ROBERTO RIGOBÓN

# Presidential Address: Macroeconomics and Online Prices

The availability of microeconomic pricing data has produced an empirical revolution in macroeconomics. Two events were the main culprits: First, national statistical offices allowed economists to study the data underlying the construction of the consumer price index (CPI), which gave our profession the chance to tackle many questions that have been open for decades. Second, scanner data from several supermarkets and merchandising companies were made available as well, offering another great opportunity for research. The earlier literature used aggregate price indexes to address questions of price rigidity, the law of one price, cost and exchange rate pass-through, international market segmentation, and so on. The aggregation and the procedures behind the construction of those indexes mask several economic phenomena. The availability of more detailed data has allowed the profession to take a closer look at old questions.[1]

The microeconomic CPI data has several advantages. The first and most important is its representativeness. Statistical offices invest in the design of the data to make sure they include a representative set of prices for the consumption basket. The second advantage is the long history. When the microeconomic price data are released, researchers generally have access to several years and sometimes decades. This feature is quite important for evaluating pass-through and relative price equilibrium deviations. The disadvantages are many, including the one indicated by Alberto Cavallo in his thesis: microeconomic CPI prices are plagued with unit values.[2] Although unit values are

Rigobon is with the Massachusetts Institute of Technology (MIT).

I want to thank El Colegio de México for organizing the 2013 LACEA conference.

1. See Klenow and Malin (2011) for a very good summary on what the literature has recently concluded on the matter.

2. See Cavallo (2012, section 3.2.1).

conceivably a very good piece of information for the computation of inflation, they are terrible for understanding price-setting dynamics, especially when evaluating price stickiness. The second disadvantage is that even though the data are representative, CPIs tend to have very few items in each sector, which means that the heterogeneity within sectors is disguised by small samples.

The scanner data resolved some of the issues in the CPI data and worsened others. The scope was clearly much bigger, and the extensive product variety allowed researchers to deal with heterogeneity much better. In addition, several scanner data sets have information on cost and quantities, as well as on the prices, which facilitates addressing questions on pass-through. The disadvantages, however, are several. First, scanner data are not representative. The data characteristics vary greatly depending on the data provider, the location, and the time period when the data were collected. Databases available for research are usually from a single retailer, which makes generalization even harder. Furthermore, the quantity data captured by scanner data can be biased. The reason is simple: a supermarket that sells milk at an unusually low price will experience unusually large sales of milk. Using the quantities to determine aggregate implications produces a massive bias. Second, scanner data also has unit values, with prices reported as a ratio of sales over quantities and averaged over a week. Not all scanner data suffer from this problem, but most of the data sets that have been used in the macroeconomic and international literature do. Observing the average price is not necessarily wrong if the question asked is about pass-through or inflation. It is, however, the incorrect data point when addressing questions of price stickiness and price dynamics.[3]

Alberto Cavallo and I started the Billion Prices Project at MIT (BPP) almost a decade ago to explore how using web methods can improve the collection of prices.[4] We use web scraping to download millions of prices every day, from hundreds of retailers in more than seventy countries. The purpose is to collect all the products sold by a store, identify the posted prices every day, and also detect information about sales, promotions, and the like, in which case we collect all prices available. There are some advantages and disadvantages to these data. First, they are not as representative as census

3. A very nice survey on the recent microeconomic pricing literature in international economics can be found in Burstein and Gopinath (2013).

4. I am using "we" quite loosely here. Alberto Cavallo does all the work downloading, cleaning, computing, and researching, while I do the tougher job of taking all the credit.

data, but they are more representative than data from a single store, which is the typical source of scanner data. In particular, every U.S. supermarket selling products through an online web page is probably in our data set, not just Safeway. Second, the data are daily and have no unit values. Hence, they are better equipped to handle price dynamics, but less useful for understanding inflation. Third, online data will not have quantities at all, which is a major disadvantage when compared with scanner data and information from the Bureau of Labor Statistics (BLS).

One immediate question is whether there is a difference between online and offline prices. For instance, are the prices posted online related to transaction prices? This is a question not only about the representativeness of the data, but it is also about its integrity, and it only be understood by checking at every store. We ran validation experiments sending people to stores and supermarkets, buying and photographing random items and then comparing them to the online prices. Also, talking to the stores proved quite useful: for example, Apple, Dell, Ikea, H&M, Zara, Mango, Lululemon, and many others have a policy of showing online exactly the same prices they have offline. But it is not the case that every price on the web is meaningful. Price aggregator pages like Kayak for tickets and Booking.com for hotels sometimes show prices that are not transactional. This severely hurts the veracity of such data sources. In all our research, we use stores that only show prices for those items they intend to sell and that provide markers when the item does not exist or is backordered, so we can code it into our database.

