

# **The Billion Prices Project: Using Online Prices for Measurement and Research**

(Preliminary and Incomplete)

Alberto Cavallo and Roberto Rigobon

Alberto Cavallo is the Douglas Drane Career Development Professor of Information Technology and Management and an Associate Professor of Applied Economics, and Roberto Rigobon is the Society of Sloan Fellows Professor of Management and a Professor of Applied Economics, both at the Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts. Cavallo is a Faculty Research Fellow and Rigobon is a Research Associate at the National Bureau of Economic Research, Cambridge, Massachusetts. Cavallo is the corresponding author and his email is [acavallo@mit.edu](mailto:acavallo@mit.edu).

**JEL-Codes:** E3, F3, F4.

**Keywords:** Online Prices, Inflation Measurement, Big Data, Scraped Internet Prices.

## **Section 1: Introduction**

New data sources, often referred to as “Big Data”, have the potential to improve economic statistics and empirical research in economics. This paper presents one example of how this can be achieved by using the vast amount of online prices being displayed on the web. We describe our work with the Billion Prices Project at MIT, and emphasize key lessons that can be used for both inflation measurement and some fundamental research questions in macro and international economics. In particular, we show how online prices can be used to construct daily price indices in multiple countries and to avoid measurement biases that distort evidence of price stickiness and international relative prices.

We start with inflation measurement. The basic procedure used in most countries to measure inflation has remained unchanged for decades. A large number of people working for National Statistical Offices (NSOs) visit hundreds of stores on a monthly or bi-monthly basis to collect prices for a pre-selected basket of goods and services. The micro data are then processed and used to construct Consumer Price Indices (CPIs) and other related indicators. Though carefully designed, this process is expensive, cumbersome, and often too slow for some users of the data. Groves (2011) describes the challenges faced by traditional survey-based methods, including growing levels of non-response. Shrinking budgets have complicated things in recent years, while crises have prompted policy-makers and other users of these statistics to demand faster and more accurate data.

Online prices have a natural appeal in this context. While the data are dispersed across hundreds of websites and thousands of pages, advances in automated “scraping” software now allow anyone to design and implement large-scale data collections on the web. Online data collection is cheap, fast, and accurate, making it an ideal complement for to traditional data collection methods, particularly in categories of goods that are

well represented online.

Our work with this type of data started back in 2007, when we (along with others) observed that in Argentina the official level of inflation that the NSO was showing did not seem to reflect the actual changes in prices. Seeking an alternative source of data, we started collecting daily prices posted on the websites of large retailers and used them to produce an alternative measure of inflation that was automatically updated on a website every day.<sup>1</sup> In Cavallo (2013) we showed that while Argentina's government announced an annual inflation rate of 8-10 percent from 2007-2011, our measure of inflation based on online-price data suggested that the actual annual inflation rate during that period was close to 25 percent on average -close to what most local macroeconomists believed the inflation rate to be . Our ability to collect prices remotely proved particularly useful in 2011, when Argentina's government started to impose fines and to threaten local economists who were collecting data independently. The manipulation of the official index continued for almost 9 years, ending in late 2015 when a new government was elected.

Argentina's statistical debacle had a positive side effect: it helped us realize the potential that online prices had for inflation measurement applications. With this idea in mind, we created the Billion Prices Project (BPP) at MIT in 2008 to extend our work to other countries, including the United States. The word "Billion" was simply meant to express our desire to collect a massive amount of prices, though we reached that number of observations in less than two year. By 2010 we were collecting 5 million prices every day from over 300 retailers in 50 countries. Half a million prices were collected every day in the US alone (by comparison, the Bureau of Labor Statistics in the US collects approximately 80 thousand prices on a monthly or bi-monthly basis). Although cheaper than traditional methods, collecting a massive amount of prices on a

---

<sup>1</sup> See [www.inflacionverdadera.com](http://www.inflacionverdadera.com). The original website had two price indices constructed with the official INDEC methodology: a "Basic Food" index and a broader "Food and Beverages" index. The website also showed the time series of prices for every good used in the index.

high-frequency level required an amount of funding that could not be sustained through grants, so in 2011 we started a company called PriceStats which nowadays collects the data and produces high-frequency indices for Central Banks and the financial sector. PriceStats greatly expanded both the quantity and quality of the data. The company currently uses about 15 million products from over 900 retailers to build daily inflation indices in 20 countries. Its micro datasets contain information from an even larger number of retailers in over 60 countries, with varying degrees of coverage. PriceStats' indices and the micro data are available to researchers working with the BPP, as we explain in Section 5.

Most of the other attempts to use “Big Data” in economics use social media or search data to *forecast* the behavior of important economic indicators such as unemployment. Our approach is different because we focus on *measurement*, not on prediction. Our goal is mainly to experiment with these new sources of information to improve the computation of traditional economic indicators, starting with the CPI. We seek to understand whether online prices have distinct dynamics, their advantages and disadvantages, and whether they could be a reliable source of information in a “production” setting (not just for a one-time research application).

We start the paper with a description of the methodology to collect online prices. A first order aspect is to realize that although the amount of data online is massive, carefully selecting the categories and retailers we sample is crucial. The goal is to obtain data that is representative of retail transactions, so we focus our data collection efforts on large multi-channel retailers such as Walmart, that sell both online and offline, instead of using online-only retailers that have a massive amount of products but a relatively small share of retail transactions. We also focus on categories of goods that are included in the official CPI baskets, for which consumer expenditure weights are available. After describing the sources of data, we discuss the advantages and

disadvantages of online data relative to other large micro-price databases (CPI data and Scanner data), and highlight the results of a large-scale validation exercise to show how online-price levels and behaviors closely resemble those that can be obtained by physically visiting offline stores.

Next, we describe the methodology used to compute online price indices and show how they closely co-move with CPIs in most countries. We emphasize two characteristics of online indices in greater detail. First, the ability to approximate hedonically-adjusted price indexes in categories with high turnover of products and a large number of overlapping price spells (such as electronics). Second, the anticipation of movements in the official CPI in many countries. This anticipation extends beyond the publication lag of the CPI and suggests that online prices often adjust sooner to aggregate shocks.

Moving on to research applications, we discuss two areas in macro and international economics where online price data can have a major impact. First, we show that online price data, collected daily, can significantly alter some key stylized facts in the price-stickiness literature. In particular, we document that online prices exhibit a very different distribution of price changes compared to micro CPI prices and scanner prices. The main reason for the difference is that online prices do not have time averages, common in Scanner Data, or imputed prices, common in CPI data, which create a large number of small spurious price changes. Second, online price data also challenges previous findings about the “Law of One Price” (LOP) as well as real exchange rate levels and dynamics across countries. The consensus in the literature is that there are large and persistent deviations from LOP, and that there is little pass-through from nominal exchange rates to relative prices, and vice-versa, causing persistent shocks to real-exchange rates that take years to dissipate. While the LOP still fails with online data most of the time, we find that it holds well across countries that

use the same currency. We also show that when real exchange rates are constructed using closely-matched goods across countries, then relative prices and exchange rates co-move much closer than previously thought, with higher pass-through rates and quicker real-exchange rate dynamics.

Both research examples illustrate how using data collected by others, with different purposes in mind, can greatly distort empirical findings. They also suggest we should not treat “Big Data” as simply large “found” datasets that solve problems simply with their size. Instead, we think one of the greatest contributions of “Big Data” for Economics is the fact that anyone can now use these new sensors, phones, web scraping, and other data collection technologies to build customized datasets designed to fit specific measurement or research needs. Instead of finding out what datasets are available, and what can be done with them, we should be asking what would be the ideal dataset to answer a particular question, and how can we use these new technologies to get it.

We end the paper by describing how the Billion Prices Project data and indices are publicly shared and by discussing why data collection is an important endeavor that macro and international economists should pursue more often.

## **Section 2: Collecting and Processing Online Price Data**

A large and growing share of retail prices all over the world are being posted online on the websites of retailers. This represents a massive untapped source of retail price information. But collecting these prices is not trivial, as they are posted on hundreds of different websites that lack a homogeneous structure and format. Retailers do not provide historical prices, so the data has to be collected continuously and consistently over time.

To collect and process these online prices we follow a “data curation” approach,

in the spirit of Stonebraker et al. (2013). It involves carefully identifying the retailers that will serve as data sources, using web-scraping software to collect the data, then cleaning, homogenizing, categorizing, and finally extracting the information so it can be used in measurement and research applications.

### *Section 2.1 The Selection of Retailers and Data Source*

The first step in our process is choosing where to collect the data. Two things are important to remember here: we do not collect all the prices that are available to us online, and we do not use all the data that we collect.

The first statement is about choosing where to sample the data. We carefully select the retailers and categories of goods that we want to focus on. These decisions are driven by our need to get prices that are representative of retail transactions. We therefore focus almost exclusively on large multi-channel retailers (those retailers that sell both offline and online, such as Walmart) and tend to ignore online-only retailers (such as Amazon.com). The reason is that multi-channel retailers still concentrate the vast majority of all retail sales, even in highly developed countries. We are also careful when we choose what categories of goods to monitor within each retailer. We concentrate on categories that are part of traditional CPI baskets. This allows us to use official CPI weights in many of our applications. We also avoid collecting too much data from categories that are over-represented online such as CDs, DVDs, and Books, (for example, Walmart sells over 300 thousand books online).

