

NBER WORKING PAPER SERIES

ARE ONLINE AND OFFLINE PRICES SIMILAR? EVIDENCE FROM LARGE MULTI-CHANNEL RETAILERS

Alberto F. Cavallo

Working Paper 22142

<http://www.nber.org/papers/w22142>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

March 2016

Financial support for this research was provided by the JFRAP at MIT Sloan and the NBER's Economics of Digitization and Copyright Initiative. I thank Maria Fazzolari for her outstanding work implementing and coordinating all the offline data collection. I also thank the MIT and Wellesley University Students that collected some of the offline data: Vivian Xi, Maurizio Boano, Sibio Wang, Descartes Holland, Sabrina Lui, Suh Yoon, Holly Zhu, Sean Bingham, Elizabeth Krauthamer, Jeffrey Zhang, William Rodriguez, Wenxin Zhang, Jake Amereno, Ivy Jiang, Diya Mo, Qi Jin, Riley Quinn, Do Yeon Park, Jung Hyun Choi, Xiaoxi Wang, Aaroshi Sahgal, Isaiah Udotong, Giulio Capolino, Tanya Bakshi, Allison Davanzo, Karen Pulido, and Bailey Tregoning. I greatly benefited from the comments and suggestions of seminar participants at the Ottawa Group Meeting in Tokyo and MIT Sloan. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w22142.ack>

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Alberto F. Cavallo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Are Online and Offline Prices Similar? Evidence from Large Multi-Channel Retailers
Alberto F. Cavallo
NBER Working Paper No. 22142
March 2016
JEL No. E31,F4,L1

ABSTRACT

Online prices are increasingly being used for a variety of inflation measurement and research applications, yet little is known about their relation to prices collected offline, where most retail transactions take place. This paper presents the results of the first large-scale comparison of online and offline prices simultaneously collected from the websites and physical stores of 56 large multi-channel retailers in 10 countries. I find that price levels are identical about 72% of the time for the products sold in both locations, with significant heterogeneity across countries, sectors, and retailers. The similarity is highest in electronics and clothing and lowest for drugstores and office-supply retailers. There is no evidence of prices varying with the location of the ip address or persistent browsing habits. Price changes are un-synchronized but have similar frequencies and average sizes. These results have implications for National Statistical Offices and researchers using online data, as well as those interested in the effect of the internet on retail prices in different countries and sectors.

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

A Data Appendix and Replication Materials is available at <http://www.mit.edu/~afc/>

1 Introduction

Online prices are increasingly being used in measurement and research applications. In 2008 we started the Billion Prices Project at MIT (BPP) with Roberto Rigobon to experiment with online price indexes in the US and other countries.¹ In recent years many National Statistical Offices (NSOs) have started to consider the use of online data in official consumer price indexes, as discussed in Horrigan (2013), Breton et al. (2015), Krsinich (2015), Nygaard (2015), Boettcher (2015), and Griffioen et al. (2014). In the context of academic research, online prices are increasingly being applied to study pricing behaviors, measure price stickiness, study international relative prices and real exchange rate dynamics.²

Despite their growing appeal, an open fundamental question about online prices is whether they are similar to the prices that can be collected offline in physical stores. The question is important because relatively few retail transactions take place online. For example, in the US online purchases are still less than 10% of all retail transactions.³

This paper provides the first large-scale comparison of online and offline prices in large multi-channel retailers designed to answer this question. Using a combination of crowdsourcing platforms, a mobile phone app, and web scraping methods, I simultaneously collected prices in both the online and offline stores of over 50 of the largest retailers in 10 countries: Argentina, Australia, Brazil, Canada, China, Germany, Japan, South Africa, UK, and the United States. These prices are used to compare price levels, the behavior of price changes, and the selection of products available for sale in the offline and online stores. I document country, sector, and retailer heterogeneity, and test whether online prices vary with ip-address locations or persistent browsing habits. The results have implications for NSOs and researchers using online data, as well as those interested in the effect of the Internet on retail prices in different countries and sectors.

¹See Cavallo and Rigobon (2016) for details. The online price indexes are computed by a private company called PriceStats, partially owned by the author. The retailers selected and sampled in this paper were chosen independently and are not necessarily the same as those sampled by PriceStats. All the data used in this paper, including the online prices, was collected independently by the BPP.

²Recent examples include Cavallo (2015), Cavallo et al. (2014), Gorodnichenko et al. (2014), and Gorodnichenko and Talavera (2014).

³See Euromonitor (2014).

The data collection effort is unprecedented in scope and size, and was carried out as part of the Billion Prices Project at MIT.⁴ I first selected the retailers to be sampled by focusing on the top 20 companies by market shares in each country that sold both online and offline (“multi-channel”), and had product ids or barcodes that could be matched across samples.⁵ Next, I used crowdsourcing platforms such as Amazon Mechanical Turk, Elance, and UpWork to hire over 370 workers to collect the offline prices. Each worker was assigned a simple task: to scan the barcodes and collect prices for a random set of 10-50 products in any physical store of a given retailer. In some cases they had to return to the same store multiple times to scan the same set of products. All workers had to use a special BPP app for android phones developed specifically to simplify and standardize the data collection process. Using the app the worker could quickly scan the product’s barcode, enter its price, take photos of the price tags, and send all the information via email to the BPP servers, where the data was cleaned and processed. Finally, a scraping software used the barcode numbers in the offline data to look for the same product in the website of each retailer, and collected the online price within a period of seven days. The matched dataset contains prices for more than 24 thousand products and 38 thousand observations sampled between December 2014 and March 2016. It is available for download at bpp.mit.edu.

The main finding in this paper is that online and offline price levels are identical about 72% of the time, with significant heterogeneity at the country, sector, and retailer level. These percentages range from 42% in Brazil to 91% in Canada and the UK. The US is close to the average, with 69%. At the sector-level, drugstores and office-product retailers have the lowest share of identical prices, with 38% and 25% respectively, while in electronics and clothing these numbers rise to 83% and 92% respectively. When there is a price difference, the online markup tends to be small, with a magnitude of just -4% in the full sample. If we pool all observations, including those with identical prices, online-offline price difference is only -1% on average.

