



MIT Open Access Articles

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

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation	Cavallo, Alberto, and Roberto Rigobon. "The Billion Prices Project: Using Online Prices for Measurement and Research." <i>Journal of Economic Perspectives</i> 30.2 (2016): 151–178.
As Published	http://dx.doi.org/10.1257/jep.30.2.151
Publisher	American Economic Association
Version	Author's final manuscript
Citable link	http://hdl.handle.net/1721.1/105176
Terms of Use	Creative Commons Attribution-Noncommercial-Share Alike
Detailed Terms	http://creativecommons.org/licenses/by-nc-sa/4.0/

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

Alberto Cavallo and Roberto Rigobon

New data-gathering techniques, often referred to as “Big Data,” have the potential to improve statistics and empirical research in economics. This paper presents one example of how this can be achieved by using the vast number of online prices 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 indexes in multiple countries and to avoid measurement biases that distort evidence of price stickiness and international relative prices.

The basic procedure used in most countries to collect inflation data has remained roughly the same for decades. A large number of people working for national statistical offices visit hundreds of stores on a monthly or bimonthly basis to collect prices for a preselected basket of goods and services. The micro data are then processed and used to construct consumer price indexes and other related indicators. This process is expensive, complex, and often too slow for some users of the data. Infrequent sampling and slow updates to the baskets can complicate

■ *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 at acavallo@mit.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at

<http://dx.doi.org/10.1257/jep.30.2.151>

doi=10.1257/jep.30.2.151

adjustments for quality changes and the introduction of new goods.¹ Groves (2011) further describes other challenges faced by traditional survey-based methods of data collection, including growing levels of nonresponse. Shrinking resources are straining the work of national statistical offices, while recent crises have prompted policymakers 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 webpages, advances in automated “scraping” software now allow anyone to design and implement large-scale data collections on the web. Detailed information can be collected for each good, and new and disappearing products can be quickly detected and accounted for. Online data collection is cheap, fast, and accurate, making it an ideal complement to traditional methods of collecting prices, particularly in categories of goods that are well-represented online.

The first use of online data to construct inflation indexes was motivated by the manipulation of inflation statistics in Argentina from 2007 to 2015. By 2007, it had become apparent that the official level of inflation reported by the national statistical office in Argentina did not reflect the actual changes in prices. Using online data collected every day from the websites of large retailers, Cavallo (2013) showed that while Argentina’s government announced an average annual inflation rate of 8 percent from 2007–2011, the online data suggested it was actually over 20 percent, in line with the estimates of some provincial governments and local economists, and consistent with the results from surveys of household inflation expectations. The online price indexes used in that paper were automatically computed and published on a website every day from March 2008 onwards.² The ability to collect prices from outside the country proved particularly useful in 2011, when Argentina’s government started to impose fines and to pressure local economists to stop collecting data independently. The manipulation of the official price index ended in December 2015 when a new government was elected.

Argentina’s statistical debacle had a positive side effect: it showed us the potential that online prices had for inflation measurement applications. With this idea in mind, we created the Billion Prices Project 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 in fact reached that number of observations in less than two years. 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 United States alone (by comparison, the US Bureau

¹For discussion of some of these measurement topics, see the “Symposium on Measuring the CPI” in the Winter 1998 issue or the “Symposium on the Consumer Price Index” in the Winter 2003 issue of this journal.

²See *InflacionVerdadera* (<http://www.inflacionverdadera.com>), which was created to provide alternative price indexes to the official ones in Argentina. The original website had two price indexes constructed with Argentina’s official National Institute of Statistics and Censuses (INDEC is the abbreviation of the Spanish translation): 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.

of Labor Statistics collects approximately 80,000 prices on a monthly or bimonthly basis). Although gathering this massive amount of prices was cheaper online than with traditional methods, it required funding that could not be sustained through grants. Thus, in 2011 we started a company called PriceStats that now collects the data and produces high-frequency indexes for central banks and financial-sector customers. 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 indexes in 20 countries. Its micro datasets contain information from an even larger number of retailers in over 60 countries, with varying degrees of coverage. The indexes and micro data from PriceStats are available to researchers working with the Billion Prices Project, as we explain later in this paper.

Many of the other attempts to use big data in economics rely on social media or search data to *forecast* the behavior of important economic indicators. Our approach is different because we focus on *measurement*, not on prediction. Our objective is to experiment with these new sources of information to improve the computation of traditional economic indicators, starting with the Consumer Price Index. 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 this paper with a description of the methodology used 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 to sample is still crucial. The goal is to obtain data that is representative of retail transactions, so we focus our data collection efforts on large multichannel retailers such as Walmart that sell both online and offline, instead of using online-only retailers that may have many products but a relatively small share of retail transactions. We also focus on categories of goods that are included in the official consumer price index 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 (including scanner data and official price-index 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 indexes and show how they co-move with consumer price indexes in most countries. We emphasize two characteristics in greater detail. First, online indexes have the ability to approximate hedonically adjusted price indexes in sectors with a large number of goods that come and go with overlapping life-cycles (as for instance, in electronics). Second, online indexes appear able to anticipate movements in the official consumer price index in many countries. This anticipation extends beyond the publication lags, which suggests that online prices often adjust sooner to aggregate shocks.

We then move on to research applications and 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 results in the price-stickiness literature. In particular, we document that online prices exhibit a very different distribution of price changes compared to prices collected for official consumer price indexes and by 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 official micro data, which create a large number of small spurious price changes. Second, using online data to test the “law of one price” (that there should not be large or persistent cross-country differences in the prices of identical goods when translated into a common currency) gives us a more nuanced picture of when and where this law works well. The existing consensus in the literature is that there are large and persistent deviations from the law of one price, with 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 deviations can also be large with online data, we find that the law of one price holds well across countries that use the same currency. We also show that, when goods are identically matched across countries, then relative prices and nominal exchange rates co-move more closely than previously thought. This implies higher pass-through rates and less persistent real-exchange rate dynamics.

Both research examples illustrate how using data collected by others, with different purposes in mind, can distort empirical findings. They also suggest we should not treat big data as simply a collection of large datasets created as a byproduct of something else. One of the greatest opportunities of big data is that anyone can now use new technologies such as web-scraping, mobile phones, satellite imaging, and all kinds of interconnected sensors to build customized datasets designed to fit specific measurement or research needs. We end the paper by describing how the Billion Prices Project data are publicly shared and by discussing why data collection is an important endeavor that macro- and international economists should pursue more often.

Collecting and Processing Online Price Data

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

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

The Selection of Retailers and Data Source

The starting point is to select the retailers and categories of goods to sample. These decisions are driven by our need to get prices that are representative of retail transactions. We therefore focus almost exclusively on large multichannel 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 multichannel retailers still are involved in the majority of all retail sales in most countries. We are also careful when we choose what categories of goods to monitor within each retailer, concentrating on those categories that are part of traditional consumer price index baskets, and avoiding categories that are overrepresented online such as CDs, DVDs, cosmetics, and books.

We make an effort to collect the data directly from each retailer's website, rather than relying on third parties such as marketplaces, price aggregators, and price comparison websites. Data collection from individual retailers is far more challenging, but it maximizes our chances of obtaining prices linked to actual transactions and prevents third-parties from filtering or altering our samples. 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 few sources of data.

