

Public Use Microdata Sample (PUMS)

Accuracy of the Data (2018-2022)

INTRODUCTION

This 5-year public use microdata sample (PUMS) for 2018-2022 is a subset of the 2018-2022 American Community Survey (ACS) and Puerto Rico Community Survey (PRCS) samples. It contains the same sample as the combined PUMS 1-year files for 2018, 2019, 2020, 2021, and 2022. Unless otherwise specified, the term “ACS” in this document will refer to both the ACS and PRCS.

This 2018-2022 ACS 5-year PUMS contains five years of data for housing units (HUs) and the population from households and the group quarters (GQ) population. The GQ population, housing units and population from households are all weighted to agree with the ACS counts, which are an average over the five year period (2018-2022). The ACS sample is selected from all counties across the nation, and all municipios in Puerto Rico.

Estimates from the PUMS file are expected to be different from the previously released ACS estimates because they are subject to additional sampling error and further data processing operations. The additional sampling error is a result of selecting the PUMS housing and person records through an additional stage of sampling. In the public use file, the basic unit is an individual housing unit, except for the sample from GQs. For the GQ sample, the basic unit is the person. The population sample is defined as all persons living in households selected in the housing unit sample, plus the persons selected from the GQ sample. Note that microdata records in this sample do not contain names, addresses, or any information that can identify a specific housing unit, GQ or person.

Users of the 2018-2022 ACS 5-year PUMS file can find detailed information on differences between the 2018-2022 files and previous PUMS files in the PUMS Users Guide and PUMS Variable Changes documents. These documents can be found at:

<https://www.census.gov/programs-surveys/acs/microdata/documentation.html>.

Table of Contents

INTRODUCTION.....	1
CONFIDENTIALITY OF THE DATA	3
TITLE 13, UNITED STATES CODE	3
DISCLOSURE AVOIDANCE	3
DATA SWAPPING	3
SYNTHETIC DATA.....	4
PUMAS	4
ADDITIONAL MEASURES	4
SAMPLE DESIGN.....	5
HOUSING UNITS.....	5
GROUP QUARTERS	6
WEIGHTING.....	6
GROUP QUARTERS PERSON WEIGHTING.....	6
HOUSING UNIT AND HOUSEHOLD PERSON WEIGHTING.....	7
ESTIMATION.....	10
ERRORS IN THE DATA.....	10
SAMPLING ERROR	10
NONSAMPLING ERROR.....	11
MEASURING SAMPLING ERROR.....	11
STANDARD ERROR.....	11
CONFIDENCE INTERVALS	12
LIMITATIONS	13
APPROXIMATING STANDARD ERRORS WITH REPLICATE WEIGHTS	13
APPROXIMATING GENERALIZED STANDARD ERRORS WITH DESIGN FACTORS	15
EXAMPLES OF STANDARD ERROR CALCULATIONS USING GENERALIZED STANDARD ERROR FORMULAS	20
WORKING WITH DOLLAR AMOUNTS	23
ADJUSTMENT FACTORS ON THE PUMS FILE	23
COMPARING PUMS FILES FROM DIFFERENT PERIODS	24
REFERENCES.....	24

CONFIDENTIALITY OF THE DATA

The Census Bureau has implemented a series of steps to protect the confidentiality of the data. Title 13 of the United States Code, Section 9, prohibits the Census Bureau from publishing results in which an individual's data can be identified.

The Census Bureau's internal Disclosure Review Board sets the confidentiality rules for all data releases.¹ A checklist approach is used to ensure that all potential risks to the confidentiality of the data are considered and addressed.

Title 13, United States Code

Title 13 of the United States Code authorizes the Census Bureau to conduct censuses and surveys. Section 9 of the same Title requires that any information collected from the public under the authority of Title 13 be maintained as confidential. Section 214 of Title 13 and Sections 3559 and 3571 of Title 18 of the United States Code provide for the imposition of penalties of up to five years in prison and up to \$250,000 in fines for wrongful disclosure of confidential census information.

Disclosure Avoidance

Disclosure avoidance is the process for protecting the confidentiality of data. A disclosure of data occurs when someone can use published statistical information to identify an individual that has provided information under a pledge of confidentiality. For data tabulations, the Census Bureau uses disclosure avoidance procedures to modify or remove the characteristics that put confidential information at risk for disclosure.

Data Swapping

Data swapping is a method of disclosure avoidance designed to protect confidentiality. Data swapping is done by editing the source data or exchanging records for a sample of cases when creating a table. A sample of households is selected and matched on a set of selected key variables with households in neighboring geographic areas that have similar characteristics (such as the same number of adults and same number of children). Because the swap often occurs within a neighboring area, there is no effect on the marginal totals for the area or for totals that include data from multiple areas. Data swapping procedures were first used in the 1990 Census.

¹ The Census Bureau's Disclosure Review Board approved the 2018-2022 PUMS 5-year data for release with DRB Clearance number CBDRB-FY24-0006.

Synthetic Data

The goals of using synthetic data are the same as the goals of data swapping, namely, to protect confidentiality. Housing units and persons are identified as being at risk for disclosure based on certain characteristics. The synthetic data technique models the values for those or other characteristics to protect the confidentiality of that housing unit or individual. One example of where this is used is home value and property tax. These variables have an increased risk of disclosure due to the availability of public records for the same or similar data. Thus, during the PUMS creation process, we model synthetic values for these variables to decrease the potential of reidentification.

