of monetary policy that lies between a very strong one, such as that of Golosov and Lucas (2007), and a weak one, such as that of Midrigan (2011).

Finally, I show that while there is a low synchronization of adjustment across products within a store, the synchronization of changes in familiar products is higher than in heterogeneous products.[1] Additionally, I find that the sizes of price changes of familiar products that adjust in the same period are identical in the majority of the cases. The retailers in my data set, and also many other consumer-goods multiproduct firms, sell numerous similar products from the same manufacturer. Hence, whether the price series of familiar products follow parallel patterns is a relevant detail to account for when integrating menu costs in macro models. In particular, I suggest that in order to reproduce this fact from the data, menu cost models with any type of product-specific cost of adjustment have to incorporate a cost curtail when a firm changes multiple familiar products.

The theoretical part of this paper builds on the main facts of price changes that I find in the micro data. I set up, calibrate, and solve the partial equilibrium of an extension of the Midrigan (2011) multiproduct menu cost model, with only regular price changes. The fundamental deviation of the model from standard price-setting theories of price adjustment is the introduction of a multidimensional cost of price adjustment.

First, the firm faces a “general” cost, ϕ^G, that has to be paid once and for all for changing *any* number of prices in a period and it is independent of the number of changed prices. This cost is analogous to the one paid for regular price changes in the Midrigan model, leading to economies of scope in price adjustment.

The second cost of price adjustment is a “product-specific” cost, ϕ^S, paid for every price that is changed. This product-specific cost rationalizes the fact that, unaffected by the number of changes, the marginal change is always costly and hence prevents my model to generate a large number of small price changes and perfect synchronization in price adjustment. Finally, the firm receives a cost curtail, ϕ^C, when it changes familiar products’ prices. This implies that the model features a high degree of economies of scope in price adjustment of similar products.

For the solution of the model, I calibrate the key parameters with the same value for all the

[1] I define a family of products as those groups of similar products obtained from the same manufacturer by the consumer-goods firm (for example, Coke Light 375ml, Coke Light 600ml, and Coke Light 1500ml).

countries available in my data, as I try to match the patterns found in the micro facts *across* countries, instead of making the model fit the facts of *each* country. I show that incorporating only one country-specific variable –the inflation rate– the model reproduces many of the facts I document, including the relationship between inflation and the main statistics, and the relatively low yet nontrivial number of small price changes.

Lastly, my model encompasses as two special cases the standard versions of uni-dimensional cost of regular price adjustment. A model with only product-specific menu costs (a multiproduct version of Golosov and Lucas (2007)) generates a bimodal distribution of the size of price changes, with null changes between the positive and negative threshold given by the size of the cost. To reproduce this type of model, $\phi^G = 0$ in my calibration. Contrarily, a model where the firm has to pay only one general cost to change any number of regular prices (à la Midrigan (2011) excluding temporary changes) predicts a bell-shaped distribution of the size of price changes, with a large mass around zero. To reproduce this type of model, I set $\phi^S = 0$ in my calibration. As Figure 1 shows, the baseline case of my model reproduces the key features of the distribution from the data in a better way than the two special cases with uni-dimensional menu costs. This suggests the necessity of incorporating a general and a product-specific cost of price adjustment in order to make an otherwise standard menu cost model fit the micro facts of price changes.

Figure 1: **Data and model simulations for the Netherlands**

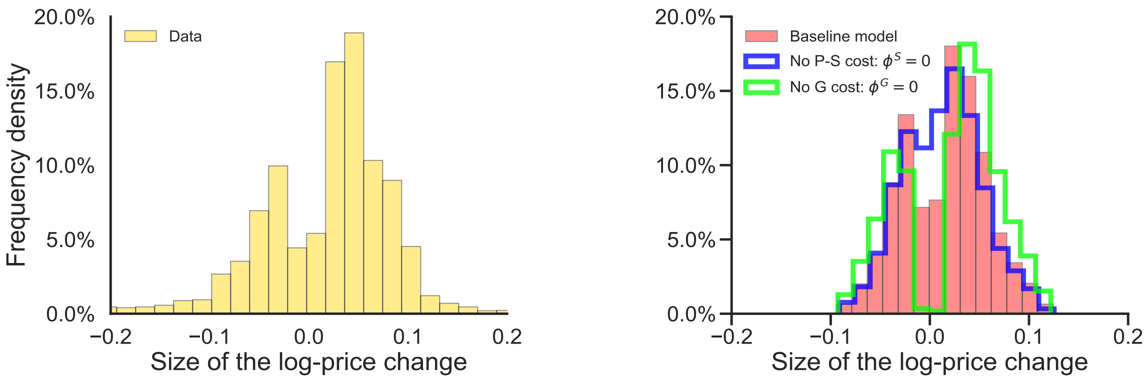

Note: The left-panel of the figure plots the distribution of the size of log-price changes found in daily data from a large retailer from the Netherlands. The right-panel plots the predicted distribution of price changes of the baseline case of the model I set up in this paper and of two special cases of it, one with only a general menu cost and the other with only a product-specific menu cost of price adjustment.

1.1. Relation to the literature on the microfoundations of nominal rigidities

This paper relates to a strand of the literature that investigates the sources of nominal rigidities and their role in macroeconomic models. Over the past decade and a half, and thanks to the availability of new micro data sets of prices, several papers exploited these new sources aiming to provide a better understanding of the characteristics of firm's price setting behavior. Bils and Klenow (2004) seminal paper is a cornerstone in this literature since they were the firsts in using data underlying the consumer price index from the United States (CPI data) to present evidence on price stickiness, taking advantage of a much broader set of information (both in terms of the number of products and products categories) than earlier papers which tended to focus on the evolution of the price of a narrow set of products, raising concerns about the representativeness of the results.

Since Bils and Klenow (2004) there was a surge in the discussion on the microfoundation of price rigidities. Three early influential papers are Dhyne et al. (2006), Klenow and Kryvtsov (2008), and Nakamura and Steinsson (2008). They use official monthly CPI data from the U.S. and from different countries of the euro area and document the main facts and characteristics of price changes.

These facts quickly became a benchmark and inspired the modelling and calibration of a new generation of sticky price models, started by Golosov and Lucas (2007) (GL). Building on their model, Midrigan (2011) presents a menu cost model that matches some aspects of the distribution of price changes found in the micro data that conflict with the predictions of GL. His model generates a large amount of small, as well as large price changes, which fits the weekly scanner price data he exploits (electronic records of transactions that firms collect as part of the operation of their businesses), and contrasts the distribution predicted by GL, which is bimodal with no changes inside an inaction band around zero. The implications of these two models are notably different: while the GL model predicts near money neutrality, in the Midrigan model a weak selection effect causes a weaker response of aggregate prices to monetary shocks, resulting in real effects of monetary policy closer to those of the Calvo model and five times greater than in GL.

Two key attributes of the set up of Midrigan (2011) are the assumption of a fat-tailed distribution of cost shocks, generating large price changes; and the combination of multiproduct firms and a single menu cost that is paid for changing any number of prices, which leads to economies of scope in price adjustment and a consequent large number of small price changes. The assumptions of multiproduct firms and economies of scope in price adjustment were widely incorporated in the debate on the microfoundations of price-setting and became a component of many recent contributions to the price stickiness literature, such as Alvarez and Lippi (2014), Kehoe and Midrigan (2015), Alvarez et al. (2016) and Karadi and Reiff (2019).

In Section 4 I set up a menu cost model that features multiproduct firms and some degree of economies of scope in price changes as well. The key departure from a standard sticky-price model is the incorporation of a three-dimension cost of price adjustment. The design of the menu cost in my model is similar to that of Bonomo et al. (2020), whose menu cost is composed of a general cost paid once for change any amount of products, and a product-specific cost paid for every additional changed price. I introduce a third component of the menu cost: a cost curtail when the prices of familiar products are changed in the same period. To the best of my knowledge, this is the first paper presenting a menu cost model aimed at matching the salient facts of price changes observed in daily data from a set of countries with heterogeneous macroeconomic conditions.

This paper is also inspired by Cavallo and Rigobon (2016) and Cavallo (2018), and contributes to the growing literature closing the gap between Big Data and Economics. For their Billion Prices Project (BPP), they created a micro data set obtaining online prices on a daily frequency using a web scraping method. Despite the many advantages of scraped daily price data,[2] the number of papers that use it is yet relatively scarce because of its limited availability. Some studies that exploit data sources with daily frequency are Alvarez et al. (2016), who use data from BPP to assess measurement errors in CPI and scanner data and to compare the estimates of some key statistics of price changes with those found in traditional data sources; Alvarez et al. (2018), who find evidence in favor of the existence of strictly positive menu costs in the relatively low fraction of small price changes observed in the BPP data; and Bonomo et al. (2020), who set up a menu cost model that matches the degree of partial synchronization of price changes that they find in a daily-frequency data set with prices of a large number of retail stores from Israel.

2. Data

For most of this paper, I use a data set of prices obtained from supermarkets' e-commerce sites with web-scraping techniques. Starting in August 2019, I collected scraped price data from a set of companies from different business industries. As of June 2020, my project comprises data from supermarkets, fashion retailers, airlines, travel agencies, and energy companies from 12 countries, obtaining more than 100.000 price points every day. In addition, I am working on making all the data I collect freely available for non-profit research projects, becoming the first free data source of online prices from different industries.

I created a program that scraps several online retailers and obtains the price of the products they sell. My software's automated task is to navigate the different websites, search for products, obtain the main characteristics (e.g. name, category, price, among others), and store it in an output file. The program is coded in Python language, though web scraping can also be done in other languages, such as R or C#. In addition to the standard packages used for data analysis, I use the libraries Selenium and BeautifulSoup. Selenium is a tool that allows the software to open a web browser, navigate and interact with it imitating human behavior. With BeautifulSoup I extract the data from the websites.

The process is as follows. I provide my program with a file containing a list of URLs of the websites I want it to scrap (the input file). Using Selenium, I access each URL, which ideally will have a list or grid of products. Websites' underlying codes are written using HTML. Then, using the Inspect tool of each website, I analyze the HTML code to find the tags of the desired elements I want to obtain. Using BeautifulSoup to parse the HTML code of each URL, I identify each of those elements, and extract the text from the website code, and I store the data in my output file, which I then append to my data set. I repeat this process for each URL in the input file.

2.1. Advantages and limitations of scraped data

Previous studies analyzing the microfoundations of price stickiness mainly use two types of data sources, namely CPI data (price data underlying official consumer price indexes) and scanner price

[2]Section 2.1 comments on the advantages and limitations of this new type of data source.

data (electronic records of transactions that establishments collect as part of the operation of their businesses). Table 1 (obtained from Cavallo (2018)) provides a comprehensive comparison of these two alternative sources and scraped price data.

Table 1: **Alternative data sources.**

	Scraped Data	CPI Data	Scanner Data
Data Frequency	Daily	Monthly	Weekly
All products in retailer	Yes	No	No
Uncensored price spells	Yes	No	Yes
Comparable data across countries	Yes	Limited	Limited
Real-time availability	Yes	No	No
Product categories covered	Few	Many	Few
Retailers covered	Few	Many	Few
Quantities sold	No	No	Yes

Source: Table 1 from Cavallo (2018).

A first major advantage of scraped data is its daily frequency. While the CPI and scanner price data sets have a frequency of one month and one week respectively, a daily frequency has several advantages as it captures all the changes in prices: missing the intraweek evolution of prices could represent a relevant loss in the objective of studying their dynamics, especially in countries with high inflation rates where price changes are more frequent, as I show in Section 3.

Second, having daily scraped data is an effective way of avoiding measurement bias. Cavallo (2018) shows that using CPI or scanner price data to study price stickiness typically comes at the cost of measurement errors. He also documents the impact that this issue has on the estimates of widely accepted statistics in the price rigidity literature, such the duration of price changes and their distributions, and the estimates of the slope of the hazard function of price changes. The sources of this bias are the imputation of missing prices in CPI data; and the measurement of prices in scanner data, which is done by calculating the average price of a product in a week weighted by sales, overestimating the number of price changes and underestimating their size.

Morover, my data set captures better the development of temporary price changes than other traditional sources. The main statistics of price changes vary to a great extent if temporary price adjustments are included in the data, typically enlarging the average size and lowering the frequency of price changes. Hence, distinguishing between short-lived temporary changes and regular price changes has relevant implications for the aggregate predictions of sticky price models. While it is possible to detect sales in the CPI and scanner price data sets using techniques such as the HP filter or recognizing a 'V' behaviour in the time series of the price of a certain product (a sudden drop returning quickly to the previous trend), those techniques are not exempt from errors. Instead, web-scraping methods allow to directly recognize sales as they are signaled by the retailer in the website. Additionally, traditional data sets only recognize discounts as changes in the listed price, while I also capture other types of sales that do not necessarily affect the listed price, such as 'buy 2 pay 1' or '25% of discount in the second unit'.

Finally, another advantage is in terms of scope. Web-scraping allows an easy incorporation of new data from a new retailer from almost every country in the world with the same collection criteria, making information from different countries comparable. Thus, I can capture the high-frequency dynamics of countries with different inflation rates, levels of market concentration and other economic determinants of the price adjustment decisions, which could later serve to account for the heterogeneity in price-setting behaviour across countries: arguably, the time, size and type of price changes are not the same for all countries under all states of their economies.

