Constructing Location-Specific Price Indexes from Scanner Data

Lun Li *

September 16, 2019

Abstract

Based on item level micro-data collected from around 35,000 U.S. supermarkets from 2006 to 2016, I compare different approaches to construct location-specific price indexes from scanner data sets, and propose a method that generates unbiased, free of chain drift indexes while staying computationally feasible. I show that the resulting indexes are consistent with official statistics released by the U.S. Bureau of Labor Statistics, and discuss ways to treat missing data in the process. My results show significant regional variations in inflation rates across U.S. locations from 2006 to 2016: cities have lower inflation rates than rural areas. This finding potentially raises new questions on the effect of central banks monetary policies on regional inflation rates.

*Researcher(s) own analyses calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher(s) and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. This work was completed in part with resources provided by the University of Chicago Research Computing Center. All errors are my own. Email: lunl@uchicago.edu.

1 Introduction

For decades, economists use the consumer price index (CPI) released by Bureau of Labor Statistics (BLS) as measures of prices and inflation rates. Recently, however, many researchers start to pay attention to alternative measures of prices, usually constructed from supermarket scanner datasets. For instance, Stroebel and Vavra (2016) construct quarterly retail price indexes at MSA and zip code level, and study the causal response of local house prices on retail prices. Argentee and Lee (2017) calculate cost-of-living indexes for each specific income group. Kaplan and Schulhofer-Wohl (2016) estimate inflation rate at household level. Beraja et al.(2016) compute regional inflation rates and chained price indexes at monthly frequency, and studied how regional and aggregate fluctuations are related. These papers are just a few examples of the expanding literature where authors uses scanner data to construct their desired price measures, instead of using the published CPI figures directly.

But what are the advantages of constructing price measures using scanner data? In fact, there are at least three reasons why these new measures are superior to traditional ones. First, price measures from scanner data provide finer granularity in terms of location and frequency. BLS only releases price indexes at national or MSA level, and often at quarterly or bi-monthly frequency only. The supermarket scanner data, on the contrary, allows researchers to construct price measures at state, county, zip code or even store level, and also with higher frequency such as monthly or even weekly. Second, scanner price measures can easily connect with household characteristics such as income and education. This important connection can help researchers better understand correlations that are otherwise hidden by traditional measures, such as the heterogeneity of household inflation rates across different income groups. Third, scanner price measures allow more frequent updating of the representative basket of goods. The base period of traditional CPI measures are often set years ahead, if not decades, and many goods in the basket already become outdated when the latest CPI figures are released. Price measures computed from scanner data, however, update the basket of goods frequently, therefore taking into account of new goods that are introduced in each period.

Of course, prices measures from scanner datasets have their own limitations, too. For instance, the official CPI measures include expenditures in more categories than scanner datasets, while the latter usually only cover food, beverages, baby products, alcohol, OTC drugs, and other types of goods commonly found in supermarket chains. Stroebel and Vavra (2016) documented that the IRI Symphony Retailer Scanner dataset only represent roughly 15% of the household spending in the Consumer Expenditure Survey. Another empirical difficulty about scanner price indexes is that they are computationally demanding and more sensitive to errors than traditional measures. In particular, chained price indexes from scanner data often exhibit the so-called ”chain drift” problem, if not being handled carefully. Chain drift typically occurs when prices and quantities oscillate from one period to the next, causing Laspreyes and Paasche price indexes to display unrealistic price changes over time. Studies including Haan and van der Grient (2010) and Diewert and Fox (2017) show that households stocking up goods during sales is among the main causes for chain drift problems. Also, products that display strong seasonality in sales can aggrevate the chain drift problem. To solve this issue, Ivancic, Diewert and Fox (2010) propose the multilateral GEKS index, which is a geometric mean of all bilateral indexes between periods[1]. According to Diewert and Fox (2010), the GEKS index displays no significant signs of chain drift. However, it is difficult to construct,

[1]The name of GEKS index comes from Gini (1931), Elteto and Koves (1964), and Szulc (1964).

especially for large datasets such as the Nielsen Retailer Scanner Dataset. Other researchers take more conventional approaches. Beraja et al. (2016), for example, construct chained price indexes by using average quantity sold for each good in the previous year as weights, thus eliminating chain drift problems.

The goal of this paper is to provide a practical guide for researchers to construct location-specific price indexes on their own. The rest of this paper is structured as follows. Section 2 reviews the literature. Section 3 describes the dataset used in this project. Section 4 introduces the approach to construct price indexes from scanner data. Section 5 presents the results. Section 6 discusses the implications of our results, as well as potential areas for future work. Section 7 concludes.

