For Online Publication:

Appendix to "The Cyclicality of Sales, Regular, and Effective Prices: Comment"

Authors: Etienne Gagnon, David López-Salido, and Jason Sockin

A.1 Introduction

Prices of items with identical characteristics often differ across outlets, even when outlets are located in the same geographic area. Thus, the price that shoppers pay for a basket of goods and services can be affected by the extent to which shoppers actively seek to take advantage of price differences across outlets. In principle, a true cost of living index should factor in quantity swings in order to account for such outlet substitutions and, more generally, to account for substitution over time, within product categories, and across product categories. In practice, the use of quantities in the construction of official price indexes is an exception rather than the norm because the gathering of quantity information along with prices is often challenging or even impossible. Even if statistical agencies could observe quantities along with item prices, the tracking of substitution across outlets might necessitate, for each market, the gathering of *many* observations of items with identical characteristics to locate where shoppers are making their purchases. Under current data collection methods, such gathering would be onerous.

Coibion, Gorodnichenko, and Hong (2015, henceforth "CGH") aim to measure the size of the bias associated with ignoring swings in item quantities across outlets. Using these authors' terminology, we define "posted price inflation" as the overall change in the prices posted by firms ignoring contemporaneous movements in quantities across outlets. Similarly, we define "effective price inflation" as the overall change in prices taking these movements across outlets into account. To the extent that a bias exists, it can drive a wedge between the two inflation measures.

Like CGH, our investigation uses the IRI marketing database to assess the size of the bias. This database tracks the prices and quantities of identical items sold at multiple retailers for each of the 50 U.S. markets in the sample. The presence of item-level quantities along with prices for several outlets is a major advantage over the micro database used to compute the CPI. However, the product and outlet coverage of the IRI dataset is far more limited than that of the CPI. Importantly for our discussion, and as is the case with official statistics, researchers face a number of practical data issues when creating price indexes, such as dealing with the presence of outliers, missing observations, or item turnover. Ultimately, the quality of the price indexes and the reliability of the econometric analysis depend on how researchers tackle these issues. In this online appendix, we discuss the ways in which the preferred methodology in our paper (henceforth "GLSS") differs from that adopted by CGH and, in our view, leads to the production of more reliable posted and effective price indexes.

A.2 Comparing CGH's and GLSS' stratum-level price indexes

We begin by illustrating what we see as anomalies in CGH's posted and effective price indexes. Figure A1 depicts the difference between posted and effective price indexes for the first market and first six product categories in the IRI sample.[1] The price indexes are constructed by cumulating monthly posted and effective (log) price inflation over the sample period. Their levels are normalized to zero in January 2001. Three series are displayed in each panel. The series labeled "CGH" are based on CGH's original methodology.[2] The series labeled "GLSS" are obtained under our preferred methodology; they include our preferred treatments for outliers, missing observations, and clearance sales, as well as proper time aggregation to the monthly level. The series labeled "CGH no censoring" replace CGH's censoring of price adjustments with our treatment of outliers but otherwise keep all other aspects of their methodology constant.

One of the ways in which CGH's price series are unsatisfactory to us is that their levels frequently suffer from large jumps. The effects of such jumps can be seen, for example, in the middle-left panel in January 2005 and in the middle-right panel in January 2009, when posted prices fell about 10 percent relative to effective prices in a single month, with the divergence in the level of the price series persisting thereafter. Another way in which CGH's price series are unsatisfactory to us is that posted and effective prices often diverge by large amounts. For example, the posted price indexes for blades, carbonated beverages, and cigarettes each grew a cumulative 15 to 20 percent less than the corresponding effective price indexes in the Atlanta market. This divergence is especially worrisome because greater store switching should have had the opposite effect: As the unemployment rate doubled in the latter years of the sample period, posted prices should have risen—not fallen—relative to effective prices. Under CGH's methodology, roughly 60 percent of the "stratums" (that is, of the combinations of a product category and a market) feature a smaller rise in posted prices than in effective prices over the sample period.

Under our preferred methodology, the difference between posted and effective price indexes stays small, at five percent or less over the sample period, for each of the stratums shown in Figure A1. This greater similarity relative to CGH's methodology is also typically found for the other stratums in the sample. Overall, the figure illustrates that our posted and effective price indexes generally do not exhibit the kind of jumps and diverging trends seen for CGH's price series.

[1] The anomalies seen for the first six stratums are illustrative of the anomalies affecting the sample as a whole.

[2] The only exception with respect to CGH's original methodology is that we define UPCs in the way outlined in the IRI documentation. This difference is immaterial for our discussion; see section A.9.

A.3 Comparing CGH's and GLSS' price indexes with the official CPI

If the issues affecting CGH's stratum-level price series were simply adding noise that washed out at the aggregation stage, then CGH's regression results might still provide consistent estimates of cyclical price responses, albeit with wider confidence bands. We argue that this situation is unlikely because CGH's stratum-level posted price series, when aggregated across markets and/or product categories, typically differ much from the price series constructed by the BLS (which are also on a posted rather than an effective basis). It would thus be misguided to treat CGH's posted price inflation series as good proxies for BLS inflation series.

To make this point formally, we match the IRI product categories with their closest disaggregated price series published by the BLS. For some IRI product categories, we find a direct match. For others, the IRI product category is only a subset of a broader BLS category. In this latter situation, we combine IRI product categories whenever possible to get price indexes that are as comparable as possible to those published by the BLS. We use yearly stratum revenues to aggregate stratum-level posted price indexes across markets and product categories, in a manner corresponding to CGH's use of market-specific weights.

Figures A2 and A3 compare the posted price series in CGH and under GLSS's preferred methodology with price series published by the BLS for product categories that have a one-to-one match in coverage between the IRI and BLS samples. As is apparent, our posted price series generally track those produced by the BLS closely whereas CGH's often do not. The figures also show posted price series under CGH's methodology when we replace their censoring with our treatment of outliers and otherwise keep their methodology unchanged. These latter series show that addressing CGH's censoring may help but is insufficient, in itself, to produce well behaved series. For example, CGH's series for coffee and cold cereals are more dissimilar to those from the BLS when censoring is not implemented, revealing the presence of other methodological differences.

The similarity between the posted price series under our preferred methodology and those published by the BLS is perhaps remarkable because the IRI and BLS samples differ in a number of ways. For example, the IRI sample excludes the largest U.S. brick-and-mortar retailer (Wal-Mart) and stores that do not pertain to a retail chain, whereas the BLS sample is representative of all points of purchase. Moreover, as in CGH, we exclude private-label items whereas the BLS does not.[3] Furthermore, the BLS observes the items' regular and

[3] Private-label items account for roughly 70 percent of all sales in the milk category, and 25 percent or less in all other product categories in the IRI sample.

posted prices, and it uses that information when making item substitutions. Nonetheless, our simple treatment for clearance sales, which we discuss below, goes a long way toward fending off potential biases due to item turnover.

For completeness, Figures A4 and A5 provide corresponding posted price indexes for all remaining IRI product categories with possibly imperfect matches to BLS disaggregated series. For example, we judge that the IRI product category "yogurt" was imperfectly matched to the BLS category "other dairy and related products" because the latter category includes, in addition to yogurt, fresh cream, sour cream, half-and-half, and a number of other dairy-based products. Similarly, the IRI categories "household cleaner" and "laundry detergent" account for only part of the BLS category "household cleaning products," which also encompasses items such as sponges, brooms, and mops. We also include "photographic supplies" among imperfect matches due to these items' obsolescence over the sample period. Even though these IRI-BLS matches are imperfect, it is clear to us that the behavior of our posted price series is more similar to that of the closest official series than is the behavior of CGH's corresponding posted price series.

Summing up, Figures A1 through A5 illustrate how, under our preferred methodology, our posted and effective price indexes are better behaved overall at the stratum and product category levels than those derived by CGH. Moreover, the closeness of our price indexes with those produced by the BLS gives us confidence that our methodological approach is sound overall. In the remainder of this appendix, we discuss the considerations that led us to choose our particular data filters.

A.4 Replacing CGH's censoring with GLSS' treatment for outliers

In the paper, we illustrate the effects of CGH's censoring on their regression estimates by raising the censoring threshold for monthly log price movements from 1 to 12. We also show that imposing a high censoring threshold is essentially the same as adopting our preferred treatment of outliers. This similarity is perhaps not surprising: As we argue below, genuine outliers are rare in the sample and raising CGH's threshold to a level within the range of values used in the literature to identify outliers makes that threshold largely irrelevant. Put differently, the quality of the IRI scanner data is sufficiently high that we could have proceeded without worrying about outliers with only limited effect on the composition of the sample and subsequent regression results.

As our paper discusses, we have two concerns about CGH's censoring procedure. One is that their censoring point is much too low to be considered a valid treatment for outliers. The other is that CGH apply their censoring inconsistently between posted and effective price inflation. For posted prices, CGH censor changes in item prices, the lowest level of disaggregation. For effective prices, CGH censor changes in the (quantity-weighted) average price across items belonging to the same market and UPC. Simply raising CGH's censoring threshold to a higher value falls short of addressing these two concerns. And if the outsize price movements are truly spurious (as some clearly are), then keeping them in the sample can pollute the results.

Our preferred methodology uses a more standard treatment to identify outliers that addresses the above concerns. Prior to computing any statistics, we drop all items featuring one or more price movements that exceed a threshold value that falls within the range of values used in the literature to identify outliers.[4] By excluding problematic items, we ensure that the same observations are used to calculate both posted and effective price inflation, thus avoiding any asymmetry in aggregation.

The series labeled "CGH no censoring" in Figure A1 illustrate how replacing CGH's censoring with our treatment for outliers (while keeping all other elements of their methodology unchanged) affects their posted and effective price series at the stratum level. Two features are worth noting. First, introducing our treatment for outliers can have a large effect on CGH's price indexes. Second, only addressing their censoring is not sufficient, in itself, to produce reasonably behaved price indexes. For example, only replacing their censoring with our treatment for outliers produces a posted price series that, in the case of coffee (lower-left panel), falls 35 percent relative to the corresponding effective price series over the sample period. Moreover, the suspiciously large jumps in the difference between posted and effective prices remain. A similar conclusion applies when we aggregate the stratum-level series to the product category level. Even after the aggregation, Figures A2 through A5 show that CGH's posted price series often remain quite different from those reported by the BLS for comparable product categories.

Table A1 shows that CGH's censoring threshold of 1 for annualized monthly log price changes (column 1) is much too low to be considered a valid treatment for outliers. Under this threshold, all price drops of 8 percent or more (column 3, measured in standard percentages to facilitate the interpretation) are censored, as are correspondingly large increases. When applied at the item level, a staggering 71 percent of all weekly nonzero price movements meet the threshold. Boosting CGH's threshold value to 5—the largest value considered by CGH in a public response to our paper—would censor all item price drops in excess of 34.1 percent and correspondingly large price increases. Under this threshold, 15.3 percent of all weekly

[4] We can afford to drop entire price histories rather than only the problematic price change observations given the large size of the IRI sample. Dropping entire price histories trims the sample by only half a percent, and has essentially no effect on the findings relative to dropping solely problematic price change observations.

nonzero price changes in the sample would meet the threshold, still a high proportion. For the threshold to capture only outliers, larger values must thus be considered. For example, if we apply a threshold of 1 to *not annualized* monthly log price changes (column 2), then we would be censoring all price drops larger than 63.2 percent and correspondingly large price increases. Our choice of a threshold for outliers—85 percent for price drops—is similar to that used by Klenow and Kryvtsov (2008). All three of these higher thresholds are met by only a tiny fraction of all weekly price change observations (columns 4 and 5) and thus lead to similar regression results when imposed. Close inspection of the observations meeting these high thresholds suggests that most are genuine outliers and thus should be dropped from the dataset rather than censored.

A.5 CGH's inconsistent time aggregation of effective prices

CGH's aggregation of weekly effective price inflation to a monthly frequency contains a mistake for months that have five weeks. To compute monthly effective price inflation, they first create a *weekly* variable that contains the average paid price *over the entire month* by dividing total revenues that month by total quantities sold that month. They next calculate the *four-week* change in this weekly variable and then take, for each month, the average of those weekly four-week changes to obtain a measure of monthly effective price inflation. A problem arises for months that have five weeks because the four-week change in the monthly average price is always zero for the fifth week of the month. As a result, CGH's measure of monthly effective price inflation understates actual inflation by a fifth in months with five weeks. This time aggregation mistake dampens the measured response of effective prices. As a result, it can lead CGH to understate the flexibility of effective prices, and thus to underestimate the bias in posted price inflation due to cyclical store switching. Table A4 provides an illustration of this error.