A limitation of scraped data compared to CPI data is that the former covers only a limited number of product categories. The scraped data I use for this paper covers between 17% and 26% of the CPI categories' expenditure basket weights, depending on the country. A shortcoming of my data compared to scanner data is that the latter also contains the amounts sold for each product. While having that information is be very valuable for certain types of analysis (for example, concerning elasticities), the lack of quantities does not affect the analysis I do in this paper.

2.2. Description of the data

My data set has more than 5 million daily prices from eight retailers of six different countries. All the retailers are large multi-channel firms selling full grocery lines and department store products. The data includes prices of supermarkets from the Netherlands, United Kingdom, Brazil, Chile and Turkey, all among the top 3 in their national market shares. There are three supermarkets from Argentina, all among the top 5 in the Argentine market share.

I started collecting scraped prices from the Argentine supermarkets in August 2019 and subsequently added other firms, thus my data set covers non-identical time spans for the different retailers. Table 2 provides relevant details of each retailer data set.

I give missing values a similar treatment to the one used in Cavallo (2018), who fills the missing prices carrying forward the last recorded value until a new price is available. My treatment of missing prices is stricter because I drop from my sample all the products whose price is not reported more than a specific threshold of periods, which I set to 20% of the total number of observed periods for each retailer.

Table 2: **Description of the data**

Country	Netherlands	UK	Chile	Brazil	Turkey	Argentina
1. Retailers	1	1	1	1	1	3
2. Observations (thousands)	1,838	112	359	388	105	3,329
3. Products	20,044	929	1,619	2,239	2,495	14,213
4. Categories	56	12	39	11	53	161
5. Start month	11:2019	02:2020	11:2019	11:2019	04:2020	08:2019
6. End month	06:2020	06:2020	06:2020	06:2020	06:2020	06:2020
8. Explicitly flag sales in the website	Yes	Yes	No	No	Yes	Yes (2)

The number of observations does not include missing values.

3. Empirics

In this section, I exploit my scraped data set to document the major salient facts of price adjustment, in order to provide a microeconomic foundation for the sticky price model I discuss later in this paper. Table 3 presents the main statistics for each retailer in my data set.

3.1. Size of price changes

Magnitude

The mean (median) absolute size of log-price changes ranges from 5.57 (4.69) to 21.97 (18.23) across different retailers.[3] One can easily observe, however, that the statistics obtained from the retailer from the United Kingdom notably differ from the other retailers. Without considering this case, the range narrows to [5.57, 13.66] ([4.69, 11.12] for the median). The case of the retailer from the U.K. represents a good case of the heterogeneity of price changes across different stores documented by Klenow and Malin (2010).

The ample range of the magnitude reflects notably different dispersions in the distributions of price changes in my sample, which is consistent with the diverseness documented in earlier papers using alternative data sources. For example, Midrigan (2011) finds an average absolute size of regular prices of 11.0 in a weekly scanner price data set from the U.S. Other studies, using CPI data, report a mean of 14.0 in the U.S. (Klenow and Kryvtsov, 2008), and a mean of 15.4 in the euro area (Dhyne et al., 2006). In an extensive daily price data set from many food stores in Israel, Bonomo et al. (2020) report an average absolute size of price changes equal to 20.8.

It does not appear to exist a clear link between the absolute size of price changes in a retailer and the inflation rate of the country where the firm operates, nor with the variance of inflation. Instead, this link seems to exist when analyzing the first moment of the distribution considering both the size and sign of price adjustments, i.e. without taking the absolute value (rows 1-2 in table 3). The average size of price changes ranges from 1.70 to 6.79. The range of the median size of price changes also narrows to [3.06, 6.79]. Panel (a) of Figure 3 plots the mean and median size of price changes and the annual inflation rate for the countries analyzed, showing a positive relation between the growth rate of the price level and the size of price changes.

Share of price increases and decreases

The coexistence of a small average size of price changes with an appreciably larger average absolute size of price changes is explained by a substantial share of price decreases in my sample. The percentage of price changes that are increases ranges from 61.1% to 71.6%. This statistic coincides with the estimates from previous studies, which fluctuate around two thirds. Moreover, it also appears to exist a link between the share of prices that are increases and the inflation rate, as shown in panel (b) of Figure 3.

[3] Unless noted otherwise, when I refer to the (log) size of price changes I mean 100 times the log difference of prices, or formally $100 \cdot \ln(P_t/P_{t-1})$, which approximates to the percentage change for close values and has the advantage of symmetry. This is also the approach taken in the literature.

Table 3: **Statistics of price changes**

Country	Netherlands	UK	Chile	Brazil	Turkey	Argentina Firm 1	Argentina Firm 2	Argentina Firm 3
Inflation rate (annual)	2.6%	1.7%	2.8%	3.7%	15.2%	53.5%	53.5%	53.5%
1. Mean	1.70	7.74	1.75	2.33	2.72	4.36	4.35	4.65
2. Median	3.15	3.92	3.06	3.52	5.72	5.88	5.89	6.79
3. 25th percentile	−2.65	−9.53	−9.44	−5.21	−8.07	−1.38	−1.42	−6.77
4. 75th percentile	5.75	28.76	11.20	7.83	12.52	13.35	9.83	11.36
5. Standard deviation	6.58	27.07	18.29	10.12	16.86	16.57	8.68	12.19
6. Absolute mean	5.57	21.97	13.61	9.12	13.66	12.98	8.28	10.99
7. Absolute median	4.69	18.23	10.11	6.83	11.12	10.15	7.85	8.55
8. Absolute 25th percentile	2.92	6.46	4.19	3.44	6.27	5.29	4.50	6.77
9. Absolute 75th percentile	7.31	32.54	18.92	12.44	18.27	17.17	10.46	13.78
10. Absolute standard deviation	3.89	17.59	12.15	7.37	10.61	11.17	5.08	7.02
11. Frequency	0.002	0.003	0.009	0.010	0.011	0.027	0.019	0.020
12. Implied duration	407	305	140	116	91	37	52	49
13. Share of price increases	0.632	0.611	0.631	0.639	0.651	0.671	0.716	0.680
14. Fraction $<\lvert 1\% \rvert$	3.54%	1.37%	1.61%	4.08%	2.06%	6.73%	3.50%	0.77%
15. Fraction $>\lvert 5\% \rvert$	44.62%	81.88%	72.24%	65.59%	82.01%	74.89%	70.44%	91.86%
16. Skewness	−1.035	0.204	−0.012	−0.149	−0.461	−0.264	−0.583	−0.327
17. Excess kurtosis	1.568	0.431	−0.302	3.212	0.476	1.612	0.030	−0.017

Notes: The size of price changes is estimated as 100 times the log difference of prices. The implied duration of price spells is expressed in days. The kurtosis of the distribution of price changes is estimated using standardized price changes at the category level. The excess kurtosis is calculated as $3 - kurtosis$.

These two facts contribute to the debate on the relevance of the shocks from different nature (i.e. aggregate and idiosyncratic) in models with nominal rigidities. On the one hand, very common price declines assign a relevant role to idiosyncratic shocks (Klenow and Malin, 2010) in driving prices away from their optimum; on the other, the novel fact revealing a positive relation between the share of price increases and inflation rate gives evidence in favor of an important role of aggregate conditions in driving price changes.

Higher moments of the distribution

Table 3 also reports the skewness and (excess) kurtosis of the distribution of the size of price changes. The skewness is negative in all the cases except for the retailer from the U.K. Two of the retailers show a moderately negative skew (between -0.5 and slightly below -1) and in the remaining cases (with the exception of the U.K.) the distributions are fairly symmetrical, yet left-skewed. This statistic has received null-to-little attention in previous papers despite providing a

Figure 3: **Inflation rate and key statistics of the size of price changes**

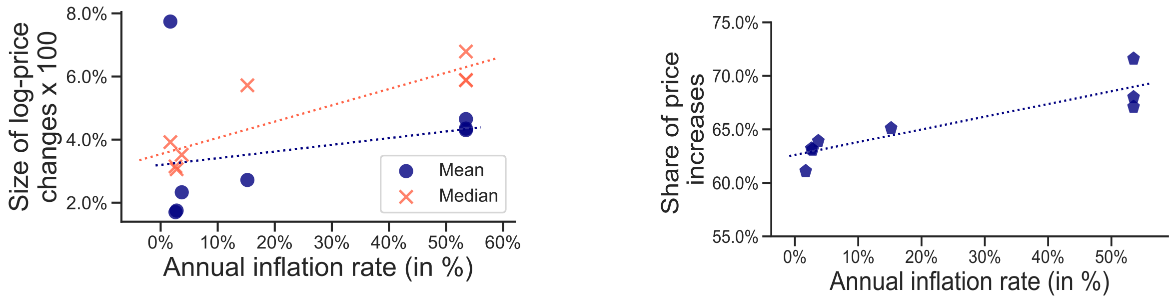

(a) Inflation rate and the average size of price changes (b) Inflation rate and share of price increases

Note: Each scatter plot represents a different retailer. The dotted lines are the plots of the best-fit linear relations between inflation rate and the different variables, and it is only added for illustration purposes. Note that the slope of this line for the relation between inflation and the mean size of price changes is flattened by the strikingly larger mean of the retailer from the UK.

helpful conceptual insight of the distribution of price changes. A negative skewness requires a larger mode than the median and the mean of the distribution. With positive mean and median, this implies that the mode of the size of price adjustment is positive and located relatively far from zero (compared to the first moment in the distribution); and, concretely, contradicts models that predict an unimodal bell-shaped distribution centered in zero.

The values of the excess kurtosis range between -0.3 and 3.2 –recall, as a reference, that the excess kurtosis of the normal distribution is 0. I estimate this statistic over the standardized log-size of price changes at the category level to prevent biased values arising from the heterogeneity across categories.[4] To obtain the standardized price change I simply subtract the mean price change of each product' category and divide by the category's standard deviation of the size of changes.

This statistic has received attention in recent papers studying the microfacts of price changes and their aggregate implications for monetary non-neutrality. Alvarez et al. (2016) find that the kurtosis of the size of price changes can be a sufficient statistic describing the real effects of monetary policy.[5] They show that in a wide variety of models, the kurtosis embodies the *selection effect* of monetary policy. The selection effect indicates that those prices that are adjusted when a firm revises them are those that are far from their optimum. A strong selection effect, such as that of the Golosov and Lucas (2007) model, indicates that price adjustments will be on average large, causing a strong response of the aggregate price level and a consequent high degree of monetary neutrality, even if the size of the monetary shock is small. Such a model displays a strongly bimodal distribution of price changes and the smallest value of excess kurtosis, which is -2. On the other extreme, in a standard Calvo model, prices adjust independently of their distance from the optimum, hence the selection effect is null and the real effects of monetary policy are relatively large. Such a model features a peaked distribution, with a large mass of small as well as large price changes, resulting in a high excess kurtosis equal to 3.

[4] The estimated kurtosis of a population composed of different elements with different variances is larger than that of each element.

[5] Their characterization of the kurtosis as a sufficient statistic remains valid on economies with zero or low inflation.

Importantly, the excess kurtosis in my data is brought down by the low density of price changes around zero, which creates a 'hole' in otherwise unimodal distributions. This feature of the data is not present in traditional data sources because of the limitations mentioned in section 2.1: time averages in scanner data and cell-relative imputation in CPI data tend to overestimate the number of small price changes. Cavallo (2018) quantifies this issue, showing that the overestimation increases the excess kurtosis from 0.96 in online data to 2.45 in CPI data.

3.2. Frequency and duration of price changes

Table 3 (rows 11-12) also reports the median frequency with which prices change and the implied duration (in days) for each price spell. I compute the frequency of price adjustment as the number of price changes divided by the number of observed prices per product. This measure is then aggregated to the category level as the mean frequency of all the products from each category.

Formally, I obtain the mean frequency for each category k as:

$$f_k = \frac{\sum_{t=1}^{T}\sum_{i=1}^{N}\varphi_{i,k;t}}{\sum_{t=1}^{T}\sum_{i=1}^{N}\mathbf{P}_{i,k;t}} \quad \text{with} \quad \varphi_{i,k;t} = \begin{cases} 1, & \text{if } P_{i,k;t} \neq P_{i,k;t-1}\ \forall i \in 1,N \wedge \forall t \in 1,T \\ 0, & \text{otherwise,} \end{cases}$$

where T is the number of periods (days), N the number of products, k an index for each category, and $\mathbf{P}_{i,k;t} = 1$ for every price observation.[6] Then, the median frequency of a retailer is $f = median(f_1, f_2, ..., f_K)$. The daily hazard rate of price changes is $\lambda = -\log(1-f)$. Therefore, defining the median implied duration (in days), d, as the inverse of the hazard rate of price changes gives:

$$d = -\frac{1}{\log(1-f)} \tag{1}$$

The range of the implied durations in my data is notably wide. The shorter duration is 37 days, for a retailer from Argentina, while the longer duration is more than one year (407 days), for the retailer from the Netherlands. These differences in the implied durations and frequencies are not a particularity of my data: previous studies have also documented very heterogeneous values for these statistics.[7] Besides potential differences in calculation methods, the estimations of the implied duration are subject to many sources of heterogeneity, such as the different types and number of goods each retailer sells, types of stores, and the role that sales play for each retailer's pricing strategy. The wide ranges of implied durations at the category level reported in Table 4 illustrate the disparity across different types of products found in my data set.