PUMAs

The Census Bureau takes further steps to prevent the identification of specific individuals, households, or housing units, on the PUMS files. The main disclosure avoidance method used is to limit the geographic detail shown in the files. The smallest geographic unit that is identified is the Public Use Microdata Area (PUMA). The current PUMAs were formed based on data and location collected in the 2020 Census and have been used by the ACS PUMS files beginning with the 2022 data year. They do not cross state boundaries.

Records on the 5-year PUMS files from 2022 onward will have the Census 2020-based PUMA codes, while records from 2021 and earlier will have Census 2010-based PUMA codes. These two PUMA codes are stored in separate variables on the 5-year PUMS files. Due to disclosure concerns, it is not possible to update the PUMA codes for the records from 2021 and earlier to 2020-based PUMAs by using their detailed geographic locations.

The Census Bureau provides maps for the PUMAs, and users can identify geographies of interest by zooming in on selected areas. These maps can be found at: https://tigerweb.geo.census.gov/tigerwebmain/TIGERweb_main.html.

For more information on the dual PUMA variables, step-by-step instructions for using TIGERweb, as well as a method to crosswalk the dual PUMA codes, see the 5-Year PUMS User Guide: <https://www.census.gov/programs-surveys/acs/microdata/documentation.html>.

Additional Measures

Other disclosure avoidance measures used in the PUMS files include top-coding, age perturbation, weight perturbation and collapsing of detail for categorical variables. The answers to open-ended questions, where an extreme value might identify an individual are top-coded (or bottom-coded). Top-coding (and bottom-coding) substitutes the value of extreme cases with the mean of the highest (or lowest) cases. Top-coded questions include age, income and housing unit value. Age perturbation disguises original data by randomly adjusting the reported ages for a subset of individuals. Weight perturbation disguises the probability of

selection for some records. Users should exercise caution when forming estimates near top-coded or bottom-coded values. More information on the variables that receive top or bottom coding in the PUMS can be found at: <https://www.census.gov/programs-surveys/acs/microdata/documentation.html>.

SAMPLE DESIGN

The 2018-2022 ACS 5-year PUMS sample is the same sample found in each of the 1-year PUMS files for the years 2018, 2019, 2020, 2021, and 2022. It contains five percent of the housing units and five percent of the GQ persons for the United States, District of Columbia, and Puerto Rico weighted to represent the average population during five years.

The PUMS GQ sample contain additional imputed records to represent the not-in-sample GQs, which effectively double the total number of GQ records.

See the Accuracy of the Data for the 2022 PUMS 1-year for a further explanation of the PUMS sampling of these imputed GQ records. By including these imputed records in the 2018-2022 ACS 5-year PUMS, the PUMS will agree better with the 2018-2022 5-year full sample ACS for population totals by state and PUMA. More details about the methodology of the large-scale whole person imputation into not-in-sample GQ facilities can be found in the 2022 ACS 1-year Accuracy of the Data at: <https://www.census.gov/programs-surveys/acs/technical-documentation/code-lists.html>.

Housing Units

The sampling for HUs (and persons from HUs) is performed on the ACS samples of HUs for each of the years 2018, 2019, 2020, 2021, and 2022 as follows:

1. Records of ACS HUs were sorted within each state by: PUMA, ACS weighting area, interview mode, type of vacant, tenure, building type, household type, householder demographics (race, Hispanic origin, sex and age), county, tract, and housing unit weight.
2. Systematic sampling is applied to ACS HUs:
 - a. Within each state, a random number is chosen between zero and the sampling interval. A counter is initialized with the random number.
 - b. At each HU record, the value of the counter is incremented by one and compared to the sampling interval.
 - i. If the counter's new value is greater than the sampling interval, the HU record is selected for the PUMS and a flag is set to 1. The counter is

decreased by the sampling interval with the new value passed to the next record.

- ii. If the counter is less than the sampling interval, the HU record is not selected for the PUMS and the value of the counter is passed to the next record without altering its value.

3. All HUs selected for PUMS were placed in the PUMS HU sample file.

The PUMS HU sample file is matched to the ACS sample of persons. All persons in selected HUs were placed in the PUMS person sample.

Group Quarters

The sampling for PUMS GQ persons is originally performed on the ACS sample of GQ persons for each of the years 2018, 2019, 2020, 2021, and 2022 as follows:

1. Interviewed GQ persons were sorted within each state by the size of their GQ facility (large vs small), the type of GQ facility, PUMA, demographics (race, Hispanic origin, sex and age), county, tract, and GQ person weight.
2. Systematic sampling is applied in the same manner as described above for HUs.
3. All selected GQ persons are added to the PUMS person sample. All imputed records derived from the selected PUMS interviews were also kept in the PUMS person sample.
4. A placeholder record is also placed in the PUMS HU file for each PUMS GQ person record.

WEIGHTING

Weights for PUMS person records are a product of the final full ACS weight, the PUMS subsampling factor, and ratio-estimate factors. The PUMS subsampling factors are the sampling intervals used to sample the PUMS HU or GQ person records within a state. The ratio-estimate factors bring the PUMS estimates into closer agreement with the published ACS estimates for several characteristics explained below.

Group Quarters Person Weighting

The procedure used to assign the weights to the GQ persons is performed independently within each state. The steps are as follows:

Initial Weight for GQ Persons

The 5-year PUMS initial weight is the product of the 1-year ACS unrounded weights for the record divided by five and the PUMS subsampling factor. Each imputed record received the same subsampling factor as its donor interview.