We make an effort to always collect the data directly from each retailer's website, rather than relying on third parties such as marketplaces, price aggregators, and price comparison websites. The data collection is far more challenging, but doing so maximizes our chances of obtaining prices linked to actual transactions and prevents third-parties from filtering or altering our samples in any way. It also gives us full control of what we choose to collect, and makes the whole process more robust, as it

does not depend on a single (or few) sources of data.

Once the data are collected, we clean them, standardize them to fit a common database schema, classify individual products using CPI categories, and start computing simple indicators to evaluate its characteristics and performance over time.

We treat each retailer as a separate sampling unit or “stratum” with potentially unique characteristics and pricing behaviors. Before including a retailer in a price index, we usually monitor its behavior for over a year to identify special characteristics in the data and know whether it is a useful and reliable source of price information.

A special characteristic of online prices is the fact that there is very little spatial differentiation. Most retailers that sell online have a single price for all shoppers in all locations within a country (though shipping costs and taxes may differ). Some retailers ask customers for zip code information to update inventory information. Grocery retailers are the only ones that tend to change prices depending on the zip code where the consumer is located. In such cases, we select a few zip codes corresponding to major cities and treat each case as an independent retailer.

The amount of data and the coverage of different categories that we can observe online vary across countries. For roughly 25 countries, our datasets currently have information on categories that cover at least 70% of CPI basket weights (services tend to account for the 30% that is not online).

### *Section 2.2 Data Collection using Web Scraping Software*

The technology to collect online prices on a large scale –called “web scraping”– is quickly improving. Just a few years ago, it required researchers to write programs in languages such as Python and PHP. See, for example, the discussion in Edelman (2012). Today, there are many simple “point-and-click” software solutions that require almost no technical expertise. Users can simply use their mouse to teach the software



what pieces of information they want to collect from a webpage. The software then creates a “robot” that is able to extract information from any other webpage with a similar structure, storing the information in a database. It identifies relevant pieces of information on a page by finding special characters of HTML code (the language that is used to create webpages) that come before and after each relevant piece of information. These “markers” are the same as long as the page does not change its look-and-feel. The challenge in web scraping is mostly to monitor the performance of the robots over time, quickly detecting errors in the data and fixing them. The robots we construct always collect the product’s id, name, description, brand, package size, category information, and a price. When available, we also collect other variables such as sale prices and stock indicators. We provide more details of the web-scraping process in the Appendix.

### *Section 2.3: Advantages and Disadvantages of Online Price Data*

To understand the strengths and weaknesses of this scraped online data for measurement and research applications, Table 1 offers a comparison with two other sources of micro-price data: traditional CPI data collected offline with survey methods by NSOs, and Scanner Data recorded from consumer purchases at cash registers by companies such as Nielsen. Detailed descriptions of these other data sources can be found in ILO (2004) and Feenstra and Shapiro (2003).

One of the most obvious advantages of online data is the low cost per observation. While not trivial, it is cheaper to use web scraping than hire people to visit physical stores or buy information from commercial scanner data providers such as Nielsen.

A second major advantage is the daily frequency. It is easier to detect errors in the data when it is collected at such high frequency. It also avoids time averages, which can generate spurious price changes as we discuss later on.

Third, online data includes detailed information for all products being sold by the sampled retailers. The cross-section of prices available is therefore much larger within categories than in traditional CPI data. In Section 3.2 we discuss how this “Big Data” feature can be used, for example, to simplify quality adjustments.

Fourth, there are no censored price spells in online data. Prices are recorded from the first day a product is offered to consumers until the day it is discontinued from the store. In contrast, in CPI data prices are often imputed or a substitute product is used when the agent surveying prices does not find the precise item from the previous time period. Knowing the full history of prices for individual goods can help to control for new good biases, make quality adjustments, and study prices and pass-through rates at time of introduction.

Fifth, online data can be collected remotely in any country. This is particularly useful in situations like the one experienced by Argentina in recent years, where the government was trying to prevent independent data collection for the computation of inflation. It also allows us to centralize the data collection and homogenize its characteristics.

Sixth, and related to the previous point, online datasets can be readily comparable across countries because prices can be collected with identical methods on matching categories of goods and time periods. This is useful in research applications that use cross-country comparisons, as we describe later on.

Finally, online data are available in real-time, without any delays to access and process the information. This is particularly useful for policy makers and anyone who needs up-to-date information.

Table 1: Alternative Micro-Price Data Sources

	Online Data	Scanner Data	CPI Data
Cost per observation	Low	Medium	High
Data Frequency	Daily	Weekly	Monthly
All Products in Retailer (Census)	Yes	No	No
Uncensored Price Spells	Yes	Yes	No
Countries with Research Data*	~50	<10	~20
Comparable Across Countries	Yes	Limited	Limited
Real-Time availability	Yes	No	No
Product Categories Covered	Few	Few	Many
Retailers Covered	Few	Few	Many
Quantities or Expenditure Weights	No	Yes	Yes

Note: \*Data from over 50 countries are currently being collected by the Billion Prices Project (bpp.mit.edu). Klenow and Malin (2010) report results with CPI data sourced from 27 papers in 23 countries.

One of the main disadvantages of online prices is that they currently cover a much smaller set of retailers and product categories than a government-run survey of consumer prices do. In particular, while the number and variety of goods whose prices are shown online is growing over time, the prices of most services are still not available on the web. The number and type of retailers is also limited compared to CPI data (though larger than in Scanner Data).

Another disadvantage is that online price data lack information on quantities sold. To obtain expenditure-weighted data, it is necessary to use weights from a government consumer expenditure survey or other sources. Scanner Data, by contrast, has detailed information on quantities sold, and could potentially be a source of high-frequency expenditure weights, although in a limited number of categories such as groceries.

### *Section 2.3: Are Online Prices Different?*

An important concern that often arises is whether online prices are different

from offline prices. This is a valid concern because most transactions still take place offline. The suspicion that online prices are different is fueled by reports that some online retailers use “dynamic pricing” strategies. See Mikians et al (2012) and Valentino-DeVries et al (2012) for some examples. In addition, many papers with “online prices” use data from online marketplaces such as Ebay or price-comparison websites such as Google Shopping. As Ellison and Ellison (2009) and Gorodnichenko et al (2014) have shown, these prices seem to change more frequently and in smaller sizes than in CPI data. However, the retailers in these datasets are mostly online-only stores participating in a fiercely competitive environment; not the type of “online data” we use.

To address some of these concerns, in Cavallo (2015a) we ran a large-scale comparison of online and offline prices for multi-channel retailers. Using a combination of a smartphone app, crowdsourced workers, and web-scraping we simultaneously collected prices on the website and at a physical store for over 20 thousand products in 50 of the largest multi-channel retailers in 10 countries. More than 500 “freelance” workers used their phones to scan barcodes in physical stores, manually enter prices, take photos of the price tags, and upload the information to our BPP servers. We then used the barcodes in the offline data to collect the prices for those exact same goods in the website of the same retailer within a 7-day time window.

This direct comparison between online and offline prices revealed a high degree of similarity in price levels, and both the frequency and size of price changes. On average about 70 percent of price levels were identical in the offline and online samples. The 30 percent that were different were mainly caused by delays in the online data collection, offline data collection errors, and sales that occur independently online and offline (as they do between different offline stores of the same retailer). The similarity was highest in retailers that sell electronics or apparel, and lowest in food retailers that

also tend price differently across different offline stores. While price changes do not have the exact same timing online and offline, they tend to have nearly identical frequency and average sizes. This suggests that the price spells for individual goods may not be synchronized online and offline, consistent with the evidence of anticipation in some online prices that we discuss in Section 3.3. Despite the general similarity, our results also revealed a great deal of heterogeneity among pricing behaviors, suggesting some validation is needed in papers with data from a limited number of retailers.

### **Section 3: Inflation Measurement**

Online prices are increasingly being used in inflation measurement applications. Besides the BPP and PriceStats, many NSOs are experimenting with the use of online data to complement traditional sources of CPI data, including the Bureau of Labor Statistics in the US (Horrigan 2013a), the Office of National Statistics in the UK (Breton et al. 2015), Statistics Netherlands (Griffioen et al. 2014), Statistics New Zealand (Krsinich 2015), Statistics Norway (Nygaard 2015), and Statistics Austria (Boettcher 2015).

In this section, we discuss the methods we have been using in the past 6 years to produce high-frequency inflation statistics with online data. We show that online price indices can closely approximate the CPI in a number of countries and settings. We then discuss how a large number of overlapping price series in the data simplify quality adjustments in categories with frequent product turnover, such as electronics. Finally, we use simple vector autoregressions to show that online price indexes can anticipate changes in the official inflation rate several months in advance—that is, beyond the one-month publication lag in official data.