This issue aside, the monthly aggregation of effective prices requires at least one pair of weekly observations—one observation in each month—separated by exactly four weeks; otherwise the four-week change will be missing for the month, resulting in a loss of information. Our procedure for handling missing weekly price observations, which is detailed in the paper and illustrated in the next two sections, helps prevent information losses in both posted and effective price movements.

A.6 Dealing with missing observations

Our paper documents that almost 40 percent of weekly observations are missing from the IRI sample. Missing observations are especially common for items with low sales volumes.

As such, the presence of missing observations is most likely to affect the reliability of CGH's unweighted regression results, which do not account for the importance of items in the dataset. Missing observations are also more likely to affect the computation of posted price inflation than effective price inflation because item-level posted price changes are more likely to be missing than UPC-market-level effective price changes.

Figure A6 provides an illustration of how missing observations affect the computation of weekly and monthly item-level inflation in CGH's analysis, and of how our preferred methodology tackles missing observations in a way that better captures variation in item-level prices. The top panel of the figure shows actual weekly price observations in the IRI sample for an item in the household cleaning product category sold in Boston.[5] In total, there are 53 price observations for this item between week 228 and week 571 in the sample. Had the price been observed every week over that period, then a total of 344 weekly observations would have been recorded. Put differently, roughly 85 percent of all weekly price observations are missing for this item in the IRI sample. From the series of 53 observed weekly prices (represented by the "x" markers), one can infer that the item's price changed at least 12 times over the period shown—and possibly more often if the price was adjusted in weeks for which there are no price observations.

CGH calculate a weekly posted price change for each of the item's 53 weekly price observations in the sample. Of these 53 observations, there are only 12 for which both the current-week and previous-week prices are observed (represented by the blue circles); in those 12 cases, CGH compute a numeric value for the weekly log price change. For the other 41 weeks with a current-week price observation but no previous-week price observation, CGH code the weekly price change with a missing value. Of note, only one of the 12 numeric values is different from zero, meaning that all but one them fail to enter CGH's regression analysis.

The middle panel shows how our imputation method for missing price observations works. Whenever the price is not observed during a week, the previous-week price—whether observed or imputed—is incremented by the amount of inflation in the stratum. This procedure ensures that the level of the resulting weekly price series (labeled "GLSS weekly price (BLS-like)") is anchored by actual weekly price observations. With inflation in the stratum being low overall, the imputed price never wanders far from the last observed price even when the price is unobserved for several consecutive weeks. This fact can be seen by comparing the "GLSS weekly price (BLS-like)" series with the "GLSS weekly price (forwarding)" series, for which we carry forward the last observed price whenever the current price is not observed. (We will return to this alternative imputation method in the section below.) Importantly

[5]The item is identified by the IRI variables SY=0, GE=1, VEND=11013, and ITEM=10020. It is an all-purpose liquid degreaser with a lemon scent.

for the reliability of our regression analysis, our weekly price series mimics observed changes in the actual price closely: Each of the 12 changes in the actual price has a counterpart in our imputed price series, meaning that we preserve the information about actual item-level price movements.

The third panel of Figure A6 contrasts CGH's time aggregation of weekly into monthly price changes, which can be inconsistent with actual price movements over time, and the time aggregation under our preferred methodology, which is consistent by design. The accompanying Table A5 shows the underlying weekly and monthly price change information. The panel and table help us illustrate two reasons why monthly price movements under CGH's approach may differ from actual monthly price movements. First, CGH compute monthly price movements by summing observed weekly price changes. As documented above, many actual weekly price changes are left out of CGH's analysis because they are missing a previous-week price in order to compute the change. For example, the item's price was observed to be \$7.49 in week 254 (which belongs to month 59) and next observed to be \$8.49 in week 261 (which belongs to month 60). Although we know that the item's price increased at some point between month 59 and month 60, there is no weekly price change between week 254 and week 261 that records the price movement under CGH's methodology. Second, whenever CGH impute all weekly price changes during a month with missing values, they also systematically impute a monthly price change equal to zero for that month.[6] For example, there is only one weekly price observation in month 60 (\$8.49 in week 261) and the corresponding weekly price change contains a missing value; therefore, CGH record a monthly price change of zero for that month. In the example shown in Figure A6, CGH code a majority of monthly inflation rates with zeros because the underlying weekly price changes contain only missing values. From a quantitative perspective, CGH's failure to capture weekly price changes because the previous-week price is missing is the largest source of revision in regression estimates when we adopt our imputation method for missing prices. By contrast, CGH's systematic imputation of monthly price changes with zeros whenever they code all underlying weekly price changes with missing values only has a marginal effect on the regression coefficients.[7]

In the third panel of Figure A6, we cumulate the monthly item-level price changes over the sample period.[8] The cumulative series under CGH's methodology captures a single monthly non-zero movement over the full sample period (in month 91), far fewer than the number

[6] The presence of such imputations and the underlying motivation are not mentioned in CGH's paper and replication materials.

[7] The small quantitative importance of this second source of bias related to missing observations was pointed out to us by CGH.

[8] We normalize the cumulative monthly log-change series to zero in the first month for which the price is observed in the sample. For ease of exposition, the panel shows only a one-year period that helps us contrast CGH's methodology with ours.

of adjustments apparent in the weekly series of actual prices. By contrast, our monthly series essentially captures all such movements, keeping track with the actual change in the weekly price over time. Many of the months for which monthly inflation is zero under CGH's procedure have all their underlying weekly price changes coded with missing values (for example, months 87–89, 93, and 96).[9]

A.7 Alternative imputation methods for missing item prices

In the paper, we adopt an imputation method for missing item prices that is similar to the method used by the BLS for the CPI. In short, for as long as an item price is missing, we impute the missing price by incrementing the last price (observed or imputed) by the amount of inflation in the stratum during the week. We choose this method in part because one of our goals is to produce indexes of posted prices that closely mimic the BLS methodology, and thus might speak to cyclical biases in official indexes. In this section, we show that our findings are robust to the consideration of two alternative imputation methods: price forwarding and linear interpolation. Like the imputation method in our paper, and in contrast with CGH's treatment of missing price observations, these two alternative methods ensure that the cumulative item-level inflation between two non-missing price observations coincides with the true change in the item's price.

Under price forwarding, each time an item's price is missing, we overwrite the missing value with the item's last observed price. Thus, the net change in an item's price over a sequence of consecutive missing observations is, in effect, attributed to the period in which the item's price is next observed. One benefit of this approach is that it is easy to implement because it requires no additional calculations and uses information available in real time. One potential downside is that it lumps all changes in periods with nonmissing price observations. To the extent that item prices are missing for extended periods or that several item prices are missing at the same time, the measured change in the price index could be delayed and more volatile relative to the true change.

Under linear interpolation, we assume that item price changes between nonmissing observations are of equal size. For example, if the price is missing for one period and then rises by 10 percent relative to the price observed two periods earlier, then we assume that the price rose 5 percent in the missing period and another 5 percent when the price was next

[9] Under CGH's methodology, months for which there are no weekly price observations do not contribute to the regression results because these months have no corresponding monthly posted price changes. In the bottom panel of Figure A6, we treat these months as if the monthly price had not changed when we cumulate monthly price movements (the months without weekly observations are 92, 94, 95, and 97). Our objective here is only to illustrate the nonzero price movements captured by CGH's methodology, not to depict the subset of monthly observations used in their regressions.

observed. One benefit of that approach is that it is agnostic about the timing of the item's price change. One possible shortcoming is that, to the extent that there is information in nonmissing observed prices, it is not making use that information. Also, the method cannot be used in real time to estimate inflation due to the lack of information about future item prices.

The results are presented in Table A6. The upper panel repeats our regression results under the preferred methodology in the paper, including the use of "BLS-like" imputation of missing price observations. The middle and bottom panels show the results under price forwarding and linear interpolation, respectively, keeping all other elements of our preferred methodology. The expenditure-weighted panel regression results (columns 5 through 8), which we emphasize throughout our paper, are little changed when we replace our imputation method with either of the two alternatives. In particular, the statistical hypothesis that effective price inflation has a larger cyclical response than posted price inflation is strongly rejected by the data. Similarly, we have strong statistical rejection for the unweighted regressions with either month fixed effects (column 3) or linear time trends (column 4). Only in the case of the unweighted regressions (columns 1 and 2) under the price forwarding method do we no longer statistically reject the hypothesis that the cyclical response of effective price inflation is larger than that of posted price inflation. However, we note the relative unreliability of these regression results, which is reflected in the point estimates being not statistically different from zero at standard significance levels for both posted and effective price inflation across all three methods shown in the table.

A.8 Controlling for clearance sales

Selection effects associated with the entry and exit of items are a well known source of bias in the official CPI over long horizons.[10] These effects can also distort the shorter-term response of price indexes to shocks, and are thus a concern for our study of the price response to local labor market conditions.[11] A fully satisfying treatment of biases due to basket turnover is beyond the scope of this paper because it would require the judgmental linking of millions of entering and exiting items. That said, we can readily mitigate one key source of bias that can potentially create a wedge in the cyclical response of posted and effective price

[10] Notably, a "new good bias" and a "quality change bias" are known to slant upward measured changes in the CPI, leading to an underestimation of the rise in living standards over long periods. See, among many contributors, the edited volume by Boskin *et al.* (1996), Bresnahan and Gordon (1996), Gordon (2006) and the conference summaries of the Ottawa Group on Price Indices.

[11] Broda and Weinstein (2010) provide evidence of cyclical creation and destruction of barcodes in scanner data. See also Berger *et al.* (2009), Nakamura and Steinsson (2012), and Gagnon, Mandel, and Vigfusson (2014) for some related evidence and theory.

inflation. "Clearance sales" is the phenomenon by which retailers sometimes offer extra discounts on items about to permanently disappear from the shelves. A failure to link the prices of disappearing items with those of their replacements can bias downward a price index because the price declines associated with the item exits would be captured but not the price increases associated with the entries of the substitutes. Posted prices are more exposed to this source of bias than effective prices because there is greater turnover in item prices than in the (quantity-weighted) average price across items belonging to the same UPC and market. We also note that, for a smaller number of observations, there is an explosion in sales volumes just before a store leaves the sample (say, with volumes reaching 10 times or more the highs achieved during large promotional sales). The source of these volume swings is unknown to us and not discussed in the IRI documentation. However, we suspect that they have to do with accounting/inventory management rather than with actual sales volumes given the implausible elasticities. Our trimming procedure automatically weeds out these anomalies.

Table A7 provides some statistics on the importance of item turnover in the IRI dataset and on the tendency of prices to fall as items are about to exit the sample. On average, 2.3 percent of items in the IRI dataset disappear from the sample every month. A slightly larger proportion of items (about 2.4 percent) join the sample each month, consistent with modest growth in stores' product offerings over time.[12] Table A7 further shows that prices are unusually low prior to item exits. On average, the price of an exiting item is over 8 percent lower than the price that prevailed a quarter before the item's exit (that is, 14 to 26 weeks earlier). For a majority of product categories, the typical price drop exceeds 10 percent.[13] The probability that an item is on sale in its final week in the sample is higher, at 30.5 percent (column 3), than for the typical item in the sample, at 23.4 percent (column 4), contributing to these relatively lower prices upon exit.

We mitigate the effects of clearance sales on our price indexes through a simple fix: We drop the last quarter (that is, 13 weeks) of every item's price trajectory. This trimming horizon is long enough to encompass the vast majority of clearance sales. At the same time,

[12] In computing these statistics, we exclude all observations in the first and last 13 weeks of a store's presence in the sample to reduce the risk of confounding store and item turnover. Some of the growth in product offering likely reflects a broadening of some product category definitions at dates when the IRI sample was extracted.

[13] We identify promotional sales using the IRI sales flag, which is extracted from price records using a proprietary algorithm. This sales flag is not directly comparable to the BLS sales flag, which is based on in-store observations by price collectors. Nonetheless, we take some comfort in that the fraction of items exiting the IRI sample (2.3 percent) multiplied by the fraction for which the IRI price-reduction flag is active upon exit (30.5 percent) is, at 0.7 percent, similar to the fraction of items subject to a clearance sale in the "processed food" and "other goods" categories of the U.S. CPI, at 0.6 percent, as reported by Nakamura and Steinsson (2008) in table 7 of their supplementary materials.

the trimming horizon is not overly long that the loss of observations alters the quality of the sample (more on this latter point below).

Figures A7 to A10 show how controlling for clearance sales affects our posted price indexes. The series labeled "GLSS" use our preferred methodology, which includes the trimming of the last quarter of data. The series labeled "GLSS (trim=7)" shorten our trimming horizon from 13 to 7 weeks. The series labeled "GLSS (no trim)" make no attempt at controlling for clearance sales, thereby using all observations in the sample. The series are compared to their nearest match among disaggregated U.S. CPI series. Because the BLS uses an extensive procedure for controlling for item substitutions, we interpret a greater proximity to official CPI series as suggestive that the series are less subject to biases.