2 Literature Review

Broadly speaking, this paper is related to two strands of economic literature that deals with inflation and price indexes. The first strand addresses the observation that inflation, as well as consumer price index (CPI), variate greatly across heterogeneous households. Michael (1979) is one of the first studies on this subject[2]. The author documents that dispersion of price changes across household exists (and is quite considerable), but the dispersion cannot be fully explained by household characteristics alone. More specifically, he find that from 1967 to 1974, the poor generally experience higher levels of inflation than the rich. But he also points out that the observed pattern does not seem stable over

[2] In fact, the idea that purchasing power cannot be measured exactly across different households can be traced back as far as Alfred Marshall, who in a 1926 article wrote that "The same change of prices affects the purchasing power of money to different persons in different ways." He then discribes how a fall in price of meat accompanied by a rise in price of bread means a lower purchasing power for the poor but vice versa for the rich.

time[3]. More recently, as computational power has become more readily available, studies such as Handbury et al. (2014), Kaplan and Schulhofer-Wohl (2017) and Argentee and Lee (2017) start to use supermarket scanner data to demonstrate the limitations of traditional aggregate CPI as a measure of consumer price changes, especially when households are heterogeneous in nature. Handbury et al. (2014), for example, claims that aggregate CPI is an imperfect indicator of inflation. From Japanese supermarket scanner data, the authors construct a ″true″ inflation measure using a Tornqvist-type price index, and compares that index with official CPI figures released by Japanese statistical agencies. They conclude that CPI is not an ideal indicator of true inflation, especially when measured inflation is low. Kaplan and Schulhofer-Wohl (2017), in their groundbreaking work, construct household level inflation rates using consumer panel data from Nielsen, and document heterogeneity in inflation rates across households with different characteristics. Their methodology differs with Michael (1979) and several other earlier studies, by relaxing the assumption that different households pay the same price for the same good and differ only in quantities of good that they purchased. Similar to Michael (1979), the authors find that poor households have higher inflation rates than richer ones. In addition, they find that most of the variances in household-level inflation rate comes from heterogeneous shocks rather than aggregate inflation – suggesting that aggregate CPI may not be a satisfactory measure to understand each household's consumption behavior.

The second strand of literature, mostly written by econometricians or statisticians, studies the ideal method to construct price indexes based on scanner data. Ivancic et al.(2010) and Haan and van der Grient (2010) are two impor-

[3]The paper also compares price changes between regions of residence (North East, North Central, South, West) but finds no significant differences.

tant studies on this matter, among many others. Because scanner data usually comes with higher frequency (e.g. daily or weekly, rather than monthly) and finer granularity (item or barcode level) than desired, aggregation is often required before an index can be computed. The standard practice adopted by BLS is to first aggregate item-level price changes into product stratum level for each area, and then take weighted average across all item stratums and areas to calculate the aggregate CPI[4]. However, as Ivancic et al.(2010) points out, aggregation choices, which usually seem innocuous when computing price indexes, can have significant effects on the final statistic. Depending on whether indexes are chained using quarterly, monthly or weekly frequency, and whether a superlative index (i.e. Tornqvist or Fisher, instead of Laspeyres or Paasche) is used, the measurement error in chained indexes can range from 0.28% to an incredible 46463.71%. This surprising result originates from the so-called "chain drift" problem, which will be discussed in greater detail in Section 4. To solve the "chain drift" problem, both Ivancic et al.(2010) and Haan and van der Grient (2010) recommend multilateral indexes (known as the GEKS index) which theoretically eliminates the problem while making maximum use of all possible observations in the scanner datasets. However, due to its complexity, studies in empirical macroeconomics often adopt their own unique ways to eliminate chain drift, if this problem is mentioned at all in the first place. For example, Beraja et al. (2016) construct monthly chained price indexes by using average monthly sales as weights, which in principle hold the consumption bundle fixed for all months in the same year, and only update this weight once each year. Argente and Lee (2017) construct income-group-specific price indexes using both

[4]The precise methodology to calculate aggregate CPI. as adopted by BLS, will be described in full detail in Section 4.

the conventional exact price index (CEPI)[5] and also an index that considers product creation and destruction, as in Broda and Weinstein (2010). Stroebel and Vavra (2016) follows a similar two-step aggregation method as BLS does in constructing their price indexes, but chained the indexes annually rather than monthly in order to avoid chain drift.

This paper contributes to the current literature in the following three ways. First, it summarizes an approach to construct price indexes from scanner datasets that is unbiased while computationally feasible. Second, this paper compares various ways to treat missing data when constructing price indexes, a problem often neglected by previous studies but nonetheless worth discussing. Third, this paper provides new evidence on the regional variations of price levels and inflation rates across the U.S., and raises new questions on the effect of monetary policies on inflation.

3 Data

The primary dataset I use in this paper is the Kilts-Nielsen Retailer Scanner Dataset (KNRS), provided by the Kilts Center of Marketing in University of Chicago, Booth School of Business. The dataset contains information on weekly price and sales data from approximately 35000 participating retail stores, and covers a time span from 2006 to 2015. In general, the database consists of over 12 billion observations each year, for over 70 million unique products (identified by 12 digit Universal Product Codes, or UPCs) in around 1000 product categories. Each observation contains weekly prices and sales quantity for each product, as well as detailed information on sales conditions (such as whether a coupon is used in transaction, for example). Overall, the entire dataset contains

[5] See Sato(1976) and Vartia(1976).