GQ Person Weighting Factors

GQ Person Post-stratification Factor

This factor adjusts the GQ person weights so that the weighted sample counts equal ACS published estimates at the state level. Due to the ACS GQ sample design and noise added for disclosure avoidance reasons, only state level PUMS GQ person estimates will agree closely with published ACS estimates. This adjustment uses the following groups:

State × Institutional/noninstitutional × Sex × Age Category

Rounding of GQ Person Weights

The final GQ person weight is rounded to an integer. Rounding is performed so that the sum of the rounded weights is within one person of the sum of the ACS total GQ person estimate for the state.

Housing Unit and Household Person Weighting

The estimation procedure used to assign the HU and person weights is performed independently within each PUMA.

Initial Weight for Persons and HUs

The 5-year PUMS initial weight is equal to the product of the ACS 1-year final weight for the record and the PUMS subsampling factor divided by five.

To accommodate the revised methodology used to create the 2020 1-year and 2018-2022 5-year data products, the initial weight for records from 2020 is equal to the weight after the first run of the person weighting, as referenced in the 2018-2022 ACS Multiyear Accuracy (located at: <https://www.census.gov/programs-surveys/acs/technical-documentation/code-lists.html>).

Person Weighting Factors

The person weights are adjusted to agree better with ACS published estimates for householders, spouses, race, Hispanic origin, sex and age by a series of two steps which are

repeated until a stopping criterion is met. This is an iterative proportional fitting or raking process. The person weights are individually adjusted at each step as described below.

The two steps are as follows:

Spouse Equalization/Householder Equalization Raking Factor

This factor is applied to individuals based on the combination of their status of being in a married-couple or unmarried-partner household and whether they are the householder. All persons are assigned to one of four groups:

1. Householder in a married-couple or unmarried-partner household
2. Spouse or unmarried partner in a married-couple or unmarried-partner household (non-householder)
3. Other householder
4. Other non-householder

The weights of persons in the first two groups are adjusted so that their sums are each equal to the ACS estimate of married-couple or unmarried-partner households using the ACS housing unit weight. The weights of persons in the third group are adjusted so that the sum is equal to the ACS estimate of occupied housing units not having a partner using the housing unit weight. The weights of persons in the fourth group are adjusted to agree with the ACS total population minus the first three groups. The goal of this step is to produce more consistent estimates of spouses or unmarried partners and married-couple and unmarried-partner households while simultaneously producing more consistent estimates of householders, occupied housing units, and households.

Demographic Raking Factor

This factor is applied to individuals based on their age, race, sex and Hispanic origin. It adjusts the person weights so that the weighted sample counts equal ACS population estimates by age, race, sex, and Hispanic origin at the PUMA level. Because of collapsing of groups in applying this factor, only total population is assured of agreeing precisely with the published ACS 2018-2022 population estimates at the PUMA level.

This uses the following groups within each PUMA (note that there are 13 Age groupings):

Race / Ethnicity (non-Hispanic White, non-Hispanic Black, non-Hispanic American Indian or Alaskan Native, non-Hispanic Asian, non-Hispanic Native Hawaiian or Pacific Islander, and Hispanic (any race)) \times Sex \times Age Groups

These two steps are repeated several times until the estimates at the PUMA level achieve their optimal consistency with regard to the spouse and householder equalization. The final Person Weighting Factor is then equal to the product of the factors from all of the iterations

of these two adjustments. The unrounded person weight is then equal to the product of Person Weighting Factor times the initial person weight.

Rounding of Person Weights

The person weight after the Person Weighting Factor has been applied is rounded to an integer. Rounding is performed so that the sum of the rounded weights is within one person of the sum of the ACS total persons from HU's estimates within state and PUMA.

Householder Adjustment Factor (HHRF)

This factor applied to occupied housing units is the same as the Person Weighting factor from the person weighting. After this stage, the weight of the housing unit is identical to the unrounded person weight of the householder after the Person Weighting Factor is applied.

Housing Unit Control Factor

This factor adjusts PUMS housing unit estimates to agree with the published ACS housing unit estimates for housing units with married couples (or partners), occupied housing units without partners and vacant housing units.

Rounding of Housing Unit Weights

The Housing Unit weight after the Housing Unit Control Factor is applied is rounded to an integer. Rounding is performed so that the sum of the rounded weights is within one housing unit of the sum of the ACS total HU's estimates within state and PUMA.

For a detailed description of how the original ACS weights are computed, see the 2018-2022 ACS Multiyear Accuracy of the Data at: <https://www.census.gov/programs-surveys/acs/technical-documentation/code-lists.html>

ESTIMATION

To produce estimates or tabulations of characteristics from the PUMS, add the weights of all persons or HUs that possess the characteristic of interest.² For instance, if the characteristic of interest is “total number of black teachers”, simply determine the race and occupation of all persons and cumulate the weights of those who match the characteristics of interest. To obtain estimates of proportions, divide the weighted estimate of persons or HUs with a given characteristic by the weighted estimate of the denominator. For example, the proportion of “black teachers” is obtained by dividing the weighted estimate of black teachers by the estimate of teachers.

PUMS estimates are expected to be different from published ACS estimates that are based on the full set of data because of the additional sampling. The exception will be characteristics controlled by the ratio-estimate factors.