Relevant insights can be obtained from the duration of price spells. First, it indicates different degrees of price stickiness across countries. In Argentina, the country with the higher inflation rate of my sample during the covered period (with an average monthly inflation rate of three percent),

[6] Since every product in my data has a price for each observed day, $\sum_{k=1}^{K}\sum_{t=1}^{T}\sum_{i=1}^{N}\mathbf{P}_{i,k;t} = TN$

[7] An interesting example of this is the difference in the reported durations by Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008) who, using the same CPI data set from the U.S. collected by the Bureau of Labor Statistics, found a median duration (in months) of 11.0 and 7.2, respectively. The time span of Nakamura and Steinsson's study is 1988:01-2005:12, while that of Klenow and Kryvtsov's is 1988:02-2005:01

Table 4: **Heterogeneity of price changes at the category and product level.**

Country						Argentina		
	Netherlands	UK	Chile	Brazil	Turkey	Firm 1	Firm 2	Firm 3
1. Range of durations across categories	[215; 1125]	[113; 1296]	[21; 695]	[47; 568]	[13; 964]	[8; 123]	[27; 99]	[27; 239]
2. Products whose prices remained unchanged	53%	57%	15%	44%	54%	1%	0%	7%

Note: The range of durations across categories for a retailer defined as the shorter and longer durations for different categories within a firm.

prices take on average one and a half month to change. On the other extreme, for the retailer from the Netherlands –the country with greater price stickiness in my data— the expected duration of a price spell is of 13.4 months.

Analyzing the entire sample, it seems to exist a link between the inflation rate and the duration of price spells. Economic reasoning suggests that if macroeconomic conditions have a predominant role in driving the deviations from the optimal price, –assuming the same degree of homogeneity between products and stores– the duration of price spells is expected to be higher with a lower rate of inflation, ceteris paribus, both in a cross-section and time-series analysis. Indeed, the duration of price spells is notably longer in countries with lower inflation than in countries with higher inflation in my data; yet –as with the previously commented connections between inflation and price change statistics– a causal relation should not be suggested without caution of missing unobserved factors, particularly because of the heterogeneity in goods and stores types mentioned above. Figure 5 plots the inflation rate and the implied duration of price changes of the firms in my data (panel a), previous findings using CPI data surveyed by Alvarez (2008) (panel b), and the values provided by Table 9 in Cavallo (2018) using scraped daily data for 31 countries (panel c).

Figure 5: **Inflation rate and duration of price spells**

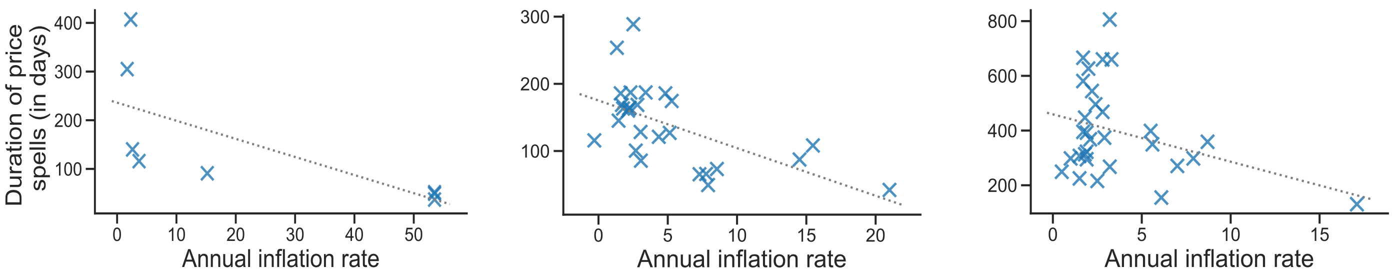

(a) Data (b) Alvarez (2008) (c) Cavallo (2018)

3.3. The behavior of temporary price changes

A segment of the recent debate on price stickiness concerns the treatment of temporary price changes, or sales. The central point of the discussion concerns the relevance and the role that is assigned to temporary changes in the estimation and modelling of price rigidities in macro models.

Proponents of not considering temporary price changes as influential to the estimations of price stickiness argue that sales are orthogonal to macroeconomic conditions, as they are unresponsive to macro shocks [Nakamura and Steinsson (2008)]. If this was the case, then sales do not contribute to the adjustment of aggregate inflation to aggregate shocks, and hence it is appropriate to exclude them when modelling price stickiness. Another rationalization for excluding sales as determinants of price stickiness is that price discounts are typically pre-established with substantial anticipation and defined by each retailer's business strategy –which includes a special budget for discounts–, and that these strategies (*price plans*) are sticky, and unresponsive to new macroeconomic information [Mankiw and Reis (2002); Burstein (2006); and Anderson et al. (2017)].

Contrarily, other papers defend the inclusion of sales in sticky price models arguing in favor of the existence of sale-type responses from firms to changes in the aggregate macroeconomic conditions [e.g. Bils and Klenow (2004), and Klenow and Kryvtsov (2008)]. Moreover, Klenow and Willis (2007) suggest that temporary price changes are in fact a source of macro price flexibility and provide evidence, using CPI data from the U.S., of significant correlation between the size of sales and the accumulated inflation since the last price change.

On the theory side, the role of sales in menu costs models is also a subject of debate, as different implications for their impact on price flexibility are obtained depending on (i) the manner they are formally incorporated into the models, and (ii) the underlying economic rationale that characterizes the decision rule resulting in a temporary change. Kehoe and Midrigan (2015) describe a model where the retailer faces a (relatively large) menu cost for adjusting regular prices and a (relatively small) menu cost for incurring into a temporary change that automatically reverts to the regular price after one period. They find that even though prices change more frequently when temporary price changes are included, their temporary nature and the fact that they automatically revert to the previous level leaves the aggregate price stickiness unaffected by the presence of temporary changes. Alternatively, Alvarez and Lippi (2020) set up a model where firms choose a price plan, defined as a set of two prices $P = \{P^L, P^H\}$, allowing the firm two move between any of the prices in the set without paying a cost (instead, a menu cost must be paid when the firm chooses another price plan). By permitting many free price reversals within a price plan and not requiring a price to automatically return to its previous value their model displays a higher flexibility of the aggregate price level relative to standard menu cost models.

Return to the previous value

A key assumption in Kehoe and Midrigan (2015) to obtain the result that sales are not a relevant measure for price flexibility is that temporary price changes automatically return to the previous value after a certain period. Hence, two relevant statistics to assess the validity of this assumption are the percentage of discounts that return to the same regular price where they departed and the concentration in the duration of sales. Table 5 shows these statistics, obtained only for the retailers that explicitly flag their sales in their websites to avoid measurement errors arising from sales-detecting algorithms.

The statistics are compatible with the assumption of Kehoe and Midrigan (2015). The number of returning-sales account for a large share of the total number of discounts (row 1). The percentage of different retailers ranges from 78% to 95%, with a mean of 85%. Also, the duration of price changes appears to be highly concentrated. The 5 durations that concentrate a larger number of sales account for more than half of all the temporary price changes of all the retailers (row 3).

Table 5: **Facts about sales**

Country				Argentina	
	Netherlands	UK	Turkey	Firm 2	Firm 3
1. Share of sales returning to previous value	95%	93%	78%	80%	81%
2. Concentration in the top 3 durations	41%	52%	45%	37%	50%
3. Concentration in the top 5 durations	52%	65%	56%	53%	65%
4. Concentration in the top 3 abs. magnitudes	74%	31%	42%	62%	62%
5. Concentration in the top 5 abs. magnitudes	79%	42%	53%	74%	79%

Concentration in the magnitude of price discounts

Another common feature of sales across different retailers is that their size is typically concentrated in a few specific values. Rows 4-5 in table 5 and Figure 13 in Appendix A.2 present this fact. This fact can be explained by two hypotheses on how firms decide to implement a temporary price change in which, importantly for the discussion on whether considering or not sales in sticky price models, temporary price changes do not contribute to aggregate price flexibility. First, the size concentration is consistent with the idea that firms have a pre-established pricing plan of sales arranged with the producers of the products, with a special budget for discounts, and with an important role of the marketing area of each retailer in the decisions concerning price discounts. The concentration of sales' sizes fits this premise since a customary discount is easier to execute (and requires less efforts devoted to the analysis of the optimal price), and because it is also more appealing for marketing communication attracting irrational consumers' behavior (for example, -15% or 'pay 6 get 7' instead of -13.48%, even if the latter was the rational agents profit maximizing discount). I label these types of changes as *marketing changes*.

Second, the concentration in the size of sales is also consistent with the idea of a relatively small physical cost of temporarily changing a price from Kehoe and Midrigan (2015). Let me consider an economy combinining small costs of making temporary price changes and the presence of marketing changes. I also assume, consistently with the evidence presented above, that sales return to their previous value after their life-period. This economy would also display a concentration of sales in a few number of sizes because of the marketing changes. However, these temporary price changes would not be responses to changes in macroeconomic conditions, since firms would prefer to do a regular (rather than a temporary) price change after a permanent aggregate shock because of their temporary nature and the fact that they automatically revert to the previous level.

3.4. Shape of the distribution of the size of price changes

One of the major salient characteristics of the distribution of the size of price changes observed in the scraped price data is the relative scarcity, yet nontrivial presence, of small price adjustments.

Table 3 reports the percentage of price changes with a size lower than 1% in absolute terms, which ranges between 0.8% and 6.7% across retailers.

Figure 7 plots a histogram of the distribution of the size of regular price changes for all the retailers in my data set. In all cases, there is a “hole” around zero reflected as a larger mass of price changes in the range (1%; 2%] and (-1%; -2%] than (0%; 1%] and (0%; -1%], respectively. Moreover, Figure 15 in Appendix A.2 shows the cumulative distribution function of the absolute size of price changes for the same retailers and compares them with the cumulative distribution function of a Gaussian distribution matching the mean and standard deviation in each case, for illustrative purposes. In all cases except for Argentina’s Firm A, the Gaussian distribution has a larger accumulated density than the data in values smaller than |3%|.

Figure 7: **Histograms of price changes distribution**

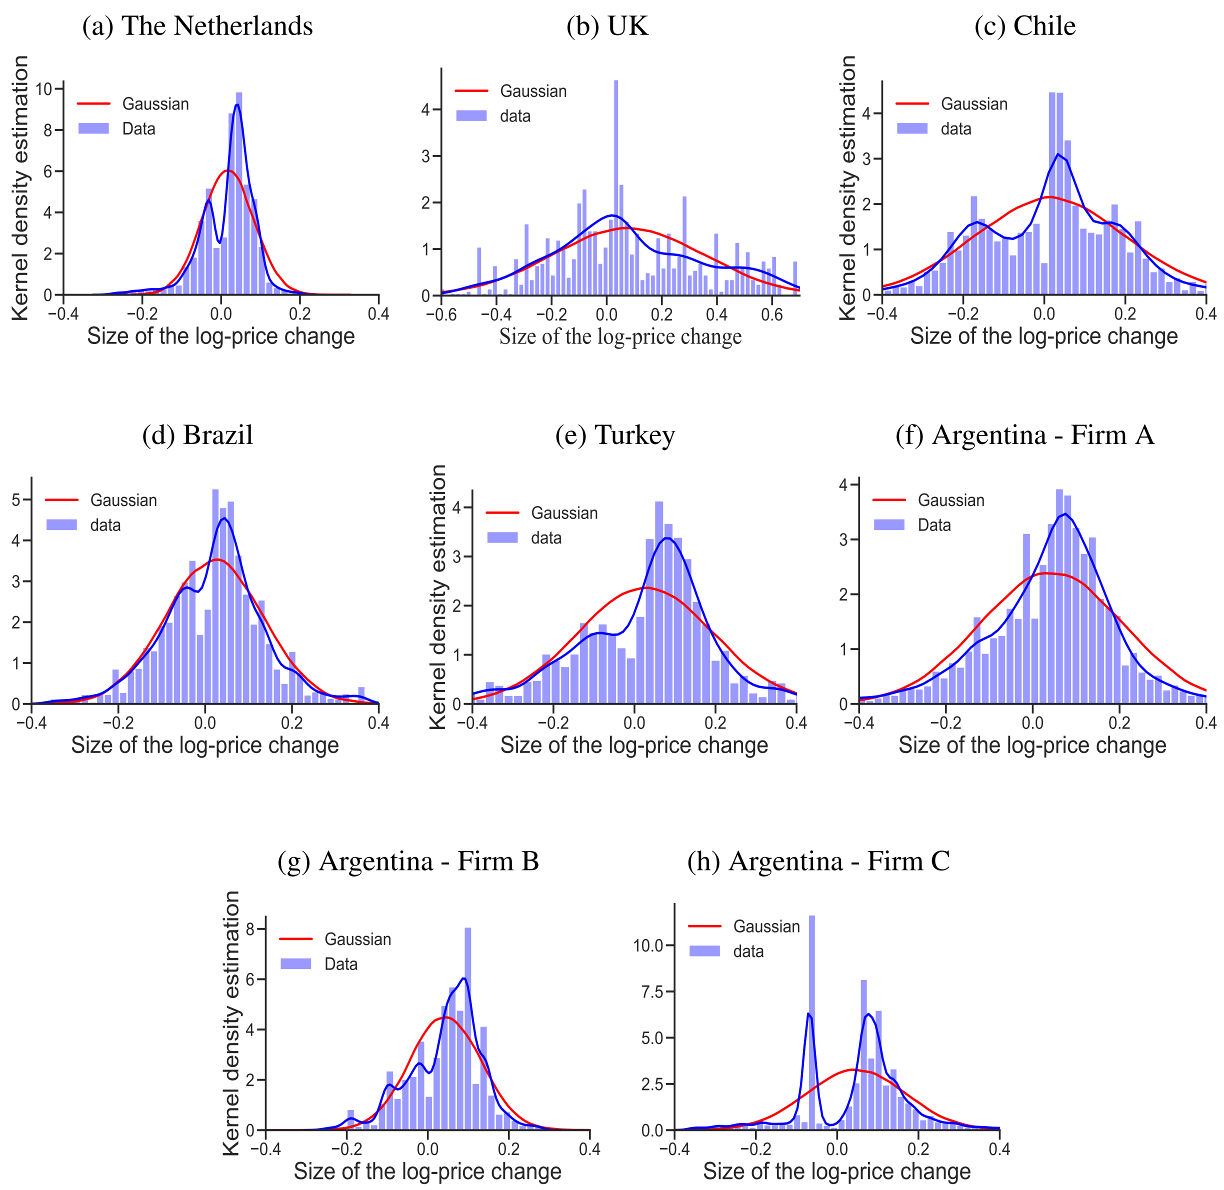

Note: The y-axis corresponds to the kernel density estimation (KDE) of each distribution –plotted in blue lines– which is a non-parametric smoothing method of obtaining a probability distribution. The height of the histograms denotes, as usual, the relative frequency of each bin.

