POINT OF PURCHASE SURVEY AND ITS USE IN THE CONSUMER PRICE INDEX REVISION

Peter A. Bounpane, U.S. Bureau of the Census
Curtis A. Jacobs, Bureau of Labor Statistics

I. Introduction

The Consumer Price Index (CPI), published monthly by the Bureau of Labor Statistics (BLS), is a measure of the changes in the retail prices of consumer goods and services. A sample of goods and services is selected for pricing in a sample of outlets.

The Bureau of Labor Statistics and the Census Bureau cooperated to conduct a series of household surveys in the early 1970's, which are being used by BLS for the revision of the CPI. This paper discusses one of those surveys-The Point of Purchase Survey: the sample design and the use of the survey results in the CPI Revision.

The Point of Purchase Survey (POPS) was designed to provide the sampling frame of outlets for food and for most commodities and services to be priced in the CPI and to provide corresponding demographic data to classify the households which reported an expenditure for an outlet. Two other surveys, the Quarterly and the Diary Expenditure Surveys, will be used to update the sample of items and weights of the CPI, but we will not directly concern ourselves with either of these surveys in this paper.

The extensive use of a household survey (POPS) to generate a list of outlets is new. Previously, existing lists of retail outlets were one method used to select the sample of outlets for pricing. For some areas, an approach similar to POPS was used.

Because of the new approach to the design of POPS, there was extensive testing prior to actual implementation. Various "hot houses" were used to test the actual questionnaire. Two full pretests (Atlanta and Cleveland) were conducted. The results of all pretesting were carefully analyzed and used as input to the development of the initial POPS.

II. Elements of the Sample Design

The development of the sample design was complex and required some assumptions because of the unique nature of the survey. The analysis of the assumptions is not fully completed. The household sample size for POPS was determined by many elements and constraints.

1. The areas within which the household survey (and thus the pricing) is made.
2. The number of households to sample per area.
3. The number of POPS categories to ask about on the questionnaire.
4. The length of the recall period for the category (which determines the number of outlets reported for the item).
5. The expected number of outlets to be selected.
6. The anticipated number and method of item selection within selected outlets.

Obviously, the decisions on the above elements of the sample size are interrelated. Many of the decisions reached were based on assumed but unknown parameters of the CPI revision scheme. Others were based on results of the POPS pretests while others were based on the opinions of the survey designers. Detailed discussions of the area design, POPS category definition, household sample selections, and the CPI item/outlet and store-item within outlet selection procedures appear in separate sections of this paper.

The Area Design

POPS was enumerated in an 85 PSU area design selected for CPI pricing. This design resulted in 27 strata with one pricing area per stratum (self-representing PSU's) and 58 nonself-representing strata. In selecting the one sample PSU for each nonself-representing stratum, a controlled selection program was used to insure the sample areas were properly distributed geographically across States. As planned, the resulting 58 area design contained at least four pricing areas within each of the 12 region-city size classes.1/ Since separate frames of outlets were required for certain individual pricing areas (PSU's), sample sizes in such PSU's had to be more than those generated by application of a uniform national sampling fraction. Hence, the POPS "national" household survey is not self-weighting. Within a PSU, however, the household sample is selected with a uniform probability.

"Sample Size" of POPS Categories

The actual categories to inquire about on the point of purchase questionnaire were selected from a larger list of potential categories. The factors in the determination of categories were:

1. Appropriate categories to allow for the outlet/item selection explained in later sections of the paper.
2. A sampling frame for certain categories was already available and hence such items were eliminated from POPS. (Example: utilities, insurance, transportation other than automobile.)
3. Emphasis on apparel categories because of previous difficulties with pricing apparel items. (Often, the outlet did not stock the specific item to price)
4. An extension of food categories based on the pretest results. (Using more categories gives better representation of food outlets other than large grocery stores.)
5. POPS was not appropriate for some items, especially those purchased rarely.