Once the data are collected, we clean them, standardize them to fit a common database schema, classify individual products using consumer price index categories, and start computing simple indicators to evaluate their 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 any special characteristics in the data so that we can know whether it is a useful and reliable source of price information.

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). Grocery retailers can sometimes show different prices for the same good depending on the zip code entered by the consumer. 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 about 25 countries, our datasets have information on categories that cover at least 70 percent of the weights in consumer price index baskets.

Data Collection Using Web Scraping Software

After selecting the sources of data, the next step is to collect the information. 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 (for an example, see the discussion in this journal by Edelman 2012). Today, there are many "point-and-click" software solutions that require almost no technical expertise. Users can simply use their mouse to teach the software what

Table 1
Alternative Micro-Price Data Sources

	<i>Online data</i>	<i>Scanner data</i>	<i>CPI data</i>
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	~60	<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

Source: Table 1 from Cavallo (2015).

Notes: The Billion Prices Project (bpp.mit.edu) datasets contain information from over 60 countries with varying degrees of sector coverage. Nielsen US scanner datasets are available at the Kilts Center for Marketing at the University of Chicago. Klenow and Malin (2010) provide stickiness results with Consumer Price Index data sources from 27 papers in 23 countries. See Cavallo (2013) for more details.

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 characters are relatively steady 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 so any errors in the data can be quickly detected and fixed. The robots we construct always collect a product identification number, the name, description, brand, package size, category information, and the 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 online Appendix available with this paper at <http://e-jep.org>.

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 consumer price index data collected offline by national statistical organizations, and scanner data recorded from consumer purchases at point-of-sale terminals by companies such as Nielsen. Detailed descriptions of these other data sources can be found in ILO et al. (2004) and Feenstra and Shapiro (2003).

One of the most obvious advantages of online data is the low cost per observation. While the cost is not trivial, it is far 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 of data collection. It is easier to detect errors in the data when it is collected at such high frequency. This approach also avoids a need to use 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 consumer price index data. Later, we discuss how this big data feature can be used to simplify quality adjustments and other traditional measurement problems.

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. Traditional data collection methods, in contrast, will typically start monitoring new goods only when the goods in the basket disappear from the stores. Knowing the full history of prices for individual goods can help to control for new-good biases, make both implicit and explicit quality adjustments, and study prices at time of product introductions.

Fifth, online data can be collected remotely. 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.

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

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, the prices of most services are still not available on the web, and the number and type of retailers is limited compared to official consumer price index data.

Another disadvantage is that online datasets lack information on quantities sold. Online prices must be combined with weights from official consumer expenditure surveys or other sources for expenditure-weighted applications. Scanner datasets, by contrast, have detailed information on quantities sold, and could potentially be a source of high-frequency expenditure weights in some categories of goods such as groceries.

Are Online Prices Different?

An important concern is whether online prices are different from offline prices; after all, 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 in which prices are varied for strategic purposes: for examples, see Mikians, Gyarmati, Erramilli, and Laoutaris (2012) and Valentino-DeVries, Singer-Vine, and Soltani (2012). In addition, many papers with “online prices” use data from online marketplaces such as Ebay or price-comparison websites such as Google Shopping. As Brynjolfsson and Smith (2000), Ellison and Ellison (2009), and Gorodnichenko, Sheremivov, and Talavera (2014) have shown, these prices seem to change more frequently and in smaller sizes than in Consumer Price Index data. However, the retailers in these datasets are mostly online-only stores participating in a fiercely competitive environment, not really the type of “online data” we use.

To better understand whether online and offline prices for multichannel retailers behave differently, Cavallo (2016) simultaneously collected prices on the websites and physical stores for over 24,000 products in 56 of the largest retailers in 10 countries. This large-scale comparison was possible thanks to the combination of a smartphone app, crowdsourced workers, and web-scraping techniques. More than 370 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 Billion Prices Project servers. We then used the barcodes in the offline data to collect the prices for those exact same goods at the website of the same retailer within a seven-day time window.

This direct comparison between online and offline prices revealed a high degree of similarity in price levels, as well as in 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 similarity was highest in retailers that sell electronics or apparel, and lowest in drugstores and office-supply retailers that also tend to price differently across offline stores. While price changes do not have the exact same timing online and offline, they tend to have similar frequency and average sizes. This suggests that the price spells for individual goods may not be synchronized online and offline, consistent with evidence to be discussed below that online prices may anticipate later price changes. Despite the general similarity between online and offline pricing, 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.

Inflation Measurement

Online prices are increasingly being used in inflation measurement applications. Besides the Billion Prices Project and PriceStats, many national statistical organizations are experimenting with the use of online data, including the US Bureau of Labor Statistics (Horrigan 2013a), the UK Office of National Statistics (Breton et al. 2015), Statistics Netherlands (Griffioen, de Haan, Willenborg 2014), Statistics New Zealand (Krsinich 2015), and Statistics Norway (Nygaard 2015).

In this section, we show that online price indexes can closely approximate the official consumer price index in a number of countries and settings. We then discuss how a large number of overlapping price series in the data can simplify quality adjustments in categories with frequent product turnover, such as electronics. Finally, we show that online price indexes can anticipate changes in the official inflation rate several months in advance.

Methodology for Comparison to Official Consumer Price Indexes

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 indexes that mimic the behavior of official consumer price indexes. The methodology for these daily indexes was based on a combination of online data with standard techniques used in official price indexes, 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 PriceStats to expand the data collection and to compute inflation measures in real time in other sectors and countries. The company is currently publishing daily price indexes in 22 countries with only a three-day lag. In Figures 1 to 4, we plot these online indexes next to the all-item nonseasonally adjusted consumer price index 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 1 illustrates the case of Argentina from 2007 to 2015. Figure 1A compares a price index produced with online data to the official consumer price index.


The fact that the two measures of inflation in Figure 1A diverge so dramatically will not surprise anyone who knows the recent story of statistics in Argentina. In February 2007, the government intervened in the National Institute of Statistics and Census (INDEC) and fired the people responsible for computing the consumer price index. The index quickly stabilized, but many local economists claimed the government was manipulating the data. Household inflation expectations increased dramatically, closely tracking some alternative estimates of inflation produced by local economists and some independent provincial governments, as shown in Cavallo, Cruces, and Perez-Truglia (2016). Suspicions were abundant, but before a measure of inflation based on online prices became available, there was no consistent way to confirm the magnitude of the discrepancy and track its evolution over time.

The manipulation in the official inflation data continued for almost nine years, ending in December 2015 when a new government was elected. During all this time, the monthly inflation rate shown in Figure 1C was consistently higher than the official data reported, with the exception of a few months in 2014 when, in response to a “motion of censure” issued by the IMF in 2013 (Rastello and Katz 2013), the Argentinian government decided to launch a new consumer price index. Unfortunately the change was temporary and the new official index quickly lost all credibility again.

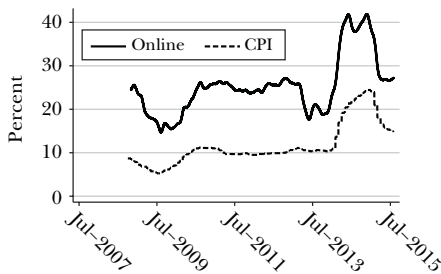
Looking only at the discrepancy in the trend of the price index or the monthly inflation rates, however, misses an important point. The online index tracked the dynamic behavior of the annual inflation rate over time, as shown in Figure 1B. The

Figure 1
Argentina

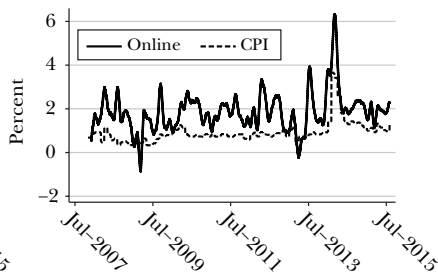
A: Price index



B: Annual inflation rate



C: Monthly inflation rate




Source: Authors using online price index computed by PriceStats and the consumer price index from the national statistical office in Argentina (INDEC).

