

NBER WORKING PAPER SERIES

SCRAPED DATA AND STICKY PRICES

Alberto Cavallo

Working Paper 21490

<http://www.nber.org/papers/w21490>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

August 2015

I am extremely grateful to Philippe Aghion, Robert Barro, and Roberto Rigobon for their advice and support with my initial data-scraping efforts. I also wish to thank Alberto Alesina, Gadi Barlevy, Michael Bordo, Jeffrey Campbell, Eduardo Cavallo, Benjamin Friedman, Gita Gopinath, Oleg Itskhoki, Alejandro Justiniano, Pete Klenow, David Laibson, Greg Mankiw, Robert Pindyck, Julio Rotemberg, Tom Stoker, several anonymous referees, and seminar participants at Harvard, MIT, UCLA, LACEA, and the Chicago Federal Reserve for their helpful comments and suggestions. Surveys of offline prices were conducted with financial support from the Warburg Fund at Harvard University, the JFRAP at MIT Sloan, and research assistance from Monica Almonacid, Enrique Escobar, Pedro Garcia, Andre Giudice de Oliveira, Andrea Albagli Iruretagoyena, and Maria Fazzolari. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Alberto Cavallo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Scraped Data and Sticky Prices
Alberto Cavallo
NBER Working Paper No. 21490
August 2015, Revised September 2015
JEL No. E30,E60

ABSTRACT

This paper introduces Scraped Data as a new source of micro-price information to measure price stickiness. Scraped data, collected from online retailers, have no time averaging or imputed prices that can affect pricing statistics in traditional sources of micro-price data. Using daily prices of 80 thousand products collected in five countries with varying degrees of inflation, including the US, I find that relative to previous findings in the literature, scraped online prices tend to be stickier, with fewer price changes close to zero percent, and with hump-shaped hazard functions that initially increase over time. I show that the sampling characteristics of the data, which minimize measurement biases, explain most of the differences with the literature. Using the cross-section of countries, I also show that only the relative frequency of price increases over decreases correlates with inflation.

Alberto Cavallo
MIT Sloan School of Management
100 Main Street, E62-512
Cambridge, MA 02142
and NBER
acavallo@mit.edu

A Data is available at:
<http://dx.doi.org/10.7910/DVN/IAH6Z6>
An Other Replication Materials is available at:
<http://www.mit.edu/~afc/data/data-page-scraped.html>

1 Introduction

Sticky prices are a fundamental element of the monetary transmission mechanism in many macroeconomic models. In the past decade, a huge empirical literature has tried to measure stickiness and understand its micro-foundations.¹ The empirical literature has produced a set of stylized facts, summarized by Klenow and Malin (2010), which have been used to motivate many theoretical papers.²

The increase in empirical work has been possible due to an unprecedented access to micro-level Consumer Price Index (CPI) data and scanner datasets in several countries. While valuable, these datasets are not collected for research purposes and their sampling characteristics can introduce measurement errors and biases that affect some of the stylized facts in the literature.³

In this paper, I use a new type of micro-level data based on online prices, called “Scraped Data”, to explicitly document the impact of measurement biases on some key stickiness statistics. In particular, I argue that two main sampling characteristics, time averaging and cell-relative imputation, can greatly affect the frequency and size of price changes in traditional data sources.

Time averaging is intrinsic in scanner datasets, such as AC Nielsen’s “Retail Scanner Data”, which report weekly averages of prices. Price imputation of temporarily unavailable products is a common characteristic of CPI datasets. In the US, Klenow and Kryvtsov (2008) report that temporarily unavailable prices are almost 7% of all items in the CPI Research Database. The Bureau of Labor Statistics (BLS) imputes these missing prices with a method called “cell-relative imputation”, which uses the average price change within related categories of goods.⁴ These two sampling characteristics, while reasonable for the purposes of the original data collection efforts, can greatly increase the number of price changes observed in the data, reducing their size and affecting related statistics such as the hazard rate of price changes.

Scraped data are not affected by these two sources of measurement bias. They are collected using specialized software that scans the websites of retailers that show prices online, finds relevant information, and stores it in a database.⁵ Once configured, the software can

¹Cecchetti (1986), Kashyap (1995), and Lach and Tsiddon (1996) provided pioneering contributions to the literature using samples of goods such as magazines and groceries. Bils and Klenow (2004) made a seminal contribution with micro-level US CPI data. They were followed by papers such as Nakamura and Steinsson (2008), Klenow and Kryvtsov (2008), Klenow and Willis (2007), Dhyne, Álvarez, Bihan, Veronese, Dias, Hoffmann, Jonker, Lünnemann, Rumler, and Vilminen (2006), Boivin, Montréal, Giannoni, and Mihov (2009), Wulfsberg (2009), Gagnon (2009), just to name a few. For a recent survey of the literature, see Nakamura and Steinsson (2013).

²Some examples are Midrigan (2011), Gorodnichenko (2008), Woodford (2009), Bonomo, Carvalho, and Garcia (2011), Costain and Nakov (2011), and Alvarez and Lippi (2014).

³For previous discussions of measurement error in the literature, see Campbell and Eden (2014), Cavallo and Rigobon (2011), and Eichenbaum, Jaimovich, Rebelo, and Smith (2014).

⁴See Bureau of Labor Statistics (2015a), page 20. Before January 2015 the BLS imputed prices using relatively broad item strata and geographic index areas. The latest methodology uses narrower elementary level items (ELIs) and metropolitan areas. This change is explained in Bureau of Labor Statistics (2015b). My results suggest that this is likely to reduce the magnitude of the imputation bias in the US CPI data in the future.

⁵The primary language used to write content on the Web is called Hyper Text Markup Language, or

be set to run automatically every day, providing high-frequency information for all goods sold by the sampled retailers in many countries. The scraped data used in this paper were collected every day between October 2007 and August 2010 for over 80 thousand individual supermarket products in five countries: Argentina, Brazil, Chile, Colombia, and the US.

The first contribution of the paper is to use the US data to document the impact of measurement bias on three common statistics in the literature: the duration of price changes, the distribution of the size of price changes, and the shape of their hazard functions.

For time-averages in scanner data, I directly compared my findings to those using data provided by AC Nielsen for the same retailer, location, and time period. I also replicate the weekly time-averaging characteristic in the scraped data, which provides a nearly perfect match to the scanner data results. As Campbell and Eden (2014) conjectured, weekly averaging makes a single price change look like two consecutive smaller changes. This creates more frequent and smaller price changes, completely altering the shape of their size distribution. It also makes the hazard rate highest on the first week after a change, producing fully downward-sloping hazard functions. Overall, time-averaging the data produces very similar results to those in papers with scanner data such as Eichenbaum, Jaimovich, and Rebelo (2011) and Midrigan (2011).

For imputation in CPI data, I replicated the cell-relative imputation in my own dataset. I show that imputing temporarily missing prices with average changes in the same category tends to increase the frequency of price changes and reduce their size, making the size distribution completely unimodal. This effect is independent of the effect of forced item substitutions, which are often removed before computing stylized facts in the literature.⁶ The bias is strongest when broader categories of goods are used as the reference for imputation, as the BLS did until January 2015. I also show that daily prices are needed to detect the initial increase in hazard rates during the first couple of months. Instead, if when cell-relative imputation is applied to monthly data, the results resemble those in papers with CPI data, such as Klenow and Kryvtsov (2008).

The second contribution of the paper is to use the data in all five countries to document three new stylized patterns. First, I study the frequency of price changes and compare it across countries. I show that only the *relative* frequency of price increases over decreases is correlated with inflation rates. This implies that discussions about the real effects of monetary policy should focus on this statistic and not the simple frequency of all changes. Second, I look at the distribution of price changes and, despite some heterogeneity across countries, it tends to be bimodal, with relatively few price changes close to zero percent. Third, I show that the hazard rates of price changes are hump-shaped in all countries, initially increasing over time. These last two results, which differ from previous findings in the literature, are mostly consistent with models that combine elements of both time and state-dependent pricing.

The findings have several uses and implications. First, they show that some stylized facts in the literature, such as the prevalence of very small price changes, are driven by the sampling characteristics in traditional data sources. Documenting and adjusting to these potential biases is key to avoid making misleading interpretations. Second, they provide new

HTML. It is written using tags, such as `<div class=price>`, which provide instructions for the browser. These tags can be used by the scraping software to locate relevant price and product information on the page.