In general, however, prices online are not exactly the same as prices offline. There is a markup, and it changes country by country: in some places online prices are higher, and in others they are lower. Cavallo shows that even in cases in which the prices are not identical, the properties regarding stickiness and inflation are consistently almost identical at the store level between online and offline prices—even in emerging markets.[5] More research is needed in this area, yet what we know so far is that the dynamic properties of online and offline prices are very similar when the prices are not identical.

The next two sections discuss papers that use online prices to reevaluate important questions that the literature has previously tackled with either CPI or scanner data. The first is about the law of one price and its deviations, while the second is about the shape of the price change distribution.

5. Cavallo (2012).

## Currency Unions, Product Introductions, and the Real Exchange Rate

In the international literature, there are massive deviations of the law of one price, even for those items narrowly defined.[6] Explanations typically focus on aspects such as transport costs, tariffs, differences in language or culture, competition, differences in income, volatility of the exchange rate, and so forth. In contrast, Cavallo, Neiman, and Rigobon find that by far the most salient determinant of price differences is the currency in which prices are quoted.[7]

We study four industry leaders: Apple, IKEA, H&M, and Zara. We analyze how the prices of these stores differ across countries for identical goods. We find significant differences across countries outside the euro area, but prices are often identical within the euro area.

Some patterns in our data are consistent with what the literature finds. When we compare products across two countries that have flexible exchange rates, price dispersion is higher than when the two countries have fixed exchange rates. This is not terribly surprising. Indeed, the average price deviation is about a third less for fixed than for flexible rates.

The surprising result appears when we compare products from fixed currencies and currency unions. Price dispersion drops by a further 33 percent! In fact, whereas identical prices almost never exist outside of currency unions, they are commonplace inside them. Put differently, when we compare Denmark to Germany—a country pair in which both countries have a very credible fixed exchange rate regime and are in the same business cycle—the prices exhibit dramatically more dispersion than when we compare Germany and France, or Germany and Greece, or Finland and Greece![8]

Formally these results can be evaluated by the average of absolute deviations or by the mass between −1 and 1 percent price deviation. Table 1 shows the dramatic impact that currency unions have on the distribution of price changes. Among floats, only 4.5 percent of the prices are ever between −1 and 1. This increases to 6.9 percent when the countries have a peg, but it bloats to 61.0 percent within currency unions!

6. See, for example, Crucini, Telmer, and Zachariadis (2005) or Gopinath and others (2011). Engel (1999) demonstrates that movement in relative prices of tradable goods across countries is the principal driver of variation in bilateral real exchange rates.

7. This section is based on Cavallo, Neiman, and Rigobon (2015).

8. This result is not a purely euro-specific phenomenon. We find the same qualitative pattern when we compare prices in the United States to dollarized countries (such as Ecuador and El Salvador) and to countries with strong pegs to the dollar (such as Jordan and Lebanon).

**TABLE 1. Price Change Distribution under Different Exchange Rate Regimes**

| *Exchange rate regime* | *All stores* | *Apple* | *IKEA* | *H&M* | *Zara* |
|---|---|---|---|---|---|
| *A. Average absolute value of log of good-level real exchange rate* | | | | | |
| Currency unions | 0.076 | 0.023 | 0.129 | 0.020 | 0.102 |
| Nominal pegs | 0.116 | 0.085 | 0.145 | 0.119 | 0.115 |
| Floats | 0.187 | 0.143 | 0.216 | 0.145 | 0.207 |
| *B. Share of absolute value of log of good-level real exchange rate that is under 1 percent* | | | | | |
| Currency unions | 0.610 | 0.681 | 0.307 | 0.911 | 0.548 |
| Nominal pegs | 0.069 | 0.140 | 0.081 | 0.069 | 0.064 |
| Floats | 0.045 | 0.049 | 0.033 | 0.062 | 0.040 |

Source: Cavallo, Neiman, and Rigobon (2014).

To visualize these results, we pool our data across goods and time and plot the logarithm of the price of each good in the listed country relative to the price in the United States in dollars (see figure 1). Figure 2 shows the exact same calculation for the relative prices between a given country and Spain (now in euros). For example, in figure 2 an $x$-axis value of zero in the bottom right histogram labeled "United States" would mean that a good is sold at the identical price in the United States and Spain after taking into account the dollar-euro exchange rate. An $x$-axis value of 0.1 would mean a good costs roughly 10 percent more in the United States than in Spain. The $y$ axes in the histograms capture the density of products corresponding to each of these relative prices.