Notes: The figure compares a price index produced with online data to a comparable official consumer price index (CPI) for the case of Argentina from 2007 to 2015. It also looks at annual and monthly inflation rates using each source of data. Monthly inflation rates for the online index are 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 the online index are computed as the percentage change in the average of the previous 30 days compared to the same average 365 days before. All price indexes are nonseasonally adjusted.

difference was mostly in the *level* of the annual inflation rate, not its movements over time. The online index also quickly reacted to aggregate shocks, such as the massive roadblocks by farmers who protested export tax hikes in 2008. This strongly suggested that online data was capable of capturing the fundamental dynamics of inflation and prompted us to collect data in other countries.³

³It 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 index by simply dividing the online inflation rate by three.

Figure 2
United States



Source: Authors using online price index computed by PriceStats and a comparable consumer price index (CPI) from the US Bureau of Labor Statistics.

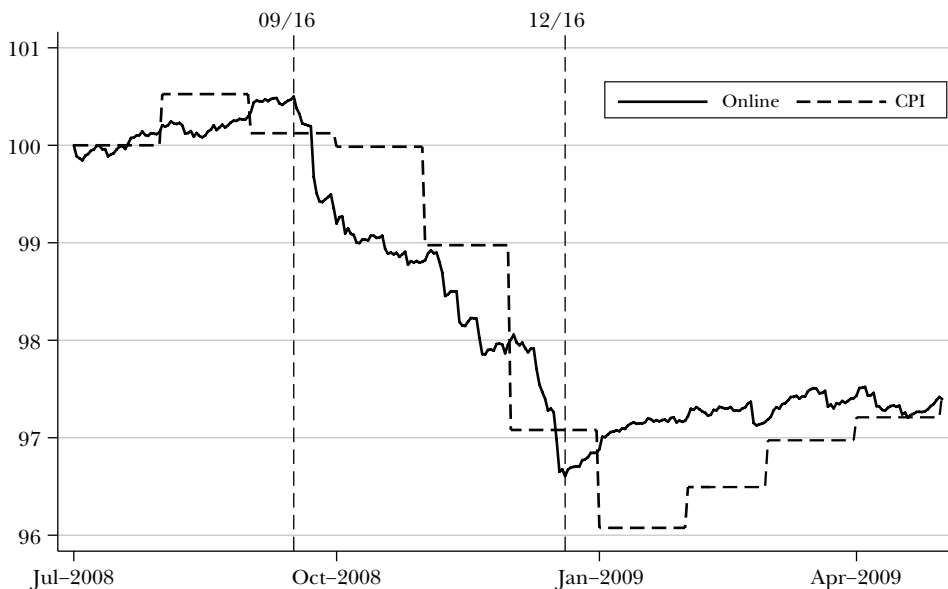
Notes: The figure compares a price index produced with online data to the US CPI for the United States January 2008 to January 2016. The monthly inflation rates for the online index are computed as the percentage change in the average of the previous 30 days compared to the same average a month before. Annual inflation rates for the online index are computed as the percentage change in the average of the previous 30 days compared to the same average 365 days before. All price indexes are non-seasonally adjusted.

The comparison of online and offline indices in other countries is completely different. The daily US index, shown in Figure 2, 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 the official Consumer Price Index for over seven years. Although there are periods where the indexes diverge, the differences are relatively small and temporary. This can also be seen in the monthly and annual inflation rates in Figures 2B and 2C.

The US online index is particularly good at anticipating major changes in inflation trends. Predicting these changes is important for participants in financial

Figure 3

US Consumer Price Index around the Bankruptcy of Lehman Brothers

Source: Authors using online price index computed by PriceStats and the Consumer Price Index from the US Bureau of Labor Statistics.

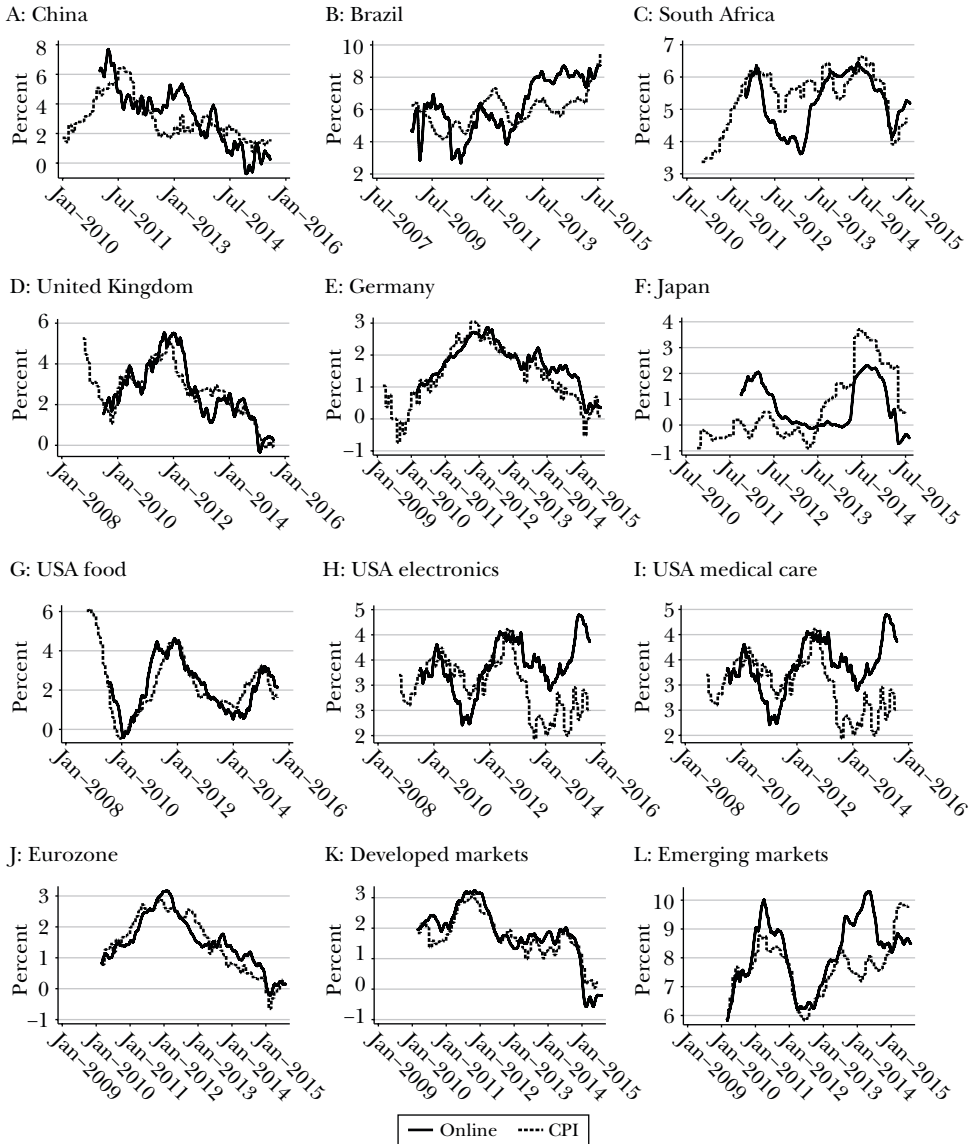
Note: The figure highlights the events around the bankruptcy of Lehman Brothers, the fourth-largest investment bank in the United States, during September 2008.