⁶See Klenow and Kryvtsov (2008), Nakamura and Steinsson (2008), and Nakamura and Steinsson (2013).

stylized facts that are robust across retailers in multiple countries and are consistent with models that combine both information and menu costs, as in Alvarez, Lippi, and Paciello (2011). Although limited here to supermarket prices, scraped data can be collected in many sectors and countries, potentially providing the statistics needed to parametrize the models in the literature. Third, they illustrate how new data collection techniques allow macroeconomists to build customized datasets, designed to address specific research questions. As pointed out by Einav and Levin (2014), the emergence of “Big Data” requires economists to develop new data management and programming capabilities, and new data collection skills may become an essential part of that process.

My paper is directly related to several others in the literature that discuss sources of measurement bias. Campbell and Eden (2014) identified prices which could not be expressed in whole cents in an AC Nielsen scanner dataset, noting that technical errors and time aggregation were likely causing them. My results with scanner data confirm their intuition. Cavallo and Rigobon (2011) discuss the potential effect of time-averaging and unit values on the distribution of size changes by running simulations with online data from multiple retailers. My papers expands the analysis to measure the effects on durations and hazard functions. It also directly compares online and scanner data from the same retailer, location, and time period, to show that time-averages are the key source of bias in scanner data. Eichenbaum, Jaimovich, Rebelo, and Smith (2014) use CPI and scanner data from multiple stores to show how unit-value prices, reported as the ratio of sales revenue of a product to the quantity sold, affect the prevalence of small price changes. While they also use daily data, their focus is on looking at the effects of unit-values using data from multiple stores on the distribution of daily price changes. Instead, I compare the weekly-averaged price change distributions to the daily price change distribution for a single store, to show that even scanner datasets which are not affected by unit-values, such as AC Nielsen’s Retail Scanner Data, still produce biases results for the frequency, size, and hazard rates of price changes. Eichenbaum, Jaimovich, Rebelo, and Smith (2014) also show the effect of unit-values and bundled goods in CPI categories such as “Electricity” and “Cellular Phone Services”. I focus instead on price imputations for temporarily missing prices, which affect nearly all CPI categories, and are particularly relevant for food categories with significant CPI weights.

My work is also related to papers that use online prices, such as Lünemann and Wintr (2011), Gorodnichenko and Talavera (2014), and Gorodnichenko, Sheremirov, and Talavera (2014). These papers find that online prices tend to be more flexible and have smaller price changes than offline prices. The difference with my results likely comes from their focus on online-only retailers that participate in price-comparison websites. As Ellison and Ellison (2009) showed, this type of retailers face a different competitive environment, which will tend to increase the frequency and reduce the size of their price changes. Instead, I use data from retailers that have an online presence but sell mostly offline. Lünemann and Wintr (2011) report that these online-offline retailers represent only 9% of all price quotes in their sample.

The paper is organized as follows. In section 2, I describe the collection methodology and characteristics of scraped data. In section 3, I use the US data to document the impact of measurement error by comparing the duration of prices, the distribution of the size of price changes, and the hazard functions with previous results in the literature, sampling simulations, and a comparable scanner dataset. Section 4 uses the scraped data to compute stickiness statistics in all countries and discuss their differences. Section 5 concludes.

2 Description of Scraped Data

2.1 The Data Collection Methodology

A large and growing share of retail prices are being posted online all over the world. Retailers show these prices either to sell online or to advertise prices for potential offline customers. This source of data provides an important opportunity for economists wanting to study price dynamics, yet it has been largely untapped because the information is widely dispersed among thousands of webpages and retailers. Furthermore, there is no historical record of these prices, so they need to be continually collected over time.

The technology to periodically record online prices on a large scale is now becoming widely available. Using a combination of web programming languages, I built an automated procedure that scans the code of publicly available web-pages every day, identifies relevant pieces of information, and stores the data in a file. This technique is commonly called “web scraping”, so I use the term *Scraped Data* to describe the information collected for this paper.

The scraping methodology works in three steps. First, at a fixed time each day, a software downloads a selected list of public web-pages where product and price information are shown. These pages are individually retrieved using the same web-address (URL) every day. Second, the underlying code is analyzed to locate each piece of relevant information. This is done by using special characters in the code that identify the start and end of each variable, and have been placed by the page programmers to give the website a particular look and feel. For example, prices are usually shown with a dollar sign in front of them and enclosed within a `<price>` and `</price>` tags. Third, the software stores the scraped information in a database that contains one record per product per day. These variables include the product’s price, the date, category information, and an indicator for whether the item was on sale or not (if available).

2.2 Advantages and Disadvantages

The main differences between Scraped Data and the two other sources of price information commonly used in studies of price dynamics, CPI and Scanner Data, are summarized in Table 1.

Table 1: Alternative Data Sources

	Scraped Data	CPI Data	Scanner Data
Data Frequency	Daily	Monthly - Bi-Monthly	Weekly
All Products in Retailer (Census)	Yes	No	No
Product Details (size, brand, sale)	Yes	Limited	Yes
Uncensored Price Spells	Yes	No	Yes
Countries Available for Research	~50*	10-15	<5
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

Notes: *Data from over 50 countries are currently being collected by the Billion Prices Project (bpp.mit.edu).

Scraped data have some important advantages. First, these datasets contain posted *daily* prices that are free from unit values, time-averaging, and imputations that can greatly affect some stickiness statistics, as is later shown in this paper. Second, detailed information can be obtained for *all* products sold by the sampled retailers, instead of just a few models or selected categories. Third, there are no censored or imputed price spells in scraped data. Prices are recorded from the first day they are offered to consumers until the day they are discontinued from the store. In CPI, by contrast, there are frequent imputations and forced substitutions when the agent surveying prices does not find the item she was looking. Fourth, scraped data can be collected remotely, in any country where information can be found online. In this paper, I include data for four *developing* countries, where scanner data are scarce and product-level CPI prices are seldom disclosed.⁷ Fifth, scraped datasets are comparable across countries, with prices that can be collected with identical technique on the same categories of goods and time periods. This makes it easier to perform simultaneous cross-country analyzes.⁸ Finally, Scraped Data are available in real-time, without any delays to access and process the information. This can be potentially used to provide estimates of

⁷The study of stickiness in developing countries is rare in the literature. An exception is Gagnon (2009), who provides a detailed analysis of sticky prices in Mexico using disaggregated CPI data manually digitalized from printed books. Another is Alvarez, Gonzalez-Rozada, Neumeyer, and Beraja (2015), who use micro-CPI data from Argentina to document the behavior of price stickiness from the hyperinflation in the late 80s to the period of low inflation in the 90s.

⁸Past cross-country comparisons in the literature have had to rely on results provided by different papers, often with different data sources, time periods, and event methods and data treatments. See, for example, Klenow and Malin (2010). An exception is Dhyne, Álvarez, Bihan, Veronese, Dias, Hoffmann, Jonker, Lünnemann, Rumler, and Vilmunen (2006), who were able to use similar data from multiple countries thanks to the coordination provided by the European Inflation Persistence Network at the European Central Bank. They describe, however, how each national statistic agency was unwilling to share the micro data, even with Eurostat, so the frequency analysis had to be conducted independently in each country by a different team, each facing a dataset with different characteristics. For example, some countries allowed researchers to discriminate price changes due to product replacement, sale prices and other discounts, while other did not.

stickiness that quickly capture changes in the underlying economic conditions.

Table 1 also shows the main disadvantages of scraped prices. First, they typically cover a much smaller set of retailers and product categories than CPI prices. In particular, the supermarket products in this paper come from a single retailer in each country and cover only 40% of all CPI expenditure weights in the four Latin American countries and 12% in the US. While this is enough to demonstrate the effect of measurement errors on pricing statistics, the quantitative findings on stickiness and size of changes shown here should not be taken as representative of the economies of these countries as a whole. Yet the number and variety of goods whose prices are shown online is growing over time, so future papers will be able to provide aggregate statistics using online data collected from a much larger share of retailers and sectors. Second, another major disadvantage of scraped data relative to scanner datasets is the lack of information on quantities sold. In the context of measuring stickiness, quantities are useful to provide weights in frequency and other related statistics.