#### *Section 3.1: Methodology for Comparison to Official CPIs*

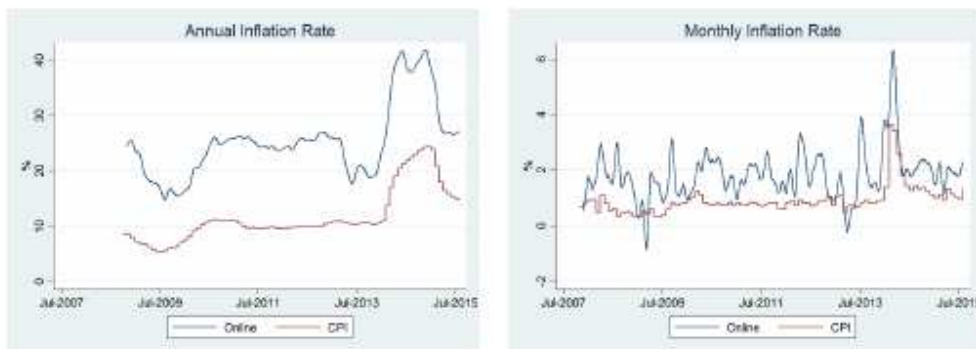
For multiple Latin American countries, Cavallo (2013) showed that online prices

could be effectively used as an alternative source of price information to construct price indices that mimic the behavior of official CPIs. The methodology for these daily indices is based on a combination of online data with standard CPI techniques, including expenditure weights for each sector where online data are available. This initial work included only data from food retailers and a handful of countries. In 2010, we founded a private company called PriceStats to expand the data collection and to start publishing inflation measures in real-time in other sectors and countries. Since March 2011, PriceStats has been publishing daily price indices in over 20 countries with only a three-day lag. In Figures 2 to 5 we plot these online indices next to the all-item non-seasonally adjusted CPI in each country. We first highlight the cases of Argentina and the United States, and then show some selected cases in a larger set of countries.

Figure 2 illustrates the case of Argentina from 2007 to 2015. Panel (a) compares a price index produced with online data to a comparable level of the official CPI.



(a) Price Index



(b) Annual Rate

(c) Monthly Rate

Figure 2: Argentina

Notes: Monthly Inflation Rates for daily online indices computed as the percentage change in the average of the previous 30 days compared to the same average a month before. Annual Inflation Rates for daily online indices computed as the percentage change in the average of the previous 30 days compared to the same average 365 days before. Online price indices computed by PriceStats. CPIs are sourced from the National Statistic Offices in each country. Detailed Sources provided in the Appendix. All price indices are non-seasonally adjusted.

The fact that the two measures of inflation in 2(a) diverge so dramatically will not surprise anyone who knows the recent story of statistics in Argentina. Back in 2007, when we started collecting online data from large supermarkets in the country, there were many suspicions that the Argentinian government, which was extremely sensitive to the accusation that its policies might be causing inflation, was keeping the official inflation estimates low. But before a measure of inflation based on online prices became

available, there was no clear way to confirm the magnitude of the discrepancy and track its evolution over time.

The manipulation of the inflation index in Argentina continued for almost 9 years, starting in February 2007 and ending in December 2015 when a new government was elected. As can be seen in Figure 2(c), the monthly inflation rate was consistently higher than the official data suggested, with the exception of a brief period in 2014 when the government launched a new CPI in an attempt to regain credibility. This was in response to the IMF's "motion of censure" issued in 2013 (See Rastello and Katz 2013). Unfortunately both indices started to diverge again within a few months and the new CPI quickly lost all credibility.

Looking at the discrepancy in the trend, however, misses an important point. The online price index tracked the dynamic behavior of the annual inflation in the official CPI, as shown Figure 2(b). There difference was mostly in the *level* of the annual inflation rate, not its movements over time. The online index also quickly reacted of large macro shocks, such as the massive road blocks by farmers in 2008. This suggested that online data was capable of capturing the fundamental dynamics of the official CPI and prompted us to collect data in other countries to see how online and offline price indices behaved in other settings<sup>2</sup>.

The results in other countries were completely different. The daily US index, shown in Figure 3, is a great example. Despite the multiple reasons why we might expect inflation indexes based on online and offline prices to deviate, the US online index has co-moved closely with official CPI for over seven years.

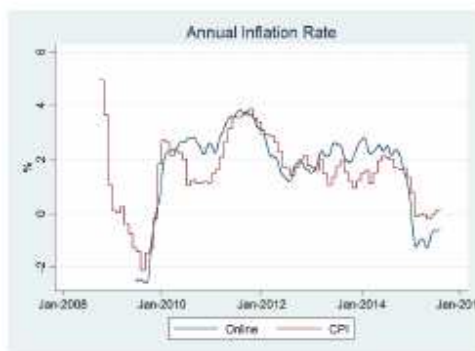
---

<sup>2</sup> It also implied that the government was not using a particularly sophisticated algorithm to change the inflation rate. In Cavallo (2013) we showed that one could closely approximate the official inflation rate by simply dividing the online inflation rate by three.





(a) Price Index



(b) Annual Rate



(c) Monthly Rate

Figure 3: USA

Notes: Monthly Inflation Rates for daily online indices computed as the percentage change in the average of the previous 30 days compared to the same average a month before. Annual Inflation Rates for daily online indices computed as the percentage change in the average of the previous 30 days compared to the same average 365 days before. Online price indices computed by PriceStats. CPIs are sourced from the National Statistic Offices in each country. Detailed Sources provided in the Appendix. All price indices are non-seasonally adjusted.

Although there are periods where the indices diverge, the differences are relatively small and temporary. This can also be seen in the monthly and annual inflation rates in Figures 3(b) and 3(c).

Interestingly, the online index seems to be particularly good at anticipating major changes in inflation trends. Predicting these changes in the official inflation rate is important for participants in financial markets, policymakers, and those economists

who monitor the economy closely.

One remarkable example of a turning point detected with online data months before it showed up in CPI data was September 16, 2008, the Tuesday after Lehman Brothers filed for bankruptcy. As Figure 4 shows, the online price index peaked that day and started falling. By October 15th, it had lost almost 1.2 percent in a single month. On October 16th, the CPI for September came out with only a 0.14 percent drop. When the October CPI numbers were published on November 19th, it has fallen another 1.01 percent. In other words, it took more than two months after Lehman's disaster for the official CPI numbers to reflect the full impact. Two months later, on December 16th 2008, the online price index stopped falling and started to increase once again. The official CPI did not show this change in the trend until the estimates for January were published on February 20th. We measure the degree of anticipation in online data more formally later on.

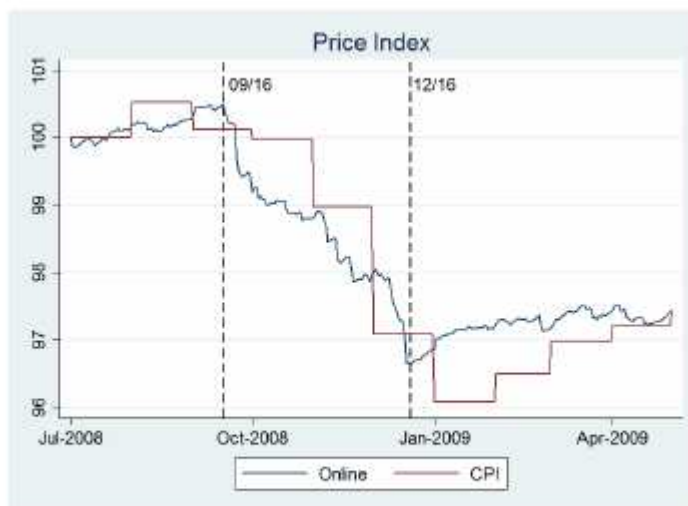


Figure 4: US Price Index around Lehman Brother's Bankruptcy

Figure 5 compares inflation as measured by online prices and by the offline prices in the official CPI for a selection of other countries and individual sectors. The main lesson of the figure is that the correspondence is reasonably close, but some more

specific insights are also possible here. First, we do not find any evidence that China is systematically holding its official inflation rate below the rate based on online prices. Second, the difference between the online price index and the official CPI appears to be smaller in developed countries like the UK and Germany, and greater for countries like Brazil or Turkey, where the online sector seems to have more independent patterns. In Japan, we observed significantly more inflation after the earthquake and an immediate impact of the sales tax changes in April 2014. While the online index in Japan does not follow the official CPI closely, it does seem to anticipate certain changes in trends.

The third row of Figure 5 shows results for a few US sectors. As one might expect, the online data matches the CPI better in sectors such as food and electronics, for which online information is widely available. By contrast, some official inflation patterns seen in the “Medical Care” sector are not well captured by online prices, probably because many services cannot be monitored online. The fourth row shows that online data can be used to provide global aggregates using country consumption weights.