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 official US Consumer Price Index data was September 16, 2008, the Tuesday after Lehman Brothers filed for bankruptcy. As Figure 3 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 Consumer Price Index for September came out with only a 0.14 percent drop. When the official October Consumer Price Index numbers were published on November 19th, it had fallen another 1.01 percent. In other words, it took more than two months after Lehman's disaster for the official Consumer Price Index numbers to reflect the full impact on price levels. Two months later, on December 16, 2008, the online price index stopped falling and started to increase once again. The Consumer Price Index did not show this change in the trend until the estimates for January were published on February 20, 2009. We measure the degree of anticipation in online data more formally later on.

Figure 4 compares inflation as measured by online prices and by the offline prices in the official consumer price index for a selection of other countries and sectors. The main lesson of the figure is that the correspondence is reasonably close,

Figure 4

Online versus Consumer Price Index (CPI) Annual Inflation Rates



Source: Authors using online price indexes computed by PriceStats and consumer price indexes sourced from the national statistical office in each country.

Notes: Figure 4 compares inflation as measured by online prices and by the offline prices in the official consumer price index for a selection of countries, sectors, and regions. Annual inflation rates for daily online price indexes are computed as the percentage change in the average of the previous 30 days compared to the same average 365 days before. The series are nonseasonally adjusted. Indexes are “all-items” 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 2010 consumption weights in each country and CPIs from official sources.

but some more specific insights are also possible. First, we do not find evidence that China has been systematically holding its official inflation rate below the rate based on online prices, though we can only compare some sectoral indexes because official expenditure weights are not publicly disclosed. Second, the difference between the online price index and the official consumer price index appears to be smaller in developed countries like the UK and Germany, and greater for countries like Brazil or South Africa, where the online sector seems to have more independent patterns. In Japan, we observed significantly more inflation after the March 2011 earthquake and an immediate impact of the sales tax changes in April 2014. While the online index in Japan does not follow its official consumer price index closely, it does seem to anticipate key changes in inflation trends.

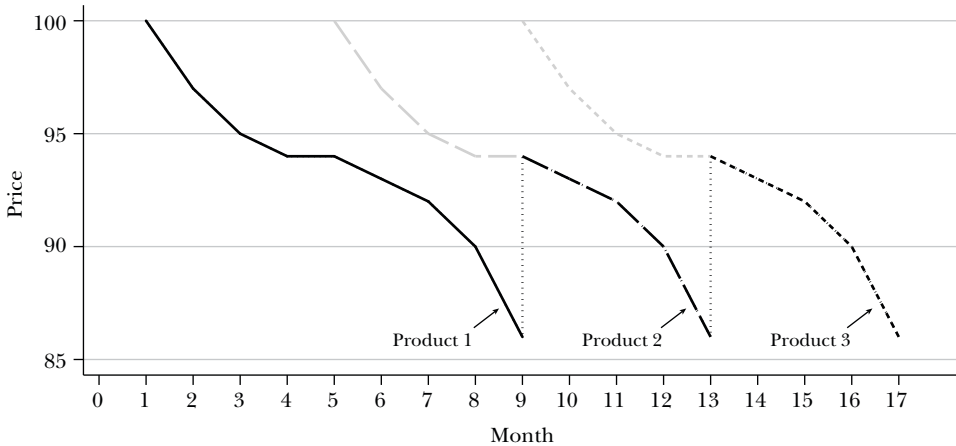
The third row of Figure 4 shows results for a few US sectors. As one might expect, the online data matches the US Consumer Price Index better in sectors such as food and electronics, for which online information is widely available. By contrast, some official inflation patterns seen in the medical care sector are not well-captured by online prices, mostly 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.

It may seem surprising to some readers that indexes based on online data have the ability to mimic official consumer price indexes 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 national statistical organizations, such as hedonic quality adjustments. We believe there are two reasons for this closer-than-expected correspondence. First, as mentioned before, we carefully design and select the data that goes into these indexes to ensure that they are representative. Second, we learned that many sampling characteristics in our data made it *simpler* to deal with some traditional measurement problems. To illustrate this point, we next discuss how online data can simplify quality adjustments by providing a large number of uncensored and overlapping price spells.

Overlapping Quality Adjustments

Quality adjustment poses a 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. National statistical organizations use different methods for quality adjustments, including seeking 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). Online datasets make it easier to deal with quality adjustments because they provide *uncensored* price spells for a large number of models and varieties of each good. With better underlying data, online price indexes can approximate the results of more sophisticated, and often impractical, hedonic-regression methods.

Figure 5
Traditional Consumer Price Index (CPI) Data with Censored Prices, and the Problem of New Products



Source: Authors.

Notes: This Figure shows hypothetical prices series for three goods. Each downward-sloping line reflects the prices of a different good. The shaded portions of these lines illustrate the data points that would be missed with traditional data collection methods.

To build some intuition for why this result holds true, consider a hypothetical example of a series of prices in Figure 5. It 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 electronic products, such as televisions, dishwashers, washing machines, and vacuum cleaners, tend to be introduced at relative high prices and then are discounted gradually over their life-cycle, with clearance sales occurring right before the product disappears from the stores (Silver and Heravi 2005).

With traditional data collection methods, it is too expensive to collect the prices for every good available for sale at each point in time within a sampled retailer. Instead, the data collector focuses on one (or a few) of the most popular models and records its price once per month until it disappears from the store. When a particular model is no longer available, the data collector starts to sample 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 new model are unknown (shown where the line is shaded more lightly on the figure). The problem is to decide how much of the price gap at that point in time is attributable to quality differences. This issue is exacerbated in goods that experience extreme price movements along their life-cycles and may have steep discounts right before disappearing from the shelves.

National statistical organizations have two main ways of dealing with this problem. One 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. While hedonic techniques have become popular in recent decades, the question of what traits should be included in the regression, how they can be measured, and what specification should be used can make hedonic techniques too data intensive and complex to implement in practice.

