



An Overview of the State-Level Weighting Procedure for the Consumer Expenditure Survey

Susan King, Taylor Wilson, and Sharon Krieger

February 2021

Consumer Expenditure Surveys Program

Introduction

This document presents an overview of the process used to create state-level weights for the Consumer Expenditure Surveys (CE) public-use microdata. Every household in the sample has a weight indicating the number of households in the population it represents. This document describes the process in which those weights are modified to make them represent the state in which they are located rather than the region of the U.S. in which they are located. The process modifies the households' base weights along with their nonresponse and calibration adjustment factors that are used to generate their final weights. That allows the modified weights to be used to make statistically valid state-level expenditure estimates. State-level weights can only be provided for states where there is a sufficient sample size and every type of primary sampling unit (PSU) in the state is adequately represented. The goal of providing state-level weights is to support state-level data analysis by public-use microdata users.

Background

The CE survey is a nationally representative household survey that collects expenditure data from a random sample of households across the U.S. The survey program produces expenditure estimates using the data collected from those households, and the data needs to be weighted to properly reflect the households' probability of being selected for CE's sample and participating in the survey. The sample design is a two-step process in which a random sample of geographic areas is selected from the U.S., and then a random sample of households is selected inside those selected areas. The geographic areas are small clusters of counties called PSUs, and the process is designed to produce unbiased expenditure estimates at the national, U.S. Census Region, and U.S. Census Division levels, but not at the state level.

The U.S. Census Bureau partitions the U.S. into four geographic "regions" (Northeast, Midwest, South, and West), and it partitions each region into two geographic "divisions" (except the South, which is partitioned into three divisions). Then a random sample of PSUs is drawn from each division, and a random sample of households is drawn from each PSU. Prior to 2017, CE's weights were calibrated to the four Census region population totals, but beginning in 2017, they are calibrated to the nine Census division population totals. This allows statistically valid expenditure estimates to be made at the division, region, and national levels, but not at geographic levels below divisions and regions (such as the state level). The gap at the state level

is filled by re-calculating the base weights from the sample of households in the state, adjusting the weights to account for nonresponse, and calibrating the weights to their state population totals. The state-level weights are statistically valid only for state-level estimates. They are not statistically valid for lower geographic areas, such as counties within a given state, or higher geographic areas such as divisions, regions, or the nation.¹

Calculating the State-Level Weights – Overview

The state-level weights are computed in a three-step process that starts with “base weights,” which are the number of households in a state’s population that a household in the sample represents. It is typically around 15,000, which means that one out of every 15,000 households in a state is selected for the survey. Then the base weights are increased by multiplying them by a nonresponse adjustment factor to account for households that were selected for the survey but did not participate in it. And finally, the weights are multiplied by a calibration adjustment factor to make the weights of all the respondent households in the state’s sample add up to the state’s true population. The final weights are typically around 50,000, which means that one out of every 50,000 households in a state participates in the survey.

CE’s Geographic Sample Design

Before describing the way base weights are calculated, CE’s geographic sample design needs to be described. First, the U.S. has 3,143 counties or county equivalents which are grouped into 1,470 PSUs. Those PSUs are the 918 “core-based statistical areas” defined by the U.S. Office of Management and Budget for urban areas of the U.S.,² and 552 additional PSUs defined by the Bureau of Labor Statistics for rural areas of the U.S.. Those PSUs are grouped or “stratified” into 91 clusters of PSUs called “strata,” and one PSU is randomly selected to represent each stratum.

The 91 strata are located in the nine divisions of the U.S., and they are divided into three size classes,³ and within each division-size class the PSUs are stratified by four variables: median household income, median household property value, latitude, and longitude. Median household income and median household property value are correlated with household expenditures, and they are designed to keep PSUs with similar expenditure levels together. Latitude and longitude are designed to keep PSUs that are geographically close to each other together.

¹ The state-level weights are not statistically valid for lower-level geographic areas such as the counties within a state because the weights are not calibrated to those levels. The weights in an area add up to the population that the area represents, not the area itself, hence in lower-level geographic areas such as counties, the weights add up to a larger population than is in the area itself. And the state-level weights are not statistically valid for higher-level geographic areas such as divisions, regions, or the nation because there are states where the CE survey does not collect data, and those states need to be accounted for. Those states are accounted for in the national weights, but not in the state weights, hence the state-level weights do not add up to the correct division, region, or national level populations.