Household Sample Size

The household sample size was determined with the following guidelines:

1. The total sample size had to be around 23,000 because of budgeting restrictions.
2. Within one of the PSU's for which separate frames for item selection were to be developed, the number of households selected was chosen to yield an expected number of outlet responses for a POPS category equal to three times the desired number of expected outlets to be selected. The desired number of outlets was four per POPS category (for example, men's shirts) per replication (half-samples: two in SR PSU's and one in each NSR PSU).
3. Another factor in determining the expected number of outlet responses is the length of the

recall period per category. Different categories had different recall periods. The final household sample size was a balance of number of selected housing units and recall periods. Pretest experience was extremely helpful in such decisions.
4. Within the other PSU's, an equal sample size per PSU was chosen to be more than a certain minimum to allow for an appropriate enumerator workload, and to keep within the 23,000 constraint.
5. Each desired sample size was inflated by 20 percent to allow for nonresponse.
6. Each desired sample size was inflated by another 25 percent to allow for having the appropriate number of outlet responses for the 80 percent "middle income" urban families. That is, only certain types of families were going to be considered eligible. It was felt that such families are about 80 percent of all families.

III. Selection of the Household Sample

With the introduction of the household survey approach, it was decided to attempt to highly cluster sample households. The assumption was that if families tend to buy in the areas where they live, the outlets given as responses to the survey would also be clustered. In order to increase the expected chance of clustering outlet responses, the household clusters were formed (where possible) around known shopping complexes. Within a cluster, it's desirable to have the selected households as dispersed as possible. Time and cost (sampling and enumeration) considerations however, led to the introduction of another stage of selection within selected clusters. Within a cluster of tracts, a sample of ED's was selected and within the selected ED's (Census enumeration districts), the sampled households were dispersed evenly. Five housing units were selected in each ED and since the total sample size per cluster was desired at 40 housing units, about eight ED's were in sample in each cluster.

To accomplish the sampling, a PSU was first divided into SSU's (secondary sampling units), an SSU being a set of contiguous census tracts around a shopping complex. The first step to form these SSU's was to spot on a tract map all known shopping complexes. These shopping complexes were the Central Business Districts, Major Retail Centers, (as defined by the 1967 Census of Business), and other shopping centers as defined by the 1970 Directory of Shopping Centers. Census tracts were then grouped into SSU's around each of these spotted shopping complexes. The guidelines for forming these SSU's were as follows:
1. All tracts containing part of the shopping complex should be in the same SSU.
2. Each SSU should ideally have only one shopping complex.
3. The SSU should have a minimum of 2,500 housing units.
4. The SSU should be no more than two miles square.

This operation did not account for all the tracts in a PSU. The remaining tracts were merely put into contiguous groups according to above guidelines 3. and 4. as well as the following:
5. If possible all tracts in an incorporated or unincorporated place outside the central city of the SMSA and greater than 2,500 population should be an SSU.
6. SSU's should not cross the central city or urbanized area boundaries.

These guidelines were just that, and so they were not always met. Only a few SSU's had fewer than 1,500 housing units and/or were larger than five miles square.

Once the PSU has been entirely subdivided into SSU's, each SSU was assigned an economic measure. This measure was a weighted average of the median family income of the tracts in the SSU as reported in the 1970 Census. The weights were the 1970 Census housing unit counts. The SSU's were then ordered as follows:
1. SSU's in the Central City were ordered from the lowest to the highest value of the economic index.
2. SSU's in the Balance of the Urbanized Area of the SMSA were ordered from the highest to the lowest value of the economic index.
3. The remaining SSU's were ordered from the lowest to the highest value of the economic index.

These orderings were used to reduce the tail-end variance of the systematic selection procedures.

With the SSU's in this order, a sample of the appropriate number was selected PPS to the 1970 housing unit count of the SSU. This cumbersome economic-geographic stratification was undertaken in order that the sample of SSU's would be a better cross-representation of the entire PSU.

Before proceeding to within SSU selection, it's necessary to introduce the concept of new construction. The scheme now being described would only lead to a selection of housing units actually listed in the 1970 Census. An additional operation was necessary to represent units constructed since the 1970 Census (new construction). It was too costly to determine all the new construction units and assign them to their appropriate SSU. Thus, the Census Bureau sampled new construction units in each PSU in their normal manner (off permit registers) without regard to SSU. Sample new construction units are, therefore, spread throughout the PSU. Since cost restraints required the total sample size in a PSU be rigidly controlled, some adjustment to the eight ED's per SSU, five housing units per ED rule was required. An added complication was that it was necessary that the number of sample SSU's be even to facilitate variance calculations.