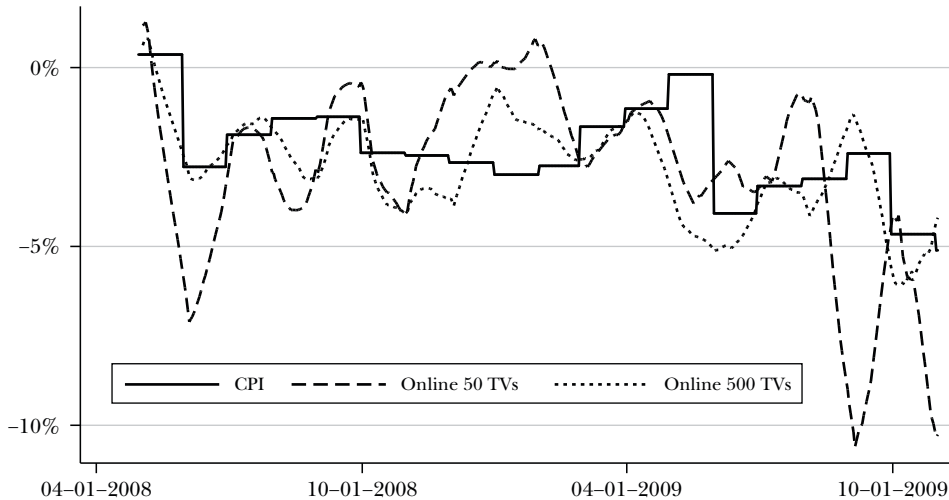
A simpler alternative method is to use “overlapping qualities.” As Armknecht and Weyback (1989) point out, if two goods coexist for some time, their overlapping prices can be used to obtain an estimate of quality change. In practice, this approach tends to assume that the price gap at the time of introduction of the new variety mostly reflects a quality difference. The problem with traditional data, however, is that the price of the new good is not observed at the time of introduction, but much later, when the old good disappears from the stores. This is noted in the *Consumer Price Index Manual* of the International Labor Organization (ILO et al. 2004, p. 27–28): “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. ... [T]he 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 this data problem by providing a large number of uncensored price spells for all models on sale at any point in time. With this type of data, a simple index using overlapping qualities can closely approximate official indexes that use complex hedonic quality-adjustment methods. Similar results were documented earlier in the price-index literature using scanner data. For example, Aizcorbe, Corrado, and Doms (2000, 2003) used scanner prices to demonstrate that, with high-frequency data, matched-model price indexes could yield results that are numerically close to those obtained using hedonic techniques in samples where product characteristics did not change much over time. More generally, the extent to which a simple matched-model price index can capture quality change will depend on several factors (Silver and Heravi 2005). First, both varieties of the product need to have a substantial degree of overlap in their prices. Second, there needs to be a large number of models so that continuing 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 6. It contains three price indexes for televisions in the United States from 2008 and 2009. The solid line shows the official Consumer Price Index for televisions 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 televisions 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 US Bureau of Labor Statistics during this time period. Intuitively, the more

Figure 6

Hedonic Consumer Price Index (CPI) versus Online Index for US Televisions
(monthly inflation rate)



Source: Authors and the US Bureau of Labor Statistics.

Notes: The solid line shows the official Consumer Price Index for televisions 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 television from a large US retailer. The line with short dashes shows an online price index with 500 models from the same source.

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.

This example illustrates one of the *size* advantages of online datasets. We may not need or want to use every single data point available in these data, but being able to extract and use uncensored spells for a large number of models can greatly simplify measurements. Even if the goal is to run a hedonic regression, online data can supply the detailed information needed to make it practical. And with more data, simpler methods can be applied. For example, Krsinich (2015) showed that online data can be used to construct a time-product dummy index that is equivalent to a fully interacted time dummy hedonic index based on all product characteristics.

Anticipation of Future Changes in the Consumer Price Index

As mentioned before, online price indexes can sometimes anticipate changes in official inflation. In this section, we document this pattern more formally and conjecture about some possible explanations.

To document the degree of anticipation, we estimate a simple autoregression equation with the US Consumer Price Index as the dependent variable and our online price index as the exogenous variable, and compute an impulse response

to see how shocks to the online index impact the official price index over time. The regression is expressed in monthly changes: specifically, we use monthly log changes in the Consumer Price Index and monthly log changes of the online index on the last day each month. We include 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.⁴

Figure 7 shows the cumulative impulse response of the Consumer Price Index to a shock in the online index over time, together with the 95 percent confidence intervals. In the United States, it takes several months for the Consumer Price Index to fully incorporate the shock to the online price inflation. At the sector level, the impact is quickest in fuel (transportation) and slowest in food and electronics (see Appendix for details). 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 other countries.

Possible reasons for why online prices can anticipate shocks in the consumer price index include delays embedded in the methodology used for the official data, differences in mixture of stores sampled, and faster adjustment of online prices in some sectors or retailers. Understanding what drives the anticipation is still an open question for future research, but the patterns in Figure 7 suggest that online data can be a useful addition to inflation forecasting models. This is explored by Aparicio and Bertolotto (2016), who show that out-of-sample inflation forecasts using online data can outperform a large number of alternative forecasting models in the US and UK economies.

Lessons for Macro and International Research

In this section, we use questions of price-stickiness and real exchange rate behaviors to illustrate how online data can change empirical results in macro and international research. Our main objective is to show how online datasets constructed to fit specific research needs can help mitigate biases and other empirical challenges that are so frequent in traditional datasets collected for other purposes.

⁴For each month t , the specification is as follows:

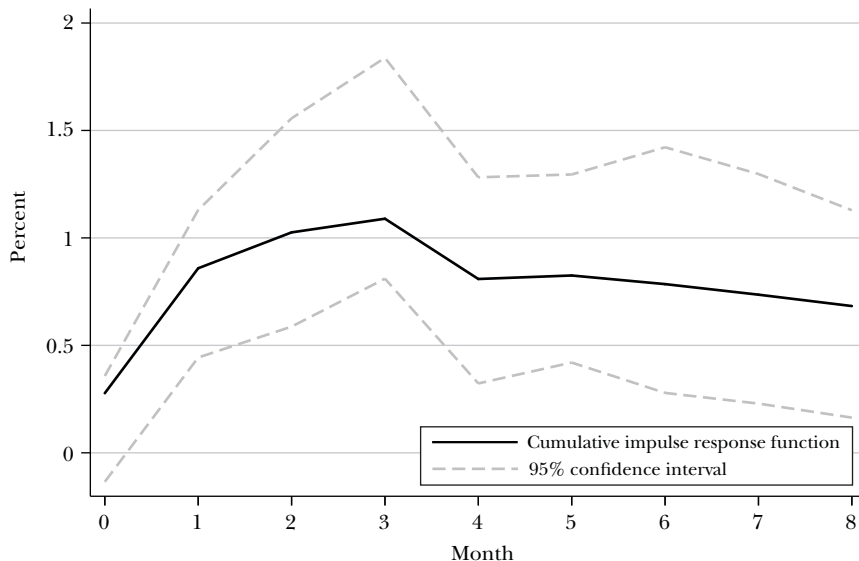
$$\Delta \ln(CPI_t) = \alpha + \beta \Delta \ln(Online_t) + \sum_{i \in [1,6]} \alpha_i \Delta \ln(CPI_{t-i}) + \sum_{i \in [1,6]} \beta_i \Delta \ln(Online_{t-i})$$

The autoregressive distributed lag (ADL) specification is equivalent to a vector auto regression (VAR) with an exclusion restriction. 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 indexes 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. The results are robust to the elimination of the contemporaneous effect of the online price index from the equation.

Figure 7

Cumulative Impulse Response of the US Consumer Price Index (CPI) to an Online Price Index Shock

(response to a 1% shock in the online index)



Source: Authors using online data computed by PriceStates and US Consumer Price Index.

