

Speaker Notes:

Things that may Affect the Estimates from the American Community Survey

Updated: 2/20/09

Slide 1

The purpose of this presentation is to provide information on issues that can affect estimates created from American Community Survey data, in particular, sampling error.

Slide 2

First we'll define sampling error, and then we'll discuss four different measures associated with sampling error, including how they are calculated and what they mean.

Then we'll talk about why measures of sampling error are important and how you can use them to help draw appropriate conclusions about ACS data.

Finally we'll talk briefly about non-sampling error and population controls.

Slide 3

Slide 4

What is Sampling Error?

Let's start our discussion with the definition:

- Sampling error is the uncertainty associated with an estimate that is based on data gathered from a **sample** of the population rather than the **full** population.

- So, why do sample estimates have uncertainty associated with them? There are two reasons:
 1. Estimates of characteristics from the sample data can differ from those that would be obtained if the entire population were surveyed.
 2. Estimates from one subset or sample of the population can differ from those based on a different sample from that same population.

Slide 5

Let's illustrate this with a simple example.

Suppose we want to estimate the average number of children per household for a population that only has 3 households. If we use the information for all households in the population, then we

would divide the total number of children, six, by the total number of households, three, to get an average of two children per household.

But, suppose we did not have enough money to collect data from all three households, the full population, and instead had to estimate the average by collecting data from only two of the three households.

Slide 6

There are three different samples of two households that we could potentially select, as shown here. If we randomly chose two households, we could get households A and B with 1 child and 2 children respectively, households B and C with 2 children and 3 children respectively, or households A and C with 1 child and 3 children respectively. If we sum the total number of children in each sample and divide by the total number of households in each sample, we will get the average number of children per household in each sample.

Sample 1 with households A and B has an average of 1.5 children per household (1 child + 2 children equals 3 total children in the sample, divided by 2 total households in the sample), while Sample 2 has an average of 2.5 children per household, and Sample 3 has an average of 2 children per household. So, these three different samples provide three different estimates of the average number of children per household for the **full** population of three households. And, the estimate from Sample 3 is the only one that is identical to the average for the full population of three households that we calculated in the last slide.

The variation between these sample estimates and the actual value for the full population is what we call sampling error. It results from collecting information from only a sample rather than the full population.

Slide 7

Slide 8

There are four key measures of sampling error that ACS data users should understand –standard error, margin of error, confidence interval, and coefficient of variation. Each of these measures are related to one another.

In the next section of this presentation, we will examine each of these measures in detail to understand how each is calculated and interpreted, and to understand how these different measures are related to each other.

Slide 9

The first measure of sampling error we will discuss is the standard error. The standard error is a measure of the variability of an estimate due to sampling. It indicates the extent to which an estimate derived from a sample survey can be expected to deviate from the population value.

The standard error for an estimate depends upon the underlying variability in the population for the characteristic as well as the sample size used for the survey. For example, if 80 percent of households in a population have two children, then the standard error of the estimate of average children per household will be smaller than in another population where there is more variation among households in the number of children.

The standard error is a foundational measure from which other sampling error measures are derived, and it is required in order to conduct tests of statistical significance. However, standard errors are not usually used for interpretation.

Slide 10

The Census Bureau provides the Margin of Error for each published ACS estimate. We'll look at margins of error in more detail a little later, but as this slide shows, the standard error for an ACS estimate can be obtained by dividing the published margin of error for the estimate by the value 1.645. For estimates from years 2005 or earlier, use the value 1.65 with the published margin of error.

Let's calculate a standard error using some of the 2007 ACS data for the city of Baltimore, Maryland. The 2007 ACS provides an estimate of 52.1 percent for males age 15 and older who live in Baltimore city and have never married. The published margin of error for this estimate is 1.7 percent. When we divide the margin of error of 1.7 percent by the value 1.645, we get a standard error of 1.033 percent.

Slide 11

The margin of error is defined as a measure of the precision of an estimate at a given level of confidence. The most commonly used confidence levels are 90%, 95% and 99%. What does the confidence level of a margin of error mean?

The confidence level of a margin of error indicates the likelihood that the difference between the population value and the sample estimate is less than or equal to the margin of error.

All ACS estimates are published with their margins of error at the 90 percent confidence level. However, it is possible to construct margins of error with higher levels of confidence, such as 95 percent or 99 percent. This is done by adjusting the published margin of error. Instructions for these adjustments can be found in the technical appendices of *The ACS Compass Products Handbooks* now available on the American Community Survey web site.

Slide 12

Even though all ACS estimates are published with margins of error at the 90 percent confidence level, it's important to understand how that measure is constructed. It will help you understand the relationship between the different measures associated with sampling error.