² CBSAs are small clusters of counties defined by OMB, consisting of one or more contiguous counties centered on an “urban core” of at least 10,000 people, where the counties’ connection to the core is measured by their residents’ commuting patterns.

³ The three size classes are S, N, and R. The “S” size class contains self-representing PSUs, which are CBSAs with over 2.5 million people. The “N” size class contains non-self-representing PSUs, which are CBSAs with under 2.5 million people. And the “R” size class contains non-self-representing PSUs which are small clusters of counties in non-CBSAs areas that are informally called “rural” PSUs.

After stratifying the PSUs, one PSU is randomly selected from each stratum with probability proportional to its size. Thus if a PSU has 20 percent of a stratum’s population, then it has a 20 percent probability of being selected for the sample; if it has 30 percent of a stratum’s population, then it has a 30 percent probability of being selected for the sample; and so on.

Calculating the State-level Base Weight

Both the national-level and state-level base weights are calculated using two probabilities: the probability of selecting a PSU, and the probability of selecting a household within the PSU. The formula is this:

$$Base\ Weight = \frac{1}{Probability\ of\ selecting\ a\ PSU} \times \frac{1}{Probability\ of\ selecting\ a\ household\ in\ the\ PSU}$$

The first term (one divided by the probability of selecting a PSU) is called the “PSU weight,” and the second term (one divided by the probability of selecting a household in the PSU) is called the “within-PSU sampling interval.” Together the product of the two terms is one divided by the probability of a household in the stratum being selected for the sample.

For example, if a PSU has 20 percent of a stratum’s population, then it has a 20 percent probability of being selected for the sample because PSUs are randomly selected with probability proportional to their size. The “PSU weight” for such a PSU would be $1/0.20 = 5.0$. That means the weight of any household in the sample needs to be multiplied by 5.0 to account for households in the other PSUs that are not selected. And if the selected PSU has 200,000 households, and 100 households are selected for the sample, then the probability of a household in the PSU being selected for the sample is $p = 100/200,000$, which means the “within-PSU sampling interval” is $1/p = 200,000/100 = 2,000$. In other words, one out of every 2,000 households in the PSU is selected for the sample. Multiplying those two factors together gives each selected household in the selected PSU a base weight of 10,000:

$$Base\ Weight = \frac{1}{0.20} \times \frac{1}{(100/200,000)} = 5 \times 2,000 = 10,000$$

That means in such a PSU, every 2,000-th household in the PSU is selected for the sample, and every 10,000-th household in the stratum is selected for the sample. Thus the base weight for each selected household in the PSU would be 10,000, meaning each selected household in the PSU’s sample represents 10,000 households in the stratum – itself plus 9,999 other households that were not selected.

The same formula is used to calculate both the national-level and state-level base weights. The only difference is how the probability of selecting a PSU is calculated. For national-level weights, the PSU’s probability of selection is based on its proportion of the division-based stratum to which it belongs, but for state-level weights, its probability of selection is based on its proportion of the state-based stratum to which it belongs. The division-based strata were defined *a priori* and used in CE’s sample selection process,

while the state-based strata were defined *a posteriori* by making adjustments to the first step in CE's sample selection process so that new state-level base weights can be calculated.

For both national-level and state-level base weights, the probability of selecting an "S" PSU is equal to 1 since "S" PSUs are selected for CE's sample with certainty. However, before the probability of selecting an "N" or "R" PSU within a state can be calculated, state-based strata must be formed by assigning every non-selected PSU in a state to one of the selected PSUs within that state, where non-selected "N" PSUs are assigned to the selected "N" PSUs and non-selected "R" PSUs are assigned to the selected "R" PSUs. The number of state-based strata in a state equals the number of selected PSUs from that state in the national sample. If there is only one selected PSU in a state, every non-selected PSU in that state is assigned to the selected PSU to form one state-based stratum. If there are multiple selected PSUs in a state, an assignment procedure is used to assign every non-selected PSU in the state to one of the selected PSUs to form multiple state-based strata. The assignment procedure ensures that the PSUs in the state-based strata are homogeneous with respect to their median household income and median household property value from five-year American Community Survey's (ACS) estimates.

Then after stratifying the PSUs into state-based strata, the probability of selecting a PSU is re-calculated at the state level, and it is equal to the population of the selected PSU divided by the population of every PSU in the state-based stratum. The "state PSU weight" is equal to one divided by the re-calculated probability of selecting a PSU, and the "within PSU sampling interval" is equal to one divided by the probability of selecting a household in the PSU which was used at the national level. Finally, the state base weight is the product of the "state PSU weight" and the "within PSU sampling interval," which is one divided by the probability of a household in the state-based stratum being selected for the sample.