2.3 Large Supermarkets in Five Countries

I built a dataset with more than 44 million supermarket prices in Argentina, Brazil, Chile, Colombia, and the US. All the data are available for download on my academic website. Table 2 provides details on each country’s database. The data come from the websites of five different supermarkets, one in each country, and were collected every day from November 2007 to August 2010 for the Latin American countries, and between June 2008 and August 2010 for the US. There are roughly 25 thousand products sampled in Argentina, Chile, and Brazil, 10 thousand in Colombia, and 30 thousand in the US.

Table 2: Database Description

	Argentina	Brazil	Chile	Colombia	USA
Total observations	10.8M	9.8M	9.7M	3.9M	10.2M
Total Products	28813	23115	24336	9526	30727
Initial date	10/2007	10/2007	10/2007	11/2007	05/2008
Final date	08/2010	08/2010	08/2010	08/2010	08/2010
Days	1041	1038	1024	1004	827
Categories	74	72	72	59	22
URLs (narrower categories)	993	319	292	122	241
Obs with sales	3%	4%	-	8%	19%
Products with sales	39%	22%	-	25%	79%
Life of goods (days, median)	540	502	634	525	495

Notes: Missing values are caused by items that go out of stock or failures in the scraping software that tend to last for a few days. I replaced missing values *within* price series for the first 90 days of the price gap with the previous price available for each product. I also removed all price changes exceeding +200% and -90%. These represent a negligible number of observations that can bias statistics related to the magnitude of price changes. See the Appendix for more details on data treatments.

All the retailers included in the dataset are market leaders in their respective countries,

with market shares of approximately 28% in Argentina, 15% in Brazil, 27% in Chile, and 30% in Colombia.⁹

To compare results for the same product categories across countries, I matched each supermarket’s classifications into 95 standardized categories containing a large variety of foods and household items.¹⁰ A narrower category indicator is also provided by the URL where the products are found, as retailers group closely similar goods in a single web-page. The retailer’s design of the website and menu pages determines the number of URLs available in each country.

2.3.1 Online vs Offline Prices

Online purchases are still a small share of transactions in most countries, so it is natural to question the representativeness of scraped data. Are online prices similar to offline prices? More specifically, if we were to physically walk into a store in these retailers and collect prices for the same products, would they be similar to the prices collected on the website at the same time?

To answer this question I conducted simultaneous surveys of *offline* and *online* prices in all the retailers included in this paper. These surveys took place in Buenos Aires, Santiago, Rio de Janeiro, Bogotá, and Washington DC with the help of five local people. They were asked to select any branch of these supermarkets and randomly buy 100 products, divided in 10 predefined categories. These categories were chosen to ensure some variety in the type of goods purchased. After the first purchase, I used the printed receipts to get unique product ids and check whether the same items were sold online or not.¹¹ Those items that could not be matched to the online database were removed from the product list for subsequent purchases. In total, four purchases took place in each supermarket, at 15-day intervals, always in the same branches. The same items were bought every time, with identical flavors and package sizes. If a product was out of stock, no price was recorded for that day, but we attempted to buy the product again in subsequent purchases.¹²

Table 3 shows the results from this data comparison exercise. The percentage of offline products that were also available online ranges from 74% in Colombia to 100% in Argentina. Most of the products that could not be matched are raw-food items, which tend to be re-packaged for online sales and have different id numbers and descriptions.

⁹The market share for the US supermarket, one of the largest in the US market, is not revealed here to prevent readers from being able to identify the specific retailer. This is strictly forbidden by the conditions of the scanner data provided by the Kilts Marketing Center at the University of Chicago Booth, used in Section 3.0.1 of the paper to compare the results with online data.

¹⁰See the Appendix for a complete list of product categories. These are based on the ELI classification used by the US Bureau of Labor Statistics.

¹¹In Argentina, Brazil, and Colombia, the matching was based exclusively on product ids. In Chile the matching was based on the item’s name, description, and package size, because the online product id did not match the offline id printed on the ticket.

¹²The offline data collection for the Latin American supermarkets took place in 2009. The US data was collected in 2015, as part of a much larger validation exercise using mobile phones and crowd-sourcing websites in 10 countries. For details see Cavallo (2015).

Table 3: Comparing Online vs. Offline Prices

	Argentina	Brazil	Chile	Colombia	USA
Matching ids	Yes	Yes	No	Yes	Yes
% Available Online	100%	80%	90%	74%	86%
PRICE LEVELS					
online=offline	18%	42%	93%	29%	79%
online>offline	78%	34%	4%	32%	5%
Price Difference* (Mean %)	5	9	2	0	-7
PRICE CHANGES					
Products with Identical Change Series**	93%	75%	94%	67%	87%
Ratio of Changes over Observations					
Offline	0.215	0.356	0.274	0.433	0.471
Online	0.215	0.411	0.249	0.433	0.436
Mean Size of Changes (%)					
Offline	1.6	4.9	1.4	8.1	25
Online	1.4	5.3	1.3	8.3	23

Notes: Data from randomly selected products in each retailer. Offline and online prices were collected within a 7-day time window. In Latin America, offline prices were collected four times every 15 days in 2009. In the US, intervals of data collection vary from a week to over a month, and prices were collected in 2015. *Excludes identical prices. **Indicator variable conditional on change: 1 if the price increased, -1 if it decreased.

I compare prices both in terms of their *levels* and the timing and size of *changes*. Even though price levels are not always the same across samples, online and offline price *changes* behave similarly in terms of timing and size of adjustments in all countries.

In the US, the share of identical prices is 79% and the price change series are highly synchronized, and both the frequency and size of price changes are similar across samples. In Chile, the matching of price levels is extremely close. 361 out of 388 comparable prices were exactly the same. The 27 price discrepancies, which averaged 2% in size, were concentrated in only 12 goods (mostly raw-food products), so that 89% of products have identical price levels across samples. In Argentina, price levels are typically higher online, yet in nearly every case there was a difference of 5% across samples. A constant markup implies that price *changes* are highly correlated, with similar frequency and size of changes.

The cases of Brazil and Colombia are more complex, but the samples still show similar price change behaviors. The evidence suggests these supermarkets may be treating their online stores as independent branches, with different price levels but similar strategies in terms of price adjustments. In Brazil, price levels are identical only 42% of the time. Most of the differences are concentrated in a small share of products, so that 75% of all goods have identical price change series across samples. The ratio of changes over total observations and the mean size of changes are very similar across samples. In Colombia, the matching of price levels, at 29%, is even lower than in Brazil, but price differences are smaller. The matching of price changes is still relatively high, with 67% of identical price changes series. In terms

of the frequency of price changes, both samples have identical ratios of changes over total observations at 0.433. Finally, the mean size of changes is also very close, with 8.1% offline and 8.2% online.

The comparison between online and offline prices is explored in great detail in a related paper, Cavallo (2015), where I simultaneously collect online and offline prices for over 40 of the largest multi-channel retailers in 10 countries. The results, consistent with those in this paper, is that on average 70% of prices are identical across samples. Price changes, while not synchronized, also have similar average frequency and size.¹³

3 How Measurement Error Affects Pricing Statistics

Measurement error has been discussed in the literature before. Campbell and Eden (2014) identified and removed prices which cannot be expressed in whole cents in an AC Nielsen scanner dataset. They noted that technical errors and time aggregation could be the cause for those “fractional prices”. Cavallo and Rigobon (2011) further discussed the potential effect of time-averaging and unit values on the distribution of size changes, and simulated the impact on the distribution of size changes using online data in a large number of countries. Eichenbaum, Jaimovich, Rebelo, and Smith (2014) used CPI and scanner data from multiple stores to show how unit-value prices, reported as the ratio of sales revenue of a product to the quantity sold, affect the prevalence of small price changes.

An advantage relative to previous papers is that I have a source of data which is not affected by unit values, time averages, or forced imputations. I am able to re-compute some classic statistics in the literature and compare them to previous results. I can simulate some of the sampling methods in Scanner and CPI data on my original data and show that they generate similar results to those in the literature. Finally, and more explicitly, I directly compare both online and scanner data from the same US Supermarket, zip code, and time period. I show that the time-averaging in scanner data basically accounts for all the differences observed with posted prices.

3.0.1 Evidence in Scanner Data

Scanner datasets have two main potential sources of measurement error. First, prices are sometimes reported as “unit values”. As Eichenbaum, Jaimovich, Rebelo, and Smith (2014) show, even with daily data, if prices are sometimes purchased with or without coupons or at different prices across stores of the same retailer, the unit values reported by some scanner datasets will tend to generate spurious small price changes. Although unit values were common in some early scanner datasets available in the literature, most scanner datasets used today, such as AC Nielsen’s “Retailer Scanner Data”, do not report prices as unit values any more. Instead, they provide the actual prices for individual goods sold by individual stores in multiple locations.