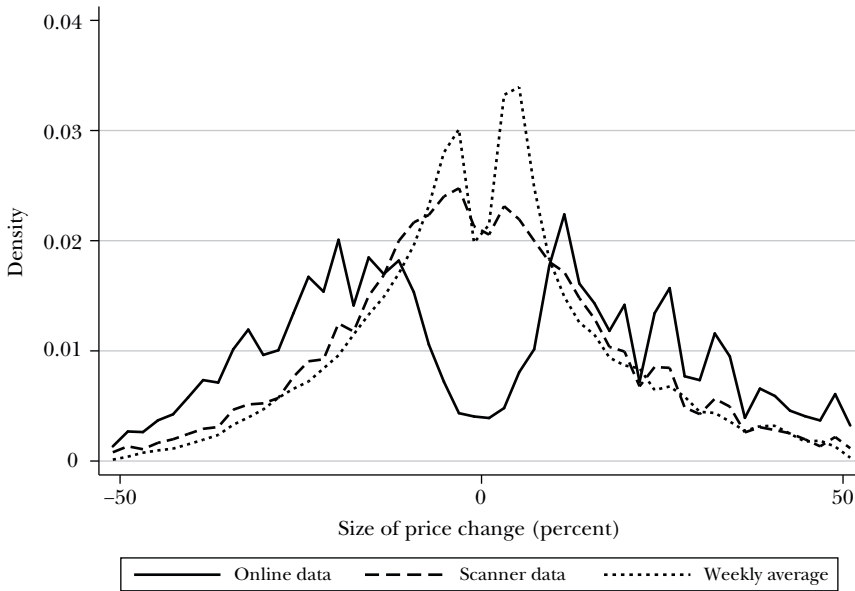
Notes: The Consumer Price Index is a US city average, all items non-seasonally adjusted from the Bureau of Labor Statistics. Data from July 2008 to January 2015.

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 microfoundations (for an example in this journal, see Dhyne et al. 2006; 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 consumer price index data and scanner datasets in several countries. Over time, the literature has settled on a set of stylized facts, summarized by Klenow and Malin (2010). In Cavallo and Rigobon (2011) and Cavallo (2015), we use online data to argue that the sampling characteristics of official consumer price index 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 scanner or consumer price index data found bell-shaped (unimodal) distributions centered around zero percent, with a significant share of small price changes, which seemed inconsistent with standard menu-cost models that predict periods

Figure 8

The Distribution of the Size of Price Changes in the United States

Source: Cavallo (2015).

Notes: Online and scanner data was collected from the same retailer, zip code, and time period. Weekly averages were computed using the daily online data. Nielsen Scanner Data provided by the Kilts Marketing Center at Chicago Booth.

of unchanging prices followed by relatively large changes (a bi-modal distribution centered around zero). This finding motivated a surge in papers trying to adapt sticky-price models to account for this fact (for example, Woodford 2009; Midrigan 2011).

However, the shape of the distribution of price changes is greatly affected by the sampling characteristics of the data. This can be seen in Figure 8, where we show a distribution of prices changes for both online and scanner data obtained from exactly the same US retailer, zip code, and time period. While the raw prices that generate these distributions are in principle the same, the results are strikingly different. The online data distribution is strongly bimodal, with very few price changes close to zero percent. There is a simple explanation for the difference. Scanner data are reported as weekly averages. As noted by Campbell and Eden (2014), this can create a large number of spurious small changes. For example, in a three-week period with a single price change in the middle of the second week, taking weekly averages would yield two small price changes: one from the first week to the second, and another from the second week to the third. These spurious changes can be seen explicitly in Figure 8, where we approximate the shape of the scanner data distribution by simply taking weekly averages of the raw online data.

Something similar happens with Consumer Price Index data, although the source of measurement bias is different in nature, as discussed in Cavallo (2015).

In particular, micro data from the Consumer Price Index will often contain imputed prices for temporarily missing items, which is a sensible thing to do when measuring inflation. This imputation is often done with the average price change of related goods, resulting in an artificial pattern of many small changes. If these imputations are not identified or removed when generating the distributions, the result is a unimodal-shaped distribution similar to those found in the literature. Furthermore, other forms of measurement biases can have a similar impact. For example, Eichenbaum, Jaimovich, Rebelo, and Smith (2014) use Consumer Price Index 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, also create a large number of spurious small changes.

Controlling for measurement bias is important, but to better understand price stickiness and its determinants, the literature also needs data with similar characteristics from multiple countries and economic settings. This is very hard with traditional data sources. For example, to obtain frequency estimates in 24 countries, Klenow and Malin (2010) had to source them from 27 different papers, each with its own particular data and methodologies. Appearing in this journal, Dhyne et al. (2006) is one of the few papers with data from multiple countries, thanks to the coordination provided by the European Inflation Persistence Network. But even in this case, each European national statistical organization was unwilling to share its 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.

Instead, online prices have the potential to provide datasets with identical sampling characteristics in a large number of countries. At the Billion Prices Project we are currently working to standardize stickiness statistics in all our data and to be able to produce them on an ongoing basis. The goal is not only to share with other researchers a range of indicators that can be used to study price stickiness, but also provide policymakers with more up-to-date information about its behavior over time.

International Prices and the Law of One Price

The global nature of online data also makes it appealing for research in international economics. In particular, the relation between relative prices and exchange rates is a classic question in international economics. A basic hypothesis is the “law of one price,” 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. When considering a group of many traded goods, the law of one price implies that exchange rates and relative prices will adjust to maintain stable purchasing power parities (“PPP”). Modest deviations from PPP are not surprising in a world with transport costs and other barriers to arbitrage. However, a huge literature documents 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 price 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 very low exchange “pass through” rates.⁵

The literature concerning the law of one price and PPP is hampered by the formidable difficulties in obtaining prices for a large number of identical goods sold simultaneously in a large number of countries, as discussed by Taylor (2001). In practice, researchers are forced to settle on having prices for identical goods from two countries (typically the US and Canada) or using price indexes 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 an index published by *The Economist* magazine based on the prices of McDonald’s Big Mac sandwiches, provide information on many countries but are limited to a single good. The World Bank’s International Comparison Program makes a worldwide effort to collect price data and to estimate PPP-adjusted GDPs in dozens of countries. But carrying out this task with traditional methods of collecting prices and adjusting for quality is so daunting that it can only be done every five years or so, severely limiting its use for research on real-exchange rate levels and dynamics.

In principle, online prices can be obtained in high frequency, for a large number of goods, in dozens of countries. The main challenge is not in the raw data collection, but rather in matching identical products across countries, as product identification codes in the data tend to be specific to the good, country, and retailer where the product is sold.

In Cavallo, Neiman, and Rigobon (2014) we addressed the matching problem by using prices collected from global retailers such as Apple, IKEA, Zara, and H&M, who sell identical goods with the same identifying information 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: for example, countries within the euro area, or countries that use the US dollar such as El Salvador and Ecuador. What really seems to matter for these global retailers is simply whether prices have to be shown to customers in the same currency, not whether countries are physically close, in a trade union, or even strongly pegging their currencies. In Cavallo, Neiman, and Rigobon (2015), 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 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. Conversely, factors that were traditionally thought to be

⁵See Rogoff (1996) for a description of the “PPP puzzle” and Taylor and Taylor (2004) for a review of the PPP literature. Burstein and Gopinath (2013) provide a review of the empirical literature on relative prices and exchange rates, and a discussion of some theoretical advances, including accounting for nontradeables or tradeables that are only locally consumed, variable markups, and pricing-to-market.

important—such as physical distance, political and tax territories, language, and culture—do not seem to matter as much. Furthermore, these patterns 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 macroeconomic models.