100 - 150 billion unique observations, in 49 states (excluding Alaska and Hawaii) and across over 2500 counties. Around 95 percent of the sales in the data comes from food, drug and mass merchandizing stores (Argentee and Lee (2017)), and the datatset represents around 30 percent of total U.S. expenditure on food and beverages (Beraja et al.(2016)). In a nutshell, the KNRS dataset provides a comprehensive and detailed overview of the U.S. retail market, and can be used as a representative sample for constructing location-specific price measures.

Table 1 shows some summary statistics of the KNRS dataset, from 2006 to 2013 [6]. The dataset contains around 10 billion unique observations each year, for a total of over 100 billion data points combined. It includes information on over 1 million unique barcodes (UPCs) in total, which belongs to over 1000 narrowly defined product modules. The dataset includes sales information in 49 states (except for Hawaii and Alaska), more than 2000 counties, and over 80 retail chains and 40 thousand stores. The vast amount of observation and detailed geographical information makes KNRS the ideal dataset to construct location-specific price indexes. Unlike the commonly studied Nielsen Homescan database that samples approximately 33000 households[7], the KNRS dataset collects weekly sales data from all participating stores in the U.S., and therefore provides a more complete landscape of the grocery retail market.

4 Methodology

In this section, I summarize the approach to construct location-specific price indexes from scanner datasets, as used by previous studies in the literature.

[6] This table is from the unpublished manuscript by Dube, Hurst, Kim and Ospina (2016).

[7] For example, see Handbury and Weinstein (2015).

4.1 Two-Step Aggregation Method

The two-step aggregation method is the most commonly used approach to construct price indexes from scanner datasets, as used by papers such as Stroebel and Vavra (2016), Beraja et al. (2016), and Kaplan and Schulhofer-Wohl (2017). This method is widely adopted the empirical macroeconomics literature, partly because it closely resembles the way BLS constructs the urban consumer chained CPI series,therefore providing a convenient benchmark to compare with.

Before beginning the "two-step" aggregation, an important preliminary step is to compute the average sales price of each item in each loacation[8], over a specific time period of our choice. In most scanner datasets, barcode-level price and quantity information is reported each week. The preliminary step to construct price indexes is aggregating weekly observations into monthly or quarterly averages:

$$p_{i,j,a,t} = \frac{\sum_{w \in t} p_{i,j,a,w} q_{i,j,a,w}}{\sum_{w \in t} q_{i,j,a,w}} \tag{1}$$

Where j stands for the product category of good i, a stands for sample area a, and w represent weeks within the desired time period t.

After the preliminary step above, the **first step** is to aggregate item-level

[8]Richardson (2000) define this entity as the "unit value". If a store has two different prices on an item within the same week, the unit value, or the weighted average of these two prices, will be reported for that week. Unit values are used here to combine weekly observations into monthly or longer time intervals, but as Richardson (2000) points out, unit values are not equivalent to prices because often no one actually pays the unit value exactly. Without a doubt, it is important to discuss whether unit values should be treated as prices in construction of indexes; however, for the purpose of this paper, I will use "unit values" and "prices" interchangeably.

price changes into product-module level:

$$\frac{p^L_{j,a,t}}{p^L_{j,a,t-1}} = \frac{\sum_i p_{i,j,a,t} q_{i,j,a,t-1}}{\sum_i p_{i,j,a,t-1} q_{i,j,a,t-1}} \quad \text{(Laspeyres)} \tag{2}$$

$$\frac{p^P_{j,a,t}}{p^P_{j,a,t-1}} = \frac{\sum_i p_{i,j,a,t} q_{i,j,a,t}}{\sum_i p_{i,j,a,t-1} q_{i,j,a,t}} \quad \text{(Paasche)} \tag{3}$$

The **second step** is to take average over all product categories for each location, to calculate the location-specific inflation rate:

$$\frac{P_{a,t}}{P_{a,t-1}} = \prod_{j=1}^{N} \left(\frac{p_{j,a,t}}{p_{j,a,t-1}} \right)^{\frac{s_{j,a,t}+s_{j,a,t-1}}{2}} \tag{4}$$

Where N is the total number of product categories, and $s_{j,a,t}$ is the revenue share of product category j in area a in time t:

$$s_{j,a,t} = \frac{TS_{j,a,t}}{\sum_{j=1}^{N} TS_{j,a,t}}$$

After choosing a base period 0, we can set $P_{a,0} = 1$ and use Equation (4) to construct a chained price index for each location a, from time 0 to any time period t, as long as data is available.