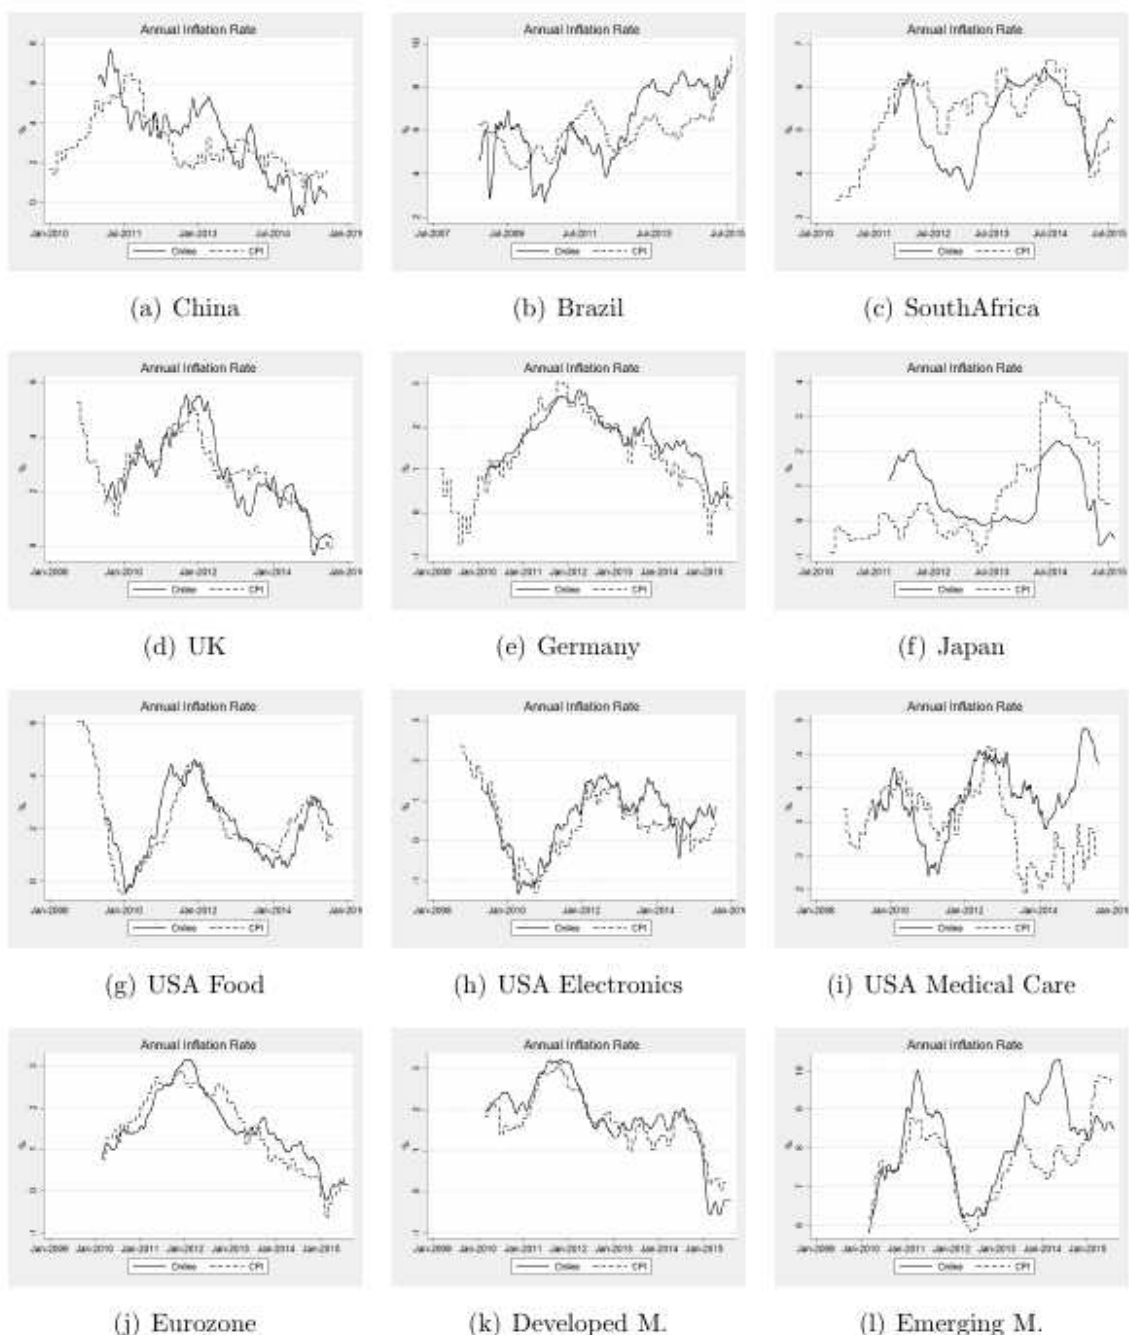


Figure 5: Online vs CPI Annual Inflation Rates

Notes: Annual Inflation Rates for daily online price indices computed as the percentage change in the average of the previous 30 days compared to the same average 365 days before. Online price indices computed by PriceStats. CPIs are sourced from the National Statistic Offices in each country. All indices are all-items non-seasonally adjusted series with the exception of China, where an online supermarket index is shown next to the official food index. Global aggregates in the last row are computed using relative consumption weights in each country and CPIs from official NSOs.

It may seem surprising to many readers how indexes based on online data have the ability to mimic official CPIs in so many cases: for large and small countries, for developed and emerging markets, and for the aggregate and sectoral data. After all, the

data differs significantly from traditional sources, and we do not apply many adjustments and methods used by NSOs, such as hedonic quality adjustments. We believe there are two reasons for this. First, we carefully design and select the data that goes into these indices to ensure that they are representative, as we mentioned in Section 2. Second, we learned that many sampling characteristics in our data made it simpler to deal with some traditional measurement problems. To illustrate this, we next discuss how online data can simplify quality adjustments by providing a large number of uncensored and overlapping price spells.

### *Section 3.2: Overlapping Quality Adjustments*

Quality adjustment poses problem for any measure of inflation: as is widely understood, if a good rises in both quality and price, then some of the price increase is presumably due to the quality changes and should not be attributed to inflation. NSOs use different methods for quality adjustments, including making efforts to find the closest comparable substitutes when a product disappears and often relying on adjustments with hedonic regressions (in which the price change is calculated while holding constant certain attributes of a good, like the memory or hard drive capabilities of a computer). It might seem that measures of inflation based on online prices would be especially susceptible to quality problems: after all, they don't adjust for quality in any formal way. But it turns out that because online data does frequent sampling of a large number of models or varieties, it can implicitly adjust for quality differences as well as seemingly more sophisticated hedonic-regression adjustments sometimes used in official CPIs.

To build some intuition for why this result holds true, consider a hypothetical example of a series of prices in Figure 6. The panel on the left illustrates the data resulting from a traditional offline data collection process. Each line represents the price of a single good over time. Many models of electronics products,

such as TVs, dishwashers, washing machines, and vacuum cleaners, that tend to be introduced at relative high prices, are discounted gradually over their life cycle and have clearance sales right before disappearing from the stores (Silver and Heravi 2005). With traditional data collection methods in an offline price index, it is not possible to collect the prices for every good available for sale at each point in time. Instead, the data collector would focus on one (or a few) of the most popular models and record its price once per month until it disappears from the store. When one particular model is no longer available, the data collector starts monitoring a different model, as shown by the vertical dashed line in the figure. But at the time of the shift, the previous prices for the model that is now being included in the index is unknown (and shown where the line is shaded more lightly on the figure). At the time when one model of a product is substituted for another, the price gap between models probably includes both a quality difference and an actual price change. The problem is compounded in goods that experience extreme price movements along their life cycles, and may have steep discounts right before disappearing from the shelves.

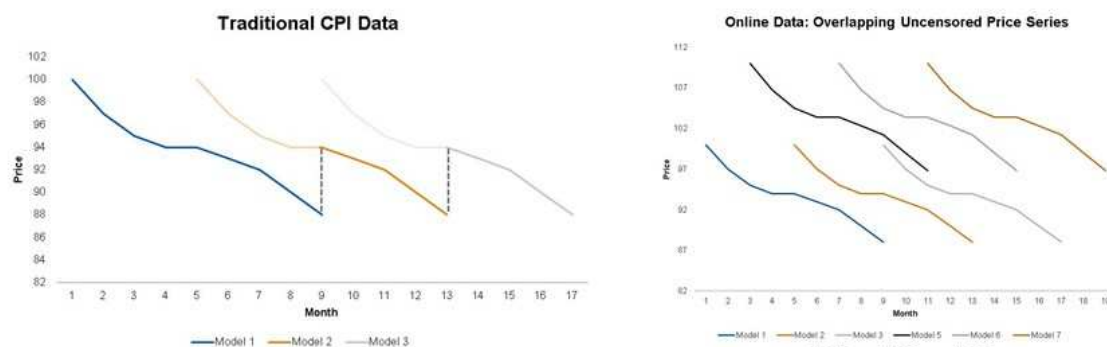


Figure 6: Overlapping Qualities

NSOs have two main ways of dealing with this problem. The preferred method is to use hedonic techniques. Again, these involve setting up a regression with the price of a good on one side and actual attributes of the good on the other side, so that future changes in the price of the good can be calculated while holding constant the attributes

of that good. While hedonic techniques have become more popular, the question of what traits should be included in the regression, how those traits should be measured, and what specification should be used can tend to make hedonic techniques very data-intensive and complex to implement.

A much simpler alternative method is to use “overlapping qualities”. As Armknecht and Weyback (1989) point out, this assumes that the price difference at the time of introduction of the new variety is mostly reflecting a quality difference. The main problem with this approach, however, is that the price of the new good is typically not observed at the time it is introduced, but much later, when the old good disappears from the stores. This is noted in the ILO CPI Manual (2004, p. 78): “When there is overlap, simple linking... may provide an acceptable solution... In practice, however, this method is not used very extensively because the requisite data are seldom available...The information needed for this...will never be available if price collectors are instructed only to introduce a new quality when an old one is dropped.”