¹³An alternative way to test the validity of scraped data is to see if the inflation dynamics obtained from this small sample of retailers can resemble those in CPI statistics, which are constructed using surveys from a large number of offline stores. In a related paper, Cavallo (2013), I show that online price indexes can closely match inflation rates in most of these countries.

The second source of measurement error is weekly-averaged prices. The effect of time-averaging in scanner data was first discussed by Campbell and Eden (2014). Their focus was not on the size of changes, but they described some complications caused by weekly averages using a simple example of a three week period with a single price change on the middle of the second week. Instead of a single price change, the weekly-averaged price data produced two price changes of smaller magnitude. The prevalence of examples like this can potentially double the frequency of changes and greatly reduce the size of price changes.

To provide evidence of time-averaging in scanner datasets, Table 4 compares results for price stickiness in the US supermarket data with other samples and papers. I use the standard methods used in the literature to compute frequency and duration. I first obtain the daily frequency per individual good by calculating the number of daily price changes over the number of total valid change observations for a particular product. Next, I calculate the *mean* frequency per good category, and finally, the *median* frequency across all categories. I then compute implied durations using $-1/\ln(1 - frequency)$, and convert them to monthly durations for comparisons across papers and samples.

Table 4: Implied Duration in US Supermarket Data

	Scraped	Scanner Large Retailer	Scanner Dominik's	Scraped Weekly Av.	Scanner Same Retailer
Period	2008-2010	2004-2006	1989-1997	2008-2010	2008-2010
Data Frequency	Daily	Weekly	Weekly	Weekly	Weekly
Duration (months)	1.4	0.6	1.0	0.8	0.8

Notes: Scanner data results in columns 3 and 4 come from Table 3 in Eichenbaum, Jaimovich, and Rebelo (2011). Scanner data results on the last column use prices from AC Nielsen provided by the Kilts Center at Chicago Booth, matching to the same retailer, zip code, and time period of the online scraped data.

The implied duration is much higher than in previous papers in the literature that have used scanner data for US Supermarkets. In particular, I include results from a “Large US Supermarket” and the other from Dominick’s Supermarket, both reported by Eichenbaum, Jaimovich, and Rebelo (2011).

While the higher duration is consistent with measurement error, there are other reasons that could cause the differences in results. For example, note that the time periods are quite different, so it is possible that goods have become stickier in recent years. Also, the retailers may not be the same, or even similar in their characteristics.

To control for time and retailer differences, I run a sampling simulation on the original scraped data by computing the weekly average price. Note that these prices are naturally free from unit values and averaging across stores, as they come from a single retailer and location. I am simply averaging prices over a week, and then calculating the weekly frequency of price changes and its implied monthly duration. The results, also reported on Table 4, show that duration fall from 1.4 to 0.8. So a simple time-averaging of daily data into weekly prices reduces the duration of prices by almost a half.

This, however, does not prove that a comparable scanner data will necessarily have the same bias, or in a similar magnitude. To make the comparison even more explicit, I purchased scanner data for the exact same retailer, location, and time period. The main challenge was to match the retailer in both samples. The scanner dataset, collected by AC Nielsen, does not explicitly identify the retailers. It only provides a supermarket chain id and the zip code of each store. However, all retailers tend to have a distinctive pattern of stores in different zip codes. By simply counting how many stores each supermarket chain in the scanner had in a given set of zip codes, I was able to find a perfect match to the retailer where my online scraped data was collected.¹⁴

The last column in Table 4 shows that the scanner data also has a duration of 0.8, identical to the effect of a simple time-averaging and the average of results reported by Eichenbaum, Jaimovich, and Rebelo (2011). Time-averages is therefore all that is needed to replicate the duration results in scanner data.

The effect on the size of price changes is even more striking. Figure 1 shows the distribution of the size of price changes in the original data, in the scanner data for the same retailer, and in the simulated weekly averaged data.

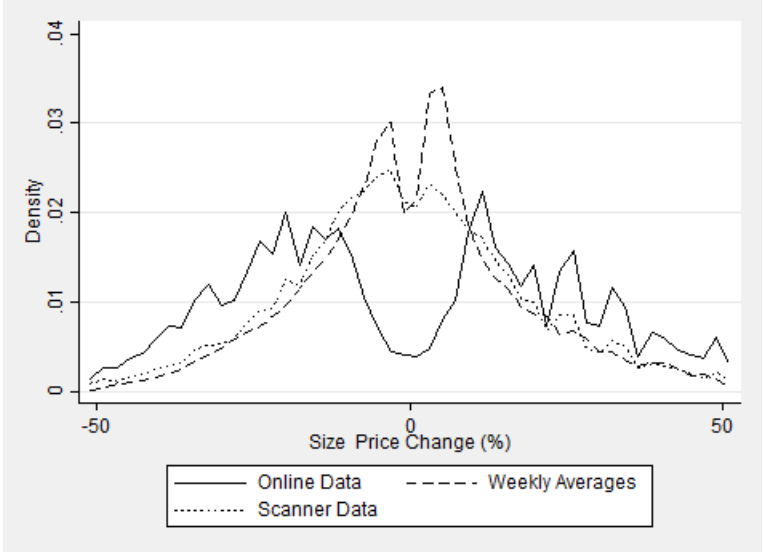


Figure 1: The Distribution of the Size of Price Changes in the US

Notes: The online and scanner data in the US was collected at the same retailer during the same time period. Scanner data was collected by Nielsen and provided by the Kilts Center at Chicago Booth.

The weekly averaging in the scanner data completely changes the shape of the distribution by turning large price changes into small ones. Interestingly, the weekly averaging and scanner dataset distribution are nearly identical, with the exception of the two spikes that remain in the weekly averaged data near zero percent. One explanation for this could be the use of coupons and loyalty cards, which would also affect the weekly averaged price even if the posted price does not change at all. This would create additional tiny price changes, further smoothing the distribution to match the actual results obtained from the scanner data.

¹⁴I obtained the distribution of stores across zip-codes in the online sample by scraping the “Find a store” form available in the website of the retailer.

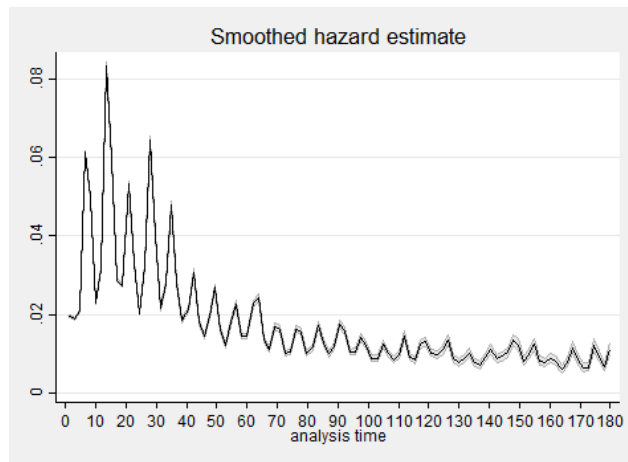
The shape of this distribution is useful to distinguish between alternative sticky price models. While the literature has mostly found unimodal distributions, the scraped data produces a distribution with very little mass near zero percent and two modes, one positive and one negative. This is consistent with models that incorporate role for an adjustment or “menu” cost that make small price changes sub-optimal. In fact, this distribution is strikingly similar to the prediction of the model in Alvarez, Lippi, and Paciello (2011), which combines both adjustment and information costs into the price-setting decision.¹⁵

Time-averaging may not only affect the frequency (duration) and size of price changes. It could also have an impact on the estimated hazard rates of price adjustment. Hazard rates measure the probability of a price change as a function of the time since the previous adjustment, and different sticky-price models will have different predictions about the shape of the hazard function over time. Adjustment-cost models, for example, tend to generate upward sloping hazards if the shocks are persistent over time. Time-dependent models, by contrast, generate spikes in the hazard function at the dates when adjustment takes place.


Figure 2 shows the daily hazard rates using the daily scraped data, the weekly averaged data, and the weekly scanner data. Details for the construction of these estimates are provided in the Appendix.