The reasons for some online-offline price level differences vary across retailers. In

⁴See Cavallo and Rigobon (2016)

⁵Market shares reported by Euromonitor International.

general, sale prices tend to create many price level differences (only 36% of sale prices are identical across samples) but they have a small impact in the aggregate results because the number of sales is relatively small (only about 11% of the total matched dataset).

Furthermore, price *changes* have different timings, but their frequency and sizes are similar in the online and offline samples. This is consistent with un-synchronized time series for individual goods and may help explain why in Cavallo and Rigobon (2016) we find that online price indices can anticipate changes in official data.

On average, about 76% of the products sampled offline were also found online by either using the automated scraping matching or by manually searching for the product description on the website. There is no evidence that retailers try to obfuscate price comparison by changing the products' id numbers. In fact, the online-offline price comparison for goods that can be automatically matched are similar to those that had to be manually matched using the product descriptions. There is also no evidence in multi-channel retailers of "dynamic pricing" strategies that can potentially cause online-offline differences. At least in the US, prices do not change with the location of the ip-address of the computer connecting to the website or when the scraping robot repeatedly browses the same webpage of a particular good for a prolonged period of time.

Nevertheless, heterogeneity in pricing behaviors is significant across retailers. Three main types of firms stand out: those with mostly identical online and offline prices, those with stable online markups (either positive or negative), and those with different prices that are not consistently higher or lower online. Some of these patterns appear to be sector-level behaviors, and in other cases are common for most retailers within a country.

For research economists using online data, these results provide evidence that most large multi-channel retailers price similarly online and offline. While there are both advantages and disadvantages of using online data, as I discuss in Cavallo (2015), the ability to collect a massive amount of prices so cheaply provides unprecedented opportunities to economic research. Retailer heterogeneity, however, implies that researchers using relatively few sources of data should be cautious to understand particular pricing patterns and control for any sampling biases.

For National Statistical Offices (NSOs), these results imply that the web can be used as an alternative data-collection technology to obtain the same prices found in offline stores. Prices collected through the web are very similar to those that can be obtained at a much higher cost by physically walking to a store. However, specific online-offline validation tests are sensible given the high degree of heterogeneity in retailer behaviors. The BPP app and methodology developed in this paper are publicly available at bpp.mit.edu to be used for these purposes.

Lastly, my findings have implications for people interested in the effects of the Internet on retail prices and their dynamics. The fact that online prices are always the same, and also similar to offline prices collected from many different zip codes, implies there is little *within-retailer* price dispersion. In practice, most retailers now seem to have a single price for each product, regardless of the location of the buyer. This suggests that while the web may not have reduced price dispersion across different retailers, it may have created incentives for firms to price identically across their own stores.

This paper is related to an IO literature that studies the behavior of online prices. An important difference is that nearly all papers in that literature focus exclusively on online-only retailers⁶. A rare exception is Brynjolfsson and Smith (2000), who compared the behavior of manually-collected prices for CDs and books in online-only and conventional retailers. They report that online prices were 9-16% lower, with smaller price changes and significant price dispersion, but their sample included both online-only and multi-channel retailers (“hybrids” in their notations). The authors note that “findings would be strengthened if we excluded hybrid retailers from our comparisons of price levels”, which implies that online and offline prices for multi-channel retailers were closer together.

Other papers in the IO and pricing literatures have used online data from “shopbots”, or price comparison websites, to study the pricing dynamics of online retailers. Examples include Brynjolfsson et al. (2003), Ellison and Ellison (2009), Lnnemann and Wintr (2011), Gorodnichenko et al. (2014), and Gorodnichenko and Talavera (2014). Although these papers do not compare prices with offline data, their results suggest that online prices

⁶See Brynjolfsson et al. (2003), Ellison and Ellison (2009), and Quan and R. (2014)

change more frequently and with smaller sizes than comparable findings in papers with offline CPI prices. This likely comes from their focus on retailers that participate in price-comparison websites, which tend to be online-only retailers. As Ellison and Ellison (2009) discuss, such retailers face a uniquely competitive environment that can significantly affect their pricing behaviors.

2 Simultaneous Online-Offline Data Collection

2.1 Multi-Channel Retailers

There are many types of “online prices”, but I focus exclusively on those obtained of large “multi-channel” retailers that sell both online and offline. When considering all retail sales, this type of retailers still concentrate the vast majority of all retail transactions, making them the most important source of price data for applications that require the use of “representative” data (such as inflation measurement). Despite its importance, this is also the type of “online prices” that has received the least attention in the IO, marketing, and internet pricing literatures due to lack of data.

Other kinds of online prices, such as those coming from online-only retailers like Amazon.com or online marketplaces like Ebay or Amazon Marketplace, are not considered in this paper. While their market shares may continue to grow in some cases, they still represent a small percentage of retail transactions in most countries. Furthermore, as pointed out by Brynjolfsson et al. (2013), technology is blurring the distinctions between physical and online retailing, making both traditional brick-and-mortar and online-only companies behave increasingly like multi-channel (“omni-channel”) retailers.

2.2 Retailer selection

The names of the retailers included in the data collection are shown in Table 1. They satisfy three conditions. First, they are a top 20 retailer by market share in a given

country.⁷ This helps ensure that I have a representative sample of the retail sector. Second, they sell both online through a website and offline through physical stores that could be visited. The majority of the retailers in the top 20 ranking satisfy this condition. Third, there has to be a way to uniquely match products online and offline. In practice, this meant that the product id number collected offline can be used to find the product on the website, something that is possible in most cases.⁸

Table 1: List of Retailers

Country	Retailers Included
Argentina	Carrefour, Coto, Easy, Sodimac, Walmart
Australia	Coles, Masters, Target, WoolWorths
Brazil	Droga Raia, Extra, Magazine Luiza, Pao de Azucar, Renner
Canada	Canadian Tire, Home Depot, The Source, Toys R Us, Walmart
China	Auchan Drive, Sams Club
Germany	Galeria Kaufhof, Obi, Real, Rewe, Saturn
Japan	Bic Camera, K's Denki, Lawson, Yamada
South Africa	Clicks, Dis-Chem Pharmacy, Mr Price, Pick n Pay, Woolworths
UK	Asda, Mark and Spencer, Sainsburys, Tesco
USA	Banana Republic, Best Buy, CVS, Forever 21, GAP, Home Depot, Lowe's, Macys, Nike, Office Max/Depot, Old Navy, Safeway, Staples, StopandShop, Target, Urban Outfitters, Walmart

Note: Results updated on 23 Mar 2016.