The distributions in figure 1 look very similar to the distributions commonly found in the literature. There is a massive dispersion of prices across countries, and price deviations of over 25 percent are common in the data, even for countries as similar and as integrated as Canada and the United States.

Figure 2 presents the exact same computation, but with Spain as the base country. It confirms the statistics described above. Countries sharing a common currency have a large mass at identical prices, corresponding to the spikes at zero. Non-euro-area countries—including Denmark, which pegs to the euro—have far more mass elsewhere.

Clearly, these international pricing patterns might not be representative of all traded goods. For example, our results are probably not very informative about the behavior of fresh food or auto prices. From my point of view, this shows that collecting the data differently and making sure that the matching of products is perfect produces a significantly different result from what the literature has found.

An immediate question is why price dispersion is so dramatically reduced among the euro area countries. We hope to answer this question in future

**FIGURE 1. Good-Level Real Exchange Rates with the United States (in logs)**

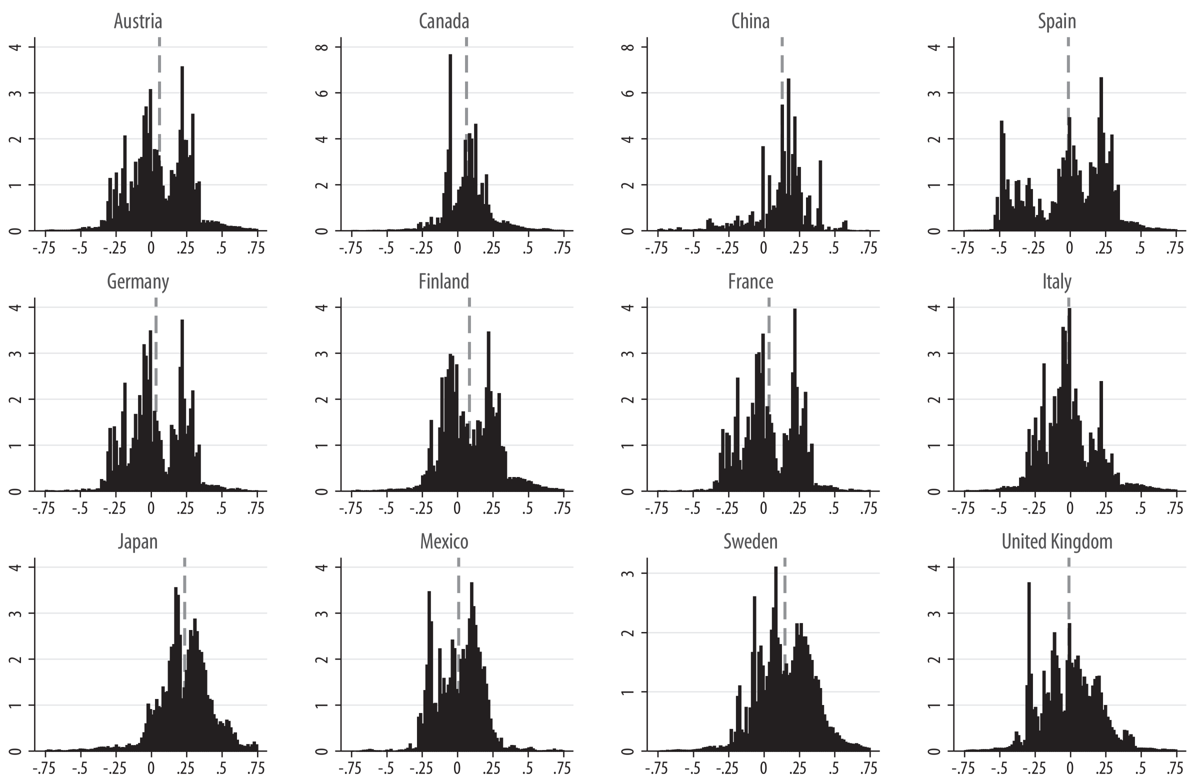

**FIGURE 2. Good-Level Real Exchange Rates with Spain (in logs)**

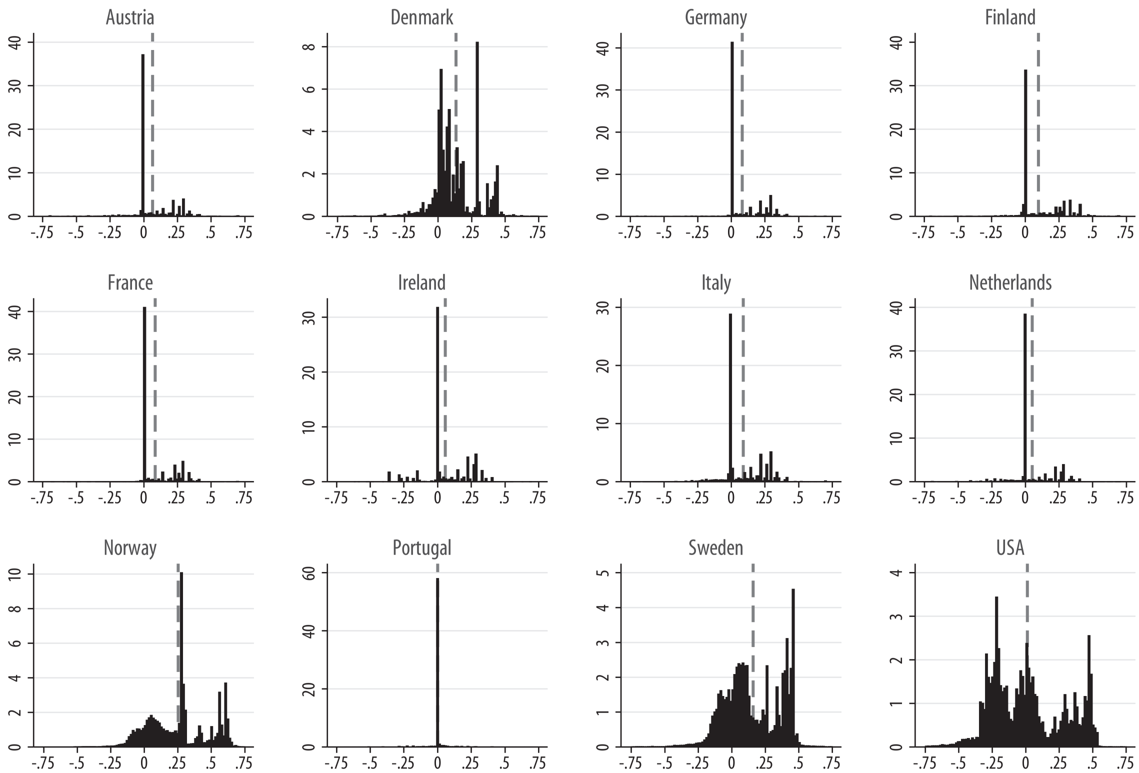

work. At this point, we can rule out several possible explanations. First, it is not merely due to all European customers being routed to the identical web page (although if firms did choose this structure, that would itself be puzzling). If a consumer in Madrid tries to access the firm's Italian web page, he or she is either rerouted to the Spanish page or is not allowed to input a shipping address outside of Italy. Second, price equalization in the euro area is not required by regulations or competition policies since all such policies exist at the European Union level as opposed to the euro area level.[9] Finally, we use the identical matching algorithm in our analysis of all bilateral country pairs, regardless of their currency regime. The fact that this same procedure identifies such small price differences in one set of countries implies that the large price differences found among others is not due to matching errors that result in price comparisons of different goods. These results perhaps suggest a greater role for consumer psychology or firm organizational structure in macroeconomic models of price determination—and the few stores I have asked actually agree with this. It is consumer anger that they are paying attention to.

