Application of Industry Specific Strata in PPI Variance Estimation

Teresa E. Hesley
U.S. Bureau of Labor Statistics
2 Massachusetts Ave NE, Washington DC, 20212

Abstract

The Producer Price Index (PPI) uses industry specific strata in their sample design. The PPI samples each of its industries independently. Establishments are typically selected at the six-digit NAICS level, though some industries are sampled as high as the three-digit level. Since PPI designs their samples on an industry by industry basis, each sample may have strata based on attributes of that industry. In the initial development of the variance estimation method, which makes use of the bootstrap method, replicates were selected over the whole of the industry. The purpose of this paper is to describe a PPI variance estimation method which takes into account the industry specific stratification.

Keywords: variance estimation, bootstrap, sample design

Note: Any opinions expressed in this paper are those of the author and do not constitute policy of the Bureau of Labor Statistics.

1. Introduction

The Producer Price Index (PPI) is a family of indexes that measures the average change over time in the selling prices received by domestic producers for their output. PPI uses data from more than 25,000 establishments, and approximately 100,000 products per month. These products, referred to as items, are organized into three different index classification structures. The industry structure uses the North American Industrial Classification System (NAICS) coding to publish indexes for over 500 industries and their aggregates. For the commodity classification structure, PPI publishes indexes for over 4,000 goods or services grouped by similarity or material composition, regardless of industry. The third structure is PPI's newly introduced Final Demand-Intermediate Demand (FD-ID) aggregation system, which recently replaced PPI's former headline Stage of Process (SOP) system. The FD-ID structure incorporates PPIs for services, construction, government purchases, and exports, covering over 75 percent of in-scope domestic production of the United States economy.

In 2009 PPI completed an analysis of several different variance estimation methods for their published 1-month and 12-month percent change statistics. The resampling variance estimation methods included Bootstrap, Jackknife, Balanced Repeated Replication (BRR) and Fay's BRR (with several different values of Fay's parameter). The study applied these methods to two years of existing PPI data for selected industries, and compared the results. Many of the methods performed well in terms of stability and confidence interval coverage. In the end, the bootstrap method of variance estimation was chosen because of these reasons as well as ease of implementation.

PPI samples each industry independently. A new sample is selected for an industry every 5-7 years. The study took into account the sampling design and selected bootstrap samples within each industry. For the majority of PPI industries, the industry-level of stratification is as detailed as the sample design becomes. However, for some industries, there are additional levels of stratification within the industry. This is done to ensure coverage of different types of establishments within the industry. When implementing the

bootstrap method, PPI refined the replicate selection process to take these additional layers of sample stratification into account.

2. PPI Sampling

PPI selects samples for approximately 25-30 industries, three times a year. A new sample is selected for an industry every 5-7 years. The reason for this methodology is that the program collects data via personal interview, and sampling in this manner distributes the burden of collection. Which industries are to be selected are based on a number of factors, including the age and attrition of the current sample, as well as the volatility of the industry itself. Once an industry is chosen for sampling, an industry study is done to identify unique characteristics of that industry and to determine if special considerations need to be taken in the sample design and/or collection of that industry's data.

The majority of industry samples will be selected from the Unemployment Insurance (UI) frame, which is a list of all businesses that pay unemployment insurance, using either employment or wages as their measure of size. However, there are quite a few industries that make use of alternative frame sources that may have size measures better suited to their data. The size measures on the frame are basically used as proxies for an establishment's revenue. These are used for weighting, and so the best sampling size measure is determined on an industry by industry basis. While each industry is itself an explicit stratum, the information from the industry study as well as the data available from the frame, determine if further stratification is necessary.

The drawing of an industry sample takes place in two stages. The first stage is the selection of establishments using PPS, and is done at the national office. The second stage happens in the field, when field economists are making personal visits to the establishments selected in stage 1. In stage 2 the field economist selects items, either goods or services, using a disaggregation method. A reporter at the establishment agrees to submit updated prices for these items on a monthly basis.