This fact represents a key difference between the usage of traditional micro data sources and scraped data for the assessment of sticky price models. Traditional data sources present a larger share of small price changes. Moreover, the relative scarcity of small changes found in my data is consistent with previous papers using scraped data [Cavallo and Rigobon (2016) and Cavallo (2018)].[8] In fact, Cavallo (2018) shows that the surfeit of small price changes found in CPI and scanner data is a consequence of measurement errors and imputation of missing prices.

This feature from the data gives rise to an important benchmark for the embodiment of menu costs in sticky price models. A standard single product menu cost model such as that of Golosov and Lucas (2007) generates a bimodal distribution of the size of price changes with a band of inaction around zero dependent on the magnitude of the menu cost, therefore failing to generate small price changes at all. Contrarily, introducing multiproduct firms but keeping a unique cost that has to be paid for adjusting any amount of regular prices as in Midrigan (2011) generates a relatively large mass of price changes around zero. Intuitively, a firm will find it optimal to change *all* its prices that differ by *any* amount from its optimal value once it has paid the adjustment cost. In Section 4, I combine both approaches in a multidimensional-menu-cost model that is able to reproduce the low density of the distribution around zero.

3.5. Synchronization of price changes of familiar products

It is very common to find similar items produced by the same manufacturer in the catalog of consumer-goods retailers. For example, a supermarket may sell different sizes and flavors of the same beverage brand, an electronic store may have different models of the same line of laptop computers, and a fashion department store may sell the same pair of jeans in different colors.

Whether the time series of the prices of similar products have akin dynamics (i.e. if they share the timing and the size of price changes) is, therefore, a relevant question for the study of price stickiness. Consider an extreme case where familiar products have equal demands and their costs (paid by the retailer to the manufacturer) have identical dynamics. Also consider two alternative data of price changes, at a given period after an aggregate shock, from the same multiproduct retailer. One has 10 identical price changes, all of familiar products, and the other has 8 price changes of heterogeneous products, with an identical average size of price changes than the data with 10 price changes. A standard approach to the study of price flexibility would consider all price changes the same way, and therefore estimate a higher degree of price flexibility in the data with the 10 changes. However, one may also argue that the fact that the familiar products share the same dynamics of price changes reduces those 10 price series to only one; and, hence, that in the case with 8 adjustments the shock prompted more changes than in the case with 10.[9]

In order to investigate the similarities in the price series of familiar products, I analyze the

[8] Figure 11 in Appendix A.2, plots the distribution of price changes as obtained in the Billion Prices Project from Cavallo and Rigobon (2016), also exhibiting a "hole" around 0%.

[9] Naturally, this illustration is highly simplistic. A complete examination requires incorporating other components to the analysis, such as consumption basket weights. However, everything else the same, the affirmation that changes in familiar products' prices with identical demands and costs imply less price flexibility than changes in heterogeneous products' prices remains valid.

timing and magnitude of their changes. I use the Levenshtein ratio as the metric to flag a family of products in my data.[10] The ratio is constructed upon the Levenshtein distance of the products' names, which measures the minimal number of changes (insertions, deletions, or substitutions of characters) necessary to transform one string into another (Schulz and Stoyan, 2002).

I set 0.75 as the threshold value for the Levenshtein ratio that defines a family of products: those set of products whose names have a ratio higher than 0.75 with all the elements of the set are then considered a family. Then, I take a random date for the retailers in my sample and analyze whether the timing and the magnitude of price changes of similar products are the same or not.

Table 6 reports some relevant statistics. Conditional on at least one product from a family adjusting its price in a period, the probability of a change in the price of a product from the same family is 0.24, denoting a greater degree of synchronization of price changes in similar products than in heterogeneous products (the unconditional probability is 0.02). Moreover, when different products from the same family adjust their prices in the same period, the size (in percentage points) of the change is identical in 71% of the cases. An immediate interpretation of this fact is that these familiar products are subject to resembling idiosyncratic shocks driving deviations from their optimal prices.

Table 6: **Similarities in the time series of familiar products' prices**

1. Unconditional probability of a price change	0.02
2. Probability of a change, conditional on at least a change in the product's family in the same period	0.24
3. Probability of two adjustments in the same family being equal in the same period	0.71

Note: The statistics were obtained for random dates from the retailers from Argentina and Turkey

4. Model

In this section I set up a tractable partial equilibrium model in which a multiproduct firm decides whether to adjust or not its prices. The firm perfectly observes the state of the economy in every period and has perfect information about the optimal price of each of its products. The model builds on a version without temporary changes of the model in Midrigan (2011), incorporating a multi-dimensional structure of the cost paid by the firm for changing prices. I also show that my model encompasses as special cases the two standard types of pure menu cost models: a Golosov-Lucas-type model with no economies of scope, and a Midrigan-type model with perfect economies of scope in price adjustment.

The model is compatible with the economic rationale of price setting decision of the firms and it is able to reproduce features observed in the scraped price data that standard menu costs models

[10] Appendix A.1 contains a formal derivation of the Levenshtein ratio and provides some illustrations of its value for different products.

are not able to match. As I commented earlier, in the Golosov-Lucas model, single product firms face a fixed cost for adjusting their prices and decide to change a price when the loss from inaction surpasses the threshold determined by the menu cost. Since firms sell only one product, the aggregate distribution of the size of price changes in their model is characterized by a band of inaction around zero bounded by two ‘peaks’ at the positive and negative threshold values. The issue of null small price changes was sorted out by Midrigan (2011) with a model featuring multiproduct firms that pay a fixed menu cost for changing any number of prices at a certain period. When a firm faces a single cost, independent of the number of prices it adjusts, it will find it optimal to change all the products’ prices that differ from its optimum –conditional on changing at least one price– no matter how small the price gap is. Then, multiproduct (single) menu cost models predict a distribution of price changes with a relatively high concentration around zero and, therefore, also fail to generate the shape of the distribution of price changes observed in the data, with a relatively low amount of small price changes.[11]

I begin defining the typical problems of the representative household and firm. Later, I characterize the decision rule of the firm facing the multi-dimensional menu cost, which is the major novelty of my model. I postulate that a model that matches the main characteristics of the distribution of price changes observed in the data has to present (i) a relatively large physical cost paid only once in every period the firm changes any price; (ii) a relatively small physical cost paid for every price that is adjusted; and (iii) a “waiver” of the small cost when another familiar product’s price is changed in the same period. In the remainder of this section, I compute the partial equilibrium of the firm’s problem, which I solve in section 5, showing that introducing this novel menu cost in an otherwise standard multiproduct menu cost model makes my model fit many of the salient features from the data.

4.1. Environment

Households. The economy contains a continuum of households with a total mass normalized to unity. The representative household maximizes her utility over consumption and time devoted to labor employment given by:

$$\max_{C_{i,z;t},L_t,B_t} E_0 \sum_{t=0}^{\infty} \beta^t U(C_t, L_t), \tag{2}$$

where β denotes the household’s discount factor, C_t denotes a composite consumption good in period t compounded of imperfectly substitutable goods, and L_t denotes labor.

The household’s budget constraint is defined as:

$$\int_0^1 \sum_{i=i}^{N} P_{i,z;t} C_{i,z;t} dz + B_t \leq (1+R_{t-1})B_{t-1} + W_t L_t + \Pi_t, \tag{3}$$

[11]The characterization of the different shocks (their distribution, specifically) plays a relevant role in the predicted distribution of price changes of these models. However, the relatively large mass around zero is present in models with typically assumed distributions of shocks, such as Gaussian or Poisson.

where $C_{i,z;t}$ indicates the household's consumption of good i produced by firm z in period t and $P_{i,z;t}$ its price. B_t is the number of one-period non-contingent bonds with nominal price equal to one held in period t that pay a nominal interest rate of R in period $t+1$. W_t is the nominal wage received for the time devoted to labor and Π_t is the nominal profit received from the household's participation in firms' ownership.

Consumption is nested into two aggregators:

$$C_t = \left(\frac{1}{N} \sum_{i=1}^{N} C_{i;t}^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}}, \qquad C_{i;t} = \left(\int_0^1 (A_{i,z;t} C_{i,z;t}^{\frac{\gamma-1}{\gamma}} dz \right)^{\frac{\gamma}{\gamma-1}} \tag{4}$$

Where z is an index over retailers, i an index over goods, C_t is an aggregator of consumption over different goods and $C_{i;t}$ is an aggregator over consumption of good i purchased from different firms. γ is the elasticity of substitution across retailers and θ is the elasticity of substitution across goods. $A_{i,z;t}$ is the quality of the good i sold by firm z. Higher $A_{i,z;t}$ increases the marginal utility of consumption for that good while it also makes that good more costly to produce, as I show in the characterization of the firm's problem.

The behaviour of the representative consumer may be regarded as the outcome of a two-stage utility maximization procedure. In the first and second stage the household optimally makes consumption decisions over products and over retailers, respectively. The resulting model is a nested version of the standard constant elasticity of substitution Dixit–Stiglitz formulation.

In the first stage the household's decision problem is:

$$\min_{C_{i;t}} \sum_{i=1}^{N} P_{i;t} C_{i;t} \qquad \text{subject to} \qquad C_t \leq \left(\frac{1}{N} \sum_{i=1}^{N} C_{i;t}^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}},$$

which gives the demand

$$C_{i;t} = \left(\frac{P_{i;t}}{P_t} \right)^{-\theta} C_t \ . \tag{5}$$

In the second stage the household's decision problem is:

$$\min_{C_{i,z;t}} \int_0^1 P_{i,z;t} C_{i,z;t} dz \qquad \text{subject to} \qquad C_{i;t} \leq \left(\int_0^1 (A_{i,z;t} C_{i,z;t})^{\frac{\gamma-1}{\gamma}} dz \right)^{\frac{\gamma}{\gamma-1}},$$

which gives the demand

$$C_{i,z;t} = \left(\frac{P_{i,z;t}}{P_{i;t}} \right)^{-\gamma} A_{i,z;t}^{\gamma-1} C_{i;t} \ . \tag{6}$$

Combining 5 and 6 yields the demand for good i produced by firm z

$$C_{i,z;t} = A_{i,z;t}^{\gamma-1} \left(\frac{P_{i,z;t}}{P_{i;t}} \right)^{-\gamma} \left(\frac{P_{i;t}}{P_t} \right)^{-\theta} C_t, \tag{7}$$

$$\text{where} \quad P_t = \left(\frac{1}{N} \sum_{i=1}^{N} P_{i;t}^{1-\theta} \right)^{\frac{1}{1-\theta}} \quad \text{and} \quad P_{i;t} = \left(\int_0^1 A_{i,z;t}^{\gamma-1} P_{i,z;t}^{1-\gamma} dz \right)^{\frac{1}{1-\gamma}},$$

are the minimum expenditure necessary to buy an unit of the final consumption good C_t and the price index for good i, respectively.

Firms. The economy also contains a continuum of identical monopolistically competitive multiproduct firms indexed by $z \in [0,1]$ that hire labor competitively from households and produce N different goods, indexed by i, according to the following technology:

$$Y_{i,z;t} = \frac{1}{A_{i,z;t}} L_{i,z;t},$$

where $L_{i,z;t}$ is firm z's demand for labor for producing good i. Note that an extra unit of product quality $A_{i,z;t}$ requires $1 / A_{i,z;t}$ more units of labor to produce the same amount of output.

The level of quality of product i sold by firm z follows an exogenous AR(1) process:

$$log(A_{i,z;t}) = \rho_A log(A_{i,z;t-1}) + \zeta^A_{i,z;t},$$

where $\zeta^A_{i,z;t}$ denotes an i.i.d. normally distributed idiosyncratic quality shock with mean zero and variance σ^2_A.

Firms take wages and the demand for their products as given and set the prices $\{P_{i,z;t}\}_{i=1}^N$ that maximize their profits. After setting their set of prices they meet the demand at their posted prices.