## The Distribution of the Size of Price Changes

One of the main stylized facts uncovered by this literature is that the distribution of price changes (conditional on a change) is close to a unimodal distribution centered at zero percent, with a large share of small price changes. This finding has also been shown to hold in scanner data sets from retailers in the United States.[10] For example, it is very common that the frequency of small price changes (between −1 and 1 percent) is high. Furthermore, it is generally higher than the frequency of price changes with absolute value between one and two. Hence, the distributions look like a normal distribution! The shape of the distribution is important because it allows us to disentangle the different theories of price stickiness.[11]

9. According to Whish and Bailey (2012), "The Court of Justice in *United Brands v. Commission* ruled that 'it was permissible for a supplier to charge whatever local conditions of supply and demand dictate, that is to say, that there is no obligation to charge a uniform price throughout the E.U.'"

10. See Midrigan (2011); Klenow and Kryvtsov (2008).

11. See Alvarez, Lippi, and Paciello (2010).

Cavallo and Rigobon study the distribution of price changes for hundreds of retailers in several countries.[12] Our first version was exclusively on supermarkets, but we now have retailers in very different sectors. We find exactly the opposite of what the literature finds: it is almost never the case that the distribution is normally distributed, and the mass of price changes close to zero is actually very small.

The original supermarket data has supermarkets in several countries (see table 2). We run the Hartigan and Silverman tests—the two most powerful tests for unimodality—and find that unimodality is rejected almost everywhere even when we concentrate on the narrow window of –5 to 5 percent price increases. In other words, even when we take out price changes that are obvious sales (10 percent, 20 percent, and more), the distribution does not look unimodal at all. Out of all the supermarkets, in only three cases is the *p* value of one mode not rejected at 2.5 percent confidence! In fact, unimodality is rejected even when estimated at the category level, where a category would be quite narrowly defined (such as organic eggs). Figure 3 presents the distributions for the supermarkets.