There are several technical details that need to be clarified here, before moving on to the next section. First, after constructing location-specific price index for a smaller location a, we can find aggregate price index for a larger location l such that $a \in l$ by taking geometric weighted averages in ways similar to Equation (4). Alternatively, we can complete the geographical aggregation in the preliminary step, by taking average prices and total quantities over the larger location l directly:

$$p_{i,j,l,t} = \frac{\sum_{w \in t, a \in l} p_{i,j,a,w} q_{i,j,a,w}}{\sum_{w \in t, a \in l} q_{i,j,a,w}} \tag{5}$$

Second, the method above is sometime modified in order to remain computationally tractable or to avoid "chain-drift". For example, Stroebel and Vavra (2016) skip the step of time aggregation in the preliminary stage, and simply take observations in the last week of each quarter to represent data in that quarter. They also do not use the Tornqvist-type index in the second step, as in Equation (4), but instead use revenue shares $s_{j,s,t}$ directly as weights:

$$\frac{P_{a,t}}{P_{a,t-1}} = \prod_{j=1}^{N} \left(\frac{p_{j,a,t}}{p_{j,a,t-1}} \right)^{s_{j,a,t}} \tag{6}$$

Beraja et al. (2016) takes the Laspeyres formula in their construction of price indexes. However, different from Equation (2), the authors use average monthly quantity $\bar{q}_{i,j,a,t}$ instead of actual sales $q_{i,j,a,t}$ in the first step of calculation, holding consumption bundle *fixed* each month and only update the basket once every year:

$$\frac{p_{j,a,t}}{p_{j,a,t-1}} = \frac{\sum_i p_{i,j,a,t} \bar{q}_{i,j,a,t-1}}{\sum_i p_{i,j,a,t-1} \bar{q}_{i,j,a,t-1}} \tag{7}$$

Kaplan and Schulhofer-Wohl (2017) construct both Laspeyres and Paasche indexes in the first step, and also compute the Fisher Index, defined as:

$$\frac{p^F_{j,a,t}}{p^F_{j,a,t-1}} = \sqrt{\frac{p^L_{j,a,t}}{p^L_{j,a,t-1}} \frac{p^P_{j,a,t}}{p^P_{j,a,t-1}}} \tag{8}$$

In this paper, I follow a similar methodology as in Beraja et al. (2016), and construct price indexes at national, state, MSA (Metropolitan Statistical Area) and county level.

5 Results

In this section, I first benchmark my price indexes with official CPI figures released by BLS, and then show summary statistics of monthly location-specific indexes that are at U.S. county level, a finer granularity than provided by U.S. statistical agencies.

5.1 Benchmark: U.S. Chained CPI

Figure 1 compares the scanner price index with official chained CPI for "food at home", at the national level[9]. Both indexes are normalized such that the price level in Jan 2006 equals to 100. From Figure 1, we see that the two series almost coincide, especially from 2006 to around 2015; from late 2015 to 2016, the two price indexes differ slightly, but are nonetheless highly correlated. Overall, the two indexes have a correlation coefficient of 0.995.

5.2 Benchmark: Local CPI in Metropolitan Areas

Figure 2 shows the scanner price indexes for 13 metropolitan statistical areas (MSAs) where official monthly CPI statistics are available. Our index performs reletively well in MSAs such as Boston, Miami, New York and Washington D.C.. However, in many other cities such as Atlanta, Chicago, Cleveland, Dallas and Houston, our scanner price index suggests lower growth rates in price levels than official statistics do.

There are 3 potential reasons why this discrepancy may occur for some MSAs. First, our scanner dataset updates the number of stores and the collection of

[9]When constructing this national scanner price index, I excluded goods that do not belong to "food at home", such as health and beauty products, non-food grocery and general merchandise. The goal is to make sure we are comparing apples to apples.

goods in each store every year, while BLS updates the sample of goods less frequently. This could mean that the two indexes are comparing price increases of different set of goods, even though both are broadly defined as "food at home". Second, the definition of MSAs also changes over time. For example, the Chicago metropolitan area changes its name from "Chicago-Gary-Kenosha, IL-IN-WI" in 1998 to "Chicago-Naperville-Elgin, IL-IN-WI" in 2018, and includes two more counties, Jasper and Newton, in Indiana. These changes could potentially cause the local scanner index and official statistics to represent different geographical areas. In this paper, I define MSAs with the latest definition as published by the Census Bureau, a definition consistent with that used by BLS. However, it is unclear whether past data from BLS uses the old definition, causing a difference between the two indexes. Finally, the problem of missing data in scanner dataset could cause unreliable upward or downward drifts of price indexes, especially when the level of aggregation is not large enough. For example, if sales records of a product are only found in January and March (but not in February), this particular product is dropped from the monthly price indexes for all three months, during the second step of aggregation. In other words, a missing price in one month affects price indexes in all adjacent months. This problem is more likely to occur when the initial level of aggregation is smaller, such as in city or county level.

To address the last issues described above, I construct two alternative price indexes in which I impute the missing prices in the dataset. Using the previous example, we can impute the missing price in February by either 1) carrying forward the price in January or 2) taking the average price of January and March as the price in Februray. From Figure 2, we see that price indexes in which missing data is imputed do not differ greatly from our benchmark indexes. For a dataset of this size and scale, imputing missing data in the construction of price indexes is not an easy task: it is both computationally demanding and time-consuming. The result in this section suggests that dropping observations with missing prices do not affect the overall price index significantly.