2.3 Collecting Offline Prices in Physical Stores

Collecting prices offline is an expensive and complicated process. NSOs rely on a large number of trained data collectors to do it correctly. Lacking the budget, I looked for alternatives using new technologies. I used popular crowd-sourcing platforms, such as Amazon Mechanical Turk, Elance, and UpWork, to find people willing to do simple

⁷The rank information was obtained from Euromonitor International's Passport Retailing Global Rankings. See www.euromonitor.com

⁸Some retailers use the same product id online and offline, either the good's UPC barcode or some internal id that is printed next to the price. Others have different online and offline codes but allow customers to search for offline UPC barcodes using the websites search forms. The Appendix provided details on the exact matching procedure for each store.

data collection tasks. To minimize the chance of data entry errors, I developed a custom mobile phone app that simplifies the data collection process.

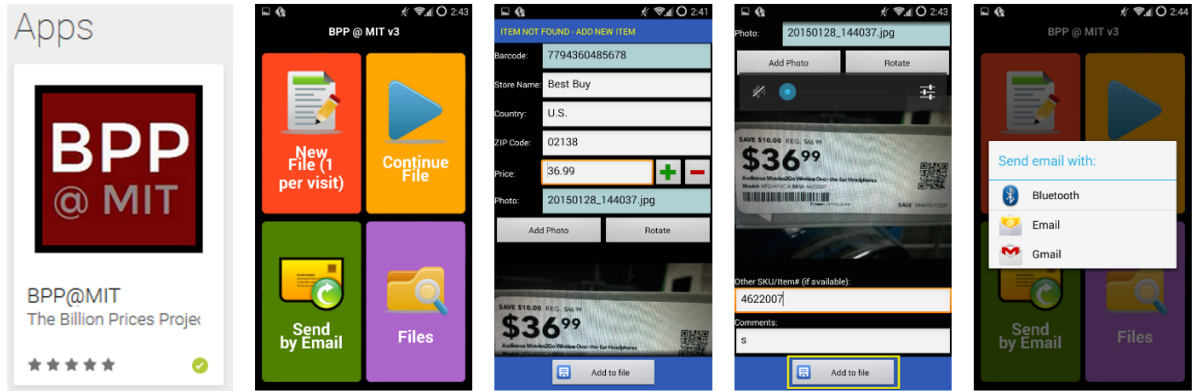
Crowd-sourcing platforms have many advantages. First, they allowed me to hire a large number of workers and reach multiple locations and cities within each country. Second, with many workers I could limit the number of individual prices that each one of them had to collect. This reduced the burden on the worker and also minimized the “show-rooming” concerns of the retailers. Showrooming is a term used to describe the practice of visiting a physical store to examine a product but later purchasing it online in another store. People who do this often use mobile apps to scan the product’s barcode and buy products online in the retailer with the lowest cost. If the data collectors spent too much time at each store, they could be required to stop and asked to leave.⁹

Two main versions of the task were posted on the crowd-sourcing websites. In the simplest case, the worker was asked to use a mobile app provided by us to scan 10 to 50 random offline products in any physical store, with some basic instructions to spread out the data collection across categories of goods. This provided the bulk of the data that I used to compare price levels across samples. A more complex version of the task required the worker to return to the same store every week for a full month and scan the same items. This gave me the panel of prices that I use to study price changes in Section 4.

The mobile app was custom-built to simplify and standardize the data collection process. It is an app for android phones called “BPP at MIT”, available for download at Google’s Play store.¹⁰ Every time a worker visits a store, she clicks on a button to open a new file. For the first product she has to enter the store’s name, zip code, and country. Then she scans the UPC barcode of the product (or the barcode on the price tag, depending on the particular retailer instructions provided), manually enters the price shown in the price tag next to the product (including all sales displayed), marks the price as

⁹I tried to conduct a similar large-scale offline data collection with MIT students in the Boston area in 2011, but most of them were forced to stop and leave the stores after a some time. It did not happen in 2015, not only because crowd-sourced workers spent just a few minutes scanning, but also because it is now quite common to see people using smartphones inside stores. FitzGerald (2013) reports that fear of showrooming has faded for many US retailers. See Balakrishnan et al. (2013) for an economic analysis of showrooming practices.

¹⁰See <https://play.google.com/store/apps/details?id=com.mit.bpp>



App available for download at the Google Play Store: <https://play.google.com/store/apps/details?id=com.mit.bpp>

Figure 1: Screenshots from BPP App for Android Phones

“regular” or “sale”, and takes a photograph of the price tag (which can be used to detect errors and validate the data). All products are scanned in a loop which makes the process quick and simple. When done, the worker presses another button to email the data to the BPP servers. A member of the BPP team verifies the submitted data and pays the worker.

Every few hours, the BPP servers automatically processed the incoming offline files to clean and consolidate the data for each retailer. I then used the offline barcode information to collect the online price in the retailer’s website, as described below.

2.4 Collecting Online Prices on each Retailer’s Website

To collect online prices I built a custom scraping “robot” for each retailer. These robots are specialized software that are programmed to use the product id (barcode) to query the retailer’s website and collect the online price and other available information. In most cases, the robot was designed to use the website’s search box to enter the product id obtained online, and then scraped the price for the product displayed by the retailer. For general details on the BPP’s online scraping methods, see Cavallo and Rigobon (2016).