The results are two fold. First, our treatment for clearance sales reduces the tendency of some posted price series to fall below the level of the BLS series over time. For some categories, such as "Breakfast cereals," "Household paper products," and "Household cleaning products," our trimming procedure reduces the gap in cumulative inflation 15 percentage points or more—a substantial improvement. We thus see our trimming procedure as helping to reduce biases related to clearance sales appreciably, and believe that our results are more reliable when the filter is included than when it is not.

Second, while the trimming of 13 weekly observations results in a loss of usable data, this loss is not so severe as to affect the quality of our price series. Indeed, when we reduce the number of trimmed observations by half (by trimming only the last seven weeks), we recover essentially the same price indexes: For most panels, the lines "GLSS" and "GLSS (trim=7)" are indistinguishable. When minor differences arise, it is not obvious to us which trimming point is preferable. In any case, the choice of trimming 7 or 13 weeks is inconsequential for our regression results, which are reported in Table A8 for reference.

A.9 Stitching of the 2001-2007 and 2008-2011 subsamples

The vast majority, although not all, of our item identifiers coincide with those in CGH's analysis. A small number of differences arise because CGH make a small departure from the procedure outlined in the IRI documentation to identify UPCs, whereas we follow exactly that procedure. We provide a brief explanation of why we believe the concerns that led CGH to depart from the IRI procedure are unwarranted. That said, the number of differing identifiers is arguably too small to make any material difference in our respective results.

The IRI documentation defines UPCs as unique combinations of four variables found in the main data files containing the weekly price and quantity observations. These variables are a vendor number (VEND), an item number (ITEM), a generation number (GE), and a system number (SY). For each UPC thus defined, IRI provides a description—stored in a separate set of files—that notably includes the UPC's brand, producer, and a volume-equivalent measure. Separate sets of UPC descriptions are provided for the 2001–2007 and the 2008–2011 subsamples.

In our paper, we adopt the definition of UPCs outlined in the IRI documentation. CGH's identification makes a small departure from this definition by adding a requirement that there be no change to a UPC's volume-equivalent measure in the UPC description files. Whenever a discrepancy in the volume-equivalent measure arises, CGH create separate UPC identifiers for each subsample, with one UPC identifier ending in 2007 and another one starting in 2008.

CGH shared with us their UPC identifiers for two product categories: mayonnaise and laundry detergent. Our UPC identifiers coincide with CGH's for the vast majority of weekly observations (99.2 percent) in these product categories. The remaining observations pertain to UPCs featuring a (typically small) change in their volume-equivalent measure in 2008. For example, the reported volume-equivalent measure may have declined from 33 ounces to 30 ounces. If these inconsistencies reflect changes in the nature of items, then one should observe a spike in the frequency of price changes in 2008 because prices before and after 2008 would pertain to different items.

As Figure A11 shows, there are no such spikes in the data. To create the figure, we only use items that are potentially problematic because their volume-equivalent measure changed in 2008. Had our sample been improperly stitched, then one should have seen a large jump in the frequency of price changes in the first week of 2008 for these items. For both product categories, the frequency of price changes remains near its normal level in early 2008. We thus conclude that the procedure outlined in the IRI documentation to define UPCs is appropriate.

So why do CGH end up creating slightly more item identifiers than we do? The likely reason is that CGH's departure from the IRI methodology exposes them to reading occasional manual entry errors in the UPC definition files. Contrary to the four variables identifying UPCs, which are read by scanners and must be exactly coded otherwise consumers would be charged for items other than the ones they are purchasing, small mistakes in the UPC definition files are largely inconsequential to retailer activities. So human data entry errors in the product description files are less likely to be identified and corrected. Because CGH draw on these imperfect files, they end up identifying slightly more items than are actually present in the IRI sample.

References

[1] Australian Bureau of Statistics. *Ottawa Group Meetings*. The Ottawa Group on Price Indices.

[2] Berger, David, Jon Faust, John H. Rogers, and Kai Steverson (2009). "Border Prices and Retail Prices," International Finance Discussion Papers 972, Board of Governors of the Federal Reserve System.

[3] Bresnahan, Timothy F., and Robert J. Gordon (1996). "The Economics of New Goods," *Studies in Income and Wealth, Volume 58*, National Bureau of Economic Research.

[4] Boskin, Michael J., Ellen R. Dulberger, Robert J. Gordon, Zvi Griliches, and Dale Jorgenson (1996). "Toward a More Accurate Measure of the Cost of Living," Final Report to the Senate Finance Committee.

[5] Broda, Christian, and David E. Weinstein (2010). "Product Creation and Destruction: Evidence and Price Implications," *American Economic Review* 100(3): 691–723.

[6] Bronnenberg, Bart J., Michael W. Kruger, and Carl F. Mela (2008). "Database Paper: The IRI Marketing Data Set," *Marketing Science* 27(4): 745–748.

[7] Bureau of Labor Statistics, Handbook of Methods, Washington DC, 2009.

[8] Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong (2015). "The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications," *American Economic Review* 105(3): 993–1029.

[9] Diewert, W. Erwin (1999). "Axiomatic and Economic Approaches to International Comparisons," NBER Chapters, in: *International and Interarea Comparisons of Income, Output, and Prices:* 13–107, National Bureau of Economic Research, Inc.

[10] Gagnon, Etienne and David López-Salido (2014). "Small Price Responses to Large Demand Shocks," Finance and Economics Discussion Series 2014-18, Board of Governors of the Federal Reserve System.

[11] Gagnon, Etienne, Benjamin R. Mandel, and Robert J. Vigfusson (2014). "Missing Import Price Changes and Low Exchange Rate Pass-Through." *American Economic Journal: Macroeconomics* 6(2): 156–206.

[12] Gordon, Robert J. (2006). "The Boskin Commission Report: A Retrospective One Decade Later," *International Productivity Monitor* 12: 7–22.

[13] Klenow, Peter J. and Oleksiy Kryvtsov (2008). “State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?” *Quarterly Journal of Economics* 123(3): 863–904.

[14] Nakamura, Emi and Jón Steinsson (2012). “Lost in Transit: Product Replacement Bias and Pricing to Market,” *American Economic Review* 102(7): 3277–3316.

Table A1: Thresholds for censoring and identification of outliers

	Threshold (monthly log changes)		**Price drop equivalent**	**Percent of price change obs. meeting threshold**	
	Annual.	Not annual.	(regular percent)	All obs.	Nonzero obs.
	(1)	(2)	(3)	(4)	(5)
CGH (2015)	1.00	0.08	8.0	21.45	70.65
CGH's reply	5.00	0.42	34.1	4.64	15.28
Not annualized thresh.	12.00	1.00	63.2	0.16	0.52
GLSS	22.77	1.90	85.0	0.02	0.07
Klenow-Kryvtsov (2008)	27.63	2.30	90.0	0.02	0.06

Source: Authors' calculations using IRI data.
Notes: CGH apply their censoring threshold to monthly price change data whereas GLSS apply their outlier threshold to weekly price change data. The statistics in columns 4 and 5 use weekly price change data; the price changes are computed using the items' current weekly prices and their last observed weekly prices.

Table A2: Inflation response to a 1-percentage-point rise in local unemployment rate under alternative censoring thresholds

	Uniformly-weighted items and UPCs				Expenditure-weighted items and UPCs			
					Market-specific	Common	Market-specific	Common
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
CGH's original estimates (threshold=1)								
Inflation								
Posted prices	-0.084^{**}	-0.087	-0.061^{***}	-0.164^{**}	-0.077^{**}	-0.075^{**}	-0.052^{**}	-0.059^{**}
	(0.041)	(0.053)	(0.017)	(0.067)	(0.021)	(0.023)	(0.026)	(0.029)
Effective prices	-0.120^{*}	-0.126	-0.219^{***}	-0.288^{***}	-0.201^{***}	-0.205^{***}	-0.136^{***}	-0.146^{***}
	(0.067)	(0.087)	(0.024)	(0.105)	(0.031)	(0.033)	(0.027)	(0.028)
Test p-value								
$\hat{\beta}^{pos} = \hat{\beta}^{eff}$	0.246	0.332	< 0.001	0.011	< 0.001	< 0.001	0.019	0.026
Largest threshold in CGH's reply (threshold=5)								
Inflation								
Posted prices	-0.195^{***}	-0.197^{*}	-0.123^{***}	-0.260^{**}	-0.150^{***}	-0.139^{***}	-0.090^{***}	-0.124^{***}
	(0.080)	(0.103)	(0.043)	(0.127)	(0.042)	(0.049)	(0.039)	(0.045)
Effective prices	-0.337^{***}	-0.364^{***}	-0.417^{***}	-0.452^{***}	-0.281^{***}	-0.288^{***}	-0.103^{***}	-0.131^{***}
	(0.104)	(0.135)	(0.031)	(0.156)	(0.033)	(0.038)	(0.041)	(0.046)
Test p-value								
$\hat{\beta}^{pos} = \hat{\beta}^{eff}$	< 0.001	0.002	< 0.001	0.003	0.011	0.004	0.947	0.889
$\hat{\beta}^{pos} < \hat{\beta}^{eff}$	< 0.001	0.001	< 0.001	0.002	0.006	0.002	0.473	0.445
Not annualized threshold (threshold=12)								
Inflation								
Posted prices	-0.234^{***}	-0.238^{**}	-0.155^{***}	-0.253^{*}	-0.193^{***}	-0.184^{***}	-0.127^{***}	-0.138^{**}
	(0.090)	(0.115)	(0.053)	(0.139)	(0.050)	(0.058)	(0.053)	(0.063)
Effective prices	-0.398^{***}	-0.428^{***}	-0.439^{***}	-0.454^{***}	-0.269^{***}	-0.277^{***}	-0.088^{**}	-0.122^{***}
	(0.107)	(0.137)	(0.033)	(0.155)	(0.030)	(0.039)	(0.042)	(0.048)
Test p-value								
$\hat{\beta}^{pos} = \hat{\beta}^{eff}$	< 0.001	< 0.001	< 0.001	0.002	0.156	0.096	0.467	0.808
$\hat{\beta}^{pos} < \hat{\beta}^{eff}$	< 0.001	< 0.001	< 0.001	0.001	0.078	0.048	0.766	0.596
Specification								
Stratum fixed effects	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	No	No	Yes	Linear trend	Yes	Yes	Yes	Yes
Weighted regressions	No	No	No	No	No	No	Yes	Yes

Source: "CGH's original estimates" are reproduced from Table 1 in Coibion, Gorodnichenko, and Hong (2015). All other numbers are the authors' calculations using IRI data.
Notes: Inflation is measured as the 12-month change in log prices. The thresholds apply to annualized monthly changes in log prices. Driscoll and Kraay (1998) standard errors are reported in parentheses below point estimates. Statistical significance at the 10, 5, and 1 percent levels is indicated with one, two, and three stars, respectively.