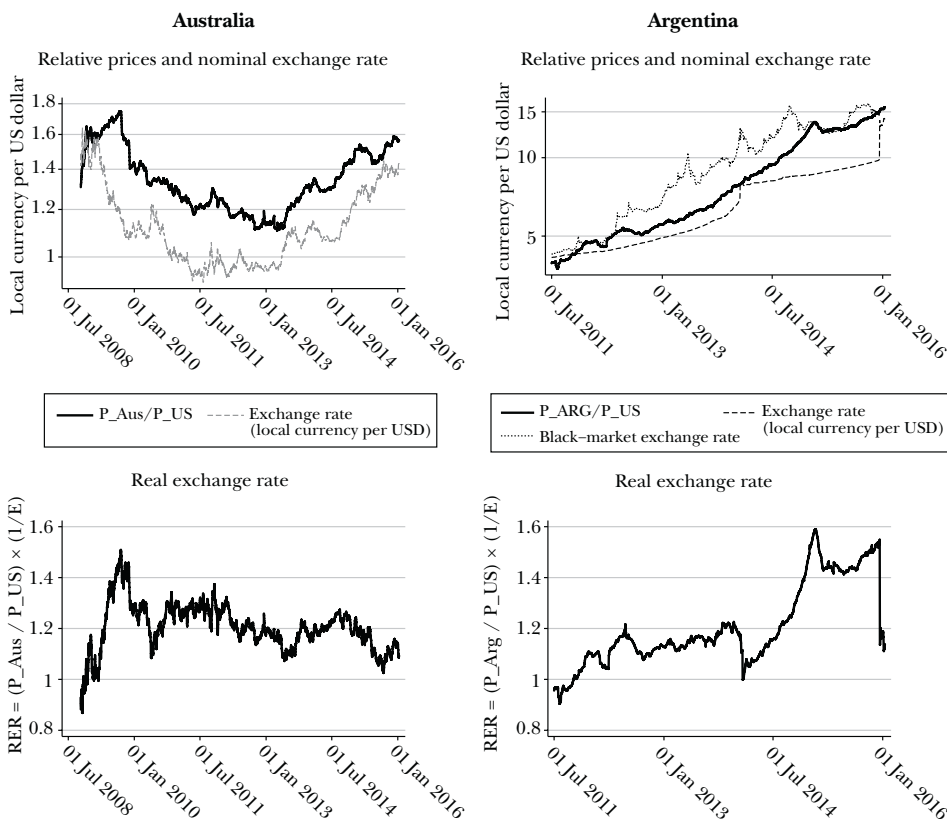
Ideally though, for some applications we need to have both global and local retailers. So since 2014, PriceStats has been expanding the product matching to include local retailers as well, classifying over 30,000 individual goods into 300 product categories. The challenge is to classify a large set of heterogeneous individual products (with varying package sizes, flavors, and retailers) into narrowly defined product categories such as “Basmati White Rice, 1kg” or “LG Basic Blu-Ray Player, 1 unit.” This is achieved by using supervised machine learning (specifically a “Naive Bayes” classifier) that trains on language-specific, hand-categorized items. The process is described in detail in Bertolotto (2016). The output resembles a collection of hundreds of “Big Mac”-type indexes for different kinds of goods.

These matched indexes can be used to study real-exchange-rate levels and dynamics, as in Cavallo and Neiman (2016). To illustrate this, Figure 9 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 United States (examples for other countries are provided in the online Appendix available with this paper at <http://e-jep.org>). The top panel shows the relative prices (in local currencies) and the nominal exchange rate (defined as local currency per US dollar). For the case of Argentina, we also plot the black-market exchange rate. The bottom panel shows the real exchange rate constructed from the other two variables (as the ratio between the relative prices and the nominal exchange rate). This is simply the relative cost of the basket when expressed in the same currency.

A common finding, present also in other countries, is that 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. In Argentina, the steady increase in relative prices was matched by the overall trend of depreciation in the currency, which is gradual in the black-market and lumpy on the official exchange rate. There are long periods where prices kept rising and the official exchange rate was held fixed by the government, causing “deviations” in the real-exchange rate, but there were sudden adjustments in the two occasions when the country devalued its currency, in January 2014 and December 2015.

This co-movement between relative prices and exchange rates implies high rates of pass-through, which can go in both directions. 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, retail price movements tend to precede nominal exchange rate adjustments.

Figure 9
Relative Prices and Exchange Rates



Source: Authors.

Notes: The top panel shows the ratio of relative prices (in local currencies, P/P_{US}) and the nominal exchange rate (E , defined as local currency per US dollar). The bottom panel shows the real exchange rate computed as $(P/P_{US}) \times (1/E)$. It is the relative cost of the basket in each country relative to the US, when expressed in the same currency. Real exchange rates and relative price series are computed by PriceStats at the product level and aggregated using a Fisher index with official expenditure weights for food, fuel, and electronics.

Another unique feature of online data is that they provide information on relative price *levels*, which are not available when using consumer price indexes. For example, the real exchange rates in Figure 9 show that the basket tends to be 20 percent more expensive in Australia relative to the United States. In Argentina, the cost is about 10 percent higher than in the United States when the currency is allowed to float. Recognizing these patterns is useful for estimating the degree of currency misalignment at different points in time, particularly in countries with managed exchange rates.

For example, in December 2015 the new government of Argentina wanted to remove all foreign-exchange market restrictions. It was unclear what the

free-market exchange rate would be, and what effect it would have on tradable prices. The nominal exchange rate implied by purchasing power parity was 14.3 pesos per dollar, suggesting that the official rate of 9.6 pesos per dollar was greatly overvalued while the black market rate of 15 pesos per dollar was slightly undervalued. When the market was freed, the new exchange rate quickly settled around 14 pesos per dollar, closely matching the implied PPP exchange rate (the ratio of relative prices). This can be seen in the jump of the official exchange rate in the top right panel of Figure 9.

While we do not expect these metrics to help predict exchange rates so closely in every country and situation, they can provide better measures of the amount of deviation of real-exchange rates from “normal” levels at a given point in time.

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

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). Most of the micro data and indexes 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 pretreatment, so other researchers can apply their own methods. We will upgrade the shared data periodically, both increasing the number of databases and retailers and also expanding the time series.

The US and Argentina inflation indexes used in this paper are published with a 30-day lag on the Billion Prices Project website, while the PPP exchange rate information discussed in the previous section are currently published with a one-year lag on the PriceStats website. The raw micro data collected by PriceStats are not publicly available but can be shared with academic researchers who collaborate with the Billion Prices Project and sign a data-access agreement.

Final Remarks

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

... [W]e 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 macro and international economists with opportunities to stop treating the data as “given” and get personally involved with data collection. We can now build datasets customized to fit specific measurement and research needs. This will help mitigate issues in empirical research such as sample selection, endogeneity, omitted variables, and error-in-variables, which are so frequent in traditional datasets.

The Billion Prices Project is just one example of the use of big data in economics.⁶ Although online price data are the focus of this paper, we hope to have convinced other 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 part of our daily lives.

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

■ *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 indexes and micro data. We also thank the large number of students who 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, CEMLA, NYU, 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 Billion Prices Project datasets can be downloaded at <http://bpp.mit.edu/datasets/>.*

⁶Einav and Levin (2014) provide a more general discussion of this topic, including new granular data sources, computational techniques such as machine learning, and the role of theory in analyzing large, unstructured datasets. In this journal, Varian (2014) describes in detail some new big data techniques that are useful to analyze large datasets.