Note that the housing unit file contains some records with zero weights. These are the GQ placeholder records.³ The housing unit weights were set to zero for these records since they are not housing units, but persons. For confidentiality reasons, the GQ data are not provided at the level of an address but only at the person-level. All of the GQ person data are included in the PUMS person file except variable for the “Yearly food stamp/Supplemental Nutrition Assistance Program (SNAP) reciprocity” variable (FS), which is included on the GQ placeholder records in the housing unit file. For food stamp reciprocity estimates of persons in GQs, you will need to match the placeholder records to the person file to obtain the person weights.

A note to GQ data users. There are limitations to the usefulness of GQ estimates at the PUMA level. The PUMS weighting controls the GQ estimates to agree with the ACS state level estimates. Depending on the application or analysis, GQ data users should consider working with state level estimates rather than PUMAs.

ERRORS IN THE DATA

Every sample survey is subject to two types of error: sampling error and nonsampling error.

Sampling Error

The data in the ACS products are estimates of the actual figures that would have been obtained by interviewing the entire population using the same methodology. The estimates from the

² Users should exercise caution when forming estimates near top-coded or bottom-coded values. More information on the variables that receive top or bottom coding in the PUMS can be found at: <https://www.census.gov/programs-surveys/acs/microdata/documentation.html>.

³ To identify HU and GQ placeholder records on the PUMS housing file, see the TYPE variable in the PUMS data dictionary: <https://www.census.gov/programs-surveys/acs/microdata/documentation.html>.

chosen sample also differ from other samples of HUs and persons within those HUs. Sampling error in data arises due to the use of probability sampling, which is necessary to ensure the integrity and representativeness of sample survey results. The implementation of statistical sampling procedures provides the basis for the statistical analysis of sample data.

Estimates made with PUMS data are subject to additional sampling error because the PUMS data consists of a subset of the full ACS sample. Thus standard errors or margins of errors of PUMS estimates can be larger than standard errors or margins of errors that would be obtained using all of the ACS microdata.

Nonsampling Error

In addition to sampling error, data users should realize that other types of errors might be introduced during any of the various complex operations used to collect and process survey data. For example, operations such as data entry from questionnaires and editing may introduce error into the estimates. These and other sources of error contribute to the nonsampling error component of the total error of survey estimates.

Nonsampling errors may affect the data in two ways. Errors that are introduced randomly increase the variability of the data. Systematic errors, which are consistent in one direction, introduce bias into the results of a sample survey. The Census Bureau protects against the effect of systematic errors on survey estimates by conducting extensive research and evaluation programs on sampling techniques, questionnaire design, and data collection and processing procedures. In addition, an important goal of the ACS is to minimize the amount of nonsampling error introduced through nonresponse for sample HUs. One way of accomplishing this is by following up on mail nonrespondents during the CAPI phase.

More information about the control of nonsampling error can be found in ACS Multiyear Accuracy of the Data at: <https://www.census.gov/programs-surveys/acs/technical-documentation/code-lists.html>.

MEASURING SAMPLING ERROR

Standard Error

A measure of the deviation of a sample estimate from the average of all possible samples. Sampling error and some types of nonsampling error, such as undercoverage and item nonresponse, are estimated by the standard error. The sample estimate and its estimated standard error permit the construction of interval estimates with a prescribed confidence that the interval includes the average result of all possible samples.

Two methods are provided for estimating the standard errors of PUMS estimates: a successive difference replicate (SDR) method using replicate weights and a generalized variance function

(GVF) method using design factors. Replicate weights have been provided with the ACS PUMS files beginning in 2005. The GVF method using design factors is used by the PUMS file produced for the Decennial Census and has also been in use by the ACS PUMS since 2000. It is important to keep in mind that standard errors produced using the SDR method will be difference from those produced using the GVF method. Generally, using the SDR method will produce a more accurate estimate than the GVF method.

Confidence Intervals

A sample estimate and its estimated standard error may be used to construct confidence intervals about the estimate. These intervals are ranges that will contain the average value of the estimated characteristic that results over all possible samples, with a known probability.

For example, if all possible samples that could result under the PUMS sample design were independently selected and surveyed under the same conditions, and if the estimate and its estimated standard error were calculated for each of these samples, then:

1. Approximately 68 percent of the intervals from one estimated standard error below the estimate to one estimated standard error above the estimate would contain the average result from all possible samples;
2. Approximately 90 percent of the intervals from 1.645 times the estimated standard error below the estimate to 1.645 times the estimated standard error above the estimate would contain the average result from all possible samples.
3. Approximately 95 percent of the intervals from 1.96 times the estimated standard error below the estimate to 1.96 times the estimated standard error above the estimate would contain the average result from all possible samples.

These intervals are referred to as 68 percent, 90 percent, and 95 percent confidence intervals, respectively. An example of how to construct a 90 percent confidence interval follows:

Add and subtract 1.645 times the standard error (SE) to the estimate to yield the lower and upper bounds of a 90% confidence interval around the estimate.

$$LB = \text{Lower bound} = \text{Estimate} - 1.645 \times SE(\text{Estimate})$$

$$UB = \text{Upper bound} = \text{Estimate} + 1.645 \times SE(\text{Estimate})$$

The 90% confidence interval is the interval (LB, UB).

Limitations

The user should be careful when computing and interpreting standard errors and confidence intervals.