Table A3: Inflation response to a 1-percentage-point rise in local unemployment rate under alternative censoring thresholds (continued)

	Uniformly-weighted items and UPCs				Expenditure-weighted items and UPCs			
					Market-specific	Common	Market-specific	Common
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
CGH's original estimates (threshold=1)								
Inflation								
Posted prices	-0.084^{**}	-0.087	-0.061^{***}	-0.164^{**}	-0.077^{**}	-0.075^{**}	-0.052^{**}	-0.059^{**}
	(0.041)	(0.053)	(0.017)	(0.067)	(0.021)	(0.023)	(0.026)	(0.029)
Effective prices	-0.120^{*}	-0.126	-0.219^{***}	-0.288^{***}	-0.201^{***}	-0.205^{***}	-0.136^{***}	-0.146^{***}
	(0.067)	(0.087)	(0.024)	(0.105)	(0.031)	(0.033)	(0.027)	(0.028)
Test p-value								
$\hat{\beta}^{pos} = \hat{\beta}^{eff}$	0.246	0.332	< 0.001	0.011	< 0.001	< 0.001	0.019	0.026
Censoring threshold equal to GLSS' trimming threshold (threshold=23)								
Inflation								
Posted prices	-0.241^{***}	-0.245^{**}	-0.156^{***}	-0.252^{*}	-0.189^{***}	-0.182^{***}	-0.123^{**}	-0.137^{**}
	(0.090)	(0.115)	(0.053)	(0.139)	(0.050)	(0.058)	(0.053)	(0.063)
Effective prices	-0.403^{***}	-0.435^{***}	-0.444^{***}	-0.452^{***}	-0.268^{***}	-0.276^{***}	-0.089^{**}	-0.123^{***}
	(0.107)	(0.137)	(0.033)	(0.154)	(0.030)	(0.038)	(0.042)	(0.048)
Test p-value								
$\hat{\beta}^{pos} = \hat{\beta}^{eff}$	< 0.001	< 0.001	< 0.001	0.002	0.140	0.091	0.523	0.820
$\hat{\beta}^{pos} < \hat{\beta}^{eff}$	< 0.001	< 0.001	< 0.001	0.001	0.070	0.045	0.738	0.590
Censoring threshold equal to Klenow and Kryvtsov's trimming threshold (threshold=28)								
Inflation								
Posted prices	-0.241^{***}	-0.246^{**}	-0.156^{***}	-0.252^{*}	-0.189^{***}	-0.182^{***}	-0.123^{**}	-0.137^{**}
	(0.090)	(0.115)	(0.053)	(0.139)	(0.050)	(0.058)	(0.053)	(0.063)
Effective prices	-0.403^{***}	-0.435^{***}	-0.444^{***}	-0.452^{***}	-0.268^{***}	-0.276^{***}	-0.089^{**}	-0.123^{***}
	(0.107)	(0.137)	(0.033)	(0.154)	(0.030)	(0.038)	(0.042)	(0.048)
Test p-value								
$\hat{\beta}^{pos} = \hat{\beta}^{eff}$	< 0.001	< 0.001	< 0.001	0.002	0.137	0.091	0.528	0.819
$\hat{\beta}^{pos} < \hat{\beta}^{eff}$	< 0.001	< 0.001	< 0.001	0.001	0.068	0.045	0.736	0.590
Specification								
Stratum fixed effects	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	No	No	Yes	Linear trend	Yes	Yes	Yes	Yes
Weighted regressions	No	No	No	No	No	No	Yes	Yes

Source: "CGH's original estimates" are reproduced from Table 1 in Coibion, Gorodnichenko, and Hong (2015). All other numbers are the authors' calculations using IRI data.
Notes: Inflation is measured as the 12-month change in log prices. The thresholds apply to annualized monthly changes in log prices. Driscoll and Kraay (1998) standard errors are reported in parentheses below point estimates. Statistical significance at the 10, 5, and 1 percent levels is indicated with one, two, and three stars, respectively.

Table A4: Illustration of CGH's time aggregation mistake for effective prices

Week	**Month**	**Monthly effective price** (1)	**Four-week change in monthly effective price** (2)	**CGH's monthly effective price inflation** (3)
1	1	1.00	.	.
2	1	1.00	.	.
3	1	1.00	.	.
4	1	1.00	.	.
5	2	1.25	0.25	0.20
6	2	1.25	0.25	0.20
7	2	1.25	0.25	0.20
8	2	1.25	0.25	0.20
9	2	1.25	0.00	0.20

Notes: The table illustrates CGH's incorrect time aggregation of weekly effective prices for months that have five weeks. We suppose that a UPC's first and second months in the sample contain four weeks and five weeks, respectively. We assume that the UPC's effective (log) price across stores in the market is 1 in the first month and 1.25 in second month (column 1), consistent with a 0.25 log increase between the two periods. To measure monthly effective price inflation, CGH first compute the four-week change in the weekly series of these monthly effective prices (column 2). For weeks 5 through 8, the four-week change coincides with the actual monthly effective price change. For week 9, the four-week change is zero because weeks 5 and 9 belong to the same month. CGH then calculate the average four-week change during the month (column 3), so they report a change of only 0.2 for month 2, instead of 0.25.

Table A5: Example of imputations of missing observations

			Imputed price adjustments (log changes)					
			CGH		**GLSS BLS-like**		**GLSS forwarding**	
Week (1)	**Month** (2)	**Price** (3)	**week** (4)	**month** (5)	**week** (6)	**month** (7)	**week** (8)	**month** (9)
367	85	8.99	.	0.000	−0.004	−0.007	0.000	0.000
368	85	8.99	0.000	0.000	0.000	−0.007	0.000	0.000
369	85	8.99	0.000	0.000	0.000	−0.007	0.000	0.000
376	87	8.99	.	0.000	−0.003	−0.015	0.000	0.000
379	88	8.99	.	0.000	0.006	0.015	0.000	0.000
385	89	6.74	.	0.000	−0.302	−0.294	−0.288	−0.288
390	90	8.49	.	0.000	0.224	0.228	0.231	0.231
391	90	8.49	0.000	0.000	0.000	0.228	0.000	0.231
392	91	8.99	0.057	0.057	0.057	0.049	0.057	0.057
395	91	8.99	.	0.057	−0.016	0.049	0.000	0.057
404	93	1.25	.	0.000	0.000	−0.007	0.000	0.000
415	96	9.99	.	0.000	0.087	0.095	0.105	0.105
424	98	9.99	.	0.000	−0.016	−0.016	0.000	0.000
425	98	9.99	0.000	0.000	0.000	−0.016	0.000	0.000

Source: Authors' calculations using IRI data.
Notes: The series labeled "GLSS (BLS-like)" use all of our data filters, including the imputation of missing weekly prices using inflation for nonmissing observations within the missing item's stratum. The series labeled "GLSS (forwarding)" impute missing item prices using the items' last observed price. The series labeled "GLSS (linear interpolation)" impute missing item prices using linear interpolations between the items' last and next observed prices.

Table A6: Inflation response to a 1-percentage-point rise in the local unemployment rate under alternative imputation methods for missing prices

	Uniformly-weighted items and UPCs				Expenditure-weighted items and UPCs			
					Market-specific	Common	Market-specific	Common
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
GLSS (BLS-like)								
Inflation								
Posted prices	−0.067	−0.058	−0.223***	−0.428**	−0.256***	−0.253***	−0.142***	−0.157***
	(0.120)	(0.157)	(0.036)	(0.189)	(0.044)	(0.051)	(0.040)	(0.047)
Effective prices	−0.083	−0.072	−0.236***	−0.416***	−0.222***	−0.210***	−0.090**	−0.098**
	(0.100)	(0.132)	(0.032)	(0.161)	(0.034)	(0.038)	(0.039)	(0.043)
Test p-value								
$\hat{\beta}^{pos} = \hat{\beta}^{eff}$	0.486	0.633	0.518	0.718	0.279	0.191	0.084	0.045
GLSS (forwarding)								
Inflation								
Posted prices	−0.049	−0.038	−0.223***	−0.411**	−0.263***	−0.258***	−0.151***	−0.164***
	(0.115)	(0.150)	(0.035)	(0.177)	(0.045)	(0.051)	(0.041)	(0.047)
Effective prices	−0.121	−0.105	−0.252***	−0.380***	−0.219***	−0.211***	−0.088***	−0.098***
	(0.103)	(0.135)	(0.034)	(0.165)	(0.034)	(0.040)	(0.039)	(0.043)
Test p-value								
$\hat{\beta}^{pos} = \hat{\beta}^{eff}$	0.000	0.013	0.198	0.224	0.238	0.155	0.072	0.073
$\hat{\beta}^{pos} < \hat{\beta}^{eff}$	0.000	0.006	0.901	0.888	0.881	0.923	0.964	0.964
GLSS (linear interpolation)								
Inflation								
Posted prices	−0.092***	−0.093**	−0.220***	−0.469**	−0.264***	−0.255***	−0.150***	−0.161***
	(0.115)	(0.150)	(0.033)	(0.178)	(0.046)	(0.050)	(0.042)	(0.047)
Effective prices	−0.121***	−0.105**	−0.252***	−0.380***	−0.219***	−0.211***	−0.088***	−0.098***
	(0.103)	(0.135)	(0.034)	(0.165)	(0.034)	(0.040)	(0.039)	(0.043)
Test p-value								
$\hat{\beta}^{pos} = \hat{\beta}^{eff}$	0.131	0.665	0.189	0.001	0.242	0.196	0.088	0.091
$\hat{\beta}^{pos} < \hat{\beta}^{eff}$	0.065	0.332	0.905	0.999	0.879	0.902	0.956	0.954
Specification								
Stratum fixed effects	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	No	No	Yes	Linear trend	Yes	Yes	Yes	Yes
Weighted regressions	No	No	No	No	No	No	Yes	Yes

Source: Authors' calculations using IRI data.
Notes: The series labeled "GLSS (BLS-like)" use all of our data filters, including the imputation of missing weekly prices using inflation for nonmissing observations within the missing item's stratum. The series labeled "GLSS (forwarding)" impute missing item prices using the items' last observed price. The series labeled "GLSS (linear interpolation)" impute missing item prices using linear interpolations between the items' last and next observed prices.

Table A7: Item turnover and clearance sales in the IRI sample

Product category	**Monthly exit rate** (percent) (1)	**Monthly entry rate** (percent) (2)	**On sale** (percent) (3)	**On sale final week** (percent) (4)	**Price drop final week** (percent) (5)
Beer	1.4	1.7	20.3	21.2	2.6
Blades	2.4	2.4	16.2	25.4	12.3
Carb. beverages	2.0	2.1	32.2	30.7	2.2
Cigarettes	2.3	2.0	6.8	4.0	−1.6
Coffee	1.8	2.2	22.1	30.7	13.3
Cold cereal	2.1	2.2	23.0	42.1	17.0
Condiments	1.3	1.2	14.8	26.7	13.8
Deodorant	2.7	2.6	22.4	34.4	18.5
Diapers	3.9	4.2	24.6	35.0	9.6
Facial tissue	2.9	2.9	23.6	33.8	11.7
Frozen dinner	2.3	2.7	32.2	44.8	14.4
Frozen pizza	1.9	2.1	33.5	38.0	7.6
Hot dogs	1.5	1.5	27.4	28.5	3.9
Household cleaners	2.0	2.9	17.8	33.6	19.3
Laundry detergent	3.1	3.2	25.6	42.6	14.2
Margarine/butter	1.4	1.3	19.2	29.5	7.6
Mayonnaise	1.6	1.7	15.6	30.4	11.7
Milk	1.8	2.1	13.9	18.8	2.9
Paper towels	4.2	4.2	19.8	32.7	9.0
Peanut butter	1.2	1.4	16.8	31.5	12.0
Photography	3.0	2.0	21.0	18.7	13.4
Razors	3.5	3.7	24.6	28.4	15.5
Salty snacks	3.7	3.8	25.1	27.3	3.3
Shampoo	3.2	3.2	25.8	33.5	16.6
Soup	1.3	1.8	20.3	37.1	20.1
Spaghetti sauce	1.3	1.3	24.8	33.9	15.4
Sugar/substitutes	1.4	1.7	12.1	28.9	19.8
Toilet tissue	3.5	3.6	21.5	35.0	8.4
Toothbrushes	2.8	3.0	21.0	29.5	20.3
Toothpaste	2.5	2.5	23.6	36.0	18.5
Yogurt	2.3	2.6	24.5	35.8	7.6
Mean					
Unweighted	2.3	2.4	21.7	30.9	11.6
Expenditures-weighted	2.3	2.4	23.4	30.5	8.2

Source: Authors' calculations using IRI data.
Notes: The statistics exclude private labels. We pool raw observations across markets and months to obtain product-category figures. "On sale" is the fraction of nonmissing weekly observations for which the IRI sales flag is activated. "On sale in final week" is the corresponding fraction using only the last observation of price trajectories. "Price drop final week" is the percent (in log changes) by which an item's last observed price is below its mean price over the previous 14 to 26 weeks. With the exception of "On sale," all statistics exclude observations within 13 weeks of a store's entry into or exit from the sample. The averaging of statistics across product categories uses either uniform weights ("Unweighted") or sample expenditures weights ("Expenditures-weighted").