Online prices offer a simple solution to the data problem by providing a large number of uncensored price spells for all models on sale at any point in time. The result is a set of data like that shown in the right-hand panel of Figure 6. With this type of data, a simple index can often closely approximate official indices that use complex hedonic quality-adjustment methods. This was documented in the price-index literature before. For example, Aizcorbe et al. (2000) and Aizcorbe et al. (2003) used scanner data to demonstrate that, with high-frequency data, matched-model price indices could yield results that are numerically close to those obtained using hedonic techniques in samples where product characteristics did not change much over time. In general, the extent to which a simple matched-model price index can capture quality change will depend on several factors (see Silver and Heravi 2005). First, both varieties of the product need to have a substantial degree of overlap in their prices (as illustrated



in the right-hand panel of Figure 6). Second, there needs to be a large number of models and relatively low degree of entry and exit, so that existing varieties can capture aggregate effects without being overly affected by idiosyncratic price movements of goods that enter and disappear from the sample.

As evidence of this effect, consider the data in Figure 7. It contains three price indices for televisions in the US from 2008 and 2009. The solid line shows the official CPI for TVs as computed by the US Bureau of Labor Statistics using hedonic methods. The line with long dashes shows an online price index based on 50 distinct models of TVs from a large US retailer. The line with short dashes shows an online price index with 500 models from the same source. As we increase the number of models included in the index, we more closely approximate the results of the hedonic price index constructed by the BLS during this time period. Intuitively, the more overlapping price series being used, the less important the extreme price movements of goods being sold at clearance prices or newly introduced will be for our price index.

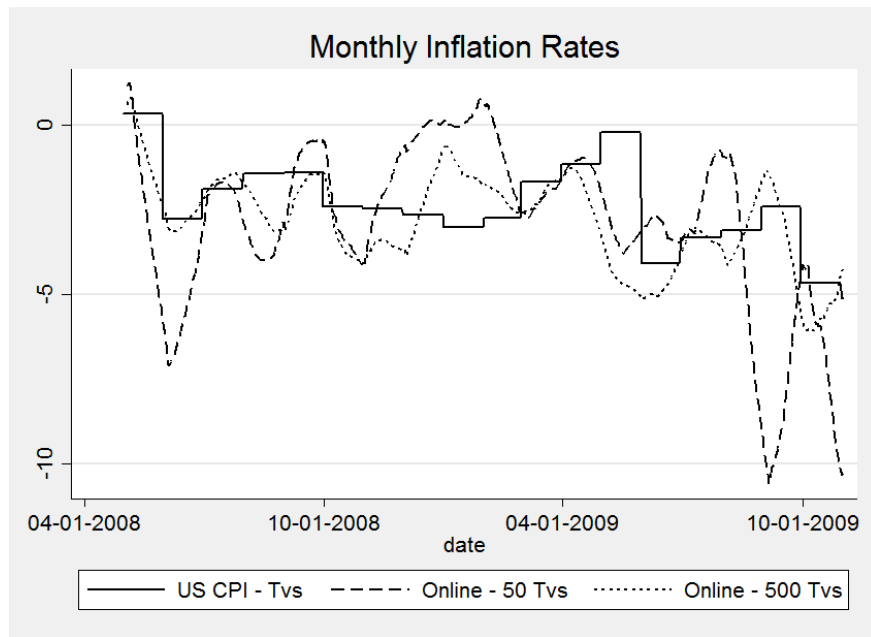


Figure 7: US CPI with Hedonics vs Online Index

This example illustrates one of the true “big data” advantages of online prices. We may not need or want to use every single data point available in these large datasets, but being able to extract and use uncensored spells for a large number of models can greatly simplify measurement in cases like these. For example, Krsinich (2015) showed that online data can be used to construct a time-product dummy index that it is equivalent to a fully-interacted time dummy hedonic index based on all product characteristics. And even if the goal is to run a full hedonic regression, online data can potentially supply a large number of models and detailed characteristics needed to do so.

### *Section 3.3: Anticipation of CPI*

As suggested before, another important characteristic of online price indices is that they can provide significant anticipation to changes in CPI inflation. In this section we document this pattern and conjecture about some possible explanations.

To more formally document the degree of anticipation, we estimate a simple vector autoregression with the official CPI as the dependent variable and our online price index is the exogenous variable, and compute an impulse response to see how shocks to the online price index impact the CPI over time. The regression is expressed in monthly changes (we use monthly log changes in the CPI and monthly log of the last day each month for our online index). The regression includes six lags of each variable, plus the contemporaneous value of the online price index to account for the early availability of the online price information<sup>3</sup>.

Figure 8 shows the cumulative impulse response of the CPI to a shock in the online index over time, together with the 95 percent confidence intervals. The top left panel is for the aggregate index, and the rest for the sectoral indices.

---

<sup>3</sup> For each month  $t$ , the specification is as follows:

$$\Delta \ln(CPI_t) = \alpha + \beta \Delta \ln(Online_t) + \sum_{i \in [1,6]} \alpha_i \Delta \ln(CPI_{t-i}) + \sum_{i \in [1,6]} \beta_i \Delta \ln(Online_{t-i})$$

The confidence bands are computed by bootstrapping in blocks. This specification gives the online price index the highest chance to explain the observed variation. There is, however, no unambiguous way of identifying the system given that under the null hypothesis both indices are valid measures of the underlying inflation. We chose this specification because it matches the actual availability of data at the end of each month: the online index is immediately available, while the CPI has a publication lag of 15 days in most countries.

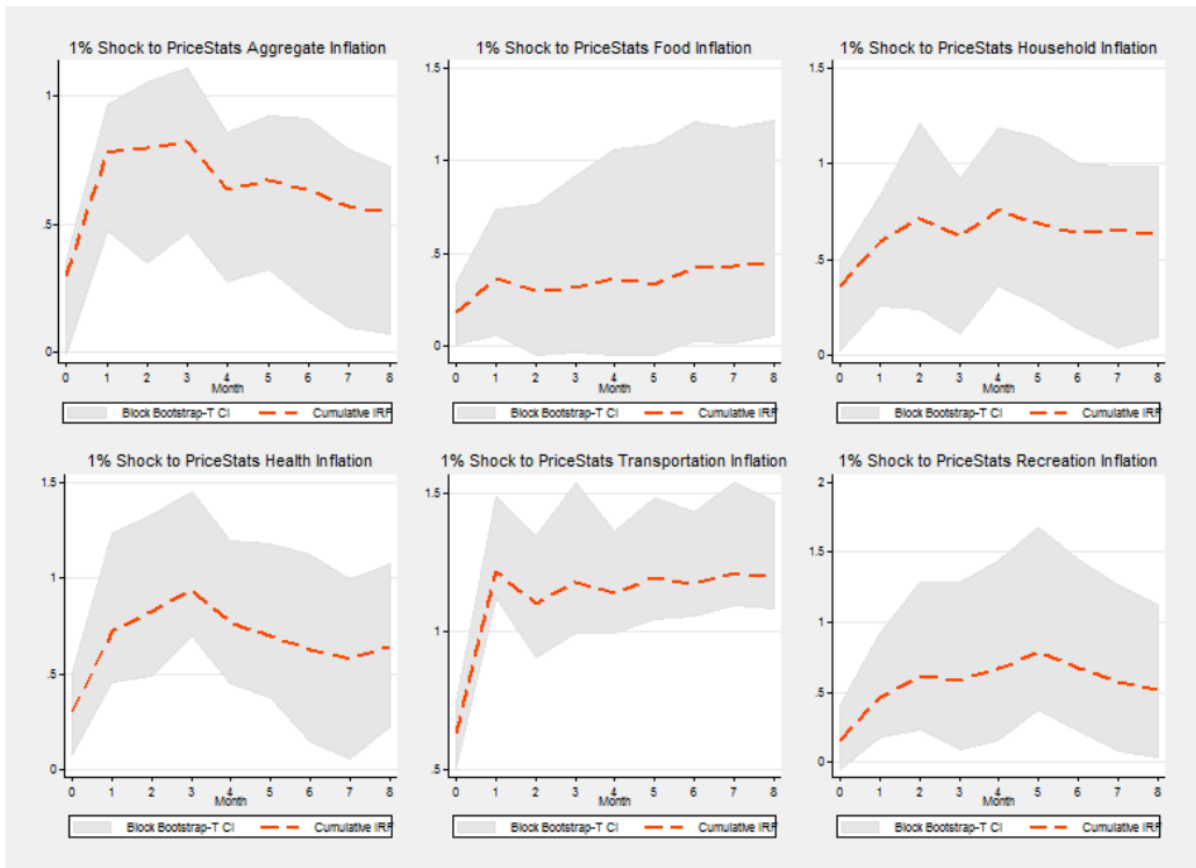


Figure 8: Impulse Response of the US CPI to a Online Price Index shock

Notes: CPI-U, US city average, all items, non-seasonally adjusted from the Bureau of Labor Statistis. The US Online Price Index is computed by PriceStats.