Nonsampling Error

The estimated standard errors included in this data product do not include all portions of the variability due to nonsampling error that may be present in the data. In particular, the standard errors do not reflect the effect of correlated errors introduced by interviewers, coders, or other field or processing personnel. Nor do they reflect the error from imputed values due to missing responses. Thus, the standard errors calculated represent a lower bound of the total error. As a result, confidence intervals formed using these estimated standard errors may not meet the stated levels of confidence (i.e., 68, 90, or 95 percent). Thus, some care must be exercised in the interpretation of the data in this data product based on the estimated standard errors.

Very Small (Zero) or Very Large Estimates

The value of almost all PUMS characteristics is greater than or equal to zero by definition. For zero or small estimates, use of the method given previously for calculating confidence intervals relies on large sample theory, and may result in negative values which, for most characteristics, are not admissible. In this case the lower limit of the confidence interval is set to zero by default. A similar caution holds for estimates of totals close to a control total or estimated proportions near one, where the upper limit of the confidence interval is set to its largest admissible value. In these situations the level of confidence of the adjusted range of values is less than the prescribed confidence level.

Approximating Standard Errors with Replicate Weights

The standard error may be calculated using the successive difference replicate (SDR) method using the replicate weights provided in the PUMS file. SDR standard errors will often be more accurate than GVF standard errors, although they may be more inconvenient for some users to calculate. The advantage of using the SDR method is that a single formula is used to calculate the standard error of many types of estimates.

Each PUMS housing unit and person record contains 80 PUMS replicate weights. These replicate weights are based on the ACS replicate weights adjusted for PUMS subsampling. For any estimate X , 80 replicate estimates are also computed using the replicate weights. For this discussion, we refer to X as the 'full sample estimate.' The first replicate estimate, X_1 , is computed using the first replicate weight, the second replicate estimate, X_2 , is computed using the second replicate weight, and so on. Each replicate estimate is computed using the replicate

weights in the same way that the full sample estimate X is computed, as described in the Estimation section.

NOTE: When programming the replicate weight standard errors, users will find the 80 replicate weights can be positive, zero or negative. The negative replicate weights are due to the addition of the Group Quarters (GQ) population to the full ACS weighting process. Within a weighting cell, GQ estimates were subtracted from population totals, sometimes resulting in negative values for the cell. The cells were collapsed in such a way as to prevent a final cell from being zero or negative for the full sample weights. The full sample weights always have a value of at least one. This restriction is not placed on the replicate weights since their only purpose is to represent the variability of the sample. PUMS replicate weights are based on ACS replicate weights so negative values may occur. Keep in mind that the replicate weights are only to be used to estimate standard errors.

The standard error of X can be approximated after the replicate estimates X_1 through X_{80} are computed. The standard error is estimated using the sum of squared differences between each replicate estimate X_r and the full sample estimate X . The standard error formula is:

$$SE(X) = \sqrt{\frac{4}{80} \sum_{r=1}^{80} (X_r - X)^2}$$

If X is zero, then use the generalized variance method for zero estimates given in Standard Errors for Totals and Percentages, to approximate the standard error.

Data users who wish to see worked examples may consult the documentation for the ACS Variance Replicate Tables, located here: <https://www.census.gov/programs-surveys/acs/data/variance-tables.html>.

The standard error can be used to form a 90% confidence interval around the estimate (X) as follows:

$$LB = \text{Lower bound} = X - 1.645 \times SE(X)$$

$$UB = \text{Upper bound} = X + 1.645 \times SE(X)$$

The 90% confidence interval is the interval (LB, UB).

As mentioned previously, we consider the SDR SEs using the replicate weight to be more accurate than the GVF SEs using the design factor. For exceptions, please note the following:

When using SDR SEs, occasionally the SE is zero for an estimate. Except for controlled estimates, all PUMS estimates are based on a sample of the population and should not have a SE of zero. If the SDR method using the replicate weight gives a SE of zero for an estimate that is not controlled, use the GVF method with the design factors to obtain a SE for that estimate. If the estimate is controlled, then a SE of zero is appropriate. Data users may check if an estimate is controlled by finding the equivalent ACS estimate on data.census.gov. If the ACS estimate is controlled it will have a MOE of five asterisks (*****).

If your estimate is a median, the SDR method using replicate weights may yield a SE of zero. Medians should always have a non-zero SE. A median with a SE of zero may occur when several records in the middle of the distribution were rounded to the same value, or when the characteristic contains few records, such as a median based on less than five records. Rounding by respondents, as well as rounding by PUMS edits may mask the variability in the median. In order to yield a more adequate standard error for that case, use the GVF method with the design factors to estimate the SE of a median.

Examples of PUMS estimates with their SE and MOE may be found by clicking on PUMS Estimates for User Verification at:

<https://www.census.gov/programs-surveys/acs/microdata/documentation.html>.

Users can check national and state level estimates with associated standard errors and margin of errors by comparing to values shown in these files. The SE and MOE are calculated using the SDR method with the PUMS replicate weights.

Approximating Generalized Standard Errors with Design Factors

Converting Standard Error to Margin of Error

The ACS uses a 90% confidence level margin of error. To convert from the standard error to the margin of error, multiply the standard error by 1.645. To calculate a margin of error at a different confidence level, data users should use the appropriate z-score (for example, 1.96 for a 95% confidence level).