The price collected online was the one displayed for each particular product on the retailer’s website, including any sales or discounts that apply to all customers. Whether taxes are added or not depends on the display conventions for prices in each country. For

example, US prices include sales but are typically shown without taxes, both online and offline. In all other countries, sales or VAT tax rates are usually included in the price in both locations. Shipping costs are never included in these online prices. Retailers have different ways to charge for shipping. The most common is a set of shipping fees that varies with the total amount of the sale or weight of the products. Some retailers offer free shipping fees, which could potentially mean that they also adjust their online prices to compensate. The online-offline price comparison provides information that can be used to determine when this is happening.

Nearly all of the online retailers in the sample have a single-price policy online for each product, independent of the location of the buyer. For example, someone purchasing a laptop from Bestbuy in San Francisco sees the same price as someone doing it from Boston. The only exceptions tend to be supermarkets, which sometimes require buyers to enter their zip code (or equivalent) before displaying prices. There are only five retailers that do this in my sample. I always use the same zipcode when collecting data online, independently of the one where the offline price was obtained, so this can cause some price level differences between the online and offline data for those retailers. However, in the Appendix I show that removing these retailers has very little impact on my results.

Finally, the scraping software could not always collect the online price on the same day, so for my benchmark results I allow online prices to be collected within 7 days of the offline price. The main tables also exclude sale prices. Results are similar for prices collected on the same day, or including sale prices, as shown in the Appendix.

2.5 The Online-Offline Matched Data

Table 2 shows the main characteristics of the matched datasets. I collected prices in 56 retailers for more than a year, between December 2014 and March 2016. There are more than 24 thousand products and 38 thousand observations in total. This dataset can be downloaded from <http://bpp.mit.edu>, together with the replication scripts for the results below.

The data coverage varies across countries. The effort was concentrated in the US, with

Table 2: Data by Country

Country	(1) Ret.	(2) Start	(3) End	(4) Days	(5) Workers	(6) Products	(7) Obs.
Argentina	5	02/15	08/15	71	18	2324	3699
Australia	4	03/15	08/15	70	17	3073	3797
Brazil	5	05/15	03/16	68	24	1438	1915
Canada	5	12/14	09/15	90	17	2665	4031
China	2	03/15	03/16	20	7	411	513
Germany	5	01/15	03/16	56	9	1215	1604
Japan	4	04/15	03/16	68	10	1127	2186
South Africa	5	03/15	03/16	83	25	2341	3212
UK	4	02/15	05/15	43	13	1663	2094
USA	17	12/14	03/16	234	235	7902	15332
ALL	56	12/14	03/16	391	373	24132	38383

Note: Results updated on 23 Mar 2016. Column 1 has the number of retailers. Columns 2 and 3 have the start and end months of data collection. Columns 4 and 5 report the number of days with data and workers that collected offline prices. Columns 6 and 7 provide the number of products and price observations that could be matched with both online and offline information.

13 retailers and about a quarter of all observations. On the other extreme is China, with only one retailer. I was unable to expand the offline data collection in China because large retailers explicitly prohibit taking photographs and recording prices at physical locations.¹¹

3 Price Levels

Table 3 compares the price levels across the online and offline samples. Column 3 shows the percentage of observations that have identical online and offline prices up to the second decimal.

The percentage of identical prices is 72% for all pooled observations and also for the average across countries. Some countries, such as Japan, have percentages close to 50%, while other such as Canada and the UK have over 90% of all prices being identical online and offline. The US is close to the average, with 69% of identical prices.

¹¹Apparently “showrooming” is more extended in China, and therefore retailers try to prevent the use of mobile phones apps and photography in their stores. According to an IBM Report in 2013, 24% of people surveyed in China admitted to having visited a physical store to buy online, compared with only 4% in the United States. See Klena and Puleri (2013) for details.

Table 3: Country - Level Differences

Country	(1) Ret.	(2) Obs	(3) Identical (%)	(4) High On (%)	(5) Low On (%)	(6) Markup (%)	(7) Difference (%)
Argentina	5	3699	60	27	13	3	1
Australia	4	3797	74	20	5	5	1
Brazil	5	1915	42	18	40	-7	-4
Canada	5	4031	91	3	5	-5	0
China	2	513	87	7	6	3	0
Germany	5	1604	74	4	23	-8	-2
Japan	4	2186	48	7	45	-13	-7
South Africa	5	3212	85	6	9	-3	-1
UK	4	2094	91	2	7	-8	-1
USA	17	15332	69	8	22	-5	-1
ALL	56	38383	72	11	18	-4	-1

Note: Results updated 23 Mar 2016. Column 3 shows the percentage of observations that have identical online and offline prices. Column 4 has the percent of observation where prices are higher online and column 5 the percentage of price that are lower online. Column 6, is the online markup, defined as the average price difference excluding cases that are identical. Column 7 is the average price difference including identical prices.

Columns 4 and 5 show the share of prices that are either higher or lower online. Conditional on a price difference, most countries tend to have lower online prices, with the exception of Argentina and Australia. The three countries with the lowest percentages of identical prices, where differences matter the most, tend to also have different behaviors. In Argentina, non-identical prices tend to be higher online, with an average markup of 3%. In Brazil, they are lower, with a markup of -7%. Japan is a clear outlier. Prices are lower online 45% of the time, with an average markup of -13%.

The average size of the price differences is quite small. This can be seen in Columns 6 and 7, where the average price difference is computed as the online price minus the offline price. Column 6 excludes cases where prices are identical, while column 7 includes them. I use the notation “Markup” and “Difference” respectively to distinguish the two cases. The online markup tends to be small, with a magnitude of -4% in the full sample. Adding prices that are identical makes the online-offline price difference only -1% on average.

Overall, these results imply little difference between online prices collected from a website and the offline prices that can be obtained by visiting the physical store of these

retailers.

The aggregate results, however, hide important heterogeneity at the sector level. Table 4 shows similar results for retailers grouped by the type of good they sell.

