AMERICAN COMMUNITY SURVEY: Evaluating Accuracy

By Terri Ann Lowenthal

PRB Population Reference