3. PPI Estimation

The PPI is calculating variance statistics for their 1-month and 12-month percent change estimates, and so it is these estimates that are calculated for each bootstrap replicate. Below is an overview of how these estimates are calculated.

a. Calculating lowest level cell index
   o Calculate long-term price relative for each item with a price that is considered 'good'. A 'good' price essentially means an item price does not require imputation because a valid price was collected for a respondent in a given time period.   Divide the current period price of an item by the base period price of an item:

$$LTR_{k,c,t} = \frac{p_{k,c,t}}{p_{k,c,0}}$$

where $LTR_{k,c,t}$ is the long-term price relative of item k with a good price in cell c at time t, $p_{k,c,t}$ is the price of item k in cell c at time t, and $p_{k,c,0}$ is the price of item k in cell c in base period 0.

o Estimate missing price relatives. The long-term price relative for any item without a good price in a particular month will be estimated based on the long-term relative of the collected items in the cell's sample:

$$LTR_{j,c,t} = LTR_{j,c,t-1} \frac{\sum_{k=1}^{n_g} w_{k,c} LTR_{k,c,t}}{\sum_{k=1}^{n_g} w_{k,c} LTR_{k,c,t-1}}$$

Where $LTR_{j,c,t}$ is the estimated long-term price relative for missing item j $(j = n_{g+1},...,n_c)$, n_g is the number of items with a good price in cell c at time t, n_c is the total number of items in cell c at time t, $LTR_{j,c,t-1}$ is the long-term price relative of item j in cell c at time t-1, $LTR_{k,c,t}$ is the long-term price relative for item k, which has a good price, $k = 1,...,n_g$, $LTR_{k,c,t-1}$ is the long-term price relative of item k in cell c at time t-1, and $w_{k,c}$ is the weight of item k in cell c.

b. Calculate cell aggregate

o To calculate the cell aggregate, multiply the long-term price relative of each item by the weight of the item and sum across the items in the cell:

$$CA_{c,t} = \sum_{i=1}^{n_c} w_{i,c} LTR_{i,c,t}$$

Where $CA_{c,t}$ is the cell aggregate for cell c at time t, $w_{i,c}$ is the weight of item i in cell c, $LTR_{i,c,t}$ is the long-term price relative of item i in cell c at time t, and n_c is the total number of items in cell c.

c. Calculate cell index

o Divide cell aggregate for current period by cell aggregate for previous period and multiply by cell index for previous period:

$$I_{c,t} = \frac{CA_{c,t}}{CA_{c,t-1}} I_{c,t-1}$$

Where $I_{c,t}$ is the index for cell c at time t, $I_{c,t-1}$ is the index for cell c at time t-1, $CA_{c,t}$ is the cell aggregate for cell c at time t, and $CA_{c,t-1}$ is the cell aggregate for cell c at time t-1.

d. Calculate the percent change for an index

$$PC_{c,t,t-m} = \left(\frac{I_{c,t}}{I_{c,t-m}} - 1 \right) 100$$

Where $PC_{c,t,t-m}$ is the percent change of index for cell c from time t to time t-m, m is equal to 1 month or 12 months, $I_{c,t}$ is the index for cell c at time t, $I_{c,t-m}$ is the index for cell c at time t-m.

4. PPI Variances

PPI uses the bootstrap method of variance estimation using 150 replicates. This approach allowed for a fairly straightforward implementation of a variance estimation method.

A. Bootstrap Replicate Selection

For replicate selection, the goal is to draw a set of bootstrap replicates once for a new industry sample, and then carry that set of replicates forward each month for estimation. The reason being that we do not want the additional noise that the re-selection of replicates can add to our variance estimates from month to month. It is not possible, in practice, to maintain a constant set of replicates, as items and establishments may be added to or dropped from an industry after the initial selection of replicates. In these instances the core of the replicates remains constant while items are added to or dropped from them. Replicates are selected when an industry's sample is introduced into estimation. The steps for replicate selection are described below.

a. Form Variance Strata

Within each industry (or industry specific stratum), assign variance strata as follows:

1. All probability establishments are placed into a single variance stratum, with each individual establishment serving as a primary sample unit (PSU) within that variance stratum.