Firm z's nominal profit absent nominal rigidities is

$$\Pi_{z;t}(\{P_{i,z;t}\}_{i=1}^N, W_t, \{A_{i,z;t}\}_{i=1}^N, \{Y_{i,z;t})\}_{i=1}^N = \frac{1}{N}\sum_{i=1}^N(\underbrace{P_{i,z;t}}_{\text{unit price}} - \underbrace{W_t A_{i,z;t}}_{\text{marginal cost}})Y_{i,z;t} \tag{8}$$

Without capital in the model $Y_{i,z;t} = C_{i,z;t}$. Aggregate demand is therefore given by $Y_t = C_t$. Combining this last equality, equation 7, and equation 8 gives the frictionless profit:

$$\Pi_{z;t} = \sum_{i=1}^N (P_{i,z;t} - W_t A_{i,z;t}) A^{\gamma-1}_{i,z;t} \left(\frac{P_{i,z;t}}{P_{i;t}}\right)^{-\gamma} \left(\frac{P_{i;t}}{P_t}\right)^{-\theta} Y_t \tag{9}$$

Solving the first order condition of the problem yields the optimal frictionless price for product i sold by firm z:

$$[P_{i,z;t}] : Y_t A^{\gamma-1}_{i,z;t} \left(\frac{P_{i,z;t}}{P_{i;t}}\right)^{-\gamma} \left(\frac{P_{i;t}}{P_t}\right)^{-\theta} (A_{i,z;t}\gamma W_t - \gamma P_{i,z;t} + P_{i,z;t}) = 0.$$

$$P^*_{i,z;t} = \mu^* A_{i,z;t} W_t. \tag{10}$$

Where $\mu^* = \frac{\gamma}{\gamma-1}$ is the steady state frictionless markup over the marginal cost that maximizes the firm's profits. The effective markup for good i is defined as: $\mu_{i,z;t} = \mu^* \frac{P_{i,z;t}}{P^*_{i,z;t}}$.

Firm z's nominal profit and nominal revenue can then be expressed in terms of the markup $\mu_{i,z;t}$:

$$\Pi_{z;t} = \sum_{i=1}^{N} (\mu_{i,z;t} - 1)\mu_{i,z;t}^{-\gamma} W_t^{1-\gamma} P_{i;t}^{\gamma-\theta} P_t^{\theta} Y_t, \tag{11}$$

$$R_{z;t} = \sum_{i=1}^{N} \mu_{i,z;t}^{1-\gamma} W_t^{1-\gamma} P_{i;t}^{\gamma-\theta} P_t^{\theta} Y_t, \tag{12}$$

Monetary Policy. I assume that the monetary authority targets a path for nominal spending,

$$\int_0^1 \sum_{i=1}^{N} P_{i,z;t} C_{i,z;t} dz = P_t C_t = M_t,$$

making nominal spending follow a random walk with a drift in logs process: $log(M_t) = \eta + log(M_{t-1}) + \zeta_t^M$, $\zeta_{i,z;t}^M \sim N(0, \sigma_M^2)$.

Finally, I follow Midrigan (2011) and assume linear disutility in labor and set the utility weight of labor relative to market consumption equal to one, which gives households preferences of the form $U(C_t, L_t) = \log(C_t) - L_t$. This assumption involves Hansen (1985) and Rogerson (1988) conditions on indivisibility of labor and agents choosing employment probabilities by trading lotteries. Importantly, this characterization of the households' preferences makes the nominal wage proportional to the nominal nominal aggregate demand and thus to the nominal money stock, ensuring a one-for-one pass-through from monetary shocks to firms' marginal cost and therefore to the optimal price of a product, which will affect the firms' price gap.

4.2. Decision rule

The fundamental deviation of my model from standard price-setting theories of price adjustment is the introduction of a three-dimensional physical cost of price adjustment. In every period, a firm that decides to adjust its prices pays three types of menu costs. First, the firm faces a "general" cost, ϕ^G, that has to be paid once and for all for changing *any* amount of prices and it is independent of the number of changed prices, which is analogous to the one in Midrigan (2011) that leads to economies of scope in price adjustment of regular prices. I imagine this cost to be associated to organizational aspects of the firm's structure, such as printing new catalogs or price lists; communication to salespeople and costumers; approval of the changes; acquiring from a third party costly information necessary to evaluate the best price; and decision-making time of employees.[12]

The second cost of price adjustment is a "product-specific" cost, ϕ^S, for every price that is changed. This product-specific cost rationalizes the fact that, unaffected by the number of changes, the marginal change is always costly (also constant in this model) and hence will prevent my model to generate a large number of small prices and perfect synchronization in price changes. The introduction of this cost is intuitive: even in online retailers that do not have to pay the physical cost of price tags, modifying an extra price requires devoting extra resources to the analysis and

[12]Zbaracki et al. (2004) comment on the existence of "managerial thinking" costs and suggest that these costs are substantially larger than the traditional physical costs of price adjustment.

application of the change.[13]

Finally, the firm receives a *cost curtail* when it changes two or more familiar products,[14] which hints that my model features a higher degree of economies of scope in price adjustment of similar products than of unfamiliar products. To exemplify this idea, consider the case of a multiproduct supermarket that is adjusting some of its prices. This assumption implies that the cost of changing the price of 'Coke Light 600ml' and 'Coke Light 1500ml' is lower than the cost of adjusting the price of 'Coke Light 600ml' and 'Aromatic candle 200grs'. Also, this cost curtail can be interpreted as an implicit introduction of asymmetric informational costs in my model: arguably, the cost of acquiring information about the optimal price of two familiar products from the same manufacturer is lower than that of acquiring information related to two utterly different products.

To illustrate this, consider the case of a firm that sells three products, with products $i \in \{2,3\}$ being familiar products. Then, the three-dimensional cost of price adjustment faced in period t by such firm is defined as

$$\Phi_t = \begin{cases} 0 & \text{if } P_{i,z;t} = P_{i,z;t-1}\ \forall\, i \in \{1,2,3\} \\ \phi^G + \phi_1^S & \text{if } P_{i,z;t} = P_{i,z;t-1}\ \forall\, i \in \{2,3\} \text{ and } P_{1,z;t} \neq P_{1,z;t-1} \\ \phi^G + \phi_2^S & \text{if } P_{i,z;t} = P_{i,z;t-1}\ \forall\, i \in \{1,3\} \text{ and } P_{2,z;t} \neq P_{2,z;t-1} \\ \phi^G + \phi_3^S & \text{if } P_{i,z;t} = P_{i,z;t-1}\ \forall\, i \in \{1,2\} \text{ and } P_{3,z;t} \neq P_{3,z;t-1} \\ \phi^G + \phi_1^S + \phi_2^S & \text{if } P_{i,z;t} \neq P_{i,z;t-1}\ \forall\, i \in \{1,2\} \text{ and } P_{3,z;t} = P_{3,z;t-1} \\ \phi^G + \phi_1^S + \phi_3^S & \text{if } P_{i,z;t} \neq P_{i,z;t-1}\ \forall\, i \in \{1,3\} \text{ and } P_{2,z;t} = P_{2,z;t-1} \\ \phi^G + \phi_2^S + \phi_3^S - \sum_{i=2}^{3} \phi_i^C & \text{if } P_{i,z;t} \neq P_{i,z;t-1}\ \forall\, i \in \{2,3\} \text{ and } P_{1,z;t} = P_{1,z;t-1} \\ \phi^G + \sum_{i=1}^{3} \phi_i^S - \sum_{i=2}^{3} \phi_i^C & \text{if } P_{i,z;t} \neq P_{i,z;t-1}\ \forall\, i \in \{1,2,3\} \end{cases} \tag{13}$$

4.3. Solving the partial equilibrium

To compute the partial equilibrium of the model I follow Yang (2019) who expresses the costs of unoptimal price in a quadratic form by doing a second order Taylor approximation of the profit function given by equation 11 around the steady state frictionless markup $\gamma/(\gamma-1)$, à la Rotemberg (1987). Then,

$$\Pi(\{\mu_{i;t}\}_{i=1}^N) \approx \Pi(\mu^*) + \underbrace{\sum_{i=1}^N \frac{\partial \Pi_t}{\partial \mu_{i;t}}\bigg|_{\{\mu_{i;t}=\mu^*\}_{i=1}^N} (\mu_{i;t} - \mu^*)}_{=0} + \frac{1}{2}\sum_{i=1}^N \frac{\partial^2 \Pi_t}{\partial \mu_{i;t}^2}\bigg|_{\{\mu_{i;t}=\mu_*\}_{i=1}^N} (\mu_{i;t} - \mu^*)^2$$

$$\Pi(\{\mu_{i;t}\}_{i=1}^N) \approx \Pi(\mu^*) + \frac{1}{2}\sum_{i=1}^N \frac{\partial^2 \Pi_t}{\partial \mu_{i;t}^2}\bigg|_{\{\mu_{i;t}=\mu_*\}_{i=1}^N} \left(\frac{\mu_{i;t} - \mu^*}{\mu^*}\right)^2 (\mu^*)^2$$

[13] Yang (2019) rationalizes a product-specific cost setting up a model of rational inattentive producers that face a cost for acquiring product-specific (in addition to aggregate) information about the optimal price.

[14] In section 3.5 I formally define my characterization of familiar products, and provide evidence in support of the idea of a cost waiver.

$$\Pi\big(\{\mu_{i;t}\}_{i=1}^{N}\big) \approx \Pi\big(\mu^*\big) + \frac{1}{2}\sum_{i=1}^{N} \frac{\partial^2 \Pi_t}{\partial \mu_{i;t}^2}\Bigg|_{\{\mu_{i;t}=\mu*\}_{i=1}^{N}} (\hat{\mu}_{i;t})^2(\mu^*)^2 \tag{14}$$

Where $\hat{\mu}_{i;t} = \log(\mu_{i;t}/\mu^*)$ is the markup gap (log-deviation from the steady state markup).

The expected losses from deviations from the optimal prices can be expressed as:

$$\mathscr{L} = \mathbb{E}\Big[\Pi\big(\{\mu_{i;t}\}_{i=1}^{N}\big) - \Pi\big(\mu^*\big)\Big]$$

Note that absent menu costs, the firm would choose in each period a price for all its product such that the markup gap is 0, $\{\mu_{i;t}\}_{i=1}^{N} = \mu^*$. However, in presence of nominal rigidities arising from costs of price adjustment, the firm's problem becomes to choose the set of prices that minimizes the expected losses given by (as a fraction of total revenue):

$$\mathscr{L} = \mathbb{E}\left[\frac{\Pi\big(\{\mu_{i;t}\}_{i=1}^{N}\big) - \Pi\big(\mu^*\big) - \Phi_t}{R\big(\mu^*\big)}\right]. \tag{15}$$

Where Φ_t is the multidimensional menu cost, as defined in equation 13. Combining equations 14 and 15 and multiplying the first by $\Pi\big(\mu^*\big)/\Pi\big(\mu^*\big)$ gives:

$$\mathscr{L} = \mathbb{E}\left[\frac{1}{2}\frac{\Pi\big(\mu^*\big)}{R\big(\mu^*\big)}\frac{\sum_{i=1}^{N} \frac{\partial^2 \Pi_t}{\partial \mu_{i;t}^2}\Big|_{\{\mu_{i;t}=\mu*\}_{i=1}^{N}} (\hat{\mu}_{i;t})^2(\mu_j^*)^2}{\Pi\big(\mu^*\big)} - \tilde{\Phi}_t\right]. \tag{16}$$

Where $\tilde{\Phi}_t = \Phi_t/R(\mu^*)$ is the multidimensional menu cost as a fraction of steady state revenues.

Moreover, since

$$\Pi\big(\mu^*\big) = \sum_{i=1}^{N}(\mu^*-1)(\mu^*)^{-\gamma}W_t^{1-\gamma}P_{i;t}^{\gamma-\theta}P_t^{\theta}Y_t,$$

$$R\big(\mu^*\big) = \sum_{i=1}^{N}(\mu^*)^{1-\gamma}(W_t)^{1-\gamma}P_{i;t}^{\gamma-\theta}P_t^{\theta}Y_t, \quad \text{and}$$

$$\frac{\partial^2 \Pi_t}{\partial \mu_{i;t}^2}\Bigg|_{\{\mu_{i;t}=\mu*\}_{i=1}^{N}} = -\gamma(\mu^*)^{-\gamma-2}W_t^{1-\gamma}P_{i;t}^{\gamma-\theta}P_t^{\theta}Y_t,$$

equation 16 can be expressed as:

$$\mathscr{L} = \mathbb{E}\left[-\gamma\frac{1}{2}\frac{\sum_{i=1}^{N}(\mu^*-1)(\mu^*)^{-\gamma}W_t^{1-\gamma}P_{i;t}^{\gamma-\theta}P_t^{\theta}Y_t}{\sum_{i=1}^{N}(\mu^*)^{1-\gamma}(W_t)^{1-\gamma}P_{i;t}^{\gamma-\theta}P_t^{\theta}Y_t}\frac{\sum_{i=1}^{N}(\mu^*)^{-\gamma}W_t^{1-\gamma}P_{i;t}^{\gamma-\theta}P_t^{\theta}Y_t(\hat{\mu}_{i;t})^2}{\sum_{i=1}^{N}(\mu^*-1)(\mu^*)^{-\gamma}W_t^{1-\gamma}P_{i;t}^{\gamma-\theta}P_t^{\theta}Y_t} - \tilde{\Phi}_t\right]$$

Operating and cancelling terms:

$$\mathscr{L} = \mathbb{E}\left[-\gamma\frac{1}{2}\frac{1}{\mu^*}\frac{\sum_{i=1}^{N}P_{i;t}^{\gamma-\theta}(\hat{\mu}_{i,z;t})^2}{\sum_{i=1}^{N}P_{i;t}^{\gamma-\theta}} - \tilde{\Phi}_t\right]$$

Notice than from the definition of $\hat{\mu}_{i,z;t}$:

$\hat{\mu}_{i,z;t} = \log(\mu_{i,z;t}) - \log(\mu^*) = \log(P_{i,z;t}) - \log(P^*_{i,z;t}) = p_{i,z;t} - p^*_{i,z;t}$.