The margin of error is a multiple of the Standard Error (SE). The margins of error published in the ACS are calculated at the 90 percent confidence level by multiplying the Standard Error by the value 1.645. The 1.645 value is related to the 90 percent confidence level. Any other confidence level, whether 95 percent or 99 percent, will have different values associated with them. Those values will increase as we increase our level of certainty or confidence level.

Slide 13

How do we interpret the margin of error for an estimate?

At a 90 percent confidence level, the margin of error indicates that there is a 90 percent probability that the estimate and the population value differ by no more than the value of the margin of error. In other words, we can be 90 percent certain that the range established by the margin of error contains the population value.

Margins of error are useful in assessing the reliability of estimates and whether differences between estimates are significant.

Later, we'll talk more about the specific meaning of the word "significant."

Slide 14

Let's continue using our previous example about never married males living in Baltimore, Maryland. The ACS estimate shows that, of the males age 15 and over who live in Baltimore city, 52.1 percent have never married. That estimate has a margin of error of plus or minus 1.7 percent, which establishes a range of 50.4 percent to 53.8 percent. Now, let's say we actually asked every male age 15 or older who lives in Baltimore city if they have ever been married. The margin of error of the ACS estimate means that there is a 90 percent chance that the range of 50.4 percent to 53.8 percent would contain the population value.

In general, the larger the margin of error relative to the size of the estimate, the less reliable is the estimate. The margin of error provides a concise measure of the precision of an ACS sample estimate.

Slide 15

The third key measure of sampling error is the confidence interval. It is defined as a range that is expected to contain the population value of the characteristic with a known probability.

Slide 16

The confidence interval is calculated by subtracting the margin of error from the estimate, to create the lower bound, and by adding the margin of error to the estimate to create the upper bound. The range between the upper and lower bounds of the confidence interval is expected to contain the population value with a given level of confidence.

Now, let's calculate the confidence interval for the ACS estimate of the percent of never married males in Baltimore city.

Slide 17

When we subtract the margin of error from the estimate, we get a lower bound of 50.4 percent, and when we add the margin of error, we get an upper bound of 53.8 percent.

The resulting confidence interval ranges from a low of 50.4 percent of the never married males to a high of 53.8 percent of never married males.

How should this confidence interval be interpreted?

Slide 18

In simple terms, this confidence interval means that we can be 90 percent certain that the confidence interval from 50.4 percent to 53.8 percent contains the population value of never married males 15 years and older in Baltimore City.

If you'll recall during our discussion of margin of error, we said the margin of error establishes a range. The upper and lower bounds of that range define the confidence interval.

In general, for similar size estimates, larger confidence intervals indicate that estimates are less precise and reliable.

Graphing confidence intervals is a useful way to provide a quick visual comparison of the extent of sampling error for different estimates.

Slide 19

In this graph, the dots show the ACS estimate and the lines show the confidence interval. It is easy to see that the confidence intervals are much larger for the 2003 and 2004 ACS estimates for Baltimore.

The confidence intervals for the estimates for 2005 through 2007 are much smaller, indicating that these estimates have smaller margins of error and are more reliable. This makes sense because the

sample size of the ACS tripled in 2005 when the survey moved out of its testing phase and expanded to include each county in the United States and Puerto Rico in its annual sample.

Slide 20

The fourth measure of sampling error we want to cover in this presentation is the coefficient of variation, or CV. It is defined as the **relative** amount of sampling error associated with a sample estimate, and is used to assess the reliability of an estimate.

The CV is calculated as the ratio of the standard error for an estimate to the estimate itself, usually expressed as a percent, which is why it is multiplied by 100 percent in the formula. A small CV indicates that the sampling error is small relative to the estimate, and therefore the user can be more confident that the estimate is a good approximation to the population value.

Because they are expressed as percents, it is easier to compare the reliability of two different estimates using CVs than it is using margins of error.

Slide 21

Now, continuing our example, let's calculate the CV for the ACS estimate of males age 15 and over living in Baltimore who have never married.

The estimate for never married males is 52.1%, and we calculated its standard error as 1.033%. The CV equals the standard error divided by the estimate, and we get a value for this CV of 1.98%.

So, how should CVs be interpreted?

Slide 22

The size of a CV is important for its interpretation. In general, the smaller the CV, the more reliable the estimate.

In the sample calculation we just completed for Baltimore city, we found a CV of only 1.98 percent. In other words, the sampling error for this estimate is about 2 percent the size of the estimate.

However, there are no hard and fast rules about the usability of an estimate based on the size of the CV. Smaller CVs imply better estimates, but the question of "when is the CV too large?" cannot be answered in a vacuum. It all depends on the context and how the estimate will be used. For example, if you just want to know if an estimate is below a certain value, an estimate with a large CV may be reliable enough for this purpose. Someone else may need an extremely precise estimate, and even a small CV (as small as 2 percent) may be too large for their use.