Design Factors

The design factors are provided in a comma-separated value (CSV) file located at:

<https://www.census.gov/programs-surveys/acs/microdata/documentation.html>.

For 2016 and previous years, they are provided at the end of the relevant PUMS Accuracy document. They can be used to approximate the standard errors and margins of error for PUMS estimates. Design factors are given by subject for the United States, all 50 states, the District of Columbia, and Puerto Rico. The term "subject" refers to a characteristic, such as

age for persons and tenure for HUs. The design factors reflect the effects of the actual sample design and estimation procedures used for the ACS.

Totals and Percentages

To approximate the standard error, use the following formulas:

Total Formula

$$SE(Y) = DF \times \sqrt{\frac{95}{5} \times Y \left(1 - \frac{Y}{N}\right)}$$

Where:

DF = Design Factor

N = Size of the Geographic Area

Y = Estimate of Characteristic Total

Percent Formula

$$SE(p) = \sqrt{\frac{95}{5B} \times p \times (100 - p)}$$

Where:

DF = Design Factor

B = Denominator of Estimated Percentage

p = Estimated Percentage

Note that here, the percent, p, ranges from 0% to 100%.

The values of N and the design factor can be determined as follows:

1. For the value of N, obtain the number of persons, number of households or number of HUs, respectively for the geography or geographies you are interested in. If the estimate is of HUs then use the number of HUs; if the estimate is of families or households then use the number of households; otherwise use the number of persons.
2. Select the appropriate table from the comma-separated value (CSV) PUMS Design Factors file, located at: <https://www.census.gov/programs-surveys/acs/microdata/documentation.html>. Use the design factor for the United States when estimating characteristics for the United States or geographic areas that cover more than one state. Use the table for a specific state when estimating characteristics for that state or geographic areas that are contained entirely within that state.

3. Then use the selected characteristic to obtain the appropriate design factor for the characteristic; for example, educational attainment or ancestry. If the estimate is a combination of two or more characteristics, we suggest the following guideline: Use the largest design factor for this combination of characteristics. **The only exception to this is for items crossed with race or Hispanic Origin. For an item(s) crossed with race or Hispanic Origin, use the largest design factor not including the race or Hispanic Origin design factor.**

An inspection of the formulas used to calculate the standard errors under simple random sampling suggests that when dealing with zero estimates or very small estimates of totals and percentages, the standard error estimates approach zero. This is also the case for very large estimates of totals and percentages. Zero or small estimates, like any other sample estimates, are still subject to sampling variability and therefore an estimated standard error of zero or close to zero is not adequate. Use one of the following procedures below for estimates that fit the following descriptions:

1. **An estimated total is less than 425 or within 425 of the total size of the tabulation area.** Use a basic standard error of 110 multiplied by the design factor for the type of estimate.
2. **For estimated percentages that are less than 2 or greater than 98.** Use a value of 2 for the estimated percentage in the percent formula.
3. **The denominator of a percentage is zero.** There are no sample observations available to compute an estimate of a proportion or an estimate of its standard error.

Medians

The sampling variability of an estimated median depends on the form of the distribution and the size of its base. The standard error of an estimated median is approximated by constructing a 68 percent confidence interval. Estimate the 68 percent confidence limits of a median based on sample data using the following procedure.

1. Obtain the weighted frequency distribution for the selected variable.
2. Calculate the weighted total to yield the base, B.
3. Approximate the standard error of a 50 percent proportion using the formula in Standard Errors for Totals and Percentages.

$$SE(50 \text{ percent}) = DF \times \sqrt{\left(\frac{95}{5B}\right) \times 50^2}$$

4. Create the variable p_lower by subtracting the SE from 50 percent. Create p_upper by adding the SE to 50 percent.

$$p_lower = 50 - SE(50 \text{ percent})$$

$$p_upper = 50 + SE(50 \text{ percent})$$

5. Determine the categories in the distribution that contain p_lower and p_upper . If p_lower and p_upper fall in the same category, follow step 6. If p_lower and p_upper fall in different categories, go to step 7.
6. If p_lower and p_upper fall in the same category, do the following:
- Define A1 as the smallest value in that category.
 - Define A2 as the smallest value in the next (higher) category.
 - Define C1 as the cumulative percent of units strictly less than A1.
 - Define C2 as the cumulative percent of units strictly less than A2.

Use the following formulas to approximate the lower and upper bounds for a confidence interval about the median:

$$\text{Lower Bound} = \left[\frac{p_{lower} - C1}{C2 - C1} \right] \times (A2 - A1) + A1$$

$$\text{Upper Bound} = \left[\frac{p_{upper} - C1}{C2 - C1} \right] \times (A2 - A1) + A1$$

4. If p_lower and p_upper fall in different categories, do the following:
- For the category containing p_lower : Define A1, A2, C1, and C2 as described in step 6. Use these values and the formula in step 6 to obtain the lower bound.
 - For the category containing p_upper : Define new values for A1, A2, C1, and C2 as described in step 6. Use these values and the formula in step 6 to obtain the upper bound.
5. Use the lower and upper bounds approximated in steps 6 or 7 to approximate the standard error of the median.

$$SE(\text{median}) = 1/2 \times (\text{Upper Bound} - \text{Lower Bound})$$

The design factor method shown here for medians may be used if the successive difference replicate (SDR) method using the replicate weights results in a SE of zero.