Then, the loss function is:

$$\mathscr{L} = \mathbb{E}\left[-\gamma \frac{1}{2}\frac{1}{\mu^*}\frac{\sum_{i=1}^{N} P_{i;t}^{\gamma-\theta}(p_{i,z;t} - p^*_{i,z;t})^2}{\sum_{i=1}^{N} P_{i;t}^{\gamma-\theta}} - \tilde{\Phi}_t \right] \tag{17}$$

Lastly, I make three extra assumptions. First, I follow Yang (2019) and assume that the elasticity of substitution between goods equals the elasticity of substitution between firms: $\theta = \gamma$. This assumption simplifies the computational solution and makes the model more tractable, and its influence is little given the minor role that the elasticity of substitution across goods plays in this model.[15] Second, I assume that familiar products share the same cost curtail, $\tilde{\phi}_i^C$, and product-specific cost of adjustment, $\tilde{\phi}_i^S$; and that the cost curtail for changing n products from the same family in the same period is $\tilde{\phi}_i^C = [(n-1)\tilde{\phi}_i^S]/n$ per changed price, which is equivalent to assuming that the firm has to pay only one product-specific cost per family of products and receives a cost waiver of the other $n-1$ product-specific costs. For instance, this suggests that a supermarket that changes the price of 'Coke Light 375ml', 'Coke Light 600ml', and 'Coke Light 1500ml' only pays *one* product-specific cost. Finally, I assume that familiar products are subject to the same idiosyncratic quality shocks. These last two assumptions imply that whenever a product departs from its optimal price, all the products of the same family will also do so; and that when a firm takes a price-adjustment decision its policy will be the same for all the products of the same family (i.e. either change all the prices or not change any price) generating perfect economies of scope in price adjustment of familiar products.

Then, firm z's problem can be written as:

$$\max_{\{p_{i,z;t}\}_{i=1}^{N}} \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \left(-\gamma \frac{1}{2}\frac{1}{\mu^*}\sum_{i=1}^{N}(p_{i,z;t} - p^*_{i,z;t})^2 \right) - \tilde{\Phi}_t \right] \tag{18}$$

For concreteness and clarity, I consider a firm that sells $N = 3$ products with products $i \in \{2,3\}$ being from the same family. The assumptions above imply that the firms' decisions over the prices of $i \in \{2,3\}$ are identical: in every period it changes either both prices or none of them. Moreover, they imply that $\tilde{\phi}_2^S = \tilde{\phi}_3^S$; and that the cost curtails received when changing both familiar products' prices are $\tilde{\phi}_2^C = \tilde{\phi}_3^C = \frac{1}{2}\tilde{\phi}_3^S = \frac{1}{2}\tilde{\phi}_2^S$. As a result, when the firm changes the prices of the two familiar products (and, say, leaves the price of product $i = 1$ unchanged), it pays $\tilde{\Phi} = \tilde{\phi}^G + \tilde{\phi}_2^S + \tilde{\phi}_3^S - \tilde{\phi}_2^C - \tilde{\phi}_3^C = \tilde{\phi}^G + \tilde{\phi}_2^S$, leaving no role for $\tilde{\phi}_3^S$.

Let $V^{C,C,C}(P_{i,z;t-1}, A_{i,z;t}, M_t)$, $V^{C,N,N}(P_{i,z;t-1}, A_{i,z;t}, M_t)$, $V^{N,C,C}(P_{i,z;t-1}, A_{i,z;t}, M_t)$, and $V^{N,N,N}(P_{i,z;t-1}, A_{i,z;t}, M_t)$ denote the firm's value of (i) adjusting all its prices, (ii) adjusting the price of the product $i = 1$ and leaving the other prices unchanged, (iii) adjusting the price of the products $i = 2,3$ and leaving the price of product $i = 1$, and (iv) leaving all prices unchanged.

[15] Concretely, this assumption implies that the elasticities of substitution between 'Coke Light 600ml' and 'Coke Light 1500ml', and between 'Coke Light 600ml' and 'Aromatic candle wooden vanilla scented 200gr' are the same, and that their value is also identical to the elasticity of substitution between two different stores.

Letting $V = \max\{V^{C,C,C}, V^{C,N,N}, V^{N,C,C}, V^{N,N,N}\}$ be the envelope of these four options, the recursive formulation of the firm's problem is characterized by the following system:

$$\begin{aligned}
\mathrm{V}^{C,C,C}(\{P_{i,z;t-1}\}_{i=1}^{3},\{A_{i,z;t}\}_{i=1}^{3},M_t) = &\max_{\{p_{1,z;t},p_{2,z;t},p_{3,z;t}\}}\Big(-\gamma\tfrac{1}{2}\tfrac{1}{\mu^*}\textstyle\sum_{i=1}^{3}(p_{i,z;t}-p^*_{i,z;t})^2 - \\
&-\tilde{\phi}^G-\tilde{\phi}_1^S-\tilde{\phi}_2^S+\beta\mathbb{E}_t\Big[V\{P_{i,z;t}\}_{i=1}^{3},\{A_{i,z;t+1}\}_{i=1}^{3},M_{t+1})\Big]\Big)\\
\mathrm{V}^{C,N,N}(\{P_{i,z;t-1}\}_{i=1}^{3},\{A_{i,z;t}\}_{i=1}^{3},M_t) = &\max_{\{p_{1,z;t}\}}\Big(-\gamma\tfrac{1}{2}\tfrac{1}{\mu^*}\textstyle\sum_{i=1}^{3}(p_{i,z;t}-p^*_{i,z;t})^2 - \\
&-\tilde{\phi}^G-\tilde{\phi}_1^S+\beta\mathbb{E}_t\Big[V\{P_{i,z;t}\}_{i=1}^{3},\{A_{i,z;t+1}\}_{i=1}^{3},M_{t+1})\Big]\Big)\\
\mathrm{V}^{N,C,C}(\{P_{i,z;t-1}\}_{i=1}^{3},\{A_{i,z;t}\}_{i=1}^{3},M_t) = &\max_{\{p_{2,z;t},p_{3,z;t}\}}\Big(-\gamma\tfrac{1}{2}\tfrac{1}{\mu^*}\textstyle\sum_{i=1}^{3}(p_{i,z;t}-p^*_{i,z;t})^2 - \\
&-\tilde{\phi}^G-\tilde{\phi}_2^S+\beta\mathbb{E}_t\Big[V(\{P_{i,z;t}\}_{i=1}^{3},\{A_{i,z;t+1}\}_{i=1}^{3},M_{t+1})\Big]\Big)\\
\mathrm{V}^{N,N,N}(\{P_{i,z;t-1}\}_{i=1}^{3},\{A_{i,z;t}\}_{i=1}^{3},M_t) = &\Big(-\gamma\tfrac{1}{2}\tfrac{1}{\mu^*}\textstyle\sum_{i=1}^{3}(p_{i,z;t}-p^*_{i,z;t})^2 + \\
&+\beta\mathbb{E}_t\Big[V(\{P_{i,z;t}\}_{i=1}^{3},\{A_{i,z;t+1}\}_{i=1}^{3},M_{t+1})\Big]\Big)
\end{aligned} \tag{19}$$

5. Quantitative solution and simulation results

5.1. Computational procedures

I solve the recursive problem given above maintaining the nature of the firm (it sells $N = 3$ products with products $i \in \{2,3\}$ being from the same family). To do so, I adapt to my model the procedure of Nakamura and Steinsson (2010), who solve a single-product menu cost model with a value function iteration method. The main difference with their solution is that the state in my problem is of a higher dimension because of the multiproduct nature of the firms in it.

In Section 4.3 I assumed that familiar products have identical processes for their quality $A_{i,z}$. This implies that their optimal prices are the same as well (see equation 10), and thus that familiar products have an identical markup gap in every period. Then, I can transform the three-products model into a two-products model in the numerical solution, by treating the two familiar products as one unique product with a two times greater slope of the loss function.

As the value of state variables (the quality and the inherited price) of the familiar products is the same, I can drop a set of $\{A_{i,z;t}, P_{i,z;t}\}$ from the state vector. This simplifies the computational burden to a great extent. Indeed, there are four state variables in the problem, which define the prior price gaps of the firm: two posted prices inherited from the last period and two qualities. I solve the firms' problem and the value functions, and obtain the optimal policy function using the algorithm described hereunder:

1. Choose a relative error tolerance level, τ.
2. Discretize the state space by constructing a grid for the inherited prices and the qualities: $p^i = \{p^i_1, p^i_2, p^i_3, \ldots, p^i_{np}\}_{i=1}^{2}$, and $a^i = \{a^i_1, a^i_2, a^i_3, \ldots, a^i_{na}\}_{i=1}^{2}$

3. Make an initial guess of the termination value function, $V^{(0)}$. This is a $(n_p^2 \times n_a^2)-$matrix. I choose the mean expected loss from inaction as the initial guess.
4. Function iteration: use the Bellman equation for the firm's problem to iterate on the expected value of the different value functions, and update the termination value $V^{(1)}$.
5. Compute the distance between the previous and the updated value as:

$$d = \frac{\beta}{1-\beta}\left[\max\left(V^{(1)} - V^{(0)}\right) - \min\left(V^{(1)} - V^{(0)}\right)\right]$$

6. If the distance is within the error tolerance, $d \leq \tau$, the value function has converged and one can obtain a numerical estimate of the termination value. If $d > \tau$, return to step (4), replacing the initial guess with the updated value, $V^{(0)} = V^{(1)}$. Repeat steps (4)-(6) until the value function has converged.
7. Once the value function has converged, take expectations of the termination value using the transition probabilities of the prices and qualities due to inflation and idiosyncratic shocks, respectively. To obtain the transition probabilities of states I use Tauchen (1986) approximation method for autoregressive processes.
8. Obtain (i) the numerical estimate of the value and (ii) the policy function, defined –for each state– as the optimal decision (CC, CN, NC, NN) and the optimal price choice in case of price adjustment.

5.2. Calibration and parametrization

For comparison purposes with previous studies, I choose a period-length of one month for my solution. I set a monthly discount factor of $\beta = 0.96^{1/12}$ which implies an annualized real interest rate of 0.04. I set the the elasticity of substitution across firms to $\gamma = 4$ resulting in an implied frictionless markup of one third, in line with Nakamura and Steinsson (2010) and Alvarez et al. (2018). This value is within the range of those estimates reported by the industrial organization literature [e.g. Berry et al. (1995) and Nevo (2001)] and lies in between the values used by Golosov and Lucas (2007), $\gamma = 7$, and Midrigan (2011), $\gamma = 3$. The elasticity of substitution across firms directly affects the slope of the loss function, thus a greater γ generates a higher frequency of price changes in the model everything else constant.

I set the drift of the random walk process of money supply, η, to equal the monthly inflation rate of each of the countries of my analysis, and calibrate the standard deviation of monetary shocks to be directly proportional to the inflation rate, as $\sigma_M = 1.7\eta$. Finally, I follow Nakamura and Steinsson (2010) and set the speed of mean reversion of the product qualities equal to $\rho = 0.7$, and I calibrate the standard deviation of the idiosyncratic quality shocks σ_{A_1} and σ_{A_2} as 0.05225 and 0.07250 respectively.[16] The choice of different variances of the idiosyncratic shocks is supported by the multiproduct nature of the firm. Setting distinct values for this parameter makes

[16] Nakamura and Steinsson (2010) set this parameter to make their model match their data and obtain a value of $\sigma_A = 0.0425$.

the processes of the two optimal prices differ, and thus prevents my model from being a multi-(identical)-product model.