Within a selected ED, a sample of housing units was selected as follows:

Since we wanted a self-weighting sample

$$\frac{1}{k} = \text{Probability of selecting an SSU } (P_1) \times \text{Probability of selecting an ED within a PSU } (P_2) \times \text{Probability of selecting a housing unit in an ED } (P_3)$$

$$= m \frac{H_{i.}}{H} \times n \frac{H_{ij}}{H_{i.}} \times \frac{a\ r}{H_{ij}}$$

where

k = desired sampling rate

H_{ij} = housing unit count of j^{th} ED in the i^{th} SSU

$H_{i.}$ = housing unit count in i^{th} SSU

$H = \sum_i H_{i.} = \sum_i \sum_j H_{ij}$

m = desired number of sample SSU's

n = desired number of sample ED's per SSU

r = expected number of sample HU's per sample ED

a = a value to create equality

The within ED sampling interval is then one over the probability of selecting a housing unit in an ED

$$TE = \frac{H_{ij}}{a\ r}$$

where a is solved from the knowns of the above equation.

In areas when the ED selection was skipped (the number of ED's in the SSU was 10 or less), $P_2 = 1$.

This operation yielded a selfweighting sample within a PSU since the new construction cases were sampled at 1 in k. Once the sampling rate (noninteger) for an ED was selected an appropriate random start was selected. Then a systematic sample of housing units was selected from all housing units listed in the ED.

The above operation was only followed in 60 of the 85 sample PSU's. These 60 were those PSU's with 1970 population of 250,000 or more. It was felt that the remaining 25 PSU's were too small to bother with the expensive and time-consuming operation of SSU formation since we were liable to be in every SSU anyway. These 25 areas were sampled as follows.

The total sample size was split into estimates of the amount expected from the 1970 Census and the amount expected from new construction. An appropriate number of ED's was determined such that there would be as close to possible as five housing units per ED. Since it was just as difficult to allocate new construction to ED as was to allocate it to SSU, new construction was sampled across the PSU regardless of the sample ED's.

Once the number of sample ED's was determined, that number of ED's was selected from all ED's PPS to the 1970 housing unit count. The ED's had previously been ordered according to the 1970 median family income for the ED. Within ED, the same scheme as for SSU's was followed to specifically designate the housing units in sample.

$\frac{1}{k}$ = Probability of selecting an ED X Probability of selecting a housing unit in the ED

$$= \frac{nH_{ij}}{H} \times \frac{a}{H_{ij}}$$

Hence $a = \frac{H}{nk}$

and $TE = \frac{H_{ij.}}{a} = \frac{nk\ H_{ij}}{a}$

The initial plans of BLS were to expand the present CPI population coverage (urban wage and clerical worker consumer units) to include all civilian noninstitutional consumer units. The specifications for the POPS survey were predicated on this assumption. Subsequently, it was decided to continue publication of an index for the urban wage and clerical worker population with the same degree of reliability as the current index, as well as to publish the more extensive index. Hence, the sample design problem was to select two samples of items and outlets such that two indexes of at least the same reliability at the national level as the present CPI could be published and also meet cost constraints. The new CPI distinction between populations is maintained primarily by the selection of outlets reported in the POPS for the appropriate population, whereas the current CPI relies on a subjectively specified level of quality of items for the population to be priced.

The following portions of the paper describe the sampling procedures to select the items, outlets, and specific store items within outlets, for the two family indexes. In addition, we'll briefly outline the parameters developed to optimize the CPI design, as it relates to the Point of Purchase Survey components.

IV. Background of Item Selection

In order to fully understand the use and effect the POPS survey had on the CPI Revision (CPIR), it is necessary to describe the relationship between the item sampling frames (results of Quarterly and Diary Surveys) and the outlet sampling frames (results of POPS). The expenditure data for the CPI is constructed from the Diary and Quarterly Consumer Expenditure Surveys conducted by Census for BLS in 1972-73. Only the first year Diary or Quarterly data was used for selection of items, since this data was all that BLS could process prior to item selection. The Diary data was the sampling frame for items frequently purchased, such as food and personal care items. The Quarterly data was used for all other selections.