Table 4: Sector - Price Level Differences

Sector	(1) Ret.	(2) Obs	(3) Identical (%)	(4) High On (%)	(5) Low On (%)	(6) Markup (%)	(7) Difference (%)
Food	10	5953	52	32	15	3	1
Clothing	7	2534	92	5	3	3	0
Household	9	7875	79	5	16	-8	-2
Drugstore	4	3053	38	11	52	-5	-3
Electronics	5	3712	83	4	13	-9	-1
Office	2	1089	25	37	38	1	1
Multiple/Mix	18	14149	80	5	15	-9	-2

Note: Results updated 23 Mar 2016. Markup excludes identical prices. Difference includes identical prices.

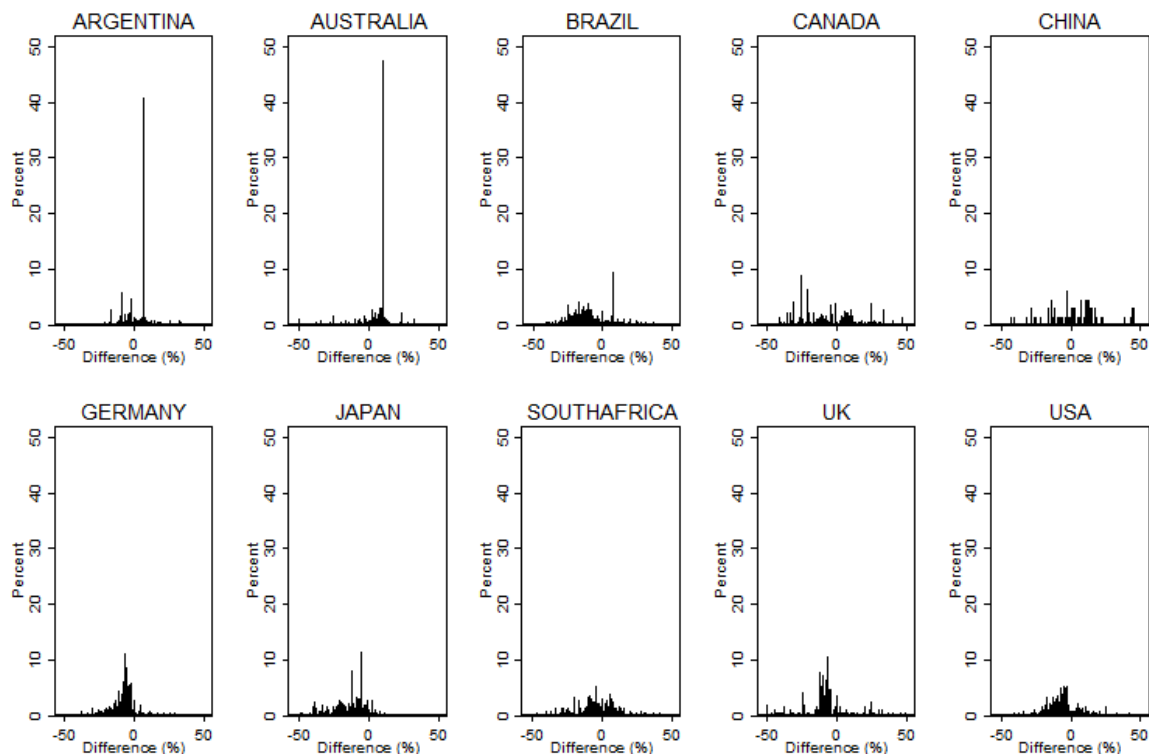
Drugstores and office-supply retailers have the lowest share of identical prices online and offline. For office products, prices are sometimes higher and sometimes lower online, without any clear patterns, as if the stores were managed independently. Drugstores, by contrast, tend to have lower prices online, possibly because they are “convenience” stores such as CVS and Walgreens in the US that can charge higher prices to offline customers.

Electronics and clothing have the highest share of identical prices. For clothing, prices are basically the same, with most of the observed differences coming from offline data collection errors and timing differences in the sampling. For electronics, prices are lower online 13% of the time, with an average markup of -9% (the highest in this sample).

Figure 2 shows the histograms for non-zero price differences in each country. The cases of Argentina and Australia stand out because there are spikes around the 5% magnitude of differences. This is caused by stable markups in online prices for some of the largest retailers. In all other countries the price differences are more dispersed in the range of -50% to 50%.

As pointed out by Nakamura and Steinsson (2008), sale events can be frequent in some countries, and the magnitude of the price changes that they generate can be large. But including or excluding sales does not alter the results significantly, as is shown in the

Figure 2: Histograms of Non-Zero Price Level Differences



Note: Results updated 23 Mar 2016. Excludes identical prices. Difference is online price - offline price. Bin width is 1%. X and Y scaled are matched for comparisons.

Appendix.

I do find that sale prices create more differences between online and offline samples, with the share of identical online and offline prices for sale observations being only 36%. But this has little impact on the full-sample results because the number of sales is small: only 11% of all matched observations have either an online sale (4.12%), an offline sale (5.03%), or both (1.92%). Unfortunately, my ability to control for sales is somewhat limited because workers could not identify offline sales with the app until October 2015, and some of the scrape jobs were not able to include online sale indicators. It is therefore likely that the main results still contain a lot of sales that I cannot control for, and the share of identical prices would rise significantly if these observations were removed.

Similarly, restricting the sample to include only prices collected on the same day (instead of allowing for a 7-day window) has little impact on the main results. The reason

is simply that prices do not typically change more than once a week. Details are provided in the Appendix.

Another potential reason for some of the price level differences is that goods have prices with similar time series that are not synchronized. I look for direct evidence of this in the next section, by comparing online and offline *changes* for a smaller sample of goods for which I have multiple weekly observations.

4 Price Changes

This section compares the behavior of price changes in the online and offline samples. A price change is defined as a non-zero difference in the price at weeks t and $t+1$. I study the timing, frequency, and size of price changes.

Table 5: Country - Price Changes

	(1) Obs.	(2) Price Changes	(3) Same Time (%)
Argentina	1392	247	35
Australia	759	73	22
Brazil	483	85	18
Canada	1427	120	32
Germany	419	16	31
Japan	1071	94	1
South Africa	882	108	15
UK	429	25	44
USA	7505	561	11
ALL	14367	1331	19

Note: Results updated 23 Mar 2016.