These impulse responses show that it takes several months for the CPI to fully incorporate the shock to the online price inflation. The impact is quickest in transportation (fuel) and slowest in food. The result is robust to the elimination of the contemporaneous effect of the official price index from the vector autoregression. In most cases, the anticipation significantly exceeds the typical publication delays in official statistics. Moreover, we find similar degrees of anticipation in all developed economies in our data.

The reasons why there is anticipation are still open questions for future research. Possible reasons include delays embedded in the methodology used for the official CPI,

difference in mixture of stores sampled by the official and online index, or different pricing dynamics by the stores.

Regardless of the reasons for this anticipation, the pattern in Figure 8 suggests that online data can be a useful addition to inflation forecasting models that mostly rely on lagged values of official data. This was explored in Bertolotto et al (2013). More recently, Aparicio and Bertolotto (2016) show that out-of-sample inflation forecasts using online data outperform a large number of alternative forecasting models in the US and the UK.

## **Section 4: Lessons for Macro and International Research**

In this section, we illustrate how online data can change empirical results in macro and international research, by focusing on price-stickiness and real-exchange rate behaviors.

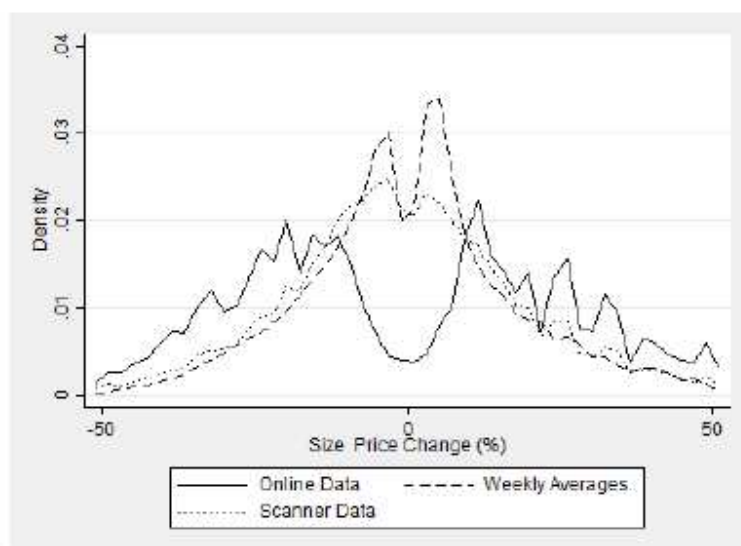
### *Section 4.1: Price Stickiness and the Distribution of Price Changes*

Sticky prices are a fundamental element of many macroeconomic models. In the past decade, a large empirical literature has tried to measure price stickiness and understand its micro-foundations (for example, Dhyne et al. (2006) in this journal; for a survey of the literature, see Nakamura and Steinsson (2013), and the references cited there). This research has been possible due to an unprecedented access to micro-level CPI data and scanner datasets in several countries. Over time, the literature settled on a set of stylized facts, summarized by Klenow and Malin (2010). In Cavallo and Rigobon (2011) and Cavallo (2015b), we use online data to argue that the sampling characteristics of official CPI and scanner data can introduce measurement biases that affect the stylized facts in the literature on patterns of price changes.

As one prominent example, a pattern that has received a lot of attention in the

literature is the shape of the distribution of the size of price changes. Most papers using CPI data found bell-shaped distributions with a significant share of small price changes, which seemed inconsistent with standard menu-cost models that predict periods of unchanging prices followed by relatively large changes, rather than a series of frequent tiny small changes. This finding motivated a surge in papers trying to adapt sticky-price models to account for this fact (for example; see Woodford 2009 and Midrigan 2011).

The shape of the distribution of price changes is, however, greatly affected by the sampling characteristics of the data. For example, scanner data are usually reported as a weekly average of sales and quantities. This time averaging can create a large number of spurious small changes, as noted by Campbell and Eden (2014). Consider, for example, a three week period with a single price change in the middle of the second week. With time averages, it would appear that there are two smaller price changes. This effect can be seen explicitly in Figure 9, where we show a distribution of prices changes for online data and scanner data collected by Nielsen from exactly the same US retailer, zip code, and time periods.



(a) Online vs Scanner

Figure 9: The Distribution of the Size of Price Changes in the US

Similarly, a growing number of papers have started to document measurement biases in the CPI and trying to adjust for them. The biases are different in nature but similar in the effects on the distribution of price changes. For example, CPI data will often include imputed prices for temporarily missing items. This imputation is often done using the average price change of related goods, resulting in an artificial pattern of many small price changes. As another example, Eichenbaum et al. (2014) also use CPI and scanner data from multiple stores to show how “unit-value prices,” which are reported as the ratio of sales revenue to the quantity sold, make it appear that small price changes are more prevalent.

Online prices can not only avoid time averages and other sampling characteristics that distort evidence of price stickiness, but also provide stickiness statistics in a large number of countries and economics settings. Until now, cross-country comparisons in the price stickiness literature had to rely on results obtained by different papers, with different data sources, time periods, methods, and data treatments. For example, Klenow and Malin

(2010) compare results for 24 countries sourced from 27 different papers. However, it is hard to know what is driven by country-level differences, or the methods and characteristics of the data. Dhyne et al. (2006) was one of the few papers that used similar data from multiple countries thanks to the coordination provided by the European Inflation Persistence Network. Even in this case, however, each European NSO was unwilling to share the micro data 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.

To better understand the determinants of price stickiness, and policies that affect them, we need detailed pricing information in multiple countries and economic settings, free from measurement biases, and with identical characteristics, treatments, and time periods.

#### *Section 4.2: International Prices and the Law of One Price*

Another research area where online data can have a significant impact is international economics. The relation between relative prices and exchange rates is a classic question in this field. A basic hypothesis is the Law of One Price (or “PPP” when considering many goods), which implies that there should not be large or persistent cross-country differences in the prices of identical goods when translated into a common currency. Modest deviations are not surprising in a world with transport costs and other barriers to arbitrage. However, a huge literature documents a surprisingly large failure of the Law of One Price for many traded goods at retail prices, resulting in significant volatility in the relative cost of consumption across countries. This failure occurs not only in levels (“Absolute PPP”), but also in changes over time (“Relative PPP”). Furthermore, nominal exchange rate shocks tend to have persistent effects on



the real-exchange rate, leading to what Kenneth Rogoff called the “PPP puzzle”. At the core of this puzzle is the fact that relative prices do not seem to adjust quickly to nominal exchange rate shocks. Many papers have documented the slow response of prices by measuring long-run exchange rate “pass-through” rates of only 20%-30%.<sup>4</sup>

The literature is hampered by the formidable difficulties in obtaining prices in a large number of identical goods sold simultaneously in a large number of countries. In practice, researchers are forced to settle on having prices for identical goods from two countries (typically the US and Canada), or use price indices from a large number of countries (constructed with different methods and baskets, and precluding any price level comparisons). Some micro sources of data, such as the Economist magazine’s Big Mac index, provide information on many countries but are limited to a single good. The World Bank’s International Comparisons Program (ICP) makes a worldwide statistical effort to collect price data in order to estimate PPP-adjusted GDPs in dozens of countries, but doing so with traditional methods is so daunting that it can only be carried out every five years or more, severely limiting its use for research on real-exchange rate levels and dynamics.

In principle, online prices can be obtained in high frequency, for large number of goods, in dozens of countries. The main challenge is not in the raw data collection, but rather in the matching of identical products using barcodes and ids that are not the same across countries or retailers. In particular, UPC codes and model numbers are often different for identical goods in different countries.

In Cavallo, Neiman and Rigobon (2014a) we solved the matching problem by using prices collected from global retailers such as Apple, Ikea, Zara, and H&M, who sell

---

<sup>4</sup> Burstein and Gopinath (2013) provide a recent review of the empirical literature and a discussion of some theoretical advances, including accounting for non-tradeables or tradeables that are only locally consumed, variable markups, and pricing-to-market.

identical goods with the same id in several dozen countries. This allowed us to directly study conditions under which the “Law of One Price” holds. Much to our surprise, we found that the law of one price only holds well in countries that share the same currency (both within the Euro area, and with countries that use US dollar, such as El Salvador and Ecuador). Countries that are physically close, or in a trade union, or even in strong pegs, do not share this property. Instead, what really seems to matter for these global retailers is simply whether prices have to be shown to customers in the same currency. In Cavallo, Neiman, and Rigobon (2014b) we used the introduction of the Euro in Latvia in January 2014 to show that the adjustment towards the law of one price can take place within a matter of days after a country joins a currency union. This type of price convergence was, after all, one of the founding objectives of the Euro.

The main implication of this line of work is that choice of currency units is far more important for defining the boundaries between markets for goods than has previously been suspected, while things that were traditionally thought to be important, such as physical distance, political and tax territories, language, and culture, do not seem to matter so much. Furthermore, the patterns we documented also point to the importance of customer psychology, organizational structure, and the Internet for price setting behavior. For example, firms may fear antagonizing customers who see prices posted on the web in the same currency across borders. Such considerations do not yet feature prominently in most standard macroeconomic models.