References

- Aizcorbe, Ana M., Carol A. Corrado, and Mark E. Doms.** 2000. "Constructing Price and Quantity Indexes for High Technology Goods." Paper prepared for the CRIW Workshop on Price Measurement at the NBER Summer Institute, July 31–August 1, 2000.
- Aizcorbe, Ana M., Carol A. Corrado, and Mark E. Doms.** 2003. "When Do Matched-Model and Hedonic Techniques Yield Similar Measures?" Federal Reserve Bank of San Francisco Working Paper, no. 2003-14.
- Aparicio, Diego, and Manuel Bertolotto.** 2016. "Forecasting Inflation with Online Prices." Unpublished paper, MIT.
- Armkecht, Paul A., and Donald Weyback.** 1989. "Adjustments for Quality Change in the US Consumer Price Index." *Journal of Official Statistics* 5(2): 107–23.
- Bertolotto, Manuel.** 2016. "Real Exchange Rates Using Online Data." Working Paper, Universidad de San Andrés.
- Breton, Robert, Gareth Clews, Liz Metcalfe, Natasha Milliken, Christopher Payne, Joe Winton, and Ainslie Woods.** 2015. "Research Indexes Using Web Scraped Data." Office for National Statistics, UK.
- Brynjolfsson, Erik, and Michael D. Smith.** 2000. "Frictionless Commerce? A Comparison of Internet and Conventional Retailers." *Management Science* 46(4): 563–85.
- Burstein, Ariel, and Gita Gopinath.** 2013. "International Prices and Exchange Rates." NBER Working Paper 18829.
- Campbell, Jeffrey R., and Benjamin Eden.** 2014. "Rigid Prices: Evidence from U.S. 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* 60(2): 152–65.
- Cavallo, Alberto.** 2015. "Scraped Data and Sticky Prices." NBER Working Paper 21490.
- Cavallo, Alberto.** 2016. "Are Online and Offline Prices Similar?" 2016 NBER Working Paper 22142.
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia.** 2016. "Learning from Potentially-Biased Statistics: Household Inflation Perceptions and Expectations in Argentina." NBER Working Paper 22103.
- Cavallo, Alberto, and Brent Neiman.** 2016. "Real Exchange Rate Behavior: Evidence from Online Retailers in Nine Countries." Unpublished Paper.
- Cavallo, Alberto, Brent Neiman, and Roberto Rigobon.** 2014. "Currency Unions, Product Introductions, and the Real Exchange Rate." *Quarterly Journal of Economics* 129(2): 529–95.
- Cavallo, Alberto, Brent Neiman, and Roberto Rigobon.** 2015. "The Price Impact of Joining a Currency Union: Evidence from Latvia." *IMF Economic Review* 63(2): 281–97.
- Cavallo, Alberto, and Roberto Rigobon.** 2011. "The Distribution of the Size of Price Changes." NBER Working Paper 16760.
- Dhyne, Emmanuel, Luis J. Álvarez, Hervé Le Bihan, Giovanni Veronese, Daniel Dias, Johannes Hoffmann, Nicole Jonker, Patrick Linnemann, 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." *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).
- 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, eds.** 2003. *Scanner Data and Price Indexes*. NBER.
- Gorodnichenko, Yuriy, Viacheslav Sheremirov, and Oleksandr Talavera.** 2014. "Price Setting in Online Markets: Does IT Click?" NBER Working Paper 20819.
- Griffioen, Robert, Jan de Haan, and Leon Willenborg.** 2014. "Collecting Clothing Data from the Internet." Proceedings of *Meeting of the Group of Experts on Consumer Price Indexes*, May 26–28. UNECE. Available at: <http://www.unece.org/stats/documents/2014.05.cpi.html#/>.
- Griliches, Zvi.** 1985. "Data and Econometricians—The Uneasy Alliance." *American Economic Review* 75(2): 196–200.
- Groves, Robert M.** 2011. "Three Eras of Survey Research." *Public Opinion Quarterly* 75(5): 861–71.
- Horrigan, Michael W.** 2013. "Big Data: A Perspective from the BLS." *Amstat News*, January 1. <http://magazine.amstat.org/blog/2013/01/01/sci-policy-jan2013/>.
- ILO, IMF, OECD, Eurostat, UNECE, and the World Bank.** 2004. *Consumer Price Index Manual: Theory and Practice*.

- Klenow, Peter J., and Benjamin A. Malin.** 2010. "Microeconomic Evidence on Price-Setting." Chap. 6 in *Handbook of Monetary Economics*, vol. 3, edited by Benjamin M. Friedman and Michael Woodford. Elsevier.
- Krsinich, Frances.** 2015. "Price Indexes from Online Data Using the Fixed-Effects Window-Splice (FEWS) Index." Paper presented at the Ottawa Group, Tokyo, Japan, May 20–22, 2015. http://www.stat.go.jp/english/info/meetings/og2015/pdf/t1s2p7_pap.pdf.
- Midrigan, Virgiliu.** 2011. "Menu Costs, Multi-product Firms, and Aggregate Fluctuations." *Econometrica* 79(4): 1139–80.
- 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. ACM. <http://doi.acm.org/10.1145/2390231.2390245>.
- Nakamura, Emi, and Jón Steinsson.** 2013. "Price Rigidity: Microeconomic Evidence and Macroeconomic Implications." *Annual Review of Economics* 5: 133–63.
- Nygaard, Ragnhild.** 2015. "The Use of Online Prices in the Norwegian Consumer Price Index." Paper prepared for the meeting of the Ottawa Group, Tokyo, Japan, May 20–22, 2015. http://www.stat.go.jp/english/info/meetings/og2015/pdf/t1s2p5_pap.pdf.
- Rastello, Sandrine, and Ian Katz.** 2013. "Argentina Is First Nation Censured by IMF for Economic Data." Bloomberg, February 2. <http://www.bloomberg.com/news/articles/2013-02-01/argentina-becomes-first-nation-censured-by-imf-on-inflation-data>.
- Rogoff, Kenneth.** 1996. "The Purchasing Power Parity Puzzle." *Journal of Economic Literature* 34(2): 647–68.
- Silver, Mick, and Saeed Heravi.** 2005. "A Failure in the Measurement of Inflation: Results from a Hedonic and Matched Experiment Using Scanner Data." *Journal of Business & Economic Statistics* 23(3): 269–81.
- Taylor, Alan M.** 2001. "Potential Pitfalls for the Purchasing-Power-Parity Puzzle? Sampling and Specification Biases in Mean-Reversion Tests of the Law of One Price." *Econometrica* 69(2): 473–98.
- Taylor, Alan M., and Mark P Taylor.** 2004. "The Purchasing Power Parity Debate." *Journal of Economic Perspectives* 18(4): 135–58.
- Valentino-DeVries, Jennifer, Jeremy Singer-Vine, and Ashkan Soltani.** 2012. "Websites Vary Prices, Deals Based on Users' Information." *Wall Street Journal*, December 24. <http://www.wsj.com/articles/SB1000142412788732377204578189391813881534>.
- Varian, Hal R.** 2014. "Big Data: New Tricks for Econometrics." *Journal of Economic Perspectives* 28(2): 3–28.
- Woodford, Michael.** 2009. "Information-Constrained State-Dependent Pricing." *Journal of Monetary Economics* 56(Supplement): S100–S124.