Table 5 focuses on the timing (only countries where there are at least 50 price change observations are shown). Column 1 shows the number of matched observations where a price change can be computed. There are fewer observations than in previous sections because I have a short time series for a limited subset of goods. Price changes can occur online, offline, or in both locations. Column 2 shows that only about 10% of observations have a price change. Column 3 reports the percentage of price changes for a given product that occur both online and offline at the same time, which I refer to as “synchronized”.

Only 19% of the 1331 price changes were synchronized across online and offline samples.

Table 6: Country - Price Change Frequency and Size

	(1) Price Changes	(2) Mean Freq. Online	(3) Mean Freq. Offline	(4) t-test p-val	(5) Mean Abs Size Online	(6) Mean Abs Size Offline	(7) t-test p-val
Argentina	247	.137	.147	.52	13.61	12.74	.67
Australia	73	.056	.092	.06	45.76	42.18	.62
Brazil	85	.167	.138	.36	10.55	9.36	.53
Canada	120	.077	.068	.48	31.11	21.71	.06
Germany	16	.035	.041	.74	27.08	16.52	.29
Japan	94	.074	.012	0	12.1	8.3	.4
South Africa	108	.1	.075	.13	23.33	16.51	.09
UK	25	.046	.07	.28	47.68	41.78	.67
USA	561	.052	.046	.28	23.78	21.18	.18
ALL	1331	.076	.068	.07	22.02	19.98	.11

Note: Results updated 23 Mar 2016.

While not synchronized, online and offline prices behave similarly in terms of how frequently they change. This can be seen in Table 6. The frequency statistics in Columns 2 and 3 are computed for each individual good first (as the share of observations with a price change), and then averaged across countries. Column 4 shows the p-value of a two-sided t-test with a null hypothesis of equal average frequencies in the online and offline samples. I can reject the null of equality with some confidence in the cases of Australia and Japan. The full sample results appear to have slightly more frequent changes online, but this is entirely driven by Japan.

In addition to similar frequencies, online and offline price changes tend to have similar sizes. This can be seen in Columns 5 and 6, where I report the mean absolute size of price changes. Column 7 is again the p-value of a two-sided t-test of equality in the online and offline means. The null hypothesis can only be rejected in Canada, where online price changes seem to be larger. In all other countries the difference is not statistically significant.

Overall, these frequency and size results suggest that the online and offline price time series behave similarly but are not synchronized. In a related paper, Cavallo and

Rigobon (2016), we find evidence that online price inflation tends to anticipate offline CPI inflation. A faster adjustment to shocks could be the reason why online price changes are not synchronized with offline changes.

5 IP-Address Location and Persistent Browsing

There has been some research suggesting that some online retailers change prices based on the browsing habits of the consumer or the location associated with the ip address of the computer being used to purchase online. See, for example, Valentino-DeVries et al. (2012), Mikians et al. (2012), and Mikians et al. (2013). If these pricing behaviors are also common for multi-channel retailers in my sample, it could help explain some of the price level differences in the data. To test whether prices vary with browsing habits or ip address, I created two experiments with special versions of the scrape robots for US retailers.

The first experiment was designed to test whether prices change based on the zip code associated with the ip address of the computer collecting the data.¹² I randomly selected 5 products in each of 10 US retailers and scraped their prices twelve times in a consecutive loop. In each loop I changed the ip address of the robots by using 12 different proxy servers in 9 US cities (Atlanta, Burbank, Charlotte, Chicago, Cleveland, Miami, Nashville, New York, and two proxies in Phoenix) and 2 international locations (Canada and UK).¹³ I did not find any evidence of this type of price discrimination. In all cases, prices were the same for a given product, regardless of what ip address was used to connect to the retailer websites.

The second experiment was designed to test if frequent visits to the webpage of a

¹²IP addresses are unique numeric identifiers for computers that are connected to a network. They are assigned by internet service providers and have an associated geographical location that is public information. For example, MIT's campus ip addresses range from 18.0.0.0 to 18.255.255.255 and are geographically linked to the 02139 zip code in Cambridge, Massachusetts. In principle, retailers could detect the ip address of the consumer visiting a site and automatically change the prices displayed based on its geolocation information.

¹³A proxy server is a computer that acts as an intermediary for the communications between two other computers in a network—in this case between the machine where the scraping software runs and the server hosting the website of the retailer—. From the retailer's website perspective, the request was coming from the ip address associated with the proxy server.

particular product could lead the retailer to change the price displayed. In this case, I scraped a single product in each retailer every five minutes for a full day. Once again, there was no evidence of price discrimination based on persistent-browsing habits: prices were always the same.

These results suggest that while these forms of price discrimination may be important in other industries (e.g. airlines) or type of retailers (e.g. online-only stores), they are not commonly used in large multi-channel retailers, at least in the US. A possible reason is that retailers may fear antagonizing their customers if reports of these tactics were to become publicized in the press, as in Valentino-DeVries et al. (2012).

6 Product Selection

The results in the previous sections would have different implications if most goods sold offline were not available online. I therefore now estimate the “overlap” in product selection across samples, defined as the share of offline goods that are also available online.¹⁴

In principle, I could use the percentage of offline product barcodes, received through the app, for which the scraping software found information online. This percentage is 63%. The problem with this number, however, is that the automated matching procedure can fail for many reason: the worker could scan the wrong barcode, the app can incorrectly read the barcode, or the scraping robot could fail while checking the website.

To better estimate the degree of product selection overlap, the BPP team manually checked a sample of 100-200 observations per retailer. They used all the information submitted by the workers, including the product description readable in the photo of the price tag, to determine how many of the offline products could also be found online.

The results, grouped by country, are reported in Table 7. Column 2 shows the number of products that were found online automatically. Column 3 are the ones that were missed by the automatic process but they were online when manually checked. Column 4 shows

¹⁴Note that, given the data characteristics, I can only estimate how many offline products are also sold online, but not the other way around. In some retailers, the online selection of goods may be larger than what can be found in a single store because online sales can be shipped from large centralized warehouses.