In more recent work, we are expanding the matching to goods sold on the largest retailers in each country, whether global or not. Without identical ids, the challenge is to classify a large set of heterogeneous varieties (different package sizes, flavors, retailers where they are sold, among others) into narrowly defined product categories such as “Basmati White Rice, 1kg” or “LG Basic Blu-Ray Player, 1 unit”. PriceStats has

been classifying over 30 thousand individual goods in 300 product categories since 2014. This is achieved by first using supervised machine learning (specifically a Naive Bayes classifier). The model trains on language specific, hand-categorized items. Then each word in the product description of a good is parsed to find explicit indicators that a product fits into a specific product definition. Next, a team of people use the set of pre-classified ids to decide whether they should be linked to each product category or not, simultaneously checking the package size details to calculate standardized unit prices across countries, as described in Bertolotto (2016). The end result resembles a collection of hundreds of “Big Mac”-type indices for different kinds of goods.

These matched indices can be used to produce high-frequency PPP metrics to study real-exchange rate levels and dynamics. For example, Figure 10 shows PPP metrics constructed by PriceStats for an average of more than 250 goods in food, electronics, and fuel, in Argentina and Australia relative to the US (other countries provided in the Appendix). The top panel shows the real exchange rate, defined as the relative cost of the basket when expressed in the same currency. This is the ratio of the relative price ratio (in local currencies) and the nominal exchange rate (defined as local currency per US dollar). Both of these variables are shown in the bottom panel. For the case of Argentina, we also included the black-market exchange rate.

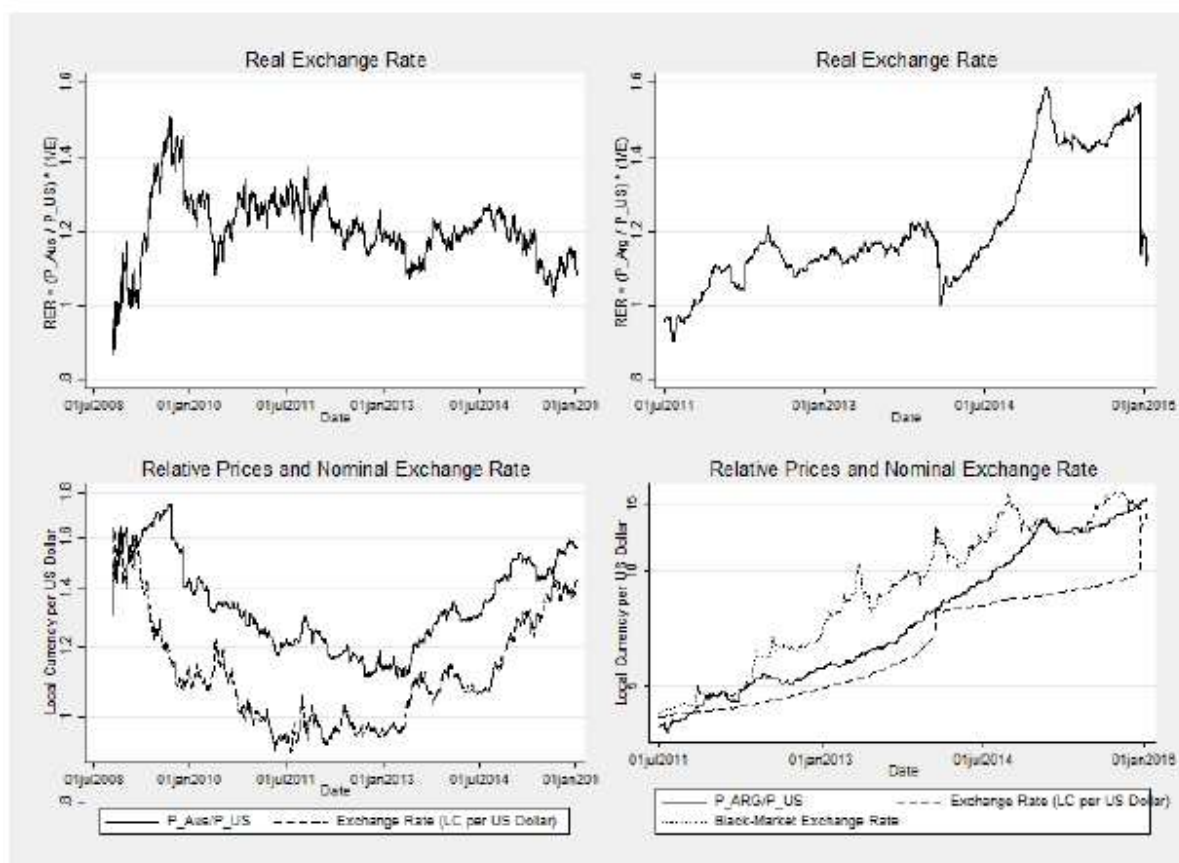


Figure 10: Relative Prices and Exchange Rates

A common finding in most countries is that shocks to the real-exchange rate seem far less persistent with online data. That is, the real-exchange rate reverts back to its average levels relatively quickly. As the bottom panel shows, this happens because relative prices co-move closely with the nominal exchange rate movements. For example, as the Australian dollar appreciated from 2008 to 2011, relative prices in Australia fell to compensate, and when the Australian dollar started to depreciate again in 2013, relative prices rose. The results in Argentina also show co-movement between these two variables. The steady increase in relative prices was matched by the overall trend of depreciation in the currency, which is gradual in the black-market but lumpy on the official exchange rate. There are long periods where prices kept rising and the official

exchange rate was fixed by the government, causing “deviations” in the real-exchange that were suddenly reversed in the two occasions when the country devalued its currency, in January 2014 and December 2015.

The co-movement between relative prices and exchange rates implies high rates of pass-through, which can go in both directions. While in Australia there is evidence that nominal exchange rate shocks affect retail prices (as the literature tries to capture in traditional “pass-through” estimates), in Argentina’s case it is relative retail price movements that tend to precede nominal exchange rate adjustments.

Perhaps more important than the dynamics is the fact that online data provides information on relative price *levels*, which are not available when using CPIs and other price indices. For example, looking at the real-exchange rate in Figure 10 we see that the basket tends to be 20% more expensive in Australia, while in Argentina the law of one price tends to hold quite well when the currency is allowed to float. Knowing this is useful to estimate the degree of currency misalignment, particularly in countries with managed foreign exchange markets. For example, in December 2015 the new government of Argentina was deciding to remove all foreign-exchange restrictions. It was unclear what the market exchange rate would become, and what effect that would have on tradable prices. Some economists argues that prices were already being set using the black-market exchange rate of 15 pesos per dollar. The official rate was instead 9.6 pesos per dollar.

The nominal exchange rate implied by PPP (the one that would make the real-exchange rate equal to one) was 14.3, suggesting that the official rate was greatly overvalued while the black market rate was slightly undervalued. When the market was liberated, the new exchange rate quickly settled around 14 pesos per dollar, matching almost perfectly the relative price ratio. This can be seen in the jump of the exchange

rate in the bottom right panel of Figure 10. While we do not expect these metrics to help predict exchange rates so closely in every country and situation, they can certainly provide better measures of the amount of deviation of real-exchange rates from “normal” levels.

So far, our micro data has only been matched for 7 countries and the time series is still too short to make strong inferences in some cases, but it is clear that some key puzzles in international economic and macroeconomics that has emerged from studies using CPI look quite different when viewed through the perspective of online data.

## **Section 5: Access to the Billion Prices Project Data**

As an academic project, we share as much data and results as possible on our webpage ([bpp.mit.edu](http://bpp.mit.edu)). Most of the micro data and indices used in our papers are currently available to download on that page, together with detailed scripts that allow others to replicate and extend our results. The micro data are posted with little pre-treatment, so researchers can apply their own methods. We plan to periodically upgrade the shared data, increasing the number of databases and retailers, and also expanding the time series as we continue to collect more prices over time.

We also publicly distribute some of the *offline* data collection tools developed at the BPP, such as an Android app used to validate online prices mentioned in Section 2. The app can be freely used to collect any kind of offline prices, and our servers will clean, consolidate, and share the data automatically with each group of users.

The US and Argentina inflation indices used in this papers are published with a 30-day lag on the BPP website, while the PPP discussed in Section 4 are currently published with a one-year lag on PriceStats’ website. The raw micro data collected by PriceStats are not publicly available but can be shared with academic researchers who collaborate with

the BPP and sign a data-sharing agreement.

## **Section 6: Final Remarks**

The Billion Prices Project is just one example of the use of “Big Data” in Economics.<sup>5</sup> To us, one of the greatest opportunities is the ability to do data collection that is customized to fit specific measurement and research needs. Although online price data are the focus of this paper, we hope to have convinced economists and perhaps a few policymakers of the benefits of experimenting with alternative data sources. Other examples include various types of scraped data, such as labor and real estate information available on the web, along with data from mobile phones, satellite images, GPS signals, and many other sensors that are increasingly becoming part of our daily lives. These new sources of information can be used to gain full control of the process of data collection, cleaning, and treatment. This will help with both the quantity and quality of the datasets used for measurement and research applications.

While many governments have been active in searching for alternative data sources, hoping both to increase the quality of statistics and to reduce cost, their use will require not only the will of policy-makers and statisticians working on the field, but also the involvement of more economists, academics, and practitioners, who can help identify the best ways to collect, process, and treat these new sources of information.