The scraped data hazard, shown in panel 2(a) has a hump-shaped pattern, initially increasing and then gradually falling over time. It is also clear that there are weekly spikes in the hazard rates.


¹⁵See figure IV in that paper.



(a) Online Data



(b) Weekly Average



(c) Scanner Data

Figure 2: Hazard Functions

Notes: Initial 180 days shown. Left-censored spells are excluded.

Panels 2(b) and 2(c) show that the weekly averaging and the scanner data produce very

different hazard function. Naturally, when a single price change is transformed into two weekly changes, most of the probability of a price change occurs in the first week after the previous change. This makes the trend of the hazard rate appear completely downward sloping from the start. In this case, since the data are weekly the spikes in these hazards are an artifact of the daily scale of the graph. Plotting these on a weekly scale would make the hazard function smooth and completely downward sloping, similar to those found in Campbell and Eden (2014).

Once again in this case, the effect of measurement error completely distorts the stylized fact. The spikes in the hazard rate are consistent with models that have information costs, as noted in Alvarez, Lippi, and Paciello (2011). But the increasing trend at the beginning, also suggests that adjustment costs play an important role, particularly because these hazard functions are also subject to survival bias. I discuss some evidence of survival bias in Section 4. For now, I simply emphasize how measurement error has a significant impact on the shape of the estimated hazards.

3.0.2 Evidence in CPI data

CPI data are collected on a monthly and sometimes bi-monthly basis. While sampling methods vary across countries, it is less affected than scanner data by time-averages or unit-values.

There is, however, one characteristic that could potentially have a significant effect: the treatment of temporarily missing prices. These missing prices occur when the person doing the data collection at the store is unable to find a particular good and considers it to be temporarily out of stock. These are independent of missing prices due to item substitutions, which have received considerably more attention in the literature. For example, Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008) emphasize how removing price changes from forced substitutions affect the frequency of price changes. Instead, I focus on temporarily missing prices that are likely to be common in any CPI dataset given the characteristics of the data collection process. Indeed, Klenow and Kryvtsov (2008) report that temporary stockouts account for almost 7% of all items in a typical month in the CPI Research Database.

In the US, the BLS uses an imputation method called cell-relative imputation for food and services. When a price is temporarily missing, it is imputed using the average observed change in the prices of goods in a similar category. Until January 2015, the BLS used item-strata, which are relatively broad product categories, as the relative “cell” used for imputation. It has now moved to using more narrow elementary level items, or ELIs.¹⁶

I simulate the sampling characteristics of cell-relative imputation to illustrate the effect on pricing statistics. To do this, I take the original scraped data and keep only the price for the 15th of each month (using the final day of the month does not change the results). Then, for each good, I impute missing prices within price spells by multiplying the previously

¹⁶See Bureau of Labor Statistics (2015b) for a description of the recent changes. In addition to the categories of goods, the imputation is applied for a given geographical aggregation level. Traditionally this was the CPI index area, and is now being replaced with the narrower “Primary Sampling Unit”. Geographical aggregation does not apply to the data of this paper, but it is potentially another reason for measurement bias in CPI micro data.

available price by the geometric average of price changes for goods in the same category. I also repeat the simulation for a narrower categorization level given by the URL.

Table 5: Implied Duration in US Supermarket Data

	Scraped Data	Scraped Category CR Imputation	Scraped URL CR imputation
Period	2008-2010	2008-2010	2008-2010
Data Frequency	Daily	Weekly	Weekly
Duration (months)	1.4	1.1	1.2

Notes: Simulation uses a single day of the month and applies cell-relative imputation to temporarily missing prices. The URL is a narrower level of categorization.

Table 5 shows that cell-relative imputation also reduces the duration of prices. It falls from 1.4 months to 1.1 months in the category cell-relative imputation. The fall is smaller (to 1.2 months) when the imputation is based on the narrower categories, as the BLS has started to do since January 2015. I am not able to compare frequency and duration results with CPI data because no paper in the literature reports this separately for supermarket data. However, the results from the simulation suggest that the magnitude of the bias on durations is not as large as the effect of time-averages in scanner data. The reason is simple: temporarily missing prices are only a small share of all observations, so the frequency of price changes is less affected.

While the impact on durations may be small, the effect on the distribution of the size of price changes can be huge, as seen in Figure 3. This is simply because no matter how rare missing prices are, imputing them with average changes of related goods will always reduce the size of any observed price change.

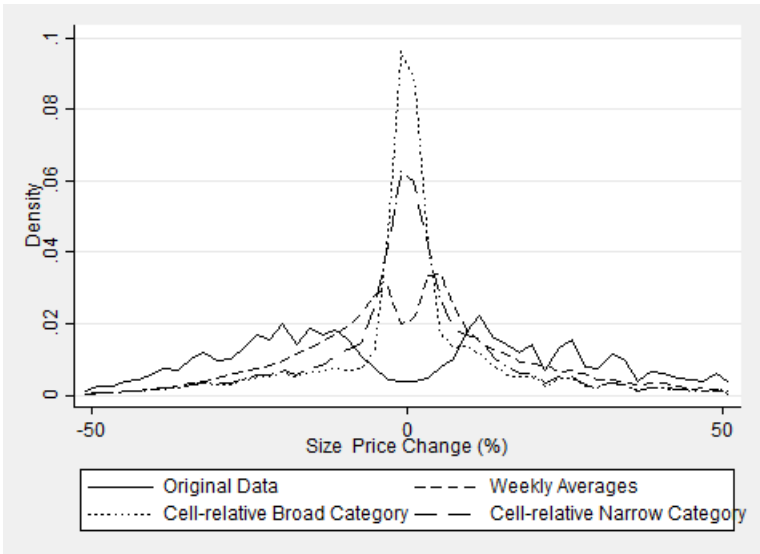



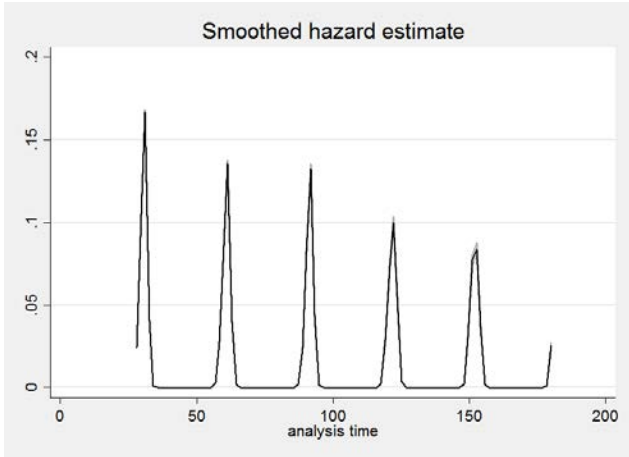
Figure 3: The Distribution of the Size of Price Changes in the US

In this case, cell-relative imputation makes the distribution completely unimodal with a large mass of price changes close to zero percent. The bias is greater than in weekly-averaged data. The share of price changes under $|1\%|$ and $|5\%|$ in absolute value, which in the original scraped data are 0.9% and 4% respectively, rise to 13% and 46% with the URL cell-relative simulation. These numbers are very close to the 11% and 40% reported by Klenow and Kryvtsov (2008) for the US CPI data.

The impact on hazard is equally important. Figure 4 shows that cell-relative imputation produces a downward sloping hazard function, similar to the effect that time-averaging introduces into scanner data.



(a) Monthly Imputation - Broad Category



(b) Monthly Imputation - Narrow Category (URL)

Figure 4: Hazard Functions

Notes: Initial 180 days shown. Left-censored spells are excluded.

The actual bias on the CPI may be lower than in my simulations. First, the number of temporarily missing observations generated by my simulation is 11%, higher than the 7% reported by Klenow and Kryvtsov (2008). Second, in practice not every single missing price may be imputed this way. Unfortunately, other sampling characteristics in CPI data may add to the bias. For example, Eichenbaum, Jaimovich, Rebelo, and Smith (2014) show

that many non-food goods are reported as unit-values (such as Telephone Services) or as composite goods (such as Airline Fares). They find that unit-values and composite good pricing account for a large share of changes smaller than 1%. And even non-missing prices can be affected by imputations, as statistical offices often correct prices for coupons, rebates, loyalty cards, bonus merchandise, and quantity discounts, depending on the share of sales volume that had these discounts during the collection period. Examples of these and other price adjustments are described in the BLS Handbook of Methods.¹⁷

The extent by which CPI data are affected by measurement bias, and the ability to control for it, will vary in different CPI research datasets coming from different countries and time periods. Nevertheless, these results suggest that distributions and hazards obtained from CPI data with imputed prices have to be treated with caution.