In the paper, we develop a test called the proportional mass, which is a very simple test to evaluate single-mode distributions at a given point. The results are consistent with the standard methods in the statistical literature.

Why the differences between our results and those found in the literature? First, some of our retailers (about one in fifty) do exhibit normal distributions for their price changes. One of those establishments is Safeway. Thus, most of the literature has been making generalizations based on what we find to be a massive exception. Second, and more importantly, most of the data are reported in unit values. This is a deeper problem than the previous point (which is mostly about sample selection).[13] The main problem is that stores actually have several prices for every item: the regular or posted price, the sales price, the price with the coupon, the price when using the loyalty card, and so on. These prices all differ by a significant amount, of 5 or 10 percent or even more especially for loyalty cards. However, the demand within each of these prices is slightly random, meaning that the average daily or weekly price is shifting by a very small number—not because the prices are moving by a small percentage, but because of the small demand changes.

12. See Cavallo and Rigobon (2011).
13. This issue is addressed in Cavallo (2012) and Eichenbaum and others (2014).

**TABLE 2. Supermarket Data**

| *Database* | *Country* | *Started* | *Days* | *Obs.* | *Products* | *# Pr P/ day* | *Pr.* | *Ch. (cc)* | *Sales* |
|---|---|---|---|---|---|---|---|---|---|
| ARGENTINA-1 | Argentina | 10/7/2007 | 876 | 13117K | 26K | 12K | 155K | 1.20% | YES |
| ARGENTINA-2 | Argentina | 23/7/2007 | 861 | 5294K | 11K | 6K | 103K | 2.00% | YES |
| AUSTRALIA-1 | Australia | 8/4/2008 | 574 | 232K | 3K | 1K | 147K | 63.40% | NO |
| AUSTRALIA-2 | Australia | 8/7/2008 | 571 | 202K | 1K | 0K | 2K | 1.00% | NO |
| AUSTRALIA-3 | Australia | 8/4/2009 | 209 | 3292K | 7K | 6K | 2K | 0.10% | NO |
| AUSTRALIA-4 | Australia | 5/3/2008 | 667 | 1967K | 18K | 4K | 46K | 2.30% | YES |
| BRAZIL-1 | Brazil | 10/10/2007 | 873 | 10780K | 22K | 11K | 260K | 2.40% | YES |
| CHILE-1 | Chile | 10/24/2007 | 859 | 12102K | 35K | 12K | 120K | 1.00% | NO |
| CHINA-1 | China | 12/5/2008 | 451 | 1101K | 7K | 3K | 6K | 0.50% | NO |
| CHINA-2 | China | 3/19/2008 | 712 | 6644K | 46K | 10K | 22K | 0.30% | NO |
| COLOMBIA-1 | Colombia | 11/13/2007 | 839 | 4186K | 9K | 5K | 77K | 1.80% | YES |
| ECUADOR-1 | Ecuador | 3/19/2009 | 347 | 667K | 3K | 2K | 6K | 0.90% | NO |
| FRANCE-1 | France | 10/29/2008 | 488 | 2806K | 10K | 5K | 11K | 0.40% | NO |
| FRANCE-2 | France | 11/18/2008 | 468 | 4878K | 17K | 10K | 18K | 0.40% | NO |
| FRANCE-3 | France | 11/5/2008 | 481 | 3102K | 21K | 6K | 33K | 1.10% | NO |
| HONGKONG-1 | Hong Kong | 5/24/2008 | 646 | 1229K | 10K | 6K | 3K | 0.30% | YES |
| IRELAND-1 | Ireland | 5/28/2008 | 642 | 11660K | 35K | 18K | 94K | 0.80% | YES |
| ITALY-1 | Italy | 11/19/2008 | 467 | 1076K | 4K | 3K | 2K | 0.20% | NO |
| ITALY-2 | Italy | 12/5/2008 | 451 | 1622K | 5K | 4K | 7K | 0.40% | YES |
| MEXICO-1 | Mexico | 5/15/2009 | 290 | 600K | 4K | 2K | 39K | 6.50% | YES |
| NETHERLANDS-1 | Netherlands | 5/2/2009 | 303 | 1485K | 10K | 8K | 4K | 0.30% | YES |
| NEWZEALAND-1 | New Zealand | 6/17/2008 | 622 | 9528K | 39K | 12K | 295K | 3.10% | NO |
| RUSSIA-1 | Russia | 2/11/2009 | 383 | 13765K | 120K | 30K | 308K | 2.20% | NO |
| SINGAPORE-1 | Singapore | 3/20/2009 | 346 | 514K | 2K | 2K | 1K | 0.10% | YES |
| SPAIN-1 | Spain | 6/27/2008 | 612 | 3017K | 11K | 5K | 28K | 0.90% | YES |
| TURKEY-1 | Turkey | 6/4/2008 | 635 | 8889K | 30K | 13K | 55K | 0.60% | YES |
| UK-1 | UK | 5/7/2008 | 663 | 8124K | 24K | 13K | 152K | 1.90% | YES |
| UK-2 | UK | 6/27/2008 | 612 | 3442K | 16K | 5K | 25K | 0.70% | NO |
| UK-3 | UK | 2/17/2009 | 377 | 494K | 6K | 4K | 5K | 1.00% | YES |
| UK-4 | UK | 10/5/2008 | 512 | 2774K | 7K | 6K | 20K | 0.70% | NO |
| UK-5 | UK | 6/18/2008 | 621 | 433K | 4K | 3K | 1K | 0.30% | NO |
| URUGUAY-1 | Uruguay | 10/23/2007 | 860 | 12297K | 46K | 10K | 79K | 0.60% | YES |
| US-1 | US | 4/11/2009 | 324 | 13484K | 57K | 35K | 486K | 3.60% | NO |
| US-2 | US | 5/6/2008 | 664 | 6309K | 14K | 10K | 35K | 0.60% | YES |
| US-3 | US | 5/8/2008 | 662 | 11868K | 29K | 15K | 262K | 2.20% | YES |
| VENEZUELA-1 | Venezuela | 5/16/2008 | 654 | 10292K | 20K | 13K | 49K | 0.50% | NO |
| Mean | | | 571 | 5236K | 20K | 8K | 80K | 2.9% | |
| Median | | | 612 | 3292K | 11K | 6K | 33K | 0.8% | |