5.3 Nielsen Price Index at County Level

One of the main contributions of this paper is to provide monthly scanner price indexes at U.S. county level, between 2006 to 2016. Figure 3 shows the average annual inflation rate for each county during the sample period. The first observation is that most counties have a mean inflation rate between 2-3 percent during this period. Also, urban areas seem to have lower inflation rate than rural areas, a result that can be partially explained by the higher income levels of urban residents[10].

Figure 4 shows the distribution of average annual inflation rates across U.S. counties. I also included the distribution of counties that belong to one of the 13 MSAs studied earlier. We observe that for all U.S. counties, the distribution centers around 2%, suggesting that the Federal Reserve is successfully targeting inflation around the 2% benchmark. However, we also find some regional variations in the inflation rate across counties, with a standard deviation of 0.72 percentage points. On the contrary, the distribution of urban counties' inflation rate have a mean of 1.66% and a standard deviation of only 0.35%. This heterogeneity in regional inflation rates may pose new challenges for the central bank: monetary policies that target a 2% national inflation rate could have heterogeneous effects across different U.S. locations.

[10] According to Argente and Lee (2017) and Kaplan Schulhofer-Wohl (2016), richer households have lower inflation rates than poorer ones.

6 Discussion

From the price index I construct from supermarket scanner datasets, there are several important features of price levels and inflation rates that deserve some further discussion here.

6.1 Regional Variation in Inflation Rates

From the result in Figure 3, we observe some significant regional variation in inflation rates across U.S. counties. Furthermore, Figure 4 seems to provide evidence that cities have *lower* inflation rates than rurual areas. What could possibly explain this pattern? Besides household characteristics such as income, wealth and preferences, there are at least two other channels working against each other that cause the regional variation in inflation rates across the U.S.. First, consumers in urban areas tend to have less time to search for bargains, and therefore inducing grocery stores to charge higher prices for each product. However, it is also true that the number of grocery stores accessible to each household is also greater in urban areas than in rural ones. This reduces the cost of urban consumers to search for cheaper bargains, which disincentivize grocery stores to charge higher prices in its products. The fact that cities have lower inflation rates provide evidence that the second channel is dominating. Of course, there could be many other reasons why inflation variation occurs, such as differences in local population, wage, housing prices, or even exposure to international trades. Whatever the true reason might be, this fact calls for more research in the future.

6.2 Dynamics of Inflation Rates

Figure 5 shows how distribution of county level inflation rates changes across time. The average inflation rates peak at 2008, hitting a level of more than 6 percent, and falls to less than 1 percent during 2009. The standard error of inflation rates is usually quite large (around 2 percentage points) and remain rather stable across time. A more interesting question is to explore whether the relative position of each county's inflation rate in the whole distribution remains constant over time. A closer look at our data suggests that a county with below-average inflation rate in a certain year is more likely to have above-average inflation in the following year. Therefore, the existence of regional variation in inflation rates may suggest that pricing decisions are made in a rather decentralized manner. Rather than observing a simultaneous increase in price levels of all U.S. counties, we find that some locations raise prices first and others follows suit later. In other words, an aggregate shock that is imposed on all counties simultaneously, such as a new monetary policy, cannot fully account for the heterogeneous behaviors of inflation rate changes across the nation. Other factors, such as local demand shocks or labor market conditions, may account for the disproportional changes in inflation rates. Maybe we should be convinced that inflation is *not* always and everywhere a monetary phenomenon, as claimed by Milton Friedman – other forces may also have non-negligible effects.

7 Conclusion

In this paper, I construct location-specific price indexes from supermarket scanner datasets, and examined their behavior at national, city and county level. The study shows considerable variaton in county level inflation rates, both cross sectional and across time, which cannot be explained solely by aggregate monetary shocks.

Appendix

A Graphs and Tables

	Individual Years								Combined	
	2006	2007	2008	2009	2010	2011	2012	2013	Total	Average
Number of Obs. (Million)	12013.1	12812.2	13037.5	12968.3	13153.4	13646.7	13618.8	13801.3	105051.0	13131.4
Number of UPCs	725224	762469	759989	753984	739768	742074	753318	769136	1487003	750745
Number of Product Modules	1085	1086	1086	1083	1085	1081	1105	1113	1113	1091
Number of Retail Chains	86	85	87	86	86	86	82	79	88	85
Number of Stores	32642	33745	34830	35343	35807	35645	36059	36316	40350	35048
Number of Zip Codes	10869	11123	11357	11476	11589	11639	11626	11553	11797	11404
Number of Counties	2385	2468	2500	2508	2519	2526	2547	2561	2593	2502
Number of States	49	49	49	49	49	49	49	49	49	49
Transaction Value (US Billion)	187.9	207.8	219.6	223.7	227.6	235.2	239.5	238.7	1779.9	222.5