The SDR method will have a SE of zero if all of the replicate medians have the same value as the median estimate. This may occur due to similar survey responses for a characteristic being rounded to the same value due to the PUMS rounding rules. The ACS calculates medians using a linear interpolation method. Using this methodology may avoid this issue. However, data users may also use the design factor method to calculate the standard errors

Means

A mean is defined here as the average quantity of some characteristic (other than the number of people, HUs, households, or families) per person, housing unit, household, or family. For example, a mean could be the average annual income of females age 25 to 34. The standard error of a mean can be approximated by the formula below. Because of the approximation used in developing this formula, the estimated standard error of the mean obtained from this formula will generally underestimate the true standard error.

$$SE(Y) = DF \times \sqrt{\left(\frac{95}{5B}\right) \times s^2}$$

Where:

B is the base (denominator) of the mean

s^2 is the sample variance of the characteristic based on weighted data.

The value of s^2 can be computed using the formula:

$$s^2 = \frac{\sum_{i=1}^n w_i y_i^2 - [(\sum_{i=1}^n w_i y_i)^2 / \sum_{i=1}^n w_i]}{(\sum_{i=1}^n w_i) - 1}$$

Where:

w_i is the weight of the i^{th} sample record

y_i is the value of the characteristic for the i^{th} sample record

n is the number of sample records

Note that $\sum_{i=1}^n w_i$ is the weighted estimate of persons/HUs in the sample (ex. the number of females age 25 to 34, and $\sum_{i=1}^n w_i Y_i$ is the weighted aggregate estimate for the characteristic of interest (ex. the aggregate income of females age 25 to 34).

Examples of Standard Error Calculations using Generalized Standard Error Formulas

We will present some examples based on the 2009-2013 PUMS 5-year data to demonstrate the use of the generalized standard error formulas.

Example 1 – Using Design Factors to Estimate the Standard Error of a Total

The estimated number of people 15 years or over who were never married is 2,136,436 from the PUMS data for the state of Virginia. To calculate the standard error, we use the total formula given in the section Standard Errors for Totals and Percentages. In this formula, Y is our estimate of 2,136,436 and N is the total PUMS population for the state of Virginia, which is 8,256,630. The design factor for “Marital Status” is 1.4.

$$SE = 1.4 \times \sqrt{\frac{95}{5} \times 2,136,436 \times \left(1 - \frac{2,136,436}{8,256,630}\right)} = 7,679.46$$

To calculate the margin of error, simply multiply 7,679.46 by 1.645 to get 12,632.72. To obtain the lower and upper bounds of the 90 percent confidence interval around 2,136,436 using the margin of error, simply add and subtract 12,632.72 from 2,136,436. Thus the 90 percent confidence interval for this estimate is [2,136,436 - 12,632.72] to [2,136,436 + 12,632.72] or 2,123,803.28 to 2,149,068.72.

Example 2 – Using Design Factors to Estimate the Standard Error of a Proportion or Percentage

The estimated percent of people 25 years or over with a bachelor’s degree or higher in Louisiana is 22.4190 = 100*(681,488/3,039,780) from the PUMS data. To calculate the standard error, we use the percent formula given in the section Standard Errors for Totals and Percentages. Use the denominator of the percentage, 3,039,780 in the formula. The design factor for “Educational Attainment” is 1.5.

$$SE = 1.5 \times \sqrt{\frac{95}{5 \times 3,039,780} \times 22.4190 \times (100 - 22.4190)} = 0.1564$$

To calculate the margin of error, multiply 0.1564 by 1.645 to get 0.2573. To obtain the lower and upper bounds of the 90 percent confidence interval around 22.4190 percent using the margin

of error, simply add and subtract 0.2573 from 22.4190. Thus the 90 percent confidence interval for this estimated percentage is $[22.4190 - 0.2573]$ to $[22.4190 + 0.2573]$ or 22.16 to 22.68.

Example 3 – Calculating the Standard Error of a Median

Users need to form a weighted frequency distribution for the variable of interest. Table 3 below shows one possible weighted frequency distribution for adjusted household income in Massachusetts.

Table 1: A Possible Distribution Frequency for Adjusted Household Income in MA

Adjusted Household Income	Frequency	Cumulative Frequency	Cumulative Percent
Less than \$10,000	153,739	153,739	6.03
\$10,000 to \$14,999	130,852	284,591	11.16
\$15,000 to \$19,999	113,550	398,141	15.62
\$20,000 to \$24,999	105,230	503,371	19.74
\$25,000 to \$29,999	95,824	599,195	23.50
\$30,000 to \$34,999	102,957	702,152	27.54
\$35,000 to \$39,999	90,972	793,124	31.11
\$40,000 to \$44,999	87,818	880,942	34.55
\$45,000 to \$49,999	83,793	964,735	37.84
\$50,000 to \$59,999	174,388	1,139,123	44.68
\$60,000 to \$74,999	231,284	1,370,407	53.75
\$75,000 to \$99,999	318,700	1,689,107	66.25
\$100,000 to \$124,999	244,795	1,933,902	75.85
\$125,000 to \$149,999	178,104	2,112,006	82.83
\$150,000 to \$199,999	202,797	2,314,803	90.79
\$200,000 or more	234,913	2,549,716	100.00

The base is the cumulative sum of the weighted frequencies, which is 2,549,716.