Table A8: Inflation response to a 1-percentage-point rise in the local unemployment rate under alternative treatments of clearance sales

	Uniformly-weighted items and UPCs				Expenditure-weighted items and UPCs			
					Market-specific	Common	Market-specific	Common
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
GLSS								
Inflation								
Posted prices	-0.067	-0.058	-0.223^{***}	-0.428^{**}	-0.256^{***}	-0.253^{***}	-0.142^{***}	-0.157^{***}
	(0.120)	(0.157)	(0.036)	(0.189)	(0.044)	(0.051)	(0.040)	(0.047)
Effective prices	-0.083	-0.072	-0.236^{***}	-0.416^{***}	-0.222^{***}	-0.210^{***}	-0.090^{**}	-0.098^{**}
	(0.100)	(0.132)	(0.032)	(0.161)	(0.034)	(0.038)	(0.039)	(0.043)
Test p-value								
$\hat{\beta}^{pos} = \hat{\beta}^{eff}$	0.486	0.633	0.507	0.718	0.248	0.193	0.084	0.045
GLSS (trim=7)								
Inflation								
Posted prices	-0.065	-0.056	-0.239^{***}	-0.433^{**}	-0.274^{***}	-0.271^{***}	-0.161^{***}	-0.174^{***}
	(0.122)	(0.159)	(0.037)	(0.188)	(0.046)	(0.053)	(0.042)	(0.049)
Effective prices	-0.083	-0.064	-0.254^{***}	-0.404^{***}	-0.231^{***}	-0.217^{***}	-0.100^{***}	-0.110^{***}
	(0.106)	(0.139)	(0.030)	(0.165)	(0.035)	(0.039)	(0.041)	(0.045)
Test p-value								
$\hat{\beta}^{pos} = \hat{\beta}^{eff}$	0.406	0.752	0.483	0.350	0.147	0.083	0.038	0.017
$\hat{\beta}^{pos} < \hat{\beta}^{eff}$	0.203	0.376	0.759	0.825	0.926	0.958	0.981	0.991
GLSS (no trimming)								
Inflation								
Posted prices	-0.353^{***}	-0.380^{**}	-0.437^{***}	-0.424^{**}	-0.364^{***}	-0.365^{***}	-0.187^{***}	-0.210^{***}
	(0.130)	(0.167)	(0.052)	(0.191)	(0.049)	(0.056)	(0.051)	(0.059)
Effective prices	-0.278^{***}	-0.284^{**}	-0.410^{***}	-0.415^{***}	-0.261^{***}	-0.253^{***}	-0.116^{***}	-0.130^{***}
	(0.099)	(0.129)	(0.040)	(0.149)	(0.035)	(0.038)	(0.042)	(0.046)
Test p-value								
$\hat{\beta}^{pos} = \hat{\beta}^{eff}$	0.042	0.039	0.366	0.875	0.002	0.003	0.057	0.019
$\hat{\beta}^{pos} < \hat{\beta}^{eff}$	0.979	0.980	0.817	0.563	0.999	0.999	0.972	0.990
Specification								
Stratum fixed effects	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	No	No	Yes	Linear trend	Yes	Yes	Yes	Yes
Weighted regressions	No	No	No	No	No	No	Yes	Yes

Source: Authors' calculations using IRI data.
Notes: The series labeled "GLSS" use all of our data filters, including the trimming of the last 13 weekly observations of each price trajectory to control for clearance sales. The series labeled "GLSS (trim=7)" shorten our trimming to 7 weekly observations. The series labeled "GLSS (no trimming)" use all of the observations in the sample, and thus do not attempt to control for clearance sales.

Figure A1: Log difference between posted and effective price indexes

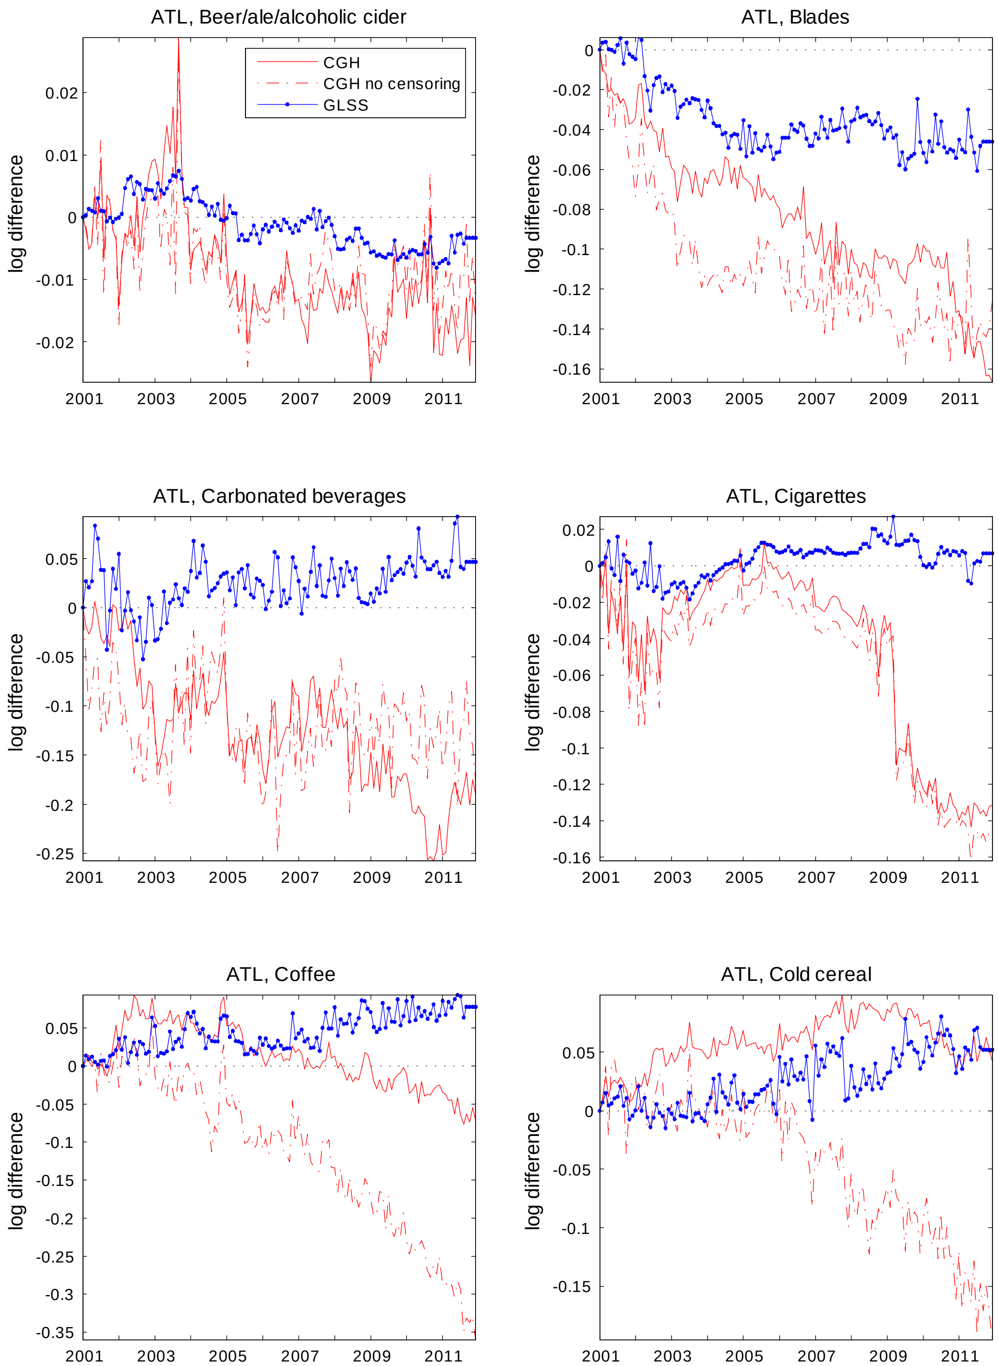

Source: Authors' calculations using IRI data.
Notes: Stratum-level posted and effective price indexes are computed by cumulating monthly inflation and are normalized to zero in January 2001. The series labeled "CGH" correspond to Coibion, Gorodnichenko, and Hong's (2015) original methodology. The series labeled "CGH no censoring" replace CGH's censoring with our treatment of outliers. The series labeled "GLSS" use all of our data filters.

Figure A2: Official and IRI-based posted price indexes: categories with direct match

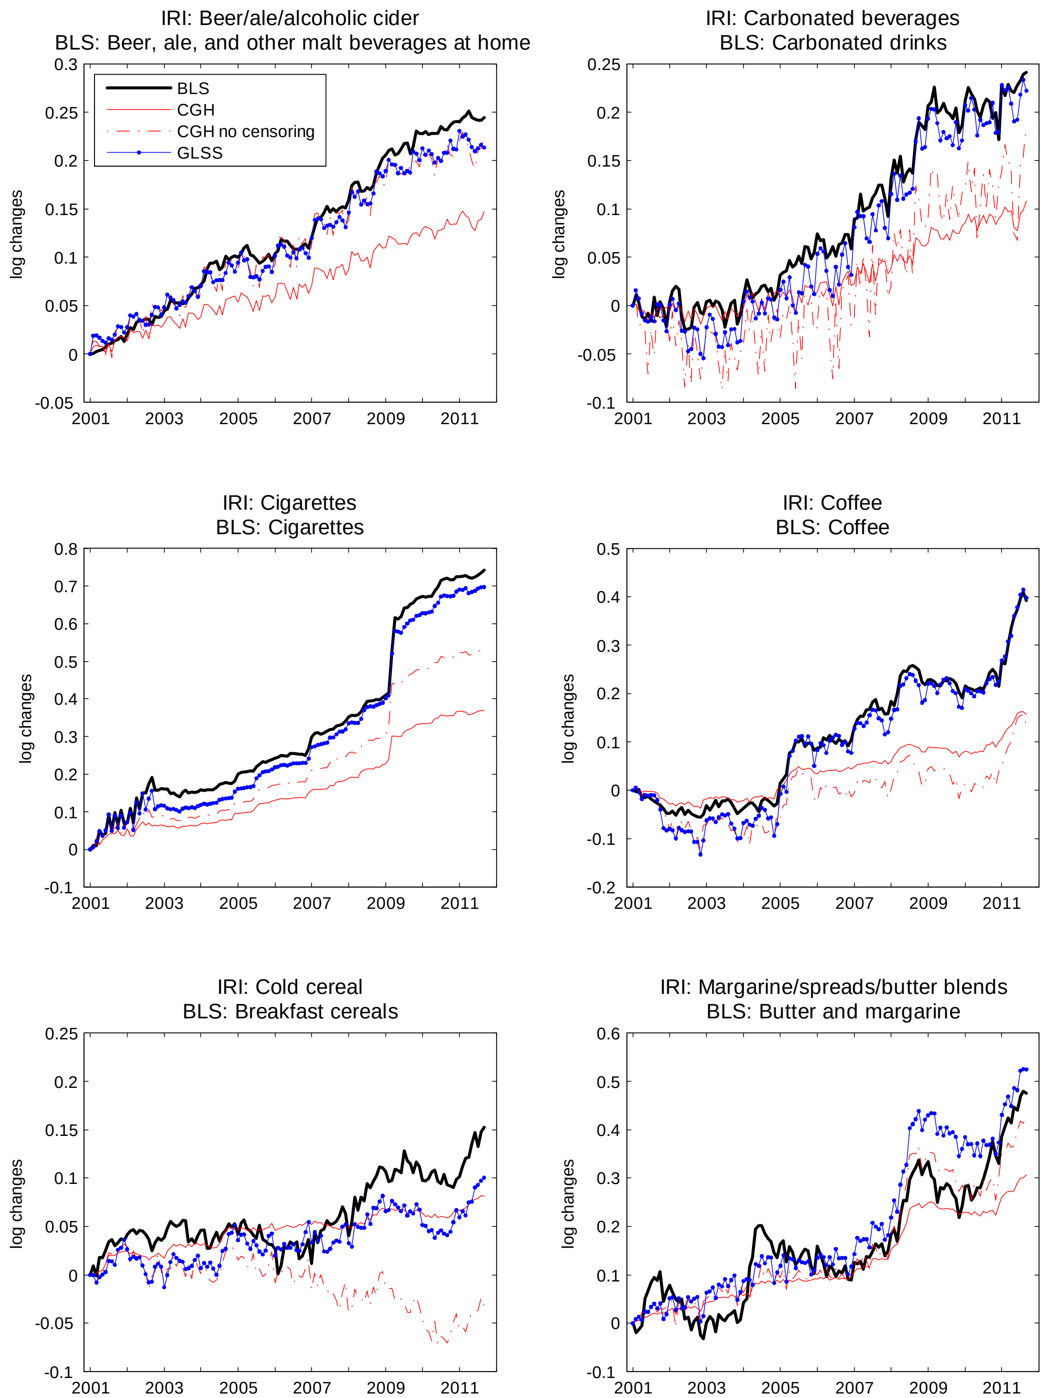

Source: Bureau of Labor Statistics and authors' calculations using IRI data.
Notes: Stratum-level posted price indexes are computed by cumulating monthly inflation, normalized to zero in January 2001, and aggregated using market-specific yearly expenditure shares. The series labeled "CGH" correspond to Coibion, Gorodnichenko, and Hong's (2015) original methodology. The series labeled "CGH no censoring" replace CGH's censoring with our treatment of outliers. The series labeled "GLSS" use all of our data filters.