2. Certainty establishments with more than one item become their own variance stratum. Within this variance stratum, each item is as a variance PSU.

3. Certainty establishments in which only one item is collected are paired with similar units, with each item being its own variance PSU. As many as three of these units can be collapsed into a single variance stratum.

   If no like unit exists, this single-item unit should be paired with a multi-item certainty unit to form a variance stratum. Again, each item will be a variance PSU.

   Finally, if there are no other certainty units in the sample, this single-item certainty unit will be included in every replicate.

An example of variance strata within an industry (industry stratum) Appendix A.

b. Select Bootstrap Replicates

After the strata are defined, there will be 150 bootstrap replicates formed from our sample. Each replicate will be selected in the following manner:

1. Of the n_h PSUs in each variance strata, a bootstrap replicate is of size $n_h - 1$. Again, the exception to this, described in section 'a.', is if there is a certainty unit containing a single item, and there were no like establishments with which to pair this certainty unit, this unit/item is the only variance PSU in its' variance strata, and is then be included in every replicate.

2. Using our sample as a frame, we then proceed to select each replicate using random sampling with replacement. This gives each PSU an equal probability of selection, and allows a PSU may be selected more than once within a single replicate.

c. Reweighting

Following Rao-Wu-Yue's (1) rescaling method, we rescale item weights for all items in each replicate:

$$w_{hik}^* = w_{hik} \left\lfloor \left(\frac{n_h}{n_h - 1} \right) m_{hi}^* \right\rfloor$$

where h is the variance stratum, i is the variance PSU, k is the item, w_{hik}^* is the bootstrap weight, w_{hik} is the original item weight, n_h is the number of original sample units in the variance stratum, and m_{hi}^* is the number of times variance PSU i in stratum h is selected in a given replicate.

B. Calculate Bootstrap Variance Estimate

This is done for each percent change estimate calculated by PPI. The replicates that are selected at the detailed industry (industry-specific stratum) level are applied to the PPI tree structures to produce estimates for all 1-month (initial, interim and final) and 12-month percent price change estimates.

$$v_{BT} = \frac{1}{B} \sum_{b=1}^{B} \left(\hat{\theta}_b - \hat{\theta}_{full} \right)^2$$

Where B is the number of bootstrap replicates (in current implementation, $B = 150$), $\hat{\theta}_b$ is the percent change estimate for a given replicate, and $\hat{\theta}_{full}$ is the percent change estimate of the original sample.

5. Difference in Study vs. Implemented Bootstrap Replicate Selection

The difference in forming variance strata between the initial variance study and the variance code that was implemented in PPI is simple, but fundamental. The initial study formed variance strata and selected bootstrap replicates across the whole of a sampled industry, as defined by the sampled NAICS code. These industry codes are explicit strata in the whole of the PPI sample, and for the majority of PPI industries, this is the level of detail for which the bootstrap samples should be selected. However, many PPI industries are further stratified by any number of factors. As the goal of this additional stratification is to reduce sampling error, these additional stratifications must be taken into account during the setup of the variance strata.

The further stratification of an industry can happen for several reasons. One reason is various establishment characteristics have been identified for a given industry, and so that industry is stratified to make sure establishments with those specific characteristics are represented. An example of this is types of ownership for fitness centers. Some establishments have their prices set at a corporate level, others are franchises, and many are individually owned. Another reason is that the sample for a given industry is composed using multiple frames. Many industries, most typically in services, will use another source of data in lieu of, or in addition to, the Unemployment Insurance (UI) file. Also, some industries are sampled at a higher NAICS level, and then further stratified by more detailed NAICS codes during the selection process. This is the

case for our wholesale industries, which are sampled at the 3-digit NAICS level and then further stratified by 4-digit NAICS during sample selection. Finally, sample strata can be assigned when one or more industry is combined to create a new industry to better reflect the economic realities we are trying to capture. No new sample is drawn in this case, and so each original industry that comprises the new industry is assigned its own sample strata code for use in the formation of variance strata. Of the 500 industries in PPI, 118 have industry specific strata.