Determine the standard error of a 50 percent proportion, using as the denominator the cumulative sum of the weighted frequencies, 2,549,716. For this example, the design factor for household income is 1.5.

$$SE(50 \text{ percent}) = 1.5 \times \sqrt{\frac{95}{5 \times 2,549,716} \times 50^2} = 0.20$$

Calculate p_{lower} and p_{upper} .

$$p_{\text{lower}} = 50 - SE(50 \text{ percent}) = 49.80$$

$$p_{\text{upper}} = 50 + SE(50 \text{ percent}) = 50.20$$

Determine the categories that contain p_{lower} and p_{upper} . The first category with a cumulative percentage that is greater than 49.80 is \$60,000 to \$74,999. The first category with a cumulative percentage that is greater than 50.20 is \$60,000 to \$74,999. Since p_{lower} and p_{upper} fall in the same category, follow the instructions given in step 5 of the section Standard Errors for Medians.

Define A1, A2, C1, and C2: A1 = 60,000, A2 = 75,000, C1 = 44.68 and C2 = 53.75. Calculate the lower bound and upper bound using these values.

$$\text{Lower Bound} = \left[\frac{49.80 - 44.68}{53.75 - 44.68} \right] \times (75,000 - 60,000) + 60,000 = 68,467.48$$

$$\text{Upper Bound} = \left[\frac{50.20 - 44.68}{53.75 - 44.68} \right] \times (75,000 - 60,000) + 60,000 = 69,129.00$$

Finally, calculate the standard error of the median:

$$SE(\text{median}) = 1/2 \times (69,129.00 - 68,467.48) = 330.76$$

Example 4 – Calculating the Standard Error of a Mean

Suppose we wish to estimate mean adjusted person income of females age 25 to 34 in Alabama. Table 4 below summarizes the computation of the terms in the formula for s^2 . The PUMS data for Alabama has 12,893 records for females age 25 to 34 that have a non-missing value for person income.

Table 2: Computations for Mean Adjusted Person Income of Females Age 25 to 34 in Alabama

Sample Record	y_i	w_i	$w_i y_i$	$w_i y_i^2$
1	21,462	5	107,309	2,303,061,466
2	21,462	13	279,004	5,987,959,811
3	9,658	32	309,051	2,984,767,660
4	65,459	21	1,374,633	89,981,762,995
.				
.				
12,893	55,169.65	17	937,884	51,742,728,026
Total	274,162,709	317,090	6,575,359,529	302,151,315,109,878

Note: The inflation-adjusted income shown in column 2 is the person income multiplied by the ADJINC variable. Unrounded adjusted income is used to compute the products in the third and fourth columns.

The mean adjusted income is:

$$\bar{Y} = \frac{6,575,359,529}{317,090} = 20,736.57$$

and s^2 is computed as follows:

$$s^2 = \frac{302,151,315,109,878 - 6,575,359,529^2/317,090}{317,090 - 1} = 522,884,429$$

The design factor for person income in Alabama, is 1.6. The standard error of the mean can now be calculated:

$$SE(\bar{Y}) = 1.6 \times \sqrt{\frac{95}{5 \times 317,090} \times 522,884,429} = 283$$

WORKING WITH DOLLAR AMOUNTS

Dollar variables must be adjusted into a common year before using them to form estimates. Generally, the older years are adjusted to the most recent year covered by the analysis. Data in the current ACS 5-year PUMS were collected over five years, so the adjustment will inflate 2017, 2018, 2019, and 2020 dollars into 2021 values.

Adjustment Factors on the PUMS File

The PUMS data dictionary for 2018-2022 describes two adjustment factors on the 5-year file that put dollar values into 2021 dollars:

ADJINC – inflation adjustment factor for income variables, such as household income, self-employment income, retirement income and wages.

ADJHSG – inflation adjustment factor for most housing dollar variables, such as utility costs, rent, food stamps, and condominium fees.

The adjustment factor for income and earnings dollar amounts (ADJINC) is applied to income and earning variables. The adjustment factor for housing dollar amounts (ADJHSG) is applied to variables related to housing costs. The adjustment factor will convert amounts into consistent 2021 dollars.

For more details, see the PUMS Data Dictionary and PUMS User Guide documents located at: <https://www.census.gov/programs-surveys/acs/microdata/documentation.html>.

Comparing PUMS files from Different Periods

When comparing dollar estimates from the 2018-2022 ACS 5-year PUMS file to estimates from other years, an additional adjustment is necessary to convert the amounts into dollars from a common year (after applying the adjustment factor described in the previous paragraphs). We use the CPI-U-RS adjustment factors from the Bureau of Labor Statistics. These factors can be found in the first table in the PDF file at:

<https://www.bls.gov/cpi/research-series/>

REFERENCES

- [1]. ACS Multiyear Accuracy of the Data (2018-2022)
<https://www.census.gov/programs-surveys/acs/technical-documentation/code-lists.html>
- [2]. Design and Methodology of the American Community Survey:
<https://census.gov/programs-surveys/acs/methodology/design-and-methodology.html>
- [3]. PUMS Accuracy of the Data for 1-Year PUMS files:
<https://www.census.gov/programs-surveys/acs/microdata/documentation.html>
- [4]. PUMS Data Dictionary:
<https://www.census.gov/programs-surveys/acs/microdata/documentation.html>
- [5]. Updated CPI-U-RS, All Items, 1977-2022
<https://www.bls.gov/cpi/research-series/r-cpi-u-rs-home.htm>