The need to get involved in data collection was eloquently pointed out many years ago in Griliches (1985) and repeated in his Presidential Address at the American Economic Association in 1994. In his words,

---

<sup>5</sup> Einav and Levin (2014) provide a more general overview of this topic, including new granular data sources, computational techniques such as machine learning, and the role of theory in analyzing large, unstructured datasets. Varian (2014) describes in detail some new “Big Data” techniques which are useful to analyze large datasets.

*“We [economists] have shown little interest in improving [the data], in getting involved in the grubby task of designing and collecting original data sets of our own. Most of our work is on “found” data, data that have been collected by somebody else, often for quite different purposes... “They” collect the data and are responsible for all their imperfections. “We” try to do the best with what we get, to find the grain of relevant information in all the chaff”*

Big Data technologies are finally providing economists, particularly in macro and international, with greater opportunities to stop treating the data as “given” and get directly involved with data collection. This will help mitigate issues in empirical research such as sample selection, endogeneity, omitted variables, and error-in-variables, which are so frequently found in traditional datasets. It will also give us a closer connection to those aspects of the economy that we seek to explain.



## Acknowledgements:

We thank Manuel Bertolotto, Maria Fazzolari, Diego Aparicio, Belen Bazano, and Augusto Ospital for outstanding research assistance at various stages of our Billion Prices Project research. We thank Pilar Iglesias and all the team at PriceStats for sharing their indices and micro data. We also thank the large number of students that helped us over the years through MIT's Undergraduate Research Opportunities Program. We thanks Chris Sims for detailed comments on an earlier version of this paper, as well as seminar participants at the MIT, IMF, IADB, Census Bureau, BEA, ECB, Macro Financial Modeling Conference, NY, Federal Reserve Bank of San Francisco, UNECE/ILO CPI Experts Conference, NBER CIRW Conference, Bureau of Labor Statistics, and the State Street Annual Research Conference. All remaining errors are ours.

Financial support for this research was provided by MIT Sloan. Publicly available BPP datasets can be downloaded at <http://bpp.mit.edu/datasets/>.

## References

- Aizcorbe, Ana, Carol Corrado, and Mark Doms. 2000. "Constructing Price and Quantity Indexes for High Technology Goods." *Industrial Output Section, Division of Research and Statistics, Board of Governors of the Federal Reserve System*, July 19.
- Aizcorbe, Ana M., Carol A. Corrado, and Mark E. Doms. 2003. "When Do Matched-Model and Hedonic Techniques Yield Similar Measures?" *FRB of San Francisco Working Paper*, no. 2003-14.
- Aparicio, Diego, and Manuel Bertolotto. 2016. "Forecasting Inflation with Online Prices." *Working Paper - MIT*.
- Armknrecht, Paul A., and Donald Weyback. 1989. "Adjustments for Quality Change in the US Consumer Price Index." *Journal of Official Statistics* 5 (2): 107.
- Bertolotto, Manuel, Alberto Cavallo, and Roberto Rigobon. 2014. "Using Online Prices to Anticipate Official CPI Inflation." UTokyo Price Project Working Paper Series 031. University of Tokyo, Graduate School of Economics.
- Bertolotto, Manuel. 2016. "Real Exchange Rate Persistence Using High-Frequency and Micro-Level Data." *Universidad de San Andres Working Paper*.
- Boettcher, Ingolf. 2015. "Automatic Data Collection on the Internet." *Statistics Austria, Working Paper*.
- Breton, Robert, Gareth Clews, Liz Metcalfe, Natasha Milliken, Christopher Payne, Joe Winton, and Ainslie Woods. 2015. "Research Indices Using Web Scraped Data." Office for National Statistics - UK.
- Burstein, Ariel, and Gita Gopinath. 2013. "International Prices and Exchange Rates." National Bureau of Economic Research.
- Campbell, Jeffrey R, and Benjamin Eden. 2014. "Rigid Prices: Evidence from US Scanner Data \*." *International Economic Review* 55 (2): 423–42.
- Cavallo, Alberto. 2013. "Online and Official Price Indexes: Measuring Argentina's Inflation." *Journal of Monetary Economics*, 152–65.
- . 2015a. "Are Online and Offline Prices Similar?" *NBER Working Paper*.
- . 2015b. "Scraped Data and Sticky Prices." *NBER Working Paper*, no. 21490.
- Cavallo, Alberto, Brent Neiman, and Roberto Rigobon. 2014a. "Currency Unions, Product Introductions, and the Real Exchange Rate." *Quarterly Journal of Economics* 129 (2).
- . 2014b. "The Price Impact of Joining a Currency Union: Evidence from Latvia." *IMF Economic Review* Forthcoming.
- Cavallo, Alberto, and Roberto Rigobon. 2011. "The Distribution of the Size of Price Changes." *NBER Working Paper* w16760: 1–40.

- Dhyne, Emmanuel, Luis J Álvarez, Hervé Le Bihan, Giovanni Veronese, Daniel Dias, Johannes Hoffmann, Nicole Jonker, Patrick Lünemann, Fabio Rumler, and Jouko 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): 171–92.
- Edelman, Benjamin. 2012. “Using Internet Data for Economic Research.” *The Journal of Economic Perspectives* 26 (2): 189–206.
- Eichenbaum, Martin, Nir Jaimovich, Sergio Rebelo, and Josephine Smith. 2014. “How Frequent Are Small Price Changes?” *American Economic Journal: Macroeconomics* 6 (2): 137–55.
- Einav, Liran, and Jonathan Levin. 2014. “Economics in the Age of Big Data.” *Science* 346 (6210): 1243089.
- Ellison, Glenn, and Sara Fisher Ellison. 2009. “Search, Obfuscation, and Price Elasticities on the Internet.” *Econometrica* 77 (2): 427–52.
- Feenstra, Robert C., and Matthew D. Shapiro. 2003. “Scanner Data and Price Indexes.” *NBER*, January.
- Goldberg, Pinelopi, and Rebecca Hellerstein. 2011. “How Rigid Are Producer Prices?” *FRB of New York Staff Report*, no. 407.
- Gorodnichenko, Yuriy, Viacheslav Sheremirov, and Oleksandr Talavera. 2014. “Price Setting in Online Markets: Does IT Click?” *NBER Working Paper Series*, no. 20819.
- Griffioen, Robert, Jan de Haan, and Leon Willenborg. 2014. “Collecting Clothing Data from the Internet.” In *Meeting of the Group of Experts on Consumer Price Indices*, 26–28.
- Griliches, Zvi. 1985. “Data and Econometricians—the Uneasy Alliance.” *The American Economic Review*, 196–200.
- Groves, Robert M. 2011. “Three Eras of Survey Research.” *Public Opinion Quarterly* 75 (5): 861–71.
- Horrigan, Michael. 2013. “Big Data: A Perspective from the BLS | Amstat News.”
- ILO, IMF, World Bank, OECD, Eurostat, and United Nations. 2004. *Consumer Price Index Manual: Theory and Practice*.
- Klenow, Peter J., and Benjamin A. Malin. 2010. “Microeconomic Evidence on Price-Setting.” In *Handbook of Monetary Economics*. Vol. 3. Elsevier.
- Krsinich, Frances. 2015. “Price Indexes from Online Data Using the Fixed-Effects Window-Splice (FEWS) Index.” *Statistics New Zealand*, Working Paper,
- Midrigan, Virgiliu. 2011. “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations.” *Econometrica* 79 (4): 1139–80. doi:10.3982/ECTA6735.
- Mikians, Jakub, László Gyarmati, Vijay Erramilli, and Nikolaos Laoutaris. 2012. “Detecting Price and Search Discrimination on the Internet.” In *Proceedings of the 11th ACM Workshop on Hot Topics in Networks*, 79–84. HotNets-XI. New York, NY, USA: ACM.

- Nakamura, Emi, and Jón Steinsson. 2013. "Price Rigidity: Microeconomic Evidence and Macroeconomic Implications." *Annual Review of Economics* 5 (1): 133–63.
- Nygaard, Ragnhild. 2015. "The Use of Online Prices in the Norwegian Consumer Price Index." *Statistics Norway*.
- Rastello, Sandrine, and Ian Katz. 2013. "Argentina Is First Nation Censured by IMF for Economic Data." *Bloomberg*. <http://www.bloomberg.com/news/articles/2013-02-01/argentina-becomes-first-nation-censured-by-imf-on-inflation-data>.
- Silver, Mick, and Saeed Heravi. 2005. "A Failure in the Measurement of Inflation." *Journal of Business & Economic Statistics* 23 (3).
- Stonebraker, Michael, Daniel Bruckner, Ihab F. Ilyas, George Beskaes, Mitch Cherniack, Stanley B. Zdonik, Alexander Pagan, and Shan Xu. 2013. "Data Curation at Scale: The Data Tamer System." In *CIDR*.
- Valentino-DeVries, Jennifer, Jeremy Singer-Vine, and Ashkan Soltani. 2012. "Websites Vary Prices, Deals Based on Users' Information." *Wall Street Journal*.
- Varian, Hal R. 2014. "Big Data: New Tricks for Econometrics." *The Journal of Economic Perspectives* 28: 25.
- Woodford, Michael. 2009. "Information-Constrained State-Dependent Pricing." *Journal of Monetary Economics* 56: S100–124.