6. Estimate Comparison

This section focuses on the difference between the standard error values calculated when the bootstrap replicates were selected over the whole of the industry and when the bootstrap replicates used industry specific strata in their formation.

Data

Variables used in analyzing the differences between the SE estimates are:
- Sampled Industry (6-digit NAICS): 118 industries affected
- Industry Sector (2-digit NAICS): 18 sectors
- Revision Month: PPI estimates have a 4 month revision period denoted N (initial estimate) through N-4 (final revision). At N-4 the index and its associated percent changes are final. In this comparison we will look at differences in the initial time period (N) and the final time period (N-4).
- Roll-up Indicator: As described above, some industry strata are created during sampling, and some are created when thinly populated detailed industries are combined into a single industry, without resampling, during estimation. This indicator is 'Y' for those strata that were the result of a roll-up vs. 'N' for strata created during sample selection.
- Industry SE estimates when replicate selection did not include industry specific strata
- Industry SE estimates when replicate selection included industry specific strata
- The percent change of the SE for the industry once the industry specific strata were used in replicate selection. This is the dependent variable.

The time period for this study is May 2013 through December 2013. This includes the N-4 final percent change estimates and associated standard errors for January of 2013 through August 2013, and the initial, N, values for May through December of 2013. There are 118 industries that use additional strata in the studied time frame.

Analysis

First, the overall chart of the changes for the final revision period. This was to get a quick snapshot of the movement of the SEs for the final PPI data. With 118 industries over 8 months, Table 1 has 938[1] observations for percent change in the standard error (PCSE).

---

[1] 118*8 = 944, but a few industries introduced new samples with strata during that time frame, and so did not have sample strata for all 8 months in the study.

Table 1 – PCSE for finalized PPI data.

Percent Change in the Standard Error (PCSE): Final Revision Jan. 2013 – Aug. 2013

The far left column, in dark green, reflects the SE estimates that showed the most improvement (decrease in SE) when including the industry specific strata in the selection of bootstrap replicates, while the far right column shows the number of SE estimates that increased with the inclusion of industry specific strata. Table two further breaks down the data showing the counts and average percent change in the SE for each of the above categories.

Table 2 – Count and Average PCSE for finalized PPI data.

	Counts	Average PCSE
Over 20% Increase in SE	45	32.13
10%-20% Increase in SE	103	14.25
5%-10% Increase in SE	112	7.38
0%-5% Increase in SE	162	2.47
0%-5% Decrease in SE	192	-2.34
5%-10% Decrease in SE	143	-7.23
10%-20% Decrease in SE	128	-14.32
Over 20% Decrease in SE	53	-33.41
Grand Total	938	-1.01

While there is a slight shift towards an improvement, it is not quite as dramatic of a change as expected. One possible reason is that although sample strata are used during sampling, it is not always with the idea

of improving the sample error, but to ensure that publication goals will be met. SE estimates are just now being used in the allocation process to help lower sample errors.

Next we further examined the impact of some of our independent variables on the dependent variable of PCSE using linear regression.   The independent variables used include Sampled Industry, Revision Month, and Roll-Up Indicator. Sampled Industry will shed light on which industries were most impacted by the use of industry specific strata in bootstrap replicate selection. Intuitively, Revision Month should have an impact, as the indexes tend to become more stable further along in the revision period. Finally, the Roll-up Indicator is included to see if industries where strata are specified during sampling differ from strata that are formed during roll-ups in estimation.

$$y_i = \alpha + \beta_1 x_{i1} + \beta_2 x_{i2} + \beta_3 x_{i3} + \epsilon_i$$

The sampled industry ended up showing strong significance, with a p-value of <.0001. The Revision Month was also significant with a p-value of 0.0201. Whether or not the industry strata was the result of a roll-up was not significant, with a p-value of 0.5134. However, examination of the revision month boxplots revealed two large outliers (Table 3).