4 Stylized Facts with Scraped Data: Cross-Country Comparisons

Section 3 uses the US data to argue that time-averaging and cell-relative imputation in traditional micro price datasets can greatly distort some of the typical statistics used in the price stickiness literature. I now use the Scraped Data, which is not affected by those types of measurement errors, to compute frequencies, durations, distributions of the size of changes, and hazard functions in all countries. The results in this section are useful to find robust “stylized facts” in the supermarket data and to compare findings across countries. The availability of comparable data in multiple countries is potentially one of the greatest advantages of scraped data, allowing us to study stickiness in economies with different levels of inflation and other conditions.

4.1 The Frequency and Size of Price Changes

Table 6 shows price change, frequency, and implied duration information for all countries. The first row reports the average annual inflation rate in the scraped data, obtained by computing a simple CPI-weighted price index in each country.¹⁸

¹⁷See Bureau of Labor Statistics (2015a), Chapter 17, pages 30 to 33.

¹⁸This is not meant to provide an accurate CPI-equivalent inflation number, but rather simply allow the comparison of inflation rates across countries when using the exact same methodology everywhere. Details for the construction of these indexes are provided in the Appendix. A related paper, Cavallo (2013) discusses how online scraped data can be used to construct indexes that are more directly comparable to the official CPIs.

Table 6: Price Changes by Country

	Argentina	Brazil	Chile	Colombia	USA
Inflation (% , average annual rate)	17.1%	5.1%	2.7%	4.2%	0.1%
Price increases (% of price changes)	68%	57%	54%	55%	52%
Price decreases (% of price changes)	32%	43%	46%	44%	48%
Size of price increases (Mean*)	13%	12%	16%	11%	29%
Size of price decreases (Mean*)	-13%	-12%	-14%	-10%	-22%
Daily Frequency	0.015	0.026	0.013	0.022	0.024
Implied Durations (days)	64	38	75	45	42
Implied Durations (months)	2.1	1.3	2.5	1.5	1.4
Frequency of Increases (Freq+)	0.010	0.015	0.007	0.012	0.012
Frequency of Decreases (Freq-)	0.005	0.011	0.006	0.010	0.012
Freq+/Freq-	2.00	1.36	1.16	1.20	1.0

Notes: * Computed as the mean within categories, and then mean across all categories.

The link between frequency and inflation has been studied before using time-series for a single country. Examples include Nakamura and Steinsson (2008), who in the US CPI data find that the frequency of price increases is correlated with inflation, but not the frequency of price decreases. Gagnon (2009) uses a time series of CPI data in Mexico and finds that, at levels of inflation below 15%, the overall frequency of price changes is not correlated with inflation because their frequency of increases rises with inflation but it is offset by a similar fall in the frequency of price decreases. In a rare example with cross-country evidence, Dhyne, Álvarez, Bihan, Veronese, Dias, Hoffmann, Jonker, Lünemann, Rumler, and Vilmunen (2006) find that inflation is positively correlated with the frequency of price increases and negatively correlated with the frequency of decreases.

As in previous studies, I find that the level of frequency is not directly correlated with inflation.¹⁹ For example, Argentina has an annual inflation rate of 17% and it is one of the stickiest countries in the sample, while the US has an inflation rate of 0.1% and is one of the most flexible. But in contrast to time-series results in the literature, the frequency of price increases and decreases are also not independently correlated with inflation. Argentina, for example, has both stickier price decreases *and* increases than the US.

What seems to matter for inflation is not whether prices are generally flexible or not, but rather how much more flexible price increases are relative to decreases. This can be seen in the last row of Table 6. The US, for example, has no inflation because the frequency of increases and decreases is perfectly balanced. In other countries, as price increases become more frequent than decreases, inflation tends to rise. In the extreme case of Argentina, price increases are twice as frequent as price decreases, consistent with a high inflation rate.

¹⁹Inflation depends on both the frequency and size of changes, so high-inflation countries could have few large price changes. However, the average size of price changes is also not correlated with the inflation level in each country.

This cross-country evidence is consistent with the time-series results in Alvarez, Gonzalez-Rozada, Neumeyer, and Beraja (2015). They use Argentine CPI data from 1988 to 1997 and show that inflation is strongly correlated with the difference between the frequency of increases and decreases at all levels of inflation.

The main implication of my findings is that discussions about the real effects of monetary policy across countries should not focus on the general frequency of price changes, but instead on the relative frequency of increases and decreases in each country. For example, Dhyne, Álvarez, Bihan, Veronese, Dias, Hoffmann, Jonker, Lünnemann, Rumler, and Vilmunen (2006) find that prices in Europe tend to be stickier than in the US. As the authors point out, this can be useful to understand how differences in the structure of the retail sector or the extent of price regulation affect stickiness. But this does not imply that monetary policy will be more effective in Europe. In principle, the same logic applies to comparisons of stickiness across goods within a single country. Much of the literature has focused on the heterogeneity in the frequency of all price changes. While this might be useful to understand how competition and other factors make price changes more likely in categories such as airline fares, it tells us nothing about the real effects of monetary policy in each sector. Instead, we should pay more attention to the relative frequencies of increases over decreases in each sector.

4.2 The Size of Price Changes

In Section 3, I showed that the US scraped data generates distributions of price changes that are bimodal with few price changes close to zero percent. Figure 5 shows that this is also the case in Argentina, Brazil, and Chile with varying degrees of bimodality and mass near zero percent. The only exception is Colombia, where the distribution is unimodal with a large number of small changes.




Figure 5: Distribution of the Size of Price Changes
 Notes: Bin size is 0.1%. Smoothed kernel density shown.

Figure 5 suggests that adjustment costs are more important in the US retailer than in the Colombian retailer. Indeed, the shape of these distributions can be used to infer the relative importance of adjustment and information costs. For example, Alvarez, Lippi, and Paciello (2011) use the coefficient of variation of the absolute size of price changes and the ratio of the mode to the mean, while Cavallo and Rigobon (2011) develop a “proportional mass” test

to measure the relative importance of small changes.

4.3 The Hazard Rate of Changes

Finally, I use the Scraped data to study the hazard rates of price changes over time. Figure 6 provide smoothed hazard rates in all countries.

A common feature across countries is the hump-shaped pattern first identified in Section 3 for the US data. With peaks at different points in time, all hazard functions are initially upward sloping. The peaks in these hazard tend to coincide with the average implied durations estimated in Table 6.²⁰

²⁰These results imply that the assumption of flat hazard rates in those estimated duration numbers is not realistic (though it may be innocuous for certain purposes).