Table 7: Retailer - Product Selection Overlap

Country	(1) Sample	(2) Found Automatically	(3) Found Manually	(4) Total Overlap (%)
ARGENTINA	500	294	52	73
AUSTRALIA	500	435	36	95
BRAZIL	400	331	12	86
CANADA	500	279	132	85
CHINA	100	50	3	53
GERMANY	400	178	23	52
JAPAN	500	329	61	74
SOUTHAFRICA	500	332	60	76
UK	500	373	59	86
USA	1600	1003	316	81
ALL (mean)	5500	3604	754	76

Note: Results updated 23 Mar 2016. Manual check with 200 products per retailer. Only a subset of retailers in each country are included.

the estimate for the total overlap $((2)+(3))$ as a percentage of (1)).

Approximately 76% of all offline products were also online. There are, once again, important differences among countries. China and Germany have the lowest overlap, while Australia, Brazil, and the UK the highest. In the US, 81% of offline products were also online.¹⁵ In the Appendix, I further show there is also considerably heterogeneity at the retailer level.

Finally, I considered the possibility that goods that could not be automatically matched might be precisely those where the online and offline prices are different. This could happen, for example, if retailers change the online id number for those goods as a way to obfuscate their price differences and prevent any comparisons. The evidence shows this is not generally the case. Both automatic and manually-matched goods produced similar results for online and offline price-level comparisons (see Appendix for details).

¹⁵These results are likely lower-bound estimates for the overlap because the manual checks often took place weeks after the original offline data was collected.

7 Retailer Heterogeneity

The country-level results in the previous sections conceal a great deal of heterogeneity across retailers in each country. Details for each retailer can be seen in Appendix Table ??, where I show price level and changes results for all retailers with at least 100 observations.

Three main types of retailers are typical. First, there are retailers where online and offline prices are identical most of the time. These are cases where the retailer explicitly chooses to have the same online and offline price. Second, there are also some retailers with a low share of identical prices, but no clear online markups. Many retailers in Brazil, for example, exhibit this pattern. These are likely cases where the online store is simply treated as another outlet, sometimes cheaper, sometimes more expensive. Third, there are retailers with a low share of identical prices and a significant online markup (either positive or negative). There are some examples in Argentina, Brazil, Japan, and the US. These patterns may reflect a desire to compensate for shipping costs or price-discriminate online consumers.

Whether each kind of retailer is useful as a source of data depends on the purpose of the paper or application. For example, using online prices for the retailer in Argentina where 79% of prices are higher online is not a problem for measuring inflation as long as the online markup is relatively constant over time, but it would bias the results if we were interested in comparing price level differences across countries. Unless a correction is applied, the online data would make prices in Argentina appear higher than what they really are. Identifying these special patterns and correcting for any biases is particularly important in papers or applications that use online data from one (or a few) retailers.

8 Conclusions

This paper shows that, for large multi-channel retailers, there is typically little difference between the online price collected from a website and the offline price obtained by visiting the physical store. Prices are identical about 72% of the time, and price changes have different timing but they are similar in terms of frequency and size. At the same

time, there is also considerable heterogeneity across countries, sectors, and retailers.

For research economists using online data for macro and international research questions, my results provide evidence that online prices are a representative source of retail prices, even if most transactions still take place offline. At a more micro level, the differences in behaviors can be used to better model the pricing dynamics and strategies of different types of retailers in various sectors and countries. However, this high degree of heterogeneity also implies that papers that use relatively few sources of data should be cautious to understand relevant pricing patterns and control for any potential sampling biases.

For National Statistical Offices (NSOs) considering the use of online data for consumer price indexes, my results show that the web can be effectively used as an alternative data-collection technology for multi-channel retailers. Particularly for products such as electronics or clothing, the price collected on the web will tend to be identical to the one that can be obtained by walking into a physical store. And online prices are not only cheaper to collect, but they also provide information for all goods sold by each retailer, with many details per product, uncensored price spells, and can be collected on a high-frequency basis without any delays.

Finally, for those interested in the effect of the Internet on retail prices, my results suggest there is little within-retailer price dispersion in multi-channel retailers. While the Internet may not have reduced dispersion across retailers, it seems to have created the incentives for retailers to price identically across their own stores. Perhaps the transparency of online prices has forced retailers to reduce geographical price discrimination for fear of antagonizing customers. This could also be affecting cross-border pricing, as suggested by Cavallo et al. (2014), where we found evidence that global firms such as Apple and Ikea tend to price identically in countries that use the same currency, where it is trivial for consumers to compare prices across countries.

Future work should try to understand why there are some observed price-level differences, and what sector and country-level characteristics could be driving them. One explanation, consistent with the results in this paper and those in Cavallo and Rigobon

(2016), is that online prices may adjust faster to shocks. Another potential cause is that location-specific sales play a larger role than I can detect in these data. Yet another possibility is that some of price differences are caused by attempts to match the pricing of online-only competitors such as Amazon.com;. To the best of my knowledge, no one has documented or quantified the effect that retailers such as Amazon.com are having on large retailers such as Walmart or Target. A direct comparison between online-only and multi-channel retailers prices would be useful to better understand how retail pricing dynamics are likely to change in the future.