The basic structure created for both the Diary and Quarterly was to define 71 Expenditure Classes (EC's) which are BLS's primary definition of publication levels of indexes. Within each ED, the expenditures defined therein were grouped into one or more item stratum. Within each item stratum, one or more substratum called Entry Level Items (ELI's) were defined. The ELI's are the ultimate sampling unit selected in Washington and are used by the data collectors as their initial level of item definition within an outlet. These ELI's are relatively broadly defined groupings of items and allow the possibility of pricing many different kinds of specifications or specific store items than a single specification as previously used in the CPI. The ELI's are mutually exclusive and account for all consumer expenditures reported in the Expenditure Surveys for an item stratum.

This structure of the Diary and Quarterly into item strata and ELI's is the same for all regions and Market Baskets; however, for sampling purposes, BLS tabulated four regional universe market baskets for each population to reflect regional differences rather than the national market basket as used in the current CPI. A single independent selection of the ELI's for each of the item stratum defines a half-sample. Within each region, BLS made eight independent selections of ELI's within each item stratum forming eight "half-

samples"[2] of ELI's. These half-samples were distributed among the CPI pricing areas (PSU's) for pricing within the region. The reason 32 half-samples were selected for each population was to reduce significantly the correlation between pricing areas for the national index. The optimum strategy would have been to make the number of selections equal to the number of half-samples in the region but workload requirements dictated that a compromise be made; however, the 32 selections provide the major portion of the possible reduction of the correlation between item samples.

Because of the requirement of publishing two family indexes, the following technique was used for item selection. Each selection of ELI's within an item stratum was made initially for the urban wage and clerical population (W) proportional to the relative expenditure of ELI's within an item stratum for the W population. Then using a technique developed by Nathan Keyfitz[3] to maximize the overlap of ELI's between populations and maintain the correct probabilities of selection, a second selection of ELI's for the all urban consumer unit population (U) was made proportional to the relative expenditures of the U population.

Each ELI was defined to be in one and only one POPS category so that the integration of the ELI sample and the outlet sample was by the ELI/POPS category concordance. For a given population/half-sample, a single selection of outlets in a POPS category was used to price up to seven ELI's for Commodities and Services (C&S) items and 14 food ELI's. Thus, the selection of the ELI identified which POPS categories were to be used for outlet selection.

V. Outlet Selection from POPS Categories

The following basic approach was used for outlet selection. Make a systematic selection of outlets reported for a given POPS category for the W population where the measure of size for each outlet was proportional to the average daily expenditure reported for the outlet by all households of the W population. The outlets for the U population were then selected by using a Keyfitzing technique to recompute the measure of size for every outlet in the universe and the sample outlets for the U population were then selected by a repeat of the systematic selection using the new measure of size.

The specific algorithm is as follows:
Let a sampling frame consist of N elements each with two measures of size U_{ik}, where i denotes the i^{th} element and k denotes the distinct measure of size, k=1, 2. Assume that

$$\sum_{i=1}^{N} U_{ik}=1 \text{ and } U_{ik} \leq \frac{1}{n}$$

where n is the number of elements in the sample to be selected. We choose a sample of size n using the measures of size U_{i1} in a systematic fashion.

Then choose a sample systematatically with measures of size w_i for the i^{th} element defined as follows:

Let $h_i=1$ if the i^{th} element was selected in the first sample, zero otherwise. Then define

$$W_i=\begin{cases} \frac{1}{n} h_i \, U_{i2}/U_{i1} \text{ if } U_{i1} \geq U_{i2} \\ \frac{1}{n}\left(h_i+\sum_{j=1}^{N} h_j \frac{(U_{j1}-U_{j2})+(U_{i2}-U_{i1})}{U_{j1}} \frac{}{\sum_j (U_{j2}-U_{j1})^+}\right) \text{ if } U_{i1}<U_{i2} \end{cases}$$

where if x is a real number, then $x^+=x$ if $x>0$, and $x^+=0$ if $x\leq 0$.