The size of the costs of adjustment are the key parameters that remain to be calibrated. As it is typically done in the sticky prices literature, I calibrate their value to allow the model to fit the microprice facts documented in Section 3. Concretely, I target the key aspects of the distribution of price changes, such as the average size, the share of price adjustments that are increases, and the fraction of small price changes. I do not set different values of the costs of adjustment for *each* country in order to match the main statistics of all the distributions of price changes, but instead I find the values of ϕ^G and ϕ^S that perform better in reproducing the main patterns of price changes *across* different countries. Hence, the only country-specific parameter in my solution is given by the inflation rate. Naturally, this aim for generality in my solution comes at the cost of more precision, but as I show in the following subsection, my general model performs good in matching the main facts of price changes. The values I set for the menu costs as a fraction of the firms' revenues are $\tilde{\phi}^G = 0.0060$ and $\tilde{\phi}^S = 0.0001$.[17]

Table 7: **Calibrated and assigned parameters**

Parameter	
1. Subjective discount factor, β	$0.96^{1/12}$
2. Elasticity of substitution across firms, γ	4
3. Implied frictionless markup, μ^*	1.33
4. Drift of the random walk process of money supply, η	
The Netherlands	0.0021
United Kingdom	0.0015
Chile	0.0023
Brazil	0.0030
Turkey	0.0094
Argentina	0.0271
5. Standard deviation of monetary policy shocks, σ_M	1.7η
6. Speed of mean reversion of the product qualities, ρ	0.7
7. Standard deviation of the idiosyncratic quality shocks	
Of product 1, σ_{A_1}	0.05225
Of product 2, σ_{A_2}	0.07250
8. General cost of price adjustment as a fraction of revenue, ϕ^G	6×10^{-3}
9. Specific cost of price adjustment as a fraction of revenue, ϕ^S	1×10^{-4}

5.3. Simulation results

I give to the tolerance parameter τ a value equal to 0.1% of the mean value from inaction. Moreover, I set the size of the grids for each price and each quality to 150 and 30 points, respectively, which implies that initially the firm's problem has more than 20 million alternative states.[18] I simulate the solution 30 times for each country. Table 8 lists the values of the main statistics of price

[17] The calibration of the menu cost parameter in previous papers tends to be dissimilar, partly depending on the micro facts the different models aim to match. For example, Midrigan (2011) sets the value of the cost (analogous to $\tilde{\phi}^G$ in his model) relative to the steady state revenue to 0.018. The cost of changing prices in Yang (2019) has a similar role as that of the Midrigan model, and it is set equal to 0.0342 in a two-products model. Karadi and Reiff (2019) calibrate the menu cost as 2.4% of steady state revenues, paid with a 6.25% probability, so the overall menu cost in their random menu cost model is 0.0015.

[18] Recall that the size of each state vector is n_p^2 and n_a^2 for prices and qualities, respectively.

changes of the model, and compares them with those found in the daily scraped data as presented in section 3.

Incorporating only one country-specific variable –the inflation rate– the model reproduces many of the facts previously documented. Both the mean and the median size of price adjustment are positively correlated with inflation, and the values for those statistics are close to those found in the data. The two cases with higher deviations between the simulation and the data are those from the United Kingdom and Argentina. As commented in Section 3, the main statistics of price changes in the data of the retailer from the United Kingdom notably differ from those of the other countries (and especially from countries with similar macroeconomic conditions, who have smaller and less dispersed price changes than the rest of the sample), and those differences cannot be explained by any aggregate country-specific variable. On the other hand, the discrepancies between the simulation and the data from Argentina arise because of the considerably higher inflation rate compared to the other countries in my sample. This reflects the generality-precision trade-off mentioned above: the parametrization of the model that best fits the distribution of low, mid and high-inflation countries overestimates the departures from the optimal price p^* arising from high inflation.

The predicted share of price increases also matches satisfactorily the data. The correlation between this statistic and inflation has a positive sign, and the values range from 62% to 67% which is also consistent with previous values reported in the literature.

Moreover, the model fails to match two statistics from the data: the dispersion of the size of price changes and the frequency of price adjustment. When it comes to the dispersion, the model predicts a greater standard deviation of price changes for higher inflation rates, since the optimal price p^* depends directly on the supply of money, which is calibrated as the rate of inflation in my numerical solution.[19] As a result, the calibration of the model that best fits the data underestimates the standard deviation of price changes of countries with low and medium inflation to prevent generating unrealistically large absolute price changes for Argentina. There are at least two potential solutions to this issue. The first one requires a characterization of the model where the pass-through from monetary policy to the optimal price is different than the 1-for-1 relation I introduce. The second one involves the introduction of Poisson shocks, instead of Gaussian, to the process driving optimal prices. This is the strategy followed by Midrigan (2011) to generate a fat-tailed distribution with highly dispersed price changes.

In addition, the predicted frequency of steady-state price changes is higher than the one in the data. The reason for this is that in the solution the firm receives many prices that are sufficiently far from their optimal values to be changed. In spite of this shortcoming, the negative relation between inflation and duration of price spells exhibited in Figure 5 is also present in the predictions of my model.[20]

[19]The lower dispersion of price changes obtained from the model is manifested in a lower standard deviation than that observed in the data. This also results in lower absolute values of the first and third quartiles; and it is pictured in Figure 9 which has less dispersed distributions than Figure 7.

[20]Recall the negative relation between the frequency of price changes and the implied duration of price spells, given by equation 1.

Finally, I focus on the relatively low, yet nontrivial, number of small price changes found in the data. Recall that this is one of the novel facts of price changes that is not present in traditional microdata sources, and that it is one of the grounds for introducing a multi-dimensional cost of price adjustment in the model. Row 8 in Table 8 compares the predictions of the model with the statistics obtained in the data, and also adds between brackets those values reported in Cavallo (2018) for robustness checks. The model produces a fraction of small price adjustments (measured as the fraction of log-price changes smaller than 0.01 in absolute terms) that is close to that observed in the data, and lies in between the predictions from the two extreme cases of menu-cost models that predict no small price changes [e.g. a multiproduct version of the Golosov-Lucas model], or a relatively large mass of small price changes [e.g. a Midrigan-type economy]. This feature of the model is also depicted in Figure 9, which shows that the predicted distribution of price changes has a dip of the density around zero, similar to that of Figure 7.

Table 8: **The model's partial equilibrium predictions against the data**

Country	Netherlands		UK		Chile		Brazil		Turkey		Argentina	
	Data	Simul.	Data	Simul.	Data	Simul.	Data	Simul.	Data	Simul.	Data[§]	Simul.
1. Mean	1.70	1.32	7.74	1.21	1.75	1.35	2.33	1.58	2.72	2.70	4.43	6.15
2. Median	3.15	1.73	3.92	1.69	3.06	1.81	3.52	1.98	5.72	2.82	6.19	6.34
3. 25th perc.	-2.65	-1.92	-9.53	-1.82	-9.44	-1.92	-5.21	-1.95	-8.07	-2.28	-3.19	-3.22
4. 75th perc.	5.75	4.09	28.76	3.81	11.20	4.22	7.83	4.83	12.52	7.00	11.51	14.96
5. Std. Dev.	6.58	4.02	27.07	3.81	18.29	4.11	10.12	4.37	16.86	7.20	12.48	15.50
6. Share of increases	0.63	0.63	0.61	0.62	0.63	0.63	0.64	0.64	0.65	0.64	0.69	0.67
7. Fraction of $\lvert\Delta p\rvert < 1\%$	3.6% (3.2%)	5.0%	1.4% (4.9%)	5.9%	1.6% (3.5%)	4.2%	4.1% (4.1%)	4.0%	2.1% (1.9%)	3.3%	3.7% (2.8%)	2.5%
8. Skewness	-1.04	-0.01	0.20	-0.1	-0.01	-0.17	-0.15	-0.19	-0.46	-0.08	-0.39	-0.13
9. Exccess kurtosis	1.57	-0.69	0.43	-0.70	-0.30	-0.68	3.21	-0.56	0.48	-0.58	0.54	-0.19

Notes: The size of price changes is estimated as 100 times the log difference of prices. The excess kurtosis is calculated as $3 - kurtosis$. The values in parentheses in row 8 are those reported in Cavallo (2018).
[§] The values for Argentina are obtained as the average of each statistic across the three retailers.

Alternative calibrations for the menu cost

The model encompasses the two extreme cases of uni-dimensional cost of regular price change. A model with only product-specific costs of price adjustment (a multiproduct version of Golosov and Lucas (2007)) generates a bimodal distribution of the size of price changes, with null changes between the positive and negative threshold given by the size of the cost. To reproduce this type of model, I set $\phi_G = 0$ in my calibration. Contrarily, a model where the firm has to pay only one general cost to change any number of regular prices (à la Midrigan (2011) excluding temporary changes) predicts a relatively high density around zero. To reproduce this type of model, I set $\phi_S = 0$ in my calibration. Figure 1 in Section 1 compares the predictions of both special cases of the model with the data from the Netherlands ($\mu = 0.002$).

Figure 9: **Histograms of price changes distribution for the model simulation**

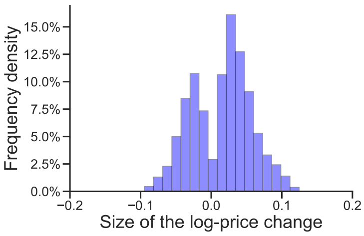

(a) The Netherlands

(b) United Kingdom

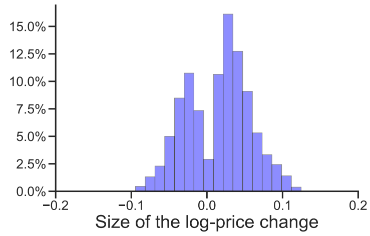

(c) Chile

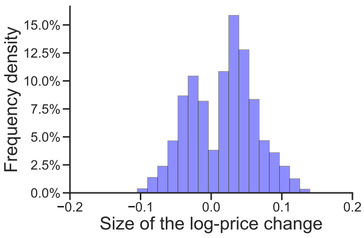

(d) Brazil

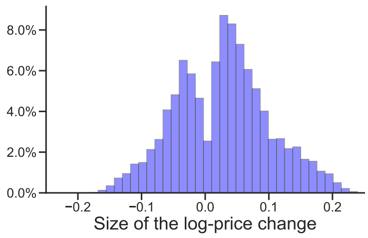

(e) Turkey

(f) Argentina

The selection effect related to the size of price changes in the baseline case of the model lies between that of the two standard menu cost models. In a Golosov-Lucas economy (with $\phi^G = 0$), those firms that adjust a price in a period are those whose price is a great distance away from its optimal value, denoting a strong selection effect. Then, when a monetary shock takes place (even a small one), prices that were close to the inaction threshold will adjust by a large amount, resulting in a large response of the aggregate level to monetary innovations and a large degree of price flexibility in such economy. Contrarily, by allowing for economies of scope in price adjustment, the selection effect in a Midrigan economy (with $\phi^S = 0$) is considerably smaller. A multiproduct firm will incur a large number of small price changes as a response to deviations from their optimal values arising from a monetary shock. As a result, the aggregate response of the price level will be smaller than in the Golosov-Lucas case, generating greater real effects of monetary shocks.

In their formalization of the relevance of the selection effect for the degree of monetary non-neutrality, Alvarez et al. (2016) give a central role to the kurtosis of price changes. They find that the kurtosis of the steady-state distribution of price changes is a sufficient statistic for the cumulative real effects of monetary shocks: for a given frequency of price adjustment, a higher kurtosis results in larger cumulative output effects of monetary policy, measured as the area under the impulse response function.

In the Golosov-Lucas model the distribution of price changes has the smallest value of (excess) kurtosis, which is -2, as all the price changes are concentrated around very large and very small values. The extreme case of my model where $\phi^G = 0$ yields a distribution with a larger, yet small, excess kurtosis of -1.2. The reason for the difference is that even in the extreme case my model generates a slightly higher steady-state dispersion than the Golosov-Lucas model. On the other hand, Midrigan (2011) model without temporary changes has an excess kurtosis close to 0 (-0.3 in the version of my model with $\phi^S = 0$). The benchmark model I set up in this paper has an excess kurtosis of between -0.7 and -0.2, varying across countries (see row 10 in Table 8). This results in a predicted larger real effect of monetary shocks than in a Golosov-Lucas economy, and smaller than in a Midrigan economy, according to Alvarez-Lippi-Le Bihan sufficient statistic approach.[21]

6. Conclusion

Firms' price-setting behavior plays a crucial role in the New Keynesian framework: without firms' capacity to adjust prices being limited these models would not display monetary non-neutrality, which is a widely accepted feature in the macroeconomic literature. Moreover, a majority of papers incorporates Calvo (1983) pricing as the source of price rigidity, partly because of its tractability.

Thanks to the recent emergence of new micro price data sets, there was an advancement in the debate on price stickiness, with various papers documenting the main facts of price changes and studying different alternatives to Calvo pricing. Many of them found that state-dependent pricing –where the timing of firms' price adjusting decision is endogenous of their profit maximization problem, such as menu cost models– performs better in matching these micro facts. However, these data sources (namely CPI and scanner price data) are not free from limitations, mainly caused by measurement and imputation errors (as showed by Cavallo (2018)), generally giving rise to an overestimation of small price changes and an underestimation of the duration of price spells.

In this paper, I follow up on this discussion by making use of a new micro price daily data set that I collected with web scraping techniques. I provide new evidence about firms' price-setting behavior and formalize my main findings in a menu cost model. I present three novel facts about price changes. First, there is a relation between the main statistics and the inflation rate of a country. Concretely, higher inflation rates are typically associated with a larger average size of price changes, with a larger share of price increases, and with a lower duration of price spells. Second, the distribution of the size of price changes has a relatively small, yet nontrivial mass around zero, which differs from the shape of the distribution found in alternative data sources and also from that predicted by the benchmark menu cost models from Golosov and Lucas (2007) and Midrigan (2011). And third, familiar products from the same manufacturer have greater similarity in the timing and magnitude of price adjustment than heterogeneous products, which suggests the necessity of a special treatment for familiar products in the incorporation of price rigidities in a multiproduct firm environment. I show that incorporating a three-dimensional cost –composed by a general cost, a product-specific cost, and a cost curtail for price changes in familiar products– makes an otherwise standard menu cost model reproduce these facts.

References

Alvarez, F., H. Le Bihan, and F. Lippi (2016). The real effects of monetary shocks in sticky price models: A sufficient statistic approach. *American Economic Review 106*(10), 2817–2851.

[21] The characterization of the kurtosis as a sufficient statistic remains valid on economies with zero or low inflation in the Alvarez-Lippi-Le Bihan framework.