Figure A3: Official and IRI-based posted price indexes: categories with direct match (continued)

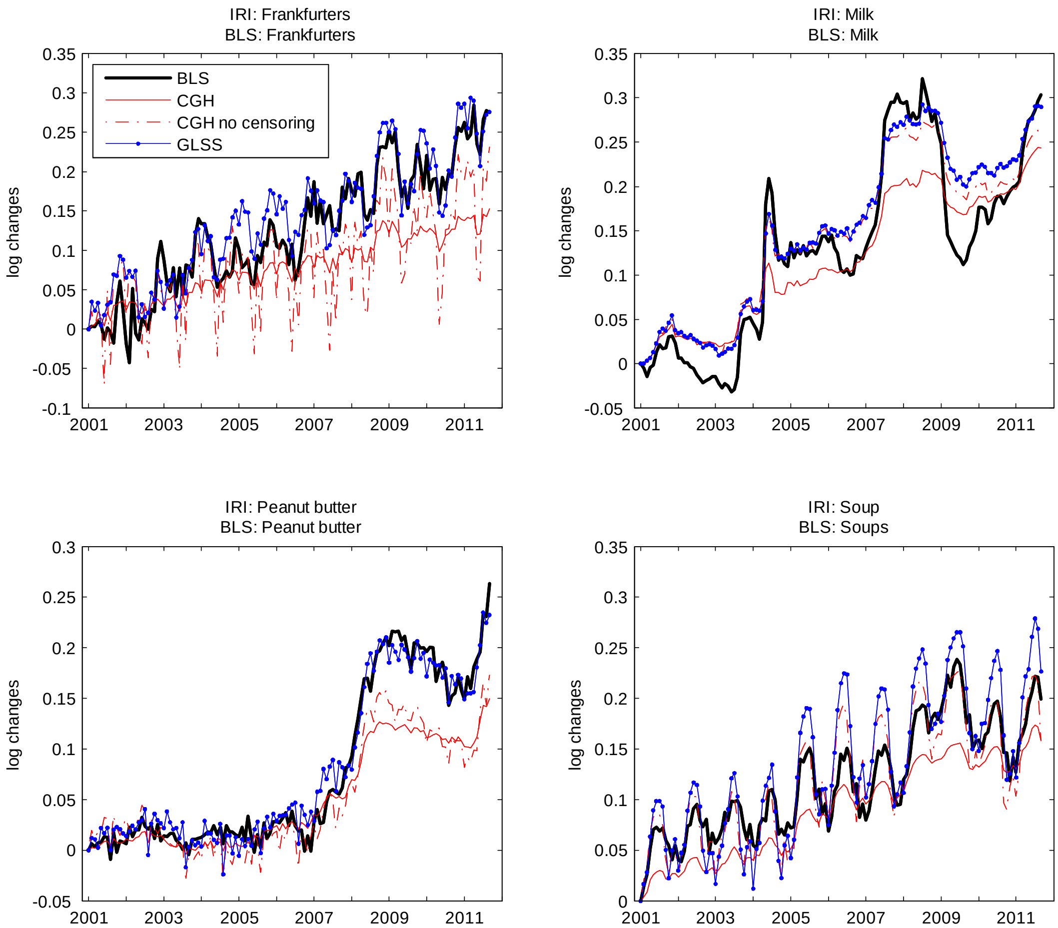

Source: Bureau of Labor Statistics and authors' calculations using IRI data.
Notes: Stratum-level posted price indexes are computed by cumulating monthly inflation, normalized to zero in January 2001, and aggregated using market-specific yearly expenditure shares. The series labeled "CGH" correspond to Coibion, Gorodnichenko, and Hong's (2015) original methodology. The series labeled "CGH no censoring" replace CGH's censoring with our treatment of outliers. The series labeled "GLSS" use all of our data filters.

Figure A4: Official and IRI-based posted price indexes: categories with imperfect match

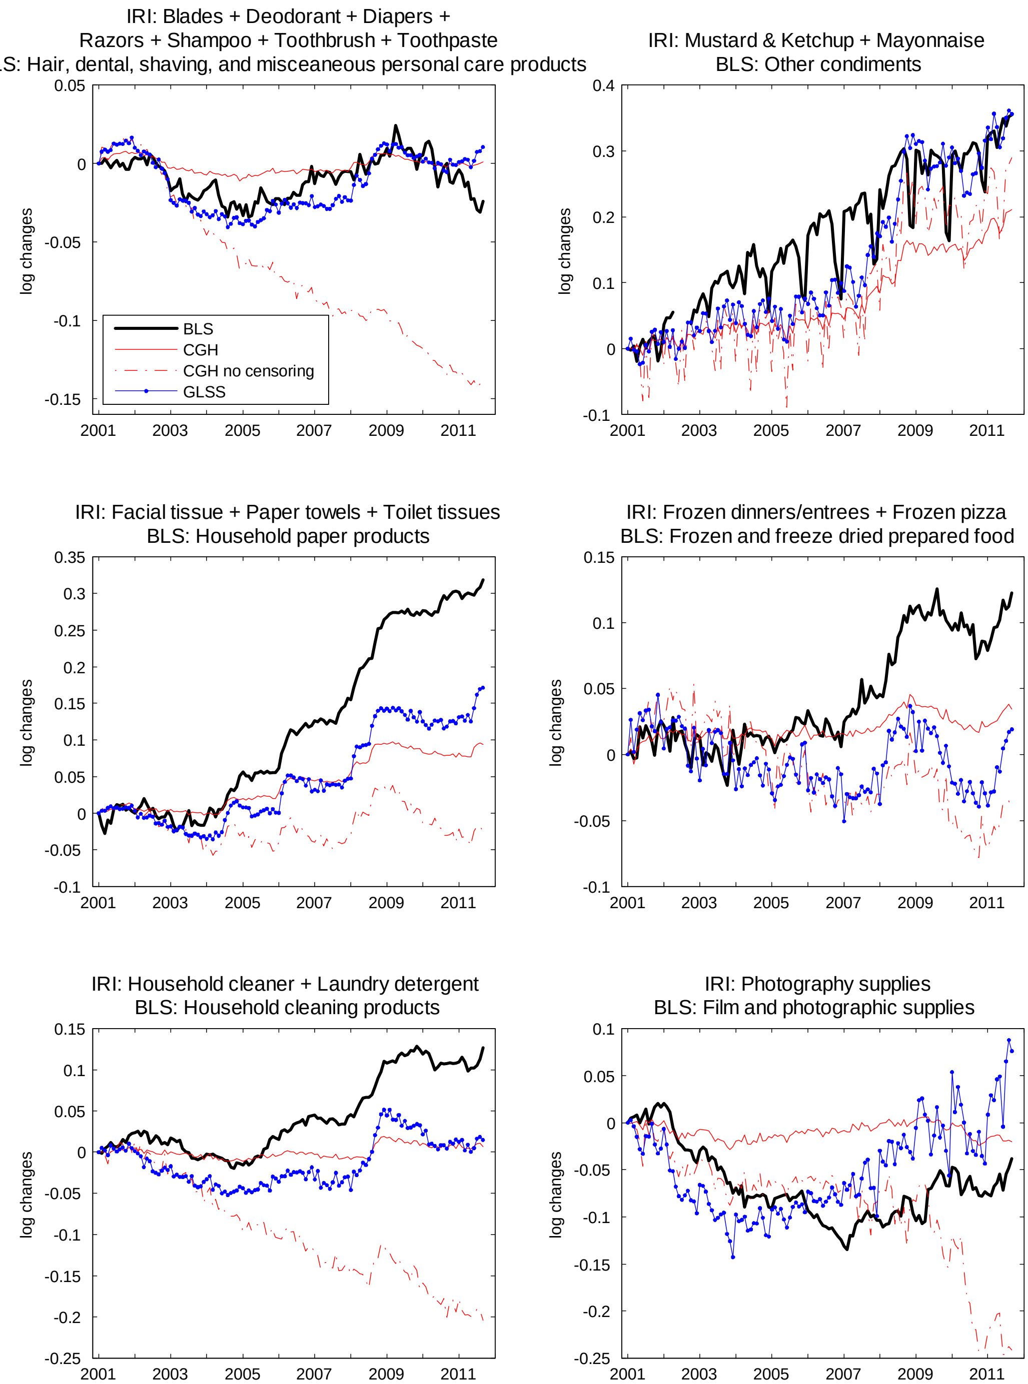

Source: Bureau of Labor Statistics and authors' calculations using IRI data.
Notes: Stratum-level posted price indexes are computed by cumulating monthly inflation, normalized to zero in January 2001, and aggregated using market-specific yearly expenditure shares. The series labeled "CGH" correspond to Coibion, Gorodnichenko, and Hong's (2015) original methodology. The series labeled "CGH no censoring" replace CGH's censoring with our treatment of outliers. The series labeled "GLSS" use all of our data filters.

Figure A5: Official and IRI-based posted price indexes: categories with imperfect match (continued)

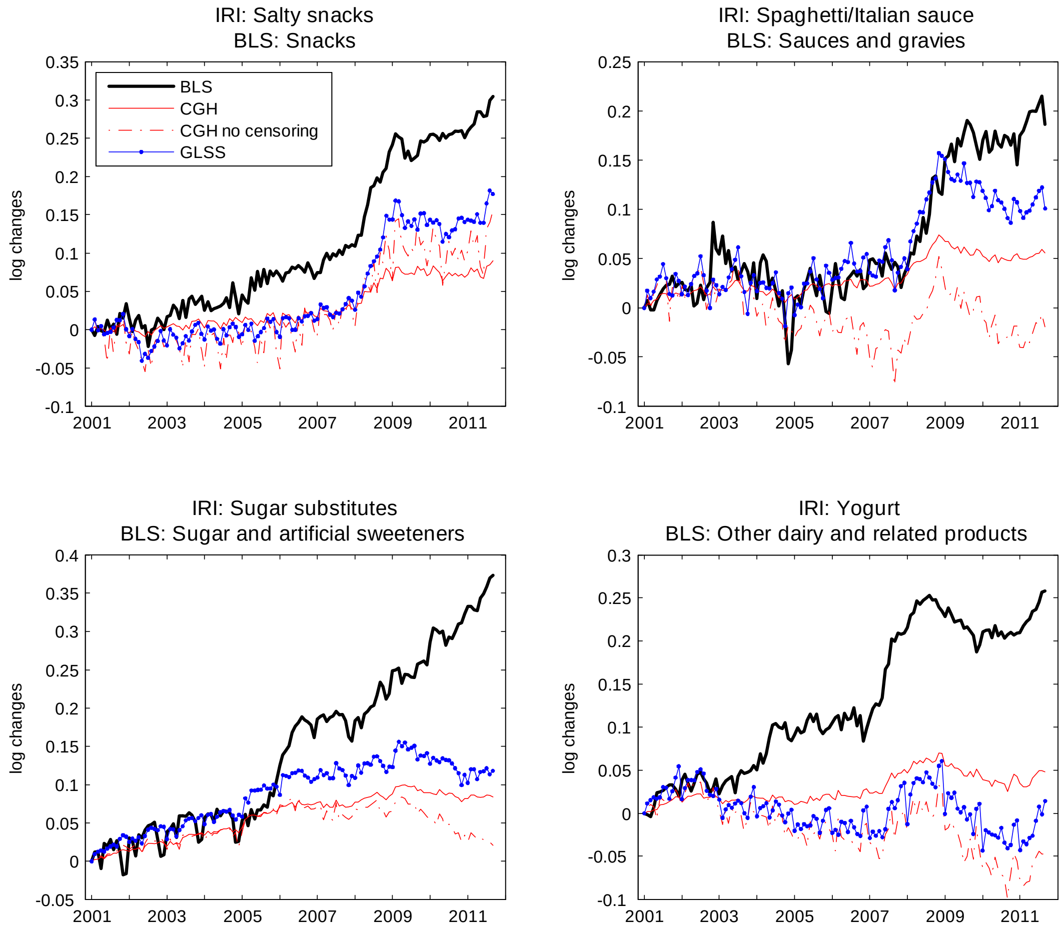

Source: Bureau of Labor Statistics and authors' calculations using IRI data.
Notes: Stratum-level posted price indexes are computed by cumulating monthly inflation, normalized to zero in January 2001, and aggregated using market-specific yearly expenditure shares. The series labeled "CGH" correspond to Coibion, Gorodnichenko, and Hong's (2015) original methodology. The series labeled "CGH no censoring" replace CGH's censoring with our treatment of outliers. The series labeled "GLSS" use all of our data filters.