**FIGURE 3. Distribution of Price Changes between −5 and 5 percent**

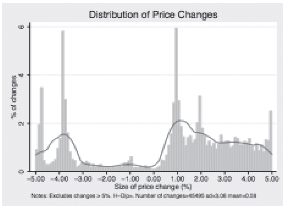


(a) ARGENTINA-1

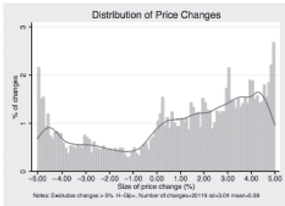


(b) ARGENTINA-2

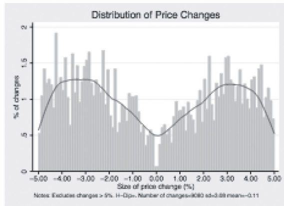


(c) AUSTRALIA-1


(d) AUSTRALIA-2

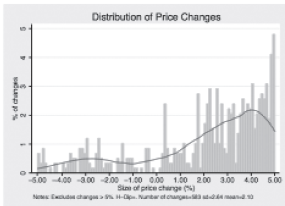


(e) AUSTRALIA-3

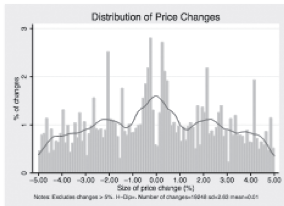


(f) AUSTRALIA-4

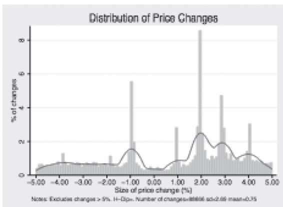


(g) BRAZIL-1

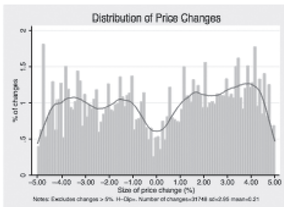


(h) CHILE-1

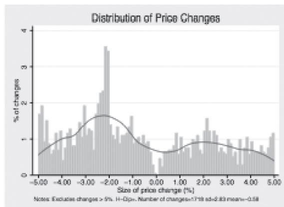


(i) CHINA-1

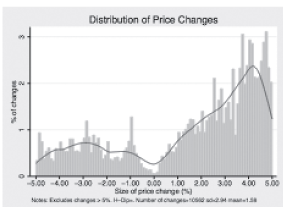


(j) CHINA-2

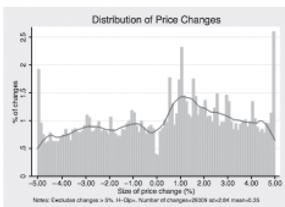


(k) COLOMBIA-1

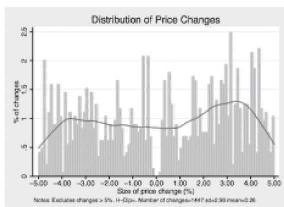


(l) ECUADOR-1

**FIGURE 3. Distribution of Price Changes between −5 and 5 percent (*Continued*)**

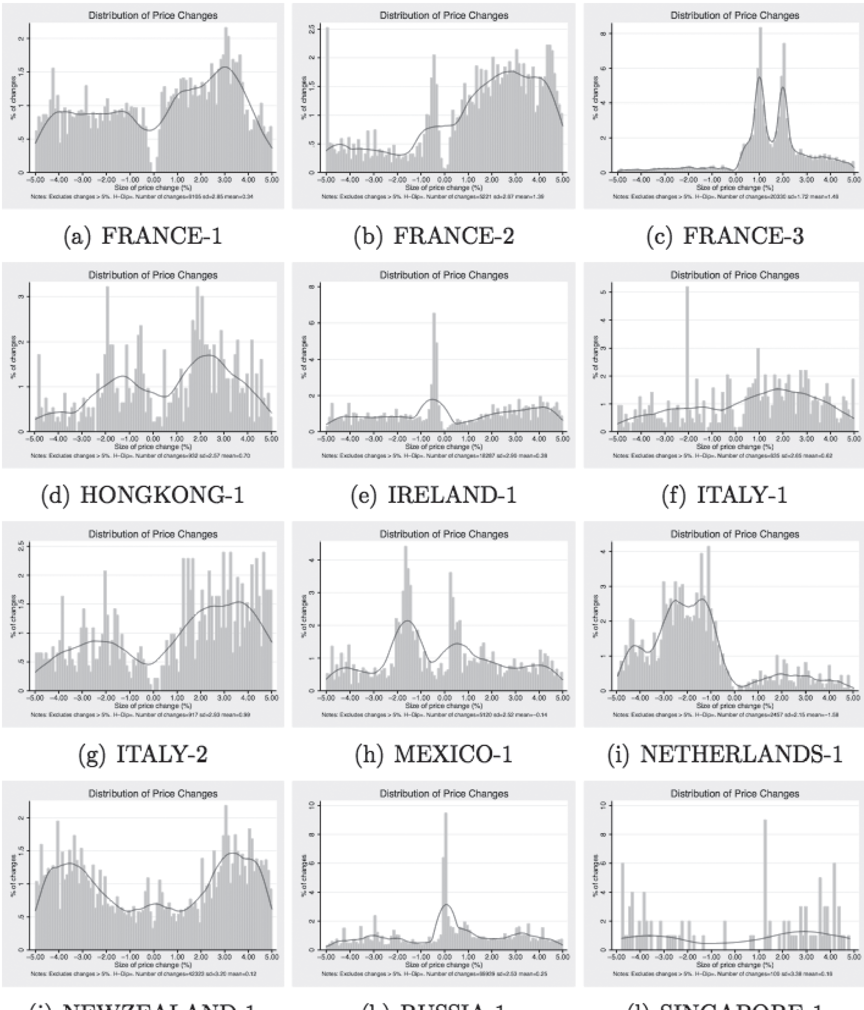


(a) FRANCE-1 (b) FRANCE-2 (c) FRANCE-3

(d) HONGKONG-1 (e) IRELAND-1 (f) ITALY-1

(g) ITALY-2 (h) MEXICO-1 (i) NETHERLANDS-1

(j) NEWZEALAND-1 (k) RUSSIA-1 (l) SINGAPORE-1

**FIGURE 3. Distribution of Price Changes between −5 and 5 percent (*Continued*)**

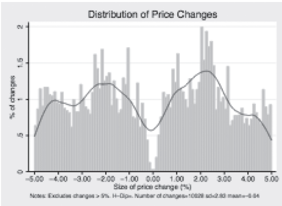


(a) SPAIN-1

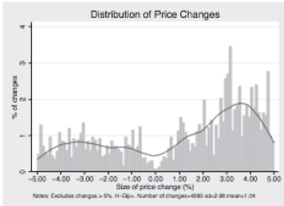


(b) TURKEY-1


(c) UK-1


(d) UK-2


(e) UK-3


(f) UK-4

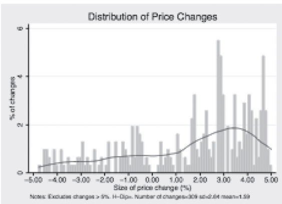


(g) UK-5


(h) URUGUAY-1


(i) US-1

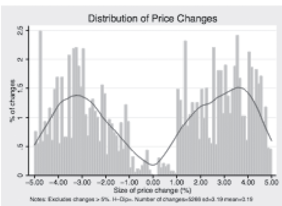


(j) US-2


(k) US-3

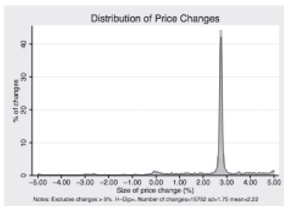


(l) VENEZUELA-1

**TABLE 3. The Effect of Using Daily versus Weekly Data on Test Scores**

| *Test* | *Daily data* | *Weekly average* |
|---|---|---|
| Mean dip (Hartigan) | 0.035 | 0.019 |
| Mean critical bandwidth (Silverman) | 1.351 | 0.799 |

Since we have the posted prices in our data, we ran a very simple exercise in which we computed the average weekly price and retested. As can be seen in table 3, the scores of all the tests drop dramatically when a simple equal weight average is computed. The typical rejection of the Silverman test drops from 1.351 to 0.799. The impact of random weights will clearly produce even more unimodal distributions. Further research is needed.