References

- Balakrishnan, Anantaram, Shankar Sundaresan, and Bo Zhang**, “Browse-and-Switch: Retail-Online Competition under Value Uncertainty,” *Production and Operations Management Society*, 2013, *23* (7), 1129–1145.
- Boettcher, Ingolf**, “Automatic data collection on the Internet,” *Statistics Austria*, 2015.
- Breton, Robert, Gareth Clews, Liz Metcalfe, Natasha Milliken, Christopher Payne, Joe Winton, and Ainslie Woods**, “Research indices using web scraped data,” 2015.
- Brynjolfsson, Erik and Michael D. Smith**, “Frictionless Commerce? A Comparison of Internet and Conventional Retailers,” *Management Science*, April 2000, *46* (4), 563–585.
- , **Astrid Andrea Dick, and Michael D. Smith**, “Search and Product Differentiation at an Internet Shopbot,” *SSRN Electronic Journal*, February 2003.
- , **Yu Jeffrey Hu, Mohammad S Rahman, and By Erik Brynjolfsson**, “Competing in the Age of Omnichannel Retailing,” *MIT Sloan Management Review*, 2013, *54* (4), 23–29.
- Cavallo, Alberto**, “Scraped Data and Sticky Prices,” *NBER Working Paper*, 2015, (21490).
- and **Roberto Rigobon**, “The Billion Prices Project: Using Online Data for Measurement and Research,” *Journal of Economic Perspectives*, 2016, (Forthcoming).
- , **Brent Neiman, and Roberto Rigobon**, “Currency Unions, Product Introductions, and the Real Exchange Rate,” *Quarterly Journal of Economics*, 2014, *129* (2).
- Eleisegui, Patricio**, “La gndola en su PC: conviene ir al sper o llenar el changuito de manera online?,” *Infobae*, 2011.
- Ellison, Glenn and Sara Fisher Ellison**, “Search, Obfuscation, and Price Elasticities on the Internet,” *Econometrica*, March 2009, *77* (2), 427–452.
- Euromonitor**, “Internet vs Store-Based Shopping: The Global Move Towards Omnichannel Retailing,” 2014.
- FitzGerald, Drew**, “Fear of ‘Showrooming’ Fades — Best Buy, Other Retailers Optimistic That Price-Matching Can Stanch Trend,” *Wall Street Journal*, November 2013.
- Fowler, Geoffrey A.**, “Do Online Grocers Beat Supermarkets?,” *Wall Street Journal*, January 2014.
- Gorodnichenko, Yuriy and Oleksandr Talavera**, “Price Setting in Online Markets: Basic Facts, International Comparisons, and Cross-border Integration,” *NBER Working Paper*, August 2014, (20406).
- , **Viacheslav Sheremirov, and Oleksandr Talavera**, “Price Setting in Online Markets: Does IT Click?,” *NBER Working Paper Series*, 2014, (20819).

- Griffioen, Robert, Jan de Haan, and Leon Willenborg**, “Collecting clothing data from the Internet,” in “meeting of the group of experts on consumer price indices” 2014, pp. 26–28.
- Groz, Martin**, “Hacer las compras online es hasta 13% ms caro que ir al supermercado,” *Clarín*, 2013.
- Hannak, Aniko, Gary Soeller, David Lazer, Alan Mislove, and Christo Wilson**, “Measuring price discrimination and steering on e-commerce web sites,” in “Proceedings of the 2014 conference on internet measurement conference” ACM 2014, pp. 305–318.
- Horrigan, Michael**, “Big Data: A Perspective from the BLS | Amstat News,” 2013.
How Amazon uses psychology to trick you into thinking everything is cheap | News | Geek.com
- How Amazon uses psychology to trick you into thinking everything is cheap | News | Geek.com.***
- Klena, Kali and Jill Puleri**, “From transactions to relationships: Connecting with a transitioning shopper,” 2013.
- Klosowski, Thorin**, “Is It Actually Cheaper to Order Everything from Amazon?”
- Krsinich, Frances**, “Price Indexes from online data using the fixed-effects window-splice (FEWS) index,” *Statistics New Zealand*, 2015.
- Lnnemann, Patrick and Ladislav Winttr**, “Price Stickiness in the US and Europe Revisited: Evidence from Internet Prices*,” *Oxford Bulletin of Economics and Statistics*, October 2011, 73 (5), 593–621.
- Mikians, Jakub, Lszl Gyarmati, Vijay Erramilli, and Nikolaos Laoutaris**, “Detecting price and search discrimination on the internet,” in “Proceedings of the 11th ACM Workshop on Hot Topics in Networks” ACM 2012, pp. 79–84. *bibtex: mikians_detecting_2012-1.*
- , – , – , and – , “Crowd-assisted search for price discrimination in e-commerce: First results,” in “Proceedings of the ninth ACM conference on Emerging networking experiments and technologies” ACM 2013, pp. 1–6.
- Nakamura, Emi and Jon Steinsson**, “Five Facts About Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics*, 2008, 123(4), 1415–1464.
- Nygaard, Ragnhild**, “The use of online prices in the Norwegian Consumer Price Index,” *Statistics Norway*, 2015.
- Passariello, Christina, Suzanne Kapner, and Manuela Mesco**, “‘Showrooming’ Hits Luxury Fashion — Late to Establish an E-Commerce Presence, Fashion Houses Lose Out as Clients Buy Elsewhere Online,” *Wall Street Journal*, April 2014.
- Quan, Thomas W. and Williams. Kevin R.**, “‘Product Variety, Across-Market Demand Heterogeneity, and the Value of Online Retail.’,” 2014.

- Rapp, Adam, Thomas L. Baker, Daniel G. Bachrach, Jessica Ogilvie, and Lauren Skinner Beitelspacher**, “Perceived customer showrooming behavior and the effect on retail salesperson self-efficacy and performance,” *Journal of Retailing*, February 2015, 91, 358–369.
- Smith, Michael D. and Erik Brynjolfsson**, “Consumer Decision-Making at an Internet Shopbot,” *SSRN Electronic Journal*, October 2001.
Study: Shopping on Amazon up to 15% cheaper than buying from stores
- Study: Shopping on Amazon up to 15% cheaper than buying from stores**, September 2014.
- Valentino-DeVries, Jennifer, Jeremy Singer-Vine, and Ashkan Soltani**, “Websites Vary Prices, Deals Based on Users’ Information,” *Wall Street Journal*, December 2012.
- Welch, David**, “Wal-Mart Beats Amazon Prices Including Glee DVD Set.”
- Wells, Rachel**, “Supermarkets charge more to buy online,” *Sydney Morning Herald*, 2013.
- Wohl, Jessica**, “In-store shoppers likely to head online next time they buy: IBM,” *Reuters*, January 2013.