Table 3

These outliers were removed from the dataset and then the regression rerun. The sampled industry is still significant (p<.0001). The revision month maintains a low p-value, but is no longer significant at the 95% test (p=0.0603). The roll-up indicator is still not significant, with p=0.3576.

Next a Tukey test was performed to see which industries differ significantly from one another. Though several industries significantly differ from one another, one stuck out by having a significantly larger average PCSE than 116 of the other 117 industries in the study. NAICS 334514, Totalizing Fluid Meter and Counting Device Manufacturing. The explanation for this is that 334514 has been in estimation for a long time and has gone through much attrition, and so its estimates are subject to more variability. This

industry has been resampled, but at the point in time of the study, the new sample was still in collection, and had not been introduced into estimation.

On the other side of the spectrum, the industry that saw the most improvement, when compared to other industries was 515210 (Cable and Other Subscription Programming). The PCSE was significantly better than 35 of the other industries. This can be attributed to favorable stratification, as the average SE estimate for this industry prior to strata application was 2.033, and saw an average reduction of nearly 20%. The range of PCSE for this industry was -5.245% to -43.4996%. This industry had been sampled using the strata National and Regional channels.

7. Conclusion

In conclusion, the inclusion of industry specific sample strata in bootstrap replicate selection provided an overall improvement in standard error estimates. Industries that saw consistent *increases* in their SE estimates could be explained by the age and attrition of the sample. Many industries saw both increases and decreases. As mentioned, this could be that sample strata were formed with the intent of covering the PPI publication structure, and not necessarily created for the purpose of reducing sampling error. The industries that saw consistent improvement could be attributed to strata definitions that helped reduce sampling error.

References

1. Spease, C., Sadler, A., Chen, H., & Keith, B. (2009). PPI Variance Estimation Study.

2. Rao, J.N.K., C.F.J. Wu & K. Yue. (1992). Some recent work on re-sampling methods for complex surveys. Survey Methodology 18 pp. 209-217.

3. Efron, B., & Tibshirani, R. J. (1993). An Introduction to the Bootstrap. Chapman & Hall/CRC.

4. Chernick, Michael R. (2008). Bootstrap Methods: A Guide for Practitioners and Researchers. 2nd edition. Hoboken, NJ, John Wiley & Sons, Inc.

5. Wolter, K. M. (2007). Introduction to Variance Estimation. 2nd edition. Chicago, IL. Springer.

Appendix A - example of variance strata within an industry (industry specific stratum):

Industry XXXXXX:

Probability Establishments (Variance Stratum #1)	
PSU = Establishment	**Items**
Establishment 1 (PSU = 1)	8 Items
Establishment 2 (PSU = 2)	8 Items
Establishment 3 (PSU = 3)	6 Items
Establishment 4 (PSU = 4)	4 Items
Establishment 5 (PSU = 5)	4 Items
Establishment 6 (PSU = 6)	4 Items

Certainty Establishment 1 (Variance Stratum #2)	
Establishment	**PSU = Items**
Certainty Establishment 1	Item 1 (PSU = 1)
	Item 2 (PSU = 2)
	Item 3 (PSU = 3)
	Item 4 (PSU = 4)
	Item 5 (PSU = 5)
	Item 6 (PSU = 6)
	Item 7 (PSU = 7)
	Item 8 (PSU = 8)

Certainty Establishment 2 (Variance Stratum #3)	
Establishment	**PSU = Items**
Certainty Establishment 2	Item 1 (PSU = 1)
	Item 2 (PSU = 2)
	Item 3 (PSU = 3)
	Item 4 (PSU = 4)
	Item 5 (PSU = 5)
	Item 6 (PSU = 6)

Certainty Establishments 3 and 4 (Variance Stratum #4)	
Establishment	**PSU = Items**
Certainty Establishment 3	Item 1 (PSU = 1)
Certainty Establishment 4	Item 1 (PSU = 2)