Figure A6: Example of missing observations and time aggregation

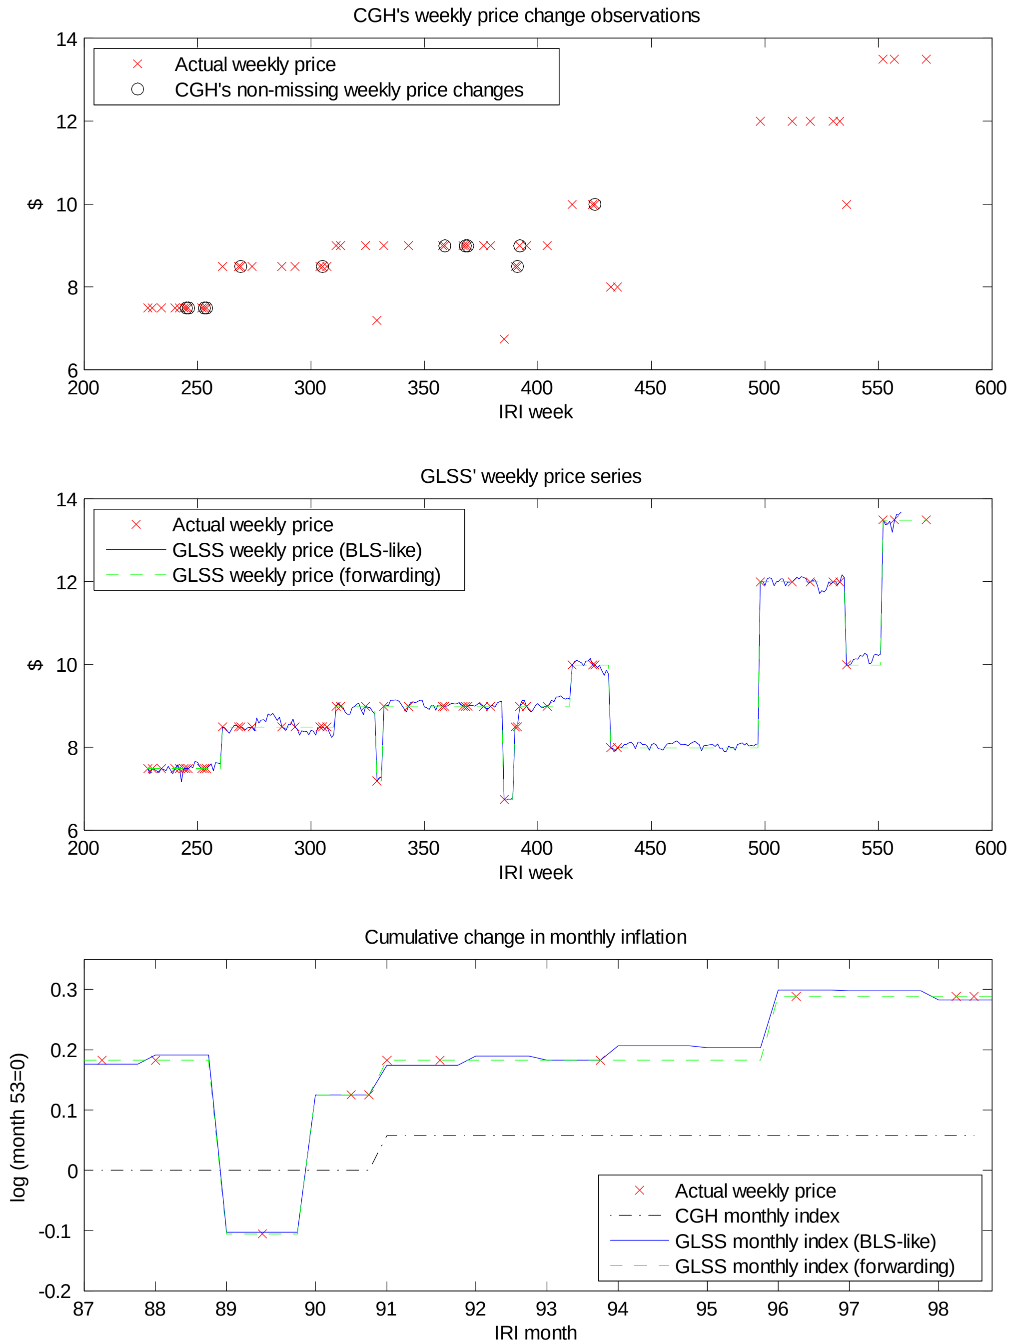

Source: Authors' calculations using IRI data.
Notes: The "x" markers indicate actual weekly observations in the IRI sample for an item in the household cleaning product category sold in Boston. The series labeled "GLSS weekly price (BLS-like)" and "GLSS weekly price (forwarding)" in the middle panel show our weekly price series when we impute the last price (observed or imputed) with inflation in the stratum or with no inflation, respectively, whenever the current price is not observed. The series labeled "CGH monthly index" , "GLSS monthly index (BLS-like)" , and "GLSS monthly index (forwarding)" in the bottom panel show cumulative monthly inflation under CGH's procedure, our BLS-like procedure, and our forwarding procedure, respectively; for CGH's procedure, we assume no inflation in months with no usable observations.

Figure A7: Controlling for clearance sales through trimming

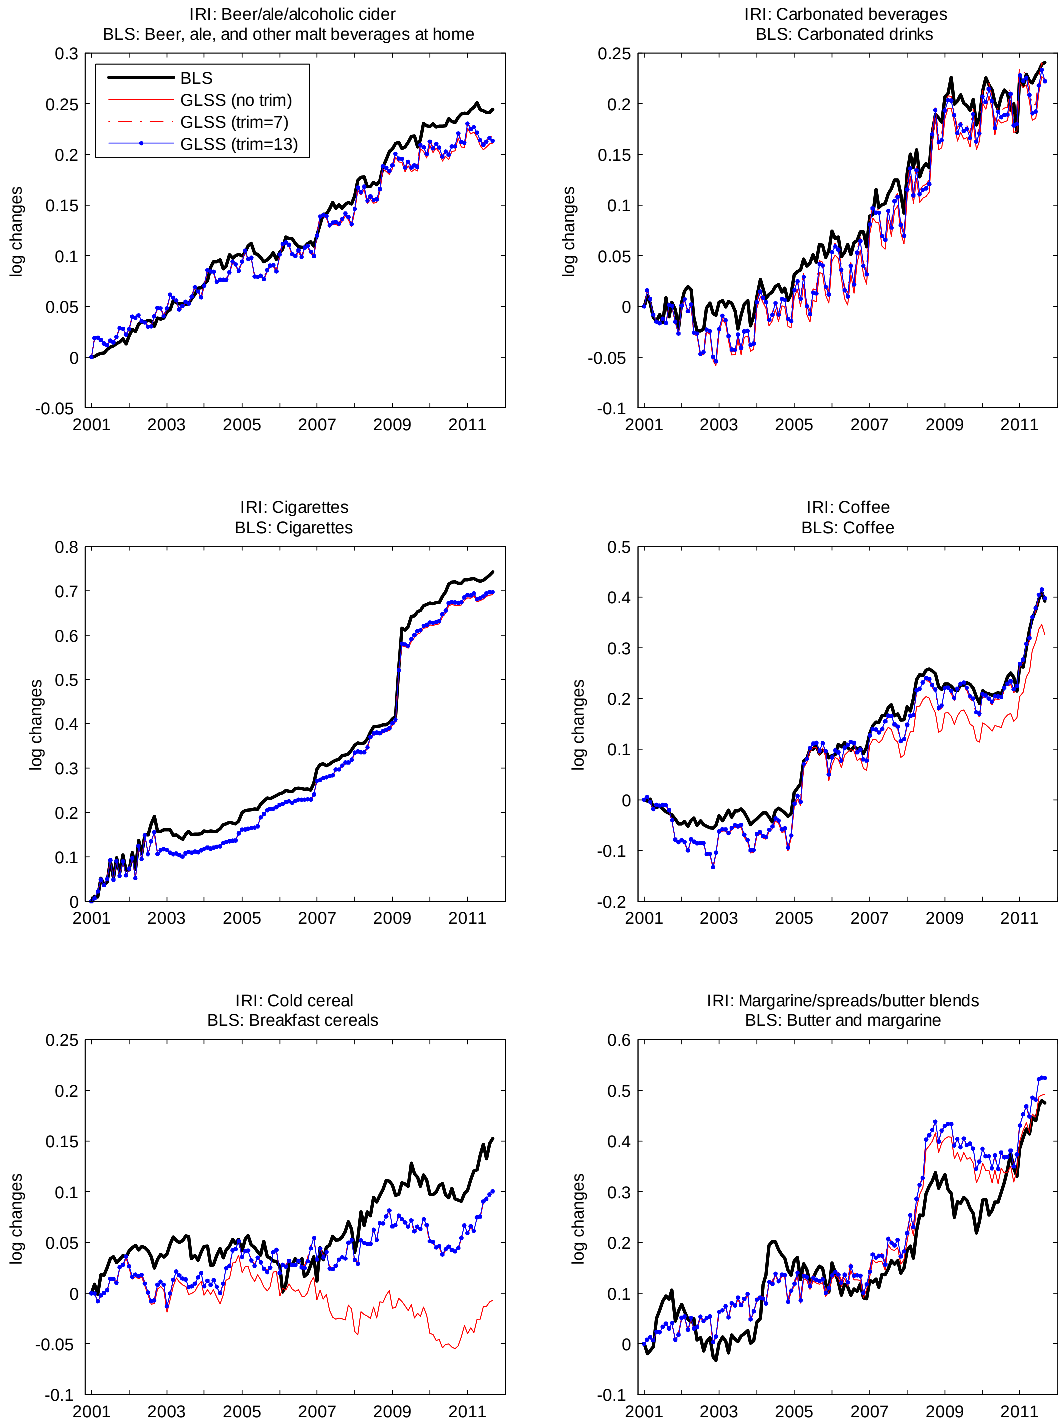

Source: Bureau of Labor Statistics and authors' calculations using IRI data.
Notes: The series labeled "GLSS (trim=13)" use all of our preferred data filters, including the trimming of the last 13 weekly observations of each price trajectory to control for clearance sales. The series labeled "GLSS (trim=7)" shorten our trimming to 7 weekly observations. The series labeled "GLSS (no trim)" use all of the observations in the sample, and thus do not attempt to control for clearance sales.

Figure A8: Controlling for clearance sales through trimming (continued)

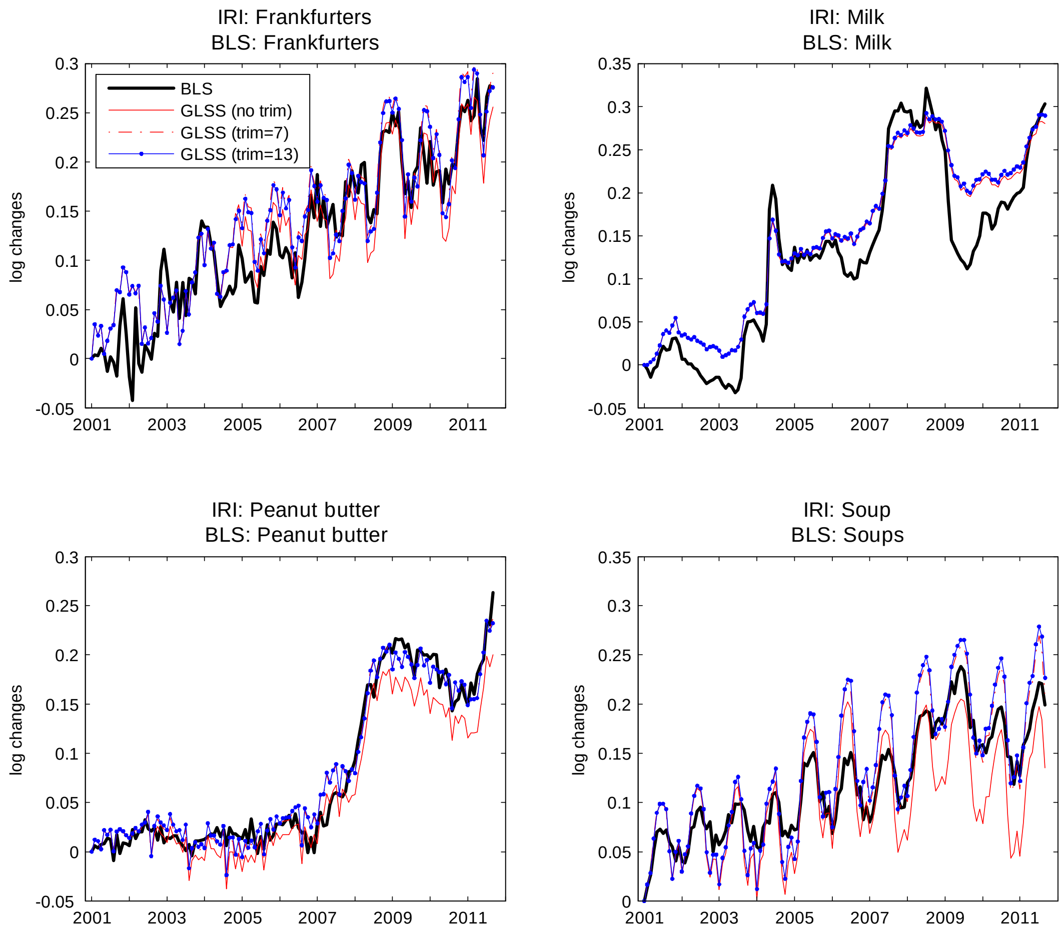

Source: Bureau of Labor Statistics and authors' calculations using IRI data.
Notes: The series labeled "GLSS (trim=13)" use all of our preferred data filters, including the trimming of the last 13 weekly observations of each price trajectory to control for clearance sales. The series labeled "GLSS (trim=7)" shorten our trimming to 7 weekly observations. The series labeled "GLSS (no trim)" use all of the observations in the sample, and thus do not attempt to control for clearance sales.