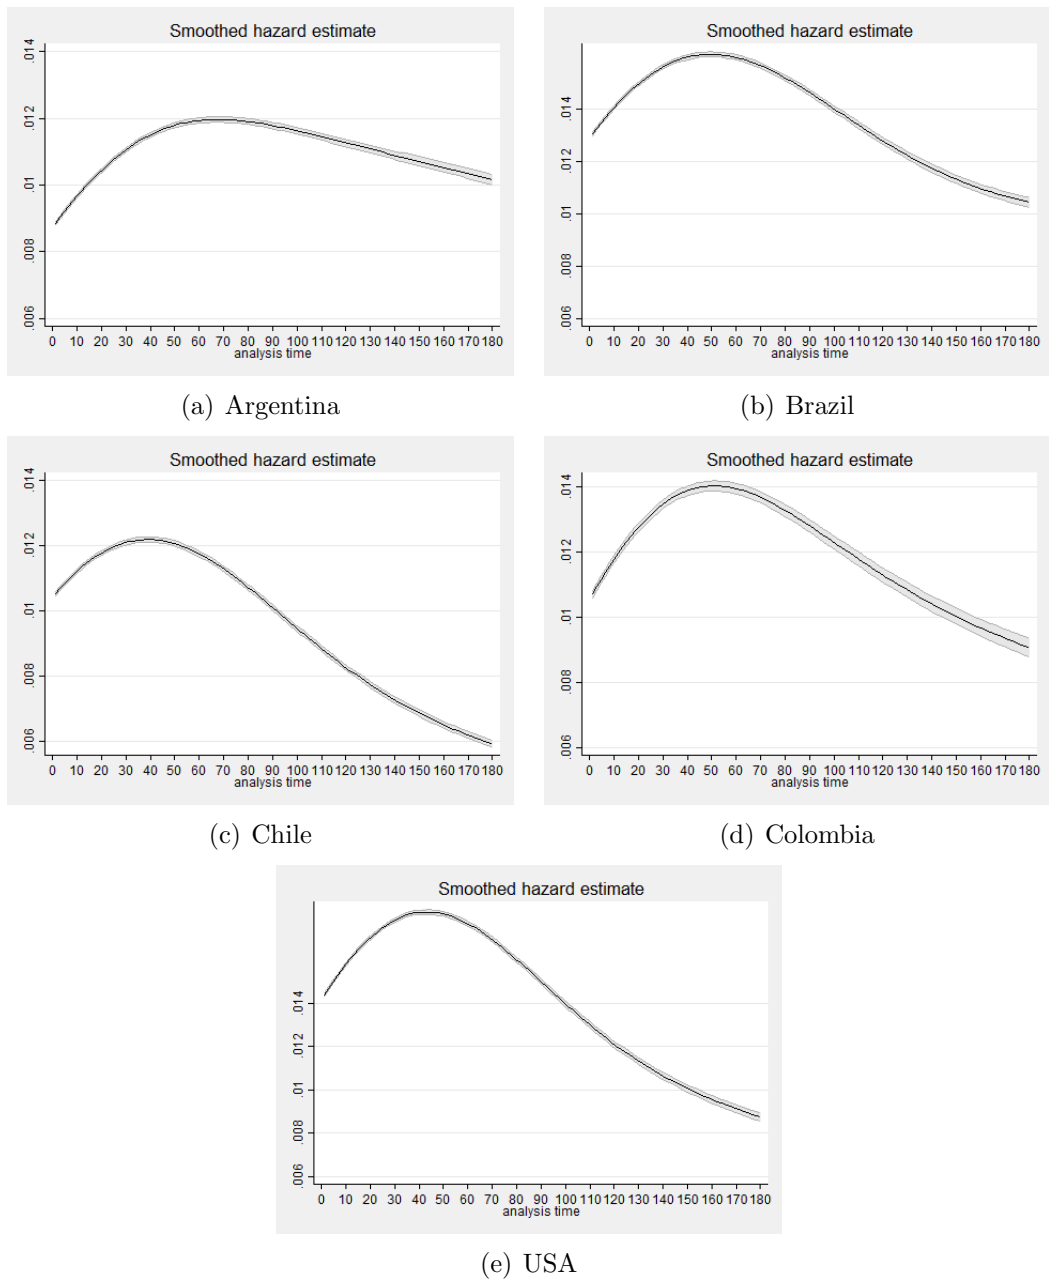


Figure 6: Smoothed Hazard Functions

Notes: Initial 180 days shown. Left-censored spells are excluded.

The differences with previous papers that found flat or downward sloping hazards is driven not only by the lack of time-averages or cell-relative imputations, but also by the fact that data are available in daily frequency and for a large and heterogeneous set of goods. The daily frequency provides the information needed to capture the initial rise in the hazard rates within the first month or two. The large set of goods provides a lot of price spells that can be used to better estimate hazard rates and try to control for problems like survival bias.

Survival bias is a well-know problem in the estimation of smoothed hazard functions. Several papers in the literature have suggested this is one of the main reasons most estimated

hazards are downward sloping.²¹ Using the granularity of scraped data, I am able to find evidence of the existence of survival bias in Figure 7, where I separate goods in terms of their average durations and re-estimated their hazard functions. The dotted line represents goods that have average durations of less than 50 days, the dashed line is for goods with average durations of 50 to 100 days, and the solid line represent stickier goods with average durations over 100 days.

²¹See Álvarez, Burriel, and Hernando (2005), Klenow and Kryvtsov (2008), Nakamura and Steinsson (2008), and Campbell and Eden (2014) among others.

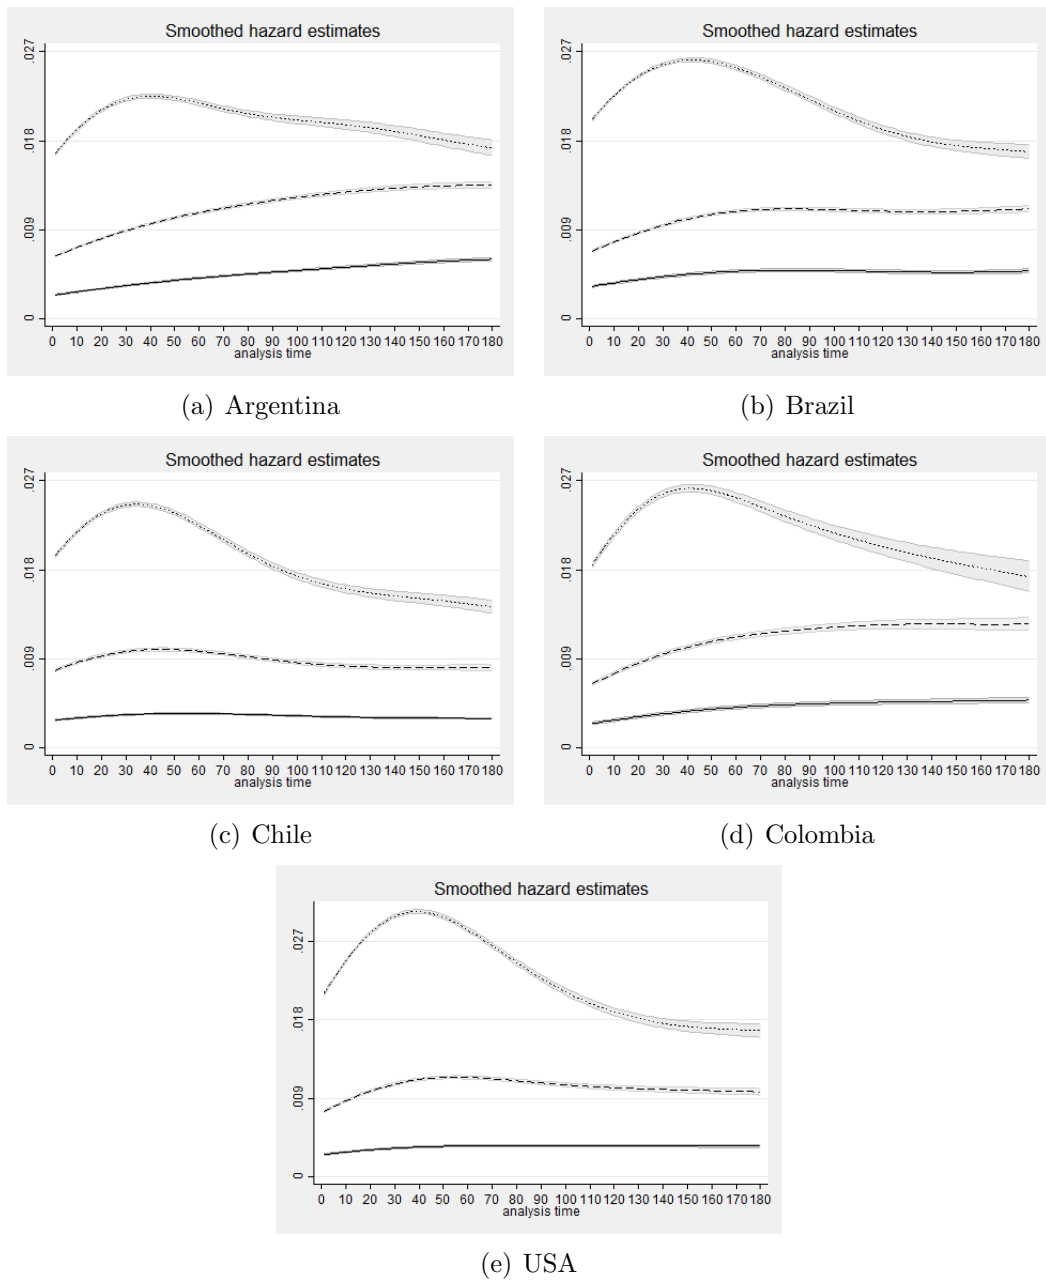


Figure 7: Hazards for Different Duration Groups

Notes: Left-censored spells are excluded. Initial 180 days shown.