Table 1: Summary Statistics of Kilts-Nielsen Retailer Scanner Dataset

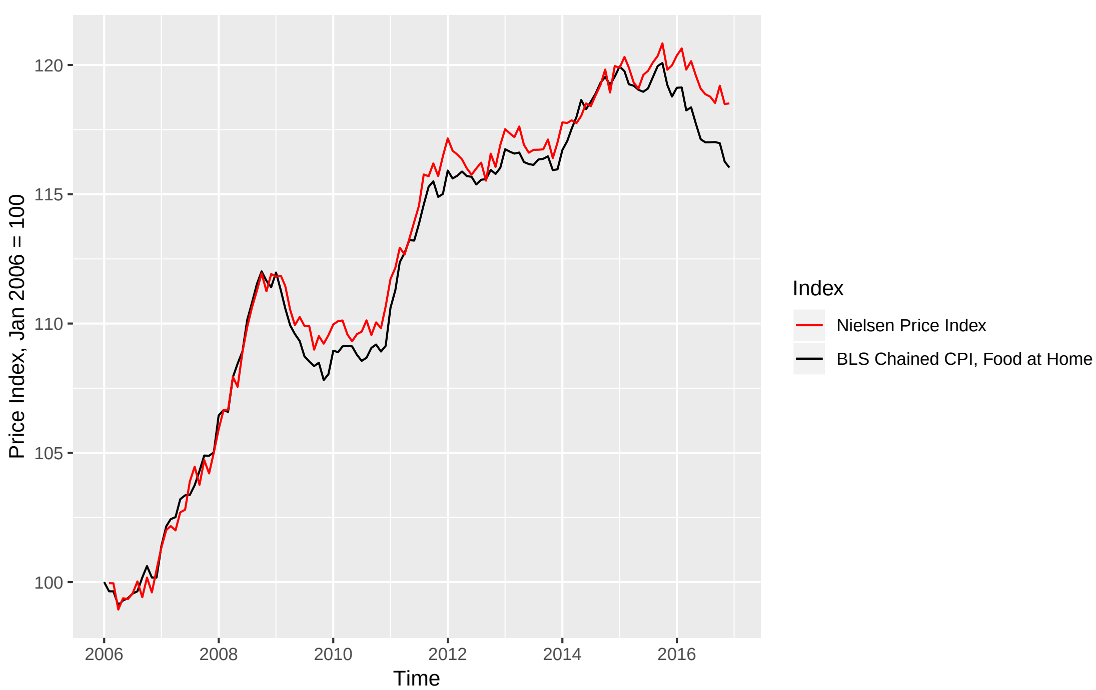

Figure 1: Comparison of Nielsen Price Index (red) with BLS chained price index (black) for food at home (series SUUR0000SAF11), at national level.

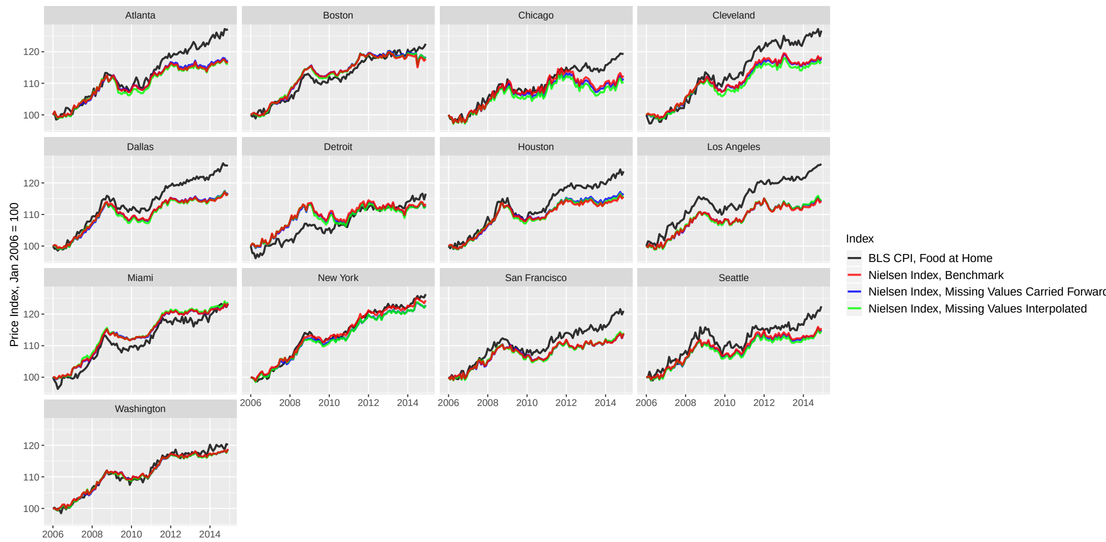

Figure 2: Comparison of Nielsen Price Index (red) with BLS chained price index (black) for food at home, at MSA level. Two indexes that address the missing data problem are also included, using carry-forward method (blue) and linear interpolation method (green).