For proportions close to zero, the CV may be unstable. In this case, data users are advised to rely on the margin of error or confidence interval rather than the CV.

In general, data users are encouraged to consider both the margin of error and the CV when evaluating the reliability of an estimate.

Slide 23

Sampling error is related to sample size. Specifically, the larger the sample size, the smaller the uncertainty or sampling error for estimates based on the survey sample.

Recall the example of estimating the average number of children per household. Now imagine instead of three households we had a full population of 100 households. It's easy to see that our estimate of average children per household generally would have much more uncertainty if we sampled only two of those 100 households, rather than say 30 of those 100 households.

This relationship between sample size and sampling error is the reason that ACS estimates for smaller geographic areas are based on multiple years of data. Combining ACS data from multiple years increases the sample size and in turn reduces the uncertainty or error associated with ACS estimates.

Sampling error is not unique to the ACS – all sample surveys have it, including the decennial census long form. Because the long form was collected from only a sample of all U.S. households, long-form estimates also have sampling error associated with them. Instructions for calculating standard errors for long form estimates can be found in the technical documentation for Summary File 3 or SF-3 on the Census Bureau's website.

Slide 24

Slide 25

So far, we have examined only one of the important ways measures of sampling error can be used – to determine whether an estimate is reliable and usable.

There are two additional ways sampling error measures can be used:

1. To make comparisons between two or more estimates, for example to conduct formal tests of statistical significance, and
2. To help data users draw appropriate conclusions about the data.

The case study we'll examine in the next section of the presentation shows how to use measures of sampling error in each of these ways.

Slide 26

Suppose that public officials in Washington, DC have implemented a series of job training initiatives they hope will increase employment and reduce poverty rates.

In 2008, after the programs have been in place for several years, officials want to assess whether there has been any change in poverty rates in the city.

Slide 27

Where can Washington, DC officials find data to address this question?

Since the resident population in Washington, DC exceeds 65,000, there are comparable data for both 2006 and 2007 from the ACS available in American FactFinder.

The 2006 ACS provides a good starting point because it provides some information about the economic characteristics of 2005. This is due to the fact that income data collected in the 2006 ACS are based on income in the “previous 12 months.”

Officials can use the ACS to examine change in the percent of people who were living in poverty in 2006 versus 2007.

Slide 28

City data analysts can use the Census Bureau’s American FactFinder website to find the estimates of people living in poverty, along with the margins of error for each estimate.

In 2006, the ACS found 19.6% of all people living in poverty in Washington DC, with a margin of error of 1.4%. The 2007 estimate for percent of people living in poverty was 16.4%, with a margin of error of 1.4%.

These estimates appear to indicate that the poverty rate has declined between 2006 and 2007. But, the analysts know they must first carefully evaluate several measures of sampling error to determine whether the estimates are reliable and usable.

Slide 29

The analysts first calculate standard errors for each of the estimates using the margins of error provided in the AFF tables. They use the denominator of 1.645 to calculate the standard errors for both estimates.

The standard error for the 2006 estimate is equal to the margin of error, 1.4%, divided by 1.645, which equals 0.85%. The CV is the standard error, 0.85%, divided by the estimate, 19.6%, which yields a CV of 4.3%. Similar calculations for 2007 give a CV of 5.2%. Both CVs are relatively small – about 5 percent in both cases.

The analysts also note that neither estimate of the percent of the population living in poverty is close to zero, so the CVs are appropriate measures of reliability.

The analysts conclude that the estimates for both 2006 and 2007 are reliable and usable for their analysis.

Slide 30

The analysts begin their comparison of these two estimates by constructing the confidence intervals around each estimate.

The analysts decide they must conduct a formal test of statistical significance.

Slide 31

When comparing survey estimates, it is very important that the comparison takes into account the sampling error associated with each estimate.

A test of statistical significance provides *statistical* evidence that indicates whether an observed difference between two estimates is likely due to chance, or likely represents a true difference that exists in the population as a whole.

Slide 32

This slide shows the formula to calculate a test of statistical significance between two estimates. In the formula, X1 and X2 represent the two estimates, and SE1 is the standard error for X1 and SE2 is the standard error for X2.

Z is the critical value for the desired confidence level for the test. For a 90% confidence level, Z is equal to 1.645.

Slide 33

The city data analysts plug in the appropriate values to determine whether the decline in the poverty rate in Washington, DC between 2006 and 2007 is statistically significant. They choose a desired confidence level of 90 percent. As indicated in the previous slide, the critical value for a 90 percent confidence level is 1.645.