## Final Remarks

The availability of micro-CPI and scanner data allowed the profession to renew its interest in the pricing dynamics literature and, more importantly, to study the macroeconomic consequences of such behavior. Although we have learned tremendously from this experience, the veracity and integrity of those data sources were never questioned. CPI prices are collected, after all, very carefully. These are extremely high quality data. However, the objective of the data collection process is the computation of inflation. I have never heard a single person at the BLS say their objective is to collect prices so researchers can study price dynamics. Yet what could constitute a very good data point for inflation could be a bad data point for studying pricing dynamics.

Online data have allowed us to evaluate the integrity of those data sources for the purpose of the research questions. We find very different results—not necessarily because the data are strange, but because the data were collected with the purpose of evaluating pricing dynamics.

The low cost of collecting online prices and online information in general is reducing our reliance on statistical offices. It is also likely to change the way statistical offices collect their information. The data are extremely high quality, but as with other sources, they are not perfect. While these collection procedures will allow more researchers to access very detailed data and to tailor the data to the question they have at hand, it is important to remember the question of representativeness. As we move through this decade and into the next, I can imagine a proliferation of data sources and the establishment of a depository where the data can be shared. I hope that LACEA becomes an integral part of that process, and initiatives like VoX LACEA are solid first steps.

## References

Alvarez, Fernando E., Francisco Lippi, and Luigi Paciello. 2010. "Optimal Price Setting with Observation and Menu Costs." Working Paper 15852. Cambridge, Mass.: National Bureau of Economic Research.

Burstein, Ariel, and Gita Gopinath. 2013. "International Prices and Exchange Rates." Working Paper 18829. Cambridge, Mass.: National Bureau of Economic Research.

Cavallo, Alberto. 2012. "Scraped Data and Sticky Prices." Working Paper 4976-12. Massachusetts Institute of Technology, Sloan School of Management.

Cavallo, Alberto, Brent Neiman, and Roberto Rigobon. 2015. "Currency Unions, Product Introductions, and the Real Exchange Rate." *Quarterly Journal of Economics* (forthcoming).

Cavallo, Alberto, and Roberto Rigobon. 2011. "The Distribution of the Size of Price Changes." Working Paper 16760. Cambridge, Mass.: National Bureau of Economic Research.

Crucini, Mario J., Chris I. Telmer, and Marios Zachariadis. 2005. "Understanding European Real Exchange Rates." *American Economic Review* 95(3): 724–38.

Eichenbaum, Marty, and others. 2014. "How Frequent Are Small Price Changes?" *American Economic Journal: Macroeconomics* 6(2): 137–55.

Engel, Charles. 1999. "Accounting for U.S. Real Exchange Rate Changes." *Journal of Political Economy* 107(3): 507–38.

Gopinath, Gita, and others. 2011. "International Prices, Costs, and Markup Differences." *American Economic Review* 101(6): 2450–86.

Klenow, Peter J., and Oleskiy Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 73(3): 863–903.

Klenow, P. J., and Benjamin A. Malin. 2011. "Microeconomic Evidence on Price-Setting." In *Handbook of Monetary Economics,* vol. 3, edited by Benjamin M. Friedman and Michael Woodford, chap. 6, pp. 231–84. Amsterdam: Elsevier.

Midrigan, Virgiliu. 2011. "Menu Costs, Multiproduct Firms, and Aggregate Fluctuations." *Econometrica* 79(4): 1139–80.

Whish, Richard, and David Bailey. 2012. *Competition Law.* 7th ed. Oxford University Press.