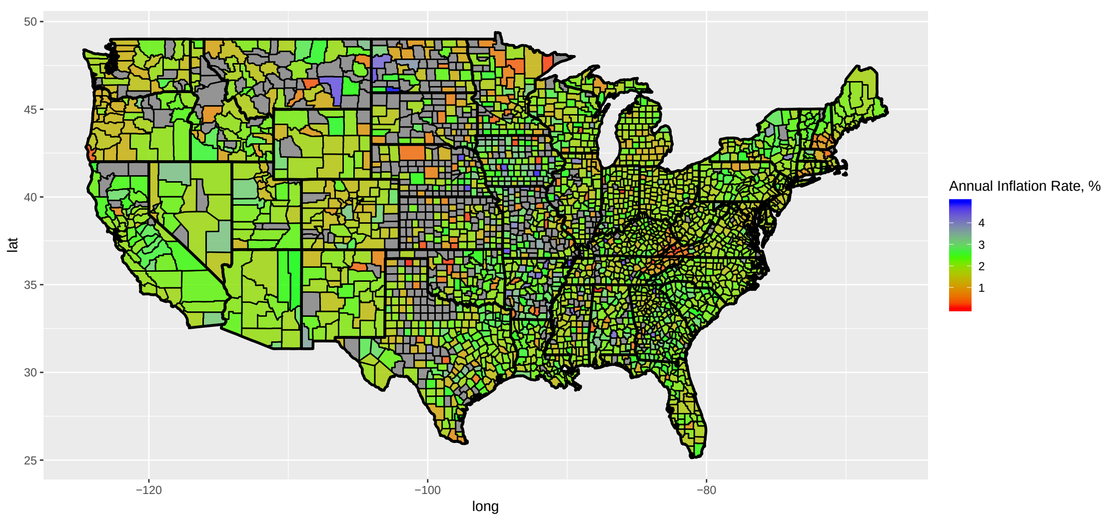

Figure 3: U.S. average annual inflation rate, 2006-2016, by county. Outliers (counties with annual inflation rate larger than 5%) are omitted for the purpose of better graphical presentation. Counties where data is not available are colored grey.

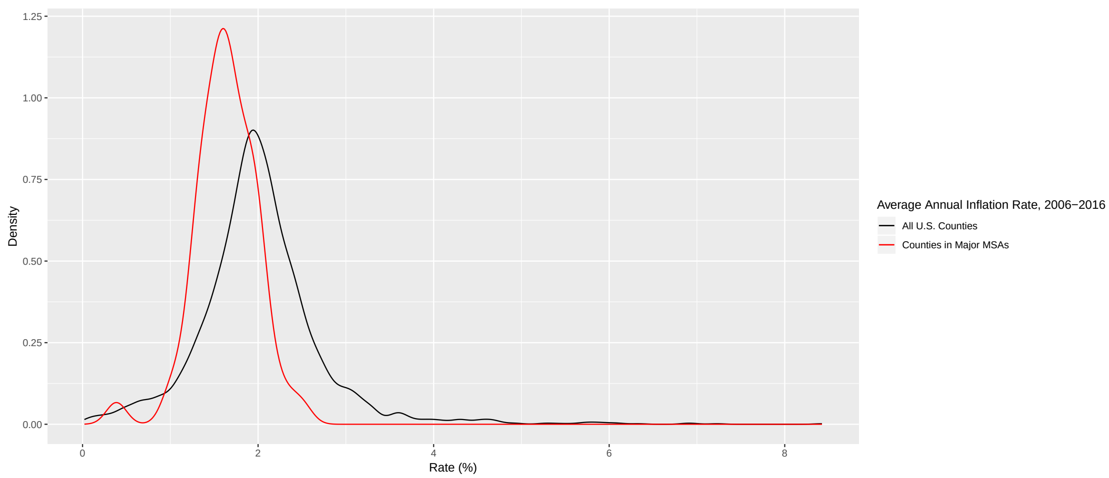

Figure 4: Distribution of average annual inflation rate between 2006 to 2016, for all U.S. counties (black) and counties in major MSAs (red). "Major MSAs" are defined as the 13 cities shown in Figure 2 (Atlanta, Boston, Chicago, Cleveland, Dallas, Detroit, Houston, Los Angeles, Miami, New York, San Francisco, Seattle and Washington D.C.).

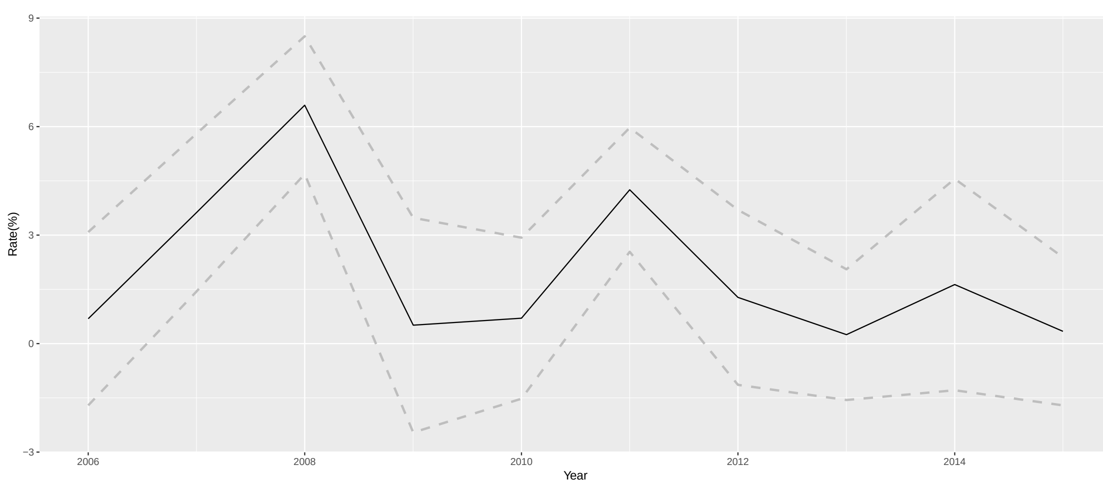

Figure 5: Change of distribution for county level price inflation rates, 2006-2015. Solid black line represents average inflation rates across U.S. counties. Dotted grey line represent the mean rates plus/minus 1 standard error.