Their calculation yields an absolute value of 2.662, which is indeed greater than 1.645. They conclude that the decline between 2006 and 2007 in the percent of people living in poverty is statistically significant at the 90 percent confidence level.

Slide 34

Although the data analysts report back to Washington, DC public officials that the decline between 2006 and 2007 was statistically significant at the 90 percent confidence level, they caution that it is too soon to tell if this is a real trend, or just a short-term fluctuation in the poverty rates. The analysts recommend that public officials track poverty rates for several more years to determine if the 2006 to 2007 decline continues.

The critical value for the test at the 95 percent confidence level is 1.96, and the critical value for a test at the 99 percent confidence level is 2.576. It can be noted that the difference between 2006 and 2007 is statistically significant at both the 95 and 99 percent confidence levels, as well.

Slide 35

Slide 36

It is also important to mention that there is an additional source of error in all censuses and surveys, including the ACS. It is called non-sampling error.

Non-sampling error refers to any error associated with a survey or census estimate other than sampling error. The important distinction is that, unlike sampling error, non-sampling error occurs in both sample surveys and censuses.

Slide 37

There are four major types of non-sampling error – non-response error, response error, processing error, and coverage error.

Non-response error occurs when some households or people refuse to participate at all in a census or survey, or when they refuse to answer some questions (which is called item non-response). If the total or item non-response rates are high, then the estimates might be biased. This will happen if the characteristics of those who don't respond are different from the characteristics of those who do respond.

Response error occurs because some respondents may intentionally or unintentionally misreport information. For example, they might not accurately remember what their income was in the last twelve months, or they may not want to report previously undisclosed income.

The third major type of non-sampling error is processing error. This could occur during data entry or coding of survey responses.

Finally, coverage error occurs when a housing unit or a person does not have a chance of being selected for a census or survey. For example, if the census or the ACS did not include newly constructed housing units, then the people who lived in those housing units would be excluded, and

this could cause coverage error in the data. In addition, coverage error can occur if a person is incorrectly included in or omitted from a household in sample.

Non-sampling error is difficult to measure directly but there are indirect measures available on the Census Bureau's website.

Slide 38

One way to reduce coverage error is the use of population controls.

Slide 39

Independent estimates from the Census Bureau's Population Estimate Program are used as controls for the ACS. Housing unit estimates are controlled at the county level for total housing, while total population estimates are controlled at the county level (or groups of counties) by age, sex, race, and Hispanic origin. The group quarters total population is controlled at the state level by major type of group quarters.

However, ACS estimates for these characteristics will not necessarily match the independent estimates. Total population for counties and states will generally match, but other estimates, such as total housing units, population by race, and total population for areas below counties, generally will not match, although they will likely be close.

The use of population controls is important to reduce sampling and non-sampling error in the ACS estimates. In particular, using controls reduces coverage error due to both missing housing units, and also missing people within housing units or group quarters.

As more current data become available, the time series of population estimates back to the preceding census are revised, while the ACS estimates for previous years are not. Therefore, some difference in the ACS estimates across time may be due to changes between the revised series.

The ACS is designed to measure the characteristics of the population and housing units. The official estimates of basic population and housing unit totals still come from the decennial census and from the Census Bureau's Population Estimates Program.

Slide 40

Slide 41

Let's recap the main topics we have covered in this presentation.

First, all surveys have sampling error and non-sampling error, including the ACS and the decennial census long form. Sampling error results from collecting information from only a sample rather

than the full population. Non-sampling error results from issues with response, processing, and coverage and is difficult to measure directly.

There are four key measures of sampling error that ACS data users should understand. They are standard error, margin of error, confidence interval, and coefficient of variation. The Census Bureau provides the MOE for most published ACS estimates, and data users can derive the standard error, confidence interval, and coefficient of variation from the MOE.

These four measures of sampling error provide important information about the reliability of ACS estimates. Data users should not use ACS estimates without first assessing their reliability.

Slide 42

In addition to providing information about the reliability of ACS estimates, sampling error measures can also be used to make comparisons between estimates and to conduct tests of statistical significance.

Finally, it is only by understanding and correctly using measures of sampling error that data users can draw appropriate conclusions about ACS data.

Slide 43

This presentation gave you an overview of sampling error and other things that may affect the estimates from the American Community Survey.

The American Community Survey staff has developed the ACS Alert, which is an e-mail newsletter giving data users the latest news about the survey. You can subscribe to the newsletter by contacting the American Community Survey staff or read past editions of the “ACS Alert” on the Internet at: <http://www.census.gov/acs/www/Special/Alerts.htm>

Please feel free to contact the Census Bureau if you have questions or need further information. If you have questions that are not answered by the Web site, please call 1-800-923-8282 or email acso.users.support@census.gov.