As I separate goods into different categories, each one of these hazards became more upward sloping. The hump-shaped patterns does not disappear completely because each one of these three hazards is itself constructed by aggregating across many goods, and therefore they are still affected by survivor bias.

Overall, my results suggest that the underlying hazard rates are even more upward sloping than what the aggregate estimates tend to reflect. They imply that, instead of focusing on theories that can generate downward sloping hazards, we should instead focus on improving the data and methods to measure them better.

5 Conclusions

This paper introduces a new way of collecting price data and applies it to study some basic stylized facts in the price stickiness literature. Scraped data, obtained directly from online retailers, provide a unique source of price information. Prices are easier to collect than in CPI and scanner data and can be obtained with daily frequency for all products sold by retailers around the world. The data are available without any delays and the collection methodology can be customized to satisfy the specific needs of sticky-price studies. More importantly for the stickiness literature, Scraped data are free from common sources of measurement error, such as time-averages and imputation methods, that can affect traditional micro-price datasets.

The paper provides two main results.

First, I use the US data to show how measurement bias affects three common stylized facts in the literature: the duration of price changes, the distribution of the size of change, and the hazard functions. I argue that scanner and CPI datasets can produce biased results for these statistics. I show this with sampling simulations in my own data, and confirm this explicitly in scanner data by comparing both online and scanner data collected from the same retailer and time period. Weekly-averaging and price imputations tend to reduce the duration of price changes (particularly in scanner data), reduce their size, make the distribution of changes unimodal, and the hazard function more downward sloping.

I then use the scraped data in five countries to document three stylized facts. First, that the *relative* frequency of price increases over decreases is the only statistic that is correlated with inflation rates across countries. This implies that discussions about the real effects of monetary policy should focus on this statistic instead of the simple frequency of changes. Second, that the distribution of price changes tends to be bimodal with relatively few price changes close to zero percent. Third, that the hazard rates of price changes are hump-shaped, initially increasing over time. These results are different from previous findings in the literature, and imply a greater role for adjustment costs. They are mostly consistent with models that combine elements of both time and state-dependent pricing, such as Alvarez, Lippi, and Paciello (2011).

My findings illustrate how scraped data can provide additional insights to the measurement of stylized facts in the stickiness literature, but the potential uses of scraped data in macroeconomics go far beyond those connected to this paper. For example, scraped prices can be used to create daily price indexes that complement official statistics, compare and test theories of international prices, and better measure exchange rate and commodity shock pass-through. To encourage their use by other people working in these and other research applications, the datasets used in this paper are publicly available for download at bpp.mit.edu.

References

- ALVAREZ, F., M. GONZALEZ-ROZADA, A. NEUMEYER, AND M. BERAJA (2015): “From hyperinflation to stable prices: Argentina’s evidence on menu cost models,” *SCID Working Paper 470*.
- ALVAREZ, F. E., AND F. LIPPI (2014): “Price Setting with menu cost for Multi-product firms,” *Econometrica*, 82, 89–135.
- ALVAREZ, F. E., F. LIPPI, AND L. PACIELLO (2011): “Optimal price setting with observation and menu costs,” *Quarterly Journal of Economics*, 126(4), 1909–1960.
- ÁLVAREZ, L. J., P. BURRIEL, AND I. HERNANDO (2005): “Do Decreasing Hazard Functions for Price Changes Make any Sense?,” *Working Paper Series - European Central Bank*, (461).
- BILS, M., AND P. J. KLENOW (2004): “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, 112, 947–985.
- BOIVIN, J., H. E. C. MONTRÉAL, M. P. GIANNONI, AND I. MIHOV (2009): “Sticky Prices and Monetary Policy : Evidence from Disaggregated U . S . Data *,” *American Economic Review*, No. 12824, 350–384.
- BONOMO, M., C. CARVALHO, AND R. GARCIA (2011): “Time- and State-Dependent Pricing: A Unified Framework,” SSRN Scholarly Paper ID 1930633, Social Science Research Network, Rochester, NY.
- BUREAU OF LABOR STATISTICS (2015a): *The Consumer Price Index*, vol. Chapter 17 of *Handbook of Methods*. BLS.
- (2015b): “New CPI Estimation System to be Introduced,” .
- CAMPBELL, J. R., AND B. EDEN (2014): “Rigid prices: evidence from u.s. scanner data *,” *International Economic Review*, 55(2), 423–442.
- CAVALLO, A. (2013): “Online and official price indexes: Measuring Argentina’s inflation,” *Journal of Monetary Economics*, pp. 152–165.
- (2015): “Are online and offline prices similar?,” *NBER Working Paper*.
- CAVALLO, A., AND R. RIGOBON (2011): “The Distribution of the Size of Price Changes,” *NBER Working Paper*, w16760, 1–40.
- CECCHETTI, S. G. (1986): “The frequency of price adjustment,” *Journal of Econometrics*, 31(3), 255–274.
- COSTAIN, J., AND A. NAKOV (2011): “Distributional dynamics under smoothly state-dependent pricing,” *Journal of Monetary Economics*, 58(6–8), 646–665.

- DHYNE, E., L. J. ÁLVAREZ, H. L. BIHAN, G. VERONESE, D. DIAS, J. HOFFMANN, N. JONKER, P. LÜNNEMANN, 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.
- EICHENBAUM, M., N. JAIMOVICH, AND S. REBELO (2011): “Reference prices, costs, and nominal rigidities,” *American Economic Review*, 101(1), 234–262.
- EICHENBAUM, M., N. JAIMOVICH, S. REBELO, AND J. SMITH (2014): “How frequent are small price changes?,” *American Economic Journal: Macroeconomics*, 6(2), 137–155.
- EINAV, L., AND J. LEVIN (2014): “Economics in the age of big data,” *Science*, 346(6210), 1243089.
- ELLISON, G., AND S. F. ELLISON (2009): “Search, Obfuscation, and Price Elasticities on the Internet,” *Econometrica*, 77(2), 427–452.
- GAGNON, E. (2009): “Price Setting During Low and High Inflation: Evidence from Mexico,” *The Quarterly Journal of Economics*, 124(3), 1221–1263.
- GORODNICHENKO, Y. (2008): “Endogenous information, menu costs and inflation persistence,” *NBER Working Paper*, (14184).
- GORODNICHENKO, Y., V. SHEREMIROV, AND O. TALAVERA (2014): “Price Setting in Online Markets: Does IT Click?,” *NBER Working Paper Series*, (20819).
- GORODNICHENKO, Y., AND O. TALAVERA (2014): “Price Setting in Online Markets: Basic Facts, International Comparisons, and Cross-border Integration,” *NBER Working Paper*, (20406).
- KASHYAP, A. K. (1995): “Sticky Prices: New Evidence from Retail Catalogues.,” *Quarterly Journal of Economics*, 110, 245–274.
- KLENOW, P. J., AND O. KRYVTSOV (2008): “State-Dependent or Time-Dependent Pricing: Does it Matter for Recent US Inflation?,” *The Quarterly Journal of Economics*, 73(3), 863–903.
- KLENOW, P. J., AND B. A. MALIN (2010): “Microeconomic evidence on price-setting,” in *Handbook of Monetary Economics*, vol. 3. Elsevier.
- KLENOW, P. J., AND J. L. WILLIS (2007): “Sticky information and sticky prices,” *Journal of Monetary Economics*, 54(SUPPL.), 79–99.
- LACH, S., AND D. TSIDDON (1996): “Staggering and Synzchronization in {Price-Setting:} Evidence from Multiproduct Firms,” *American Economic Review*, 86(5), 1175–1196.
- LÜNNEMANN, P., AND L. WINTR (2011): “Price Stickiness in the US and Europe Revisited: Evidence from Internet Prices*,” *Oxford Bulletin of Economics and Statistics*, 73(5), 593–621.

- 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,” *Quarterly Journal of Economics*, 123(4), 1415–1464.
- NAKAMURA, E., AND J. STEINSSON (2013): “Price Rigidity: Microeconomic Evidence and Macroeconomic Implications,” *Annual Review of Economics*, 5(1), 133–163.
- WOODFORD, M. (2009): “Information-constrained State-dependent Pricing,” *Journal of Monetary Economics*, 56, S100– S124.
- WULFSBERG, F. (2009): “Price adjustments and inflation - evidence from Norwegian consumer price data 1975-2004,” *Norges Bank Working Papers*.