References

Argente, David and Munseob Lee (2014) "Cost of Living Inequality during the Great Recession," *SSRN Electronic Journal*, p. 53, URL: http://dx.doi.org/10.2139/ssrn.2567357, DOI: http://dx.doi.org/10.2139/ssrn.2567357.

Beraja, Martin, Erik Hurst, and Juan Ospina (2016) "The Aggregate Implications of Regional Business Cycles," *Ssrn*, DOI: http://dx.doi.org/10.2139/ssrn.2733708.

Broda, Christian and David E. Weinstein (2010) "Product creation and destruction: Evidence and price implications," *American Economic Review*, Vol. 100, pp. 691–723, DOI: http://dx.doi.org/10.1257/aer.100.3.691.

De Haan, Jan and Heymerik A. Van Der Grient (2011) "Eliminating chain drift in price indexes based on scanner data," *Journal of Econometrics*, Vol. 161, pp. 36–46, DOI: http://dx.doi.org/10.1016/j.jeconom.2010.09.004.

DellaVigna, Stefano and Matthew Gentzkow (2017) "Uniform Pricing in U.S. Retail Chains," DOI: http://dx.doi.org/10.2139/ssrn.3067508.

Diewert, W Erwin and Kevin J Fox (2017) "Substitution Bias in Multilateral Methods for CPI Construction using Scanner Data."

Dub, Jean Pierre, Hurst Ruben, and Kim Juan (2016) "Predicting Aggregate Variables with Scanner Data," pp. 1–38.

Feenstra, Robert C. and Matthew D. Shapiro (2003) *Scanner Data and Price Indexes*: University of Chicago Press, URL: http://www.nber.org/books/feen03-1.

Fitzgerald, Jenny and Owen Shoemaker (2013) "Evaluating the Consumer Price Index Using Nielsen's Scanner Data," pp. 1033–1045.

Handbury, J. and D. E. Weinstein (2015) "Goods Prices and Availability in Cities," *The Review of Economic Studies*, Vol. 82, pp. 258–296, URL: https://academic.oup.com/restud/article-lookup/doi/10.1093/restud/rdu033, DOI: http://dx.doi.org/10.1093/restud/rdu033.

Handbury, Jessie, David E Weinstein, and Tsutomu Watanabe (2013) "How Much Do Official Price Indexes Tell Us About Inflation?," *NBER Working Papers*, DOI: http://dx.doi.org/10.3386/w19504.

Ivancic, Lorraine, W. Erwin Diewert, and Kevin J. Fox (2011) "Scanner data, time aggregation and the construction of price indexes," *Journal of Econometrics*, Vol. 161, pp. 24–35, URL: http://dx.doi.org/10.1016/j.jeconom.2010.09.003, DOI: http://dx.doi.org/10.1016/j.jeconom.2010.09.003.

Kaplan, Greg and Sam Schulhofer-Wohl (2017) "Inflation at the household level," *Journal of Monetary Economics*, Vol. 91, pp. 19–38, DOI: http://dx.doi.org/10.1016/j.jmoneco.2017.08.002.

Michael, Robert T. (1979) "Vatiation Across Households in the Rate of Inflation," *Journal of Money, Credit and Banking*, Vol. 11, pp. 32–46.

Reinsdorf, Marshall B. and Alan H. Dorfman (1999) "The Sato-Vartia index and the monotonicity axiom," *Journal of Econometrics*, Vol. 90, pp. 45–61, DOI: http://dx.doi.org/10.1016/S0304-4076(98)00027-X.

Sato, Kazuo (1976) "The Ideal Log-Change Index Number," *The Review of Economics and Statistics*, Vol. 58, p. 223, URL: https://www.jstor.org/stable/1924029?origin=crossref, DOI: http://dx.doi.org/10.2307/1924029.

Stroebel, Johannes and Joseph Vavra (2014) "House Prices, Local Demand, and Retail Prices," *Ssrn*, DOI: http://dx.doi.org/10.2139/ssrn.2500457.

Triplett, Jack E. (2003) *Using Scanner Data in Consumer Price Indexes. Some Neglected Conceptual Considerations*, pp. 151–162: University of Chicago Press, URL: http://www.nber.org/chapters/c9734.

Vartia, Yrjo O. (1976) "Ideal Log-Change Index Numbers," *Scandinavian Journal of Statistics*, Vol. 3, p. 121.