Alvarez, F. and F. Lippi (2014). Price Setting With Menu Cost for Multiproduct Firms. *Econometrica 82*(1), 89–135.

Alvarez, F. and F. Lippi (2020). Temporary Price Changes, Inflation Regimes, and the Propagation of Monetary Shocks. *American Economic Journal: Macroeconomics 12*(1), 104–152.

Alvarez, F., F. Lippi, and L. Paciello (2018). Monetary shocks in models with observation and menu costs. *Journal of the European Economic Association 16*(2), 353–382.

Alvarez, L. J. (2008). What do micro price data tell us on the validity of the new keynesian phillips curve? *Economics - The Open-Access, Open-Assessment E-Journal 2*, 1–36.

Anderson, E., B. A. Malin, E. Nakamura, D. Simester, and J. Steinsson (2017). Informational rigidities and the stickiness of temporary Sales. *Journal of Monetary Economics 90*(C), 64–83.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica 63*(4), 841–890.

Bils, M. and P. J. Klenow (2004). Some Evidence on the Importance of Sticky Prices. *Journal of Political Economy 112*(5), 947–985.

Bonomo, M., C. Carvalho, O. Kryvtsov, S. Ribon, and R. Rigato (2020). Multi-Product Pricing: Theory and Evidence from Large Retailers in Israel. Staff Working Papers 20-12, Bank of Canada.

Burstein, A. T. (2006). Inflation and output dynamics with state-dependent pricing decisions. *Journal of Monetary Economics 53*(7), 1235–1257.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics 12*(3), 383–398.

Cavallo, A. (2018). Scraped Data and Sticky Prices. *The Review of Economics and Statistics 100*(1), 105–119.

Cavallo, A. and R. Rigobon (2016). The Billion Prices Project: Using Online Prices for Measurement and Research. *Journal of Economic Perspectives 30*(2), 151–178.

Dhyne, E., L. J. Alvarez, H. Le Bihan, G. Veronese, D. Dias, J. Hoffmann, N. Jonker, P. Lunnemann, F. Rumler, and J. Vilmunen (2006). Price changes in the euro area and the united states: Some facts from individual consumer price data. *Journal of Economic Perspectives 20*(2), 171–192.

Golosov, M. and R. E. Lucas, Jr. (2007). Menu Costs and Phillips Curves. *Journal of Political Economy 115*, 171–199.

Hansen, G. D. (1985). Indivisible labor and the business cycle. *Journal of Monetary Economics 16*(3), 309–327.

Karadi, P. and A. Reiff (2019). Menu Costs, Aggregate Fluctuations, and Large Shocks. *American Economic Journal: Macroeconomics 11*(3), 111–146.

Kehoe, P. and V. Midrigan (2015). Prices are sticky after all. *Journal of Monetary Economics 75*(C), 35–53.

Klenow, P. and J. Willis (2007). Sticky information and sticky prices. *Journal of Monetary Economics 54*(Supplement 1), 79–99.

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*(3), 863–904.

Klenow, P. J. and B. A. Malin (2010). Microeconomic Evidence on Price-Setting. In B. M. Friedman and M. Woodford (Eds.), *Handbook of Monetary Economics*, Volume 3 of *Handbook of Monetary Economics*, Chapter 6, pp. 231–284. Elsevier.

Mankiw, N. G. and R. Reis (2002). Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve. *The Quarterly Journal of Economics 117*(4), 1295–1328.

Midrigan, V. (2011). Menu Costs, Multiproduct Firms, and Aggregate Fluctuations. *Econometrica 79*(4), 1139–1180.

Nakamura, E. and J. Steinsson (2008). Five Facts about Prices: A Reevaluation of Menu Cost Models. *The Quarterly Journal of Economics 123*(4), 1415–1464.

Nakamura, E. and J. Steinsson (2010). Monetary Non-neutrality in a Multisector Menu Cost Model. *The Quarterly Journal of Economics 125*(3), 961–1013.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica 69*(2), 307–342.

Rogerson, R. (1988). Indivisible labor, lotteries and equilibrium. *Journal of Monetary Economics 21*(1), 3–16.

Rotemberg, J. (1987). The new keynesian microfoundations. In *NBER Macroeconomics Annual 1987, Volume 2*, pp. 69–116. National Bureau of Economic Research.

Schulz, K. U. and M. Stoyan (2002). Fast string correction with Levenshtein automata . *International Journal on Document Analysis and Recognition 5*, 67–85.

Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters 20*(2), 177–181.

Yang, C. (2019). Rational Inattention, Menu Costs, and Multi-Product Firms: Micro Evidence and Aggregate Implications. Job market paper.

Zbaracki, M. J., M. Ritson, D. Levy, S. Dutta, and M. Bergen (2004). Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets. *The Review of Economics and Statistics 86*(2), 514–533.

Appendix A Appendix

A.1 Derivation of the Levenshtein ratio

$$lev_{v,w}(i,j) = \begin{cases} \max(i,j) & \text{if } \min(i,j)=0 \\ \min \begin{cases} \text{lev}_{v,w}(i-1,j)+1 & \text{[deletion]} \\ \text{lev}_{v,w}(i,j-1)+1 & \text{[insertion]} \\ \text{lev}_{v,w}(i-1,j-1)+\mathbf{b} & \text{[substitution]} \end{cases} & \text{otherwise} \end{cases} \tag{20}$$

Where $\mathbf{b}=0$ when $v_i = w_j$ and $\mathbf{b}=1$ otherwise, and $\text{lev}_{v,w}(i,j)$ is the distance between the first i characters of v and the first j characters of w.

Once obtained the Levenshtein distance, I can derive the Levenshtein ratio (LR), which I use to compare the similarity between two strings, as:

$$r = \frac{(|v|+|w|) - \text{lev}_{v,w}(i,j)}{|v|+|w|}.$$

By construction $r=1$ when the two strings are identical, and $r \to 0$ for completely different strings. For concreteness, consider the string ‘Coke Light 500ml’. Its LR with ‘Coke Light 375ml’, ‘Sprite 500ml’, and ‘Maple Syrup 375gr’ are 0.88, 0.57, and 0.24, respectively.

A.2 Additional tables and figures

Figure 11: **Histograms of price changes in Billion Prices Project data**

(a) Chile From 10:2007 to 03:2012 | (b) Brazil From 10:2007 to 07:2011 | (c) Argentina From 10:2007 to 03:2011

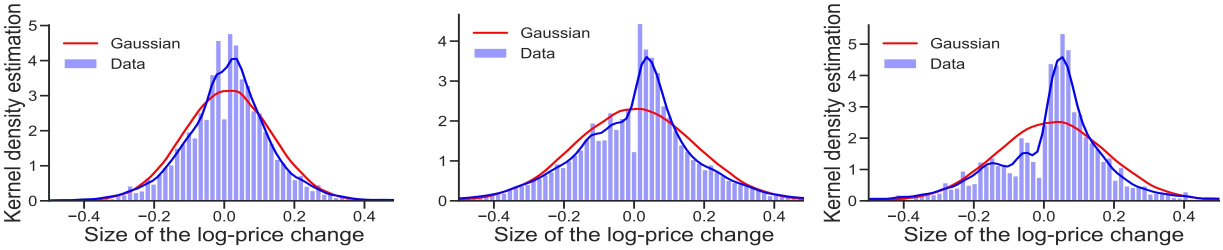

Note: The y-axis corresponds to the kernel density estimation (KDE) of each distribution –plotted in blue lines– which is a non-parametric smoothing method of obtaining a probability distribution. The height of the histograms denotes, as usual, the relative frequency of each bin.

Figure 13: **Concentration in the size of temporary price changes**

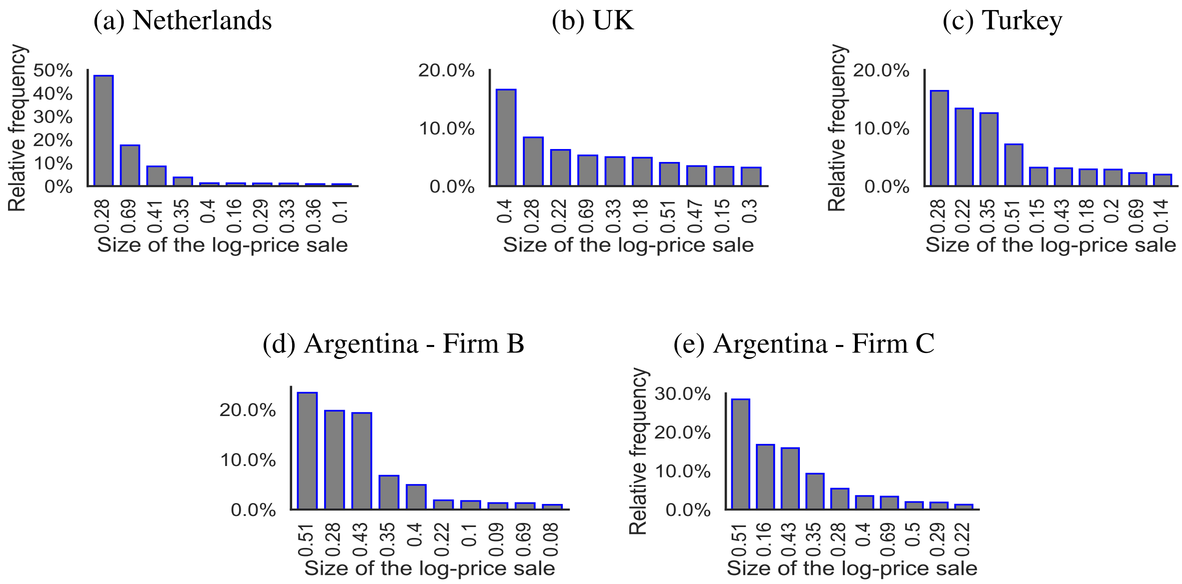

Figure 15: **Range [0, 0.10] of the cumulative distribution function of the absolute size of log-price changes**

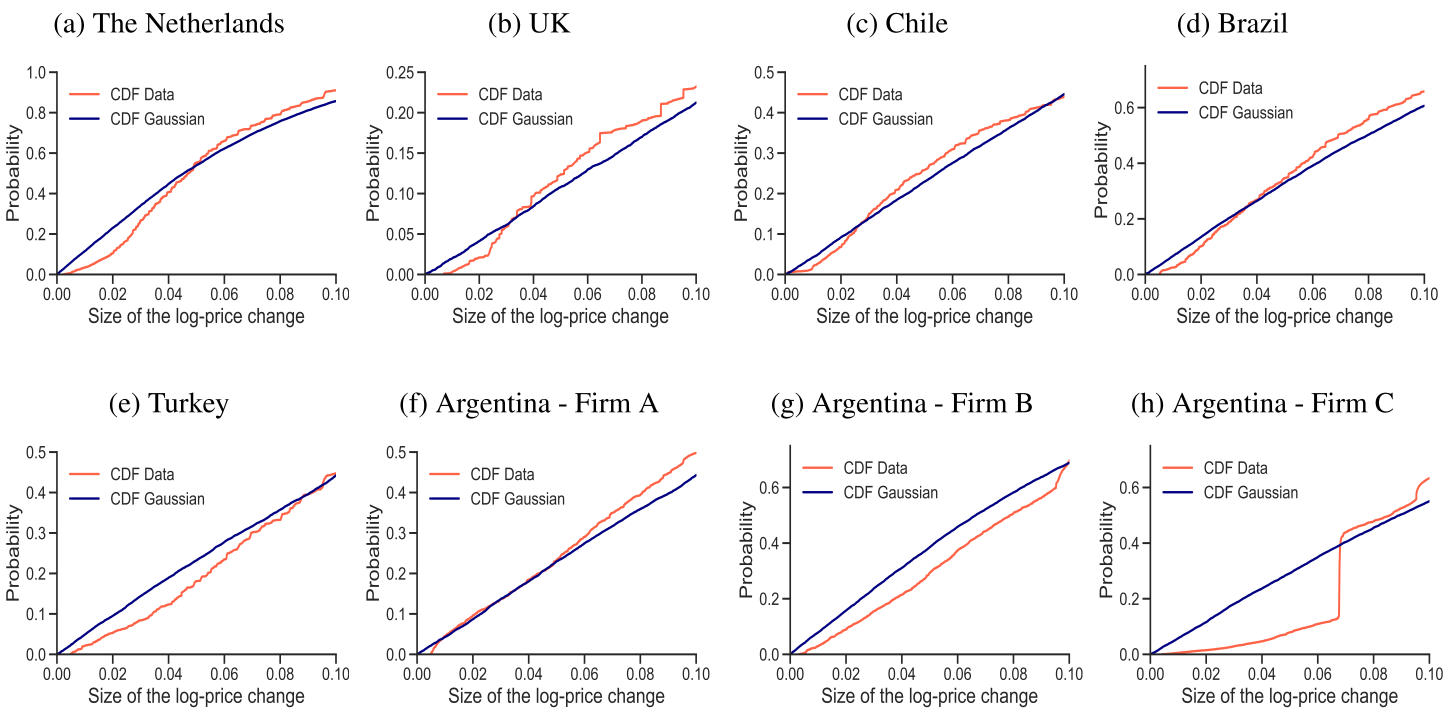

Note: The x-axis is trimmed to the range [0, 0.10] to provide a clearer picture of the number of small price changes. Also note that the scale of the vertical axis differs across retailers depending on the dispersion of the magnitude of changes, also visible in Figure 7.