The samples of ELI's and outlets were then merged and listings of the sample outlets and the sample ELI's designated within the outlet were prepared on a "facesheet" for the data collector.

VI. Within Outlet Selection for Specific Items

For each ELI, the selection of a specific store item by a data collector is done using multi-stage probability selection techniques with measures of size proportional to percentages of dollar sales usually provided by the respondent of the outlet. This procedure is new for the CPI which formerly asked for the "volume seller" for a tight specification describing a specified quality of item required for pricing.

For example, the old procedure required an outlet be selected to price a shirt of a specified grade of fabric (a tight specification as opposed to a broad ELI-shirt). Then, if more than one shirt met the required specifications, the "volume seller" brand and style was selected for pricing. In the revision, an outlet is selected for shirts. A specific shirt is chosen by sampling proportional to the percentage of sales of all the categories and types of shirts in the outlet. The detailed description of the shirt is recorded for future pricing.

To perform this operation, the data collector is provided with a checklist that includes all the price determining characteristics of items defined within the ELI. In addition, the data collector is given the definition of the ELI, suggested stages of groupings of items to aid in quickly selecing a specific store item and a series of worksheets on which to define the categories of items, post the probabilities and identify the next category within which to select the specific store item by use of the random number table on the worksheet.

In developing this procedure, it became necessary to provide the data collector with several alternatives for defining the categories and obtaining the percentage of dollar sales or approximations to those sales. The procedures developed to obtain the proportion of sales were:
1. Obtaining the proportions directly from a respondent.
2. Ranking the categories (by respondent) and then obtaining the proportions directly or using preassigned proportions.
3. Using shelf space to estimate the proportions where applicable.
4. Using equal probability if all else fails. To define the categories, direct responses from the

respondent as to what he sells or an inventory technique are used.

These procedures make possible the use of objective probability sampling techniques down to the specific source item within each outlet. They also allow broad definitions of ELI's so that the same tight specification need not be priced everywhere. The wide variety of specific items greatly reduces the within EC item variance and allows a substantial reduction in the number of quotes required to obtain the same reliability as the current index. A second important benefit from the broader ELI's, along with the POPS categories, is a significantly higher probability of finding a store item within the definition of the ELI within the sample outlet.

VII. Optimization

Since a given outlet could be reported for several POPS categories and had several chances of selection, the determination of needed sample sizes of items and outlets so as to meet reliability requirements and cost constraints was not straight-forward. Design parameters involving the amount of overlap for selection of outlets and items, and the effect of Keyfitzing the two samples (W and U) between populations, were determined empirically. These design parameters, along with estimates of the components of variance, intra-class correlations and unit costs were used to optimize the allocation of items and outlets for the portion of the CPIR covered by the POPS survey. This analysis was done separately for the food at home portion and the commodities and services (C&S) of the CPI.

Additional factors in the model accounted for percent personal visit and telephone interview, frequency of pricing (monthly, bimonthly, and quarterly pricing) and separate cost functions for initiation and for ongoing costs.

IX. Conclusion

It's too early to fully evaluate the POPS since the process of outlet and item selection is still in progress. Our aim here was to record what took place in the current POPS. Preliminary indications, however, are that the POPS is an effective method (both from an operational and cost standpoint) of item selection. Field work has begun in sixteen areas. Selection of items is scheduled for completion by March 1976.

There is some data available on one aspect of the POPS design. The outlets reported by respondents have been examined in terms of their geographical proximity. Though the analysis is incomplete, it appears that clustering households did not produce the desired clustering of outlet responses. Further examination of this assumption and the other aspects of the POPS design are worthy of further study. Such study will be helpful in any future "POPS" type survey.

This version is a short version of the presented paper. Copies of the full paper are available from the authors.

1/ The 12 classes were three city sizes (SMSA's of 400,000 or more population, other SMSA's, and Non-SMSA's) crossed by the four Census regions.

2/ For computation of variances, two selections of items are necessary. Two half-samples of ELI's are used in each self-representing PSU and one half-sample of ELI's is used in each non-self-representing PSU.

3/ Nathan Keyfitz, "Sampling with Probabilities Proportional to Size: Adjustment for Changes in the Probabilities," JASA Vol. 46 (March 1951) pp. 105-109.