Figure A9: Controlling for clearance sales through trimming (continued)

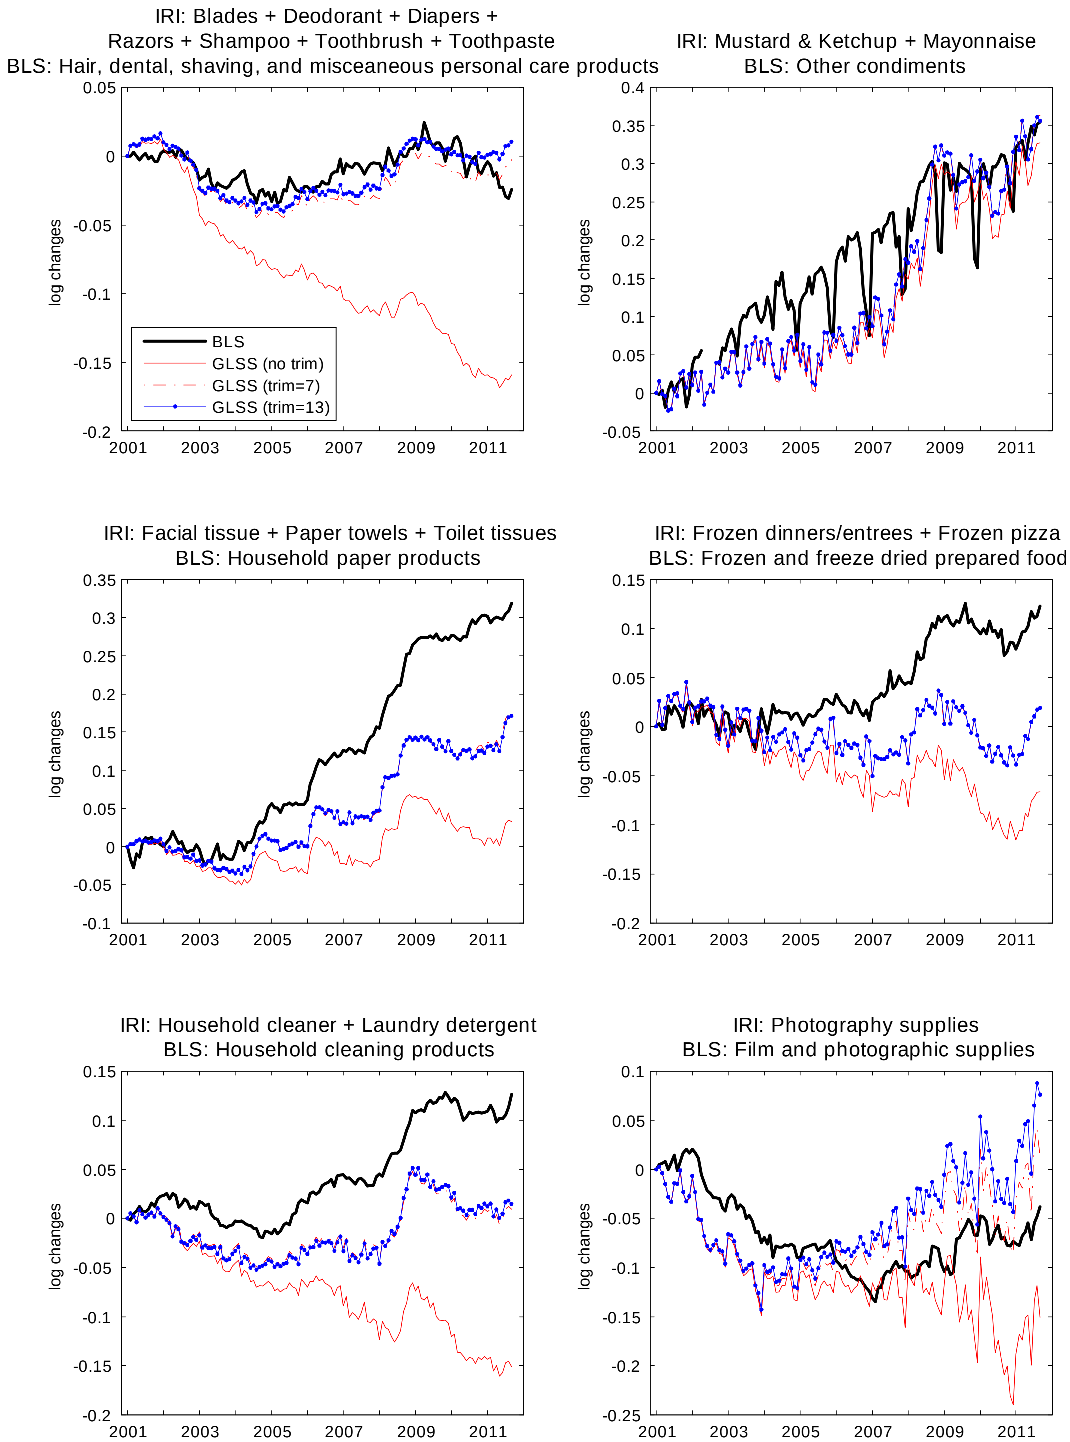

Source: Bureau of Labor Statistics and authors' calculations using IRI data.
Notes: The series labeled "GLSS (trim=13)" use all of our preferred data filters, including the trimming of the last 13 weekly observations of each price trajectory to control for clearance sales. The series labeled "GLSS (trim=7)" shorten our trimming to 7 weekly observations. The series labeled "GLSS (no trim)" use all of the observations in the sample, and thus do not attempt to control for clearance sales.

Figure A10: Controlling for clearance sales through trimming (continued)

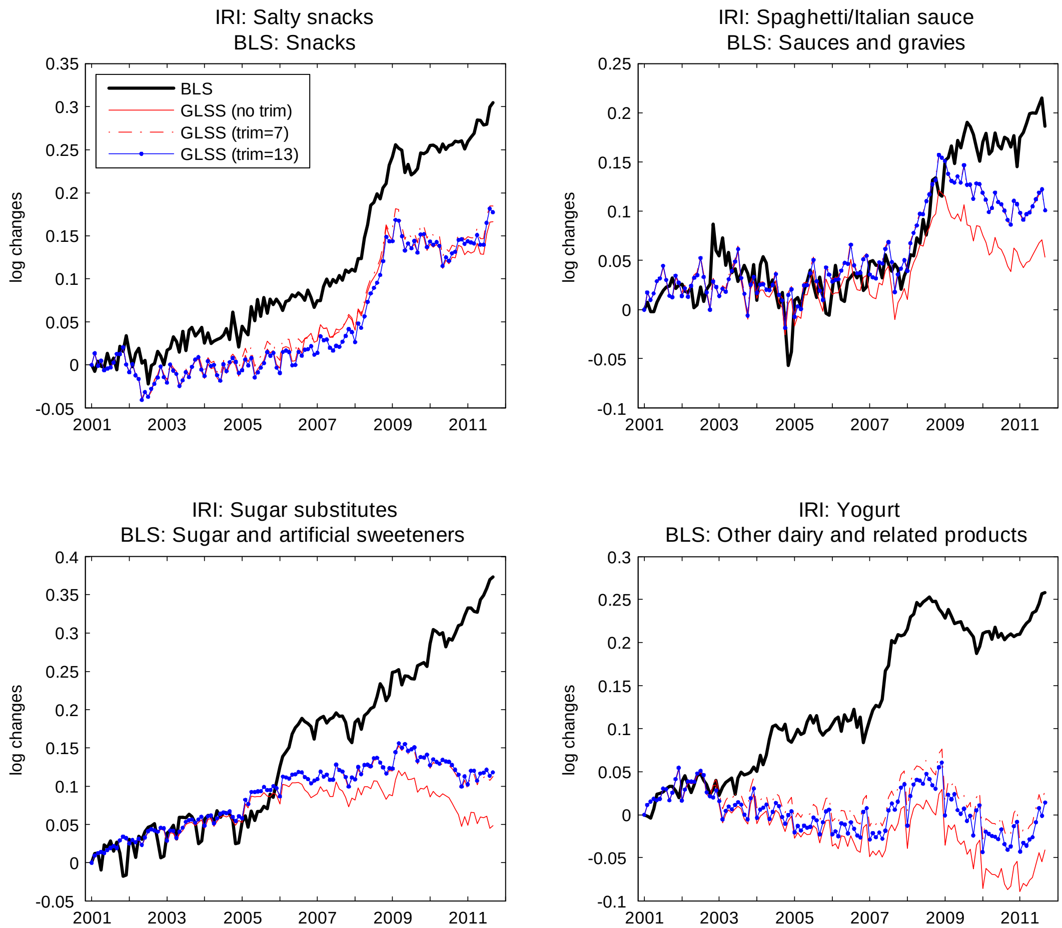

Source: Bureau of Labor Statistics and authors' calculations using IRI data.
Notes: The series labeled "GLSS (trim=13)" use all of our preferred data filters, including the trimming of the last 13 weekly observations of each price trajectory to control for clearance sales. The series labeled "GLSS (trim=7)" shorten our trimming to 7 weekly observations. The series labeled "GLSS (no trim)" use all of the observations in the sample, and thus do not attempt to control for clearance sales.

Figure A11: Frequency of price changes for items whose identifiers differ between GLSS and CGH

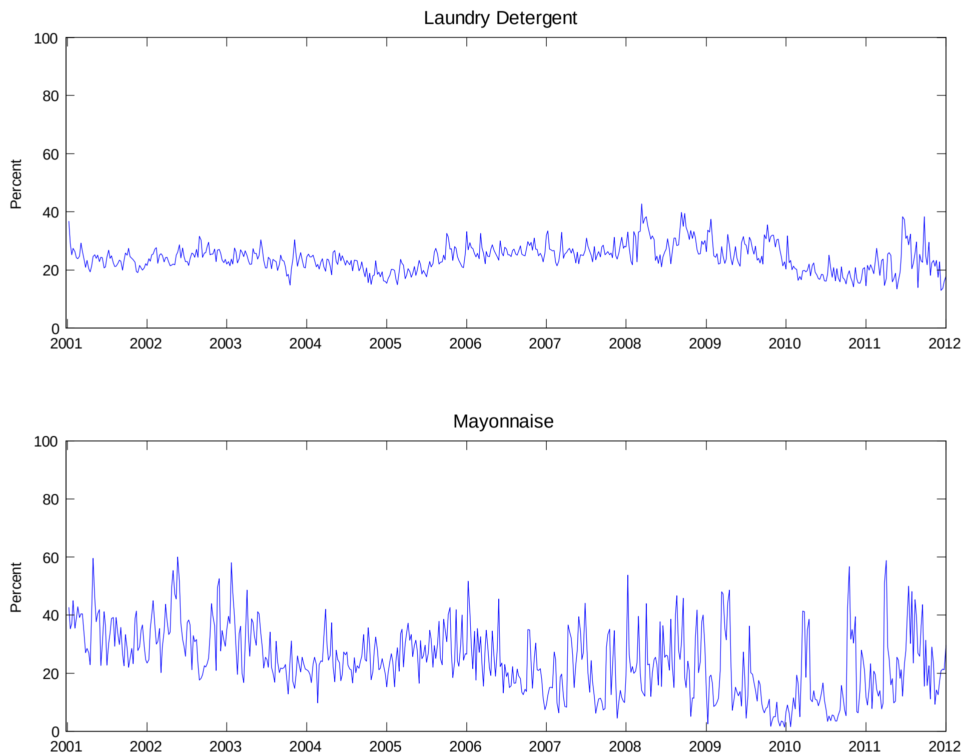

Source: Authors' calculations using IRI data and item identifiers provided by Coibion, Gorodnichenko, and Hong (2015).
Notes: The frequency of price changes is calculated using only items for which the volume-equivalent measure in the UPC definition files differ between the 2001–2007 and 2008–2011 subsamples.