The formula to calculate the "state base weight" is this:

$$\text{State Base Weight} = \text{State PSU Weight} \times \text{Within PSU Sampling Interval}$$

Subsampling Adjustment Process

After the base weights are computed, several adjustments are made to address issues that arise after the sample is drawn. The subsampling adjustment process adjusts for subsampling in the field, which occurs when a data collector visits a particular address and discovers multiple housing units where only one was expected. A "weighting control factor" is used to adjust the state base weight for these extra housing units in the survey. The need for this adjustment is very rare, hence the weighting control factor is generally 1.

Nonresponse Adjustment Process

The nonresponse adjustment process adjusts for interviews that cannot be conducted in occupied housing units due to a household's refusal to participate in the survey or an interviewer's inability to contact anyone at the housing unit despite repeated attempts. The "nonresponse adjustment factor" is used to adjust the weight of respondent households to account for the nonrespondent households and is equal to the sum of the base weights of the occupied housing units in the sample (both respondents and nonrespondents) divided by the sum of the base weights of the respondent households in the sample. This adjustment uses household information that is known for both respondent and nonrespondent

households, including each household's region of the U.S., number of household members, number of contact attempts, and the average adjusted gross income in the household's zip code according to a publicly available database from the Internal Revenue Service. Currently, the nonresponse adjustment factor for the national estimates is being used as a proxy for the nonresponse adjustment factor for the state estimates.

The state base weight for respondent households is multiplied by the weighting control factor and the nonresponse adjustment factor to create the "initial weight," which is the input variable for the calibration process. The formula to calculate the initial weight is this:

$$\text{Initial Weight} = \text{State Base Weight} \times \text{Weighting Control Factor} \times \text{Nonresponse Adjustment Factor}$$

Calibration Process

The calibration process adjusts for the coverage error that results from frame undercoverage. It uses the initial weights from respondent households as input for calculating the "calibration adjustment factor," which is the final weight adjustment. The initial weights from respondent households are calibrated quarterly to ten "known" population totals for the state to account for frame undercoverage. The process uses nonlinear programming to select a set of weights that minimize the amount of change made to the initial weights so that the calibrated weights multiplied by the number of household members in each of the demographic groups sum to the population totals. The ten population totals are from the Current Population Survey (CPS) and they are: the total number of households in the state; the total number of homeowning households in the state, the total number of Hispanic people in the state, and the total number of people in each of the seven age categories in the state.

The same process is used when calculating the calibration adjustment factor at the national and state levels, but for the nation, it calibrates to 35 known U.S. population totals, and for the state-level, it calibrates to ten known state population totals. The calibration adjustment factor for the state-level weights is based on only ten population totals rather than all 35 population totals because of the low number of respondent households available as input to the state-level process.

After adjusting the initial weights by the calibration adjustment factor, the final state weight is this:

$$\text{State Weight} = \text{Initial Weight} \times \text{Calibration Adjustment Factor}$$

Additional Information

Presently, the state-level weighting process is in an experimental phase. State-level weights are being made available to Consumer Expenditure Survey microdata users to gauge interest and usefulness. For further details, questions, or comments on the procedures or the weights themselves, please contact the Division of Consumer Expenditure Surveys at CEXINFO@bls.gov.

References

- [1] Neiman, D, King, S., Swanson, D., Ash, S., Enriquez, J., and Rosenbaum, J. (2015). Review of the 2010 Sample Redesign of the Consumer Expenditure Survey. Proceedings of the American Statistical Association, Section on Survey Methodology, 2522-2536.
- [2] King, S. (2017). Estimating State Expenditures for New Jersey, Florida, and California. <https://www.bls.gov/cex/ws17-estimating-state-expenditures-nj-fl-ca.pdf>
- [3] King, S. (2017). Calculating State Weights for the Consumer Expenditure Survey. Proceedings of the American Statistical Association.
- [4] Bureau of Labor Statistics, U.S. Department of Labor (2018). U.S. Bureau of Labor Statistics Handbook of Methods, Consumer Expenditures and Income <https://www.bls.gov/opub/hom/cex/pdf/cex.pdf>
- [5] Bureau of Labor Statistics, U.S. Department of Labor (2019). Consumer Expenditure Public Use Microdata Documentation. https://www.bls.gov/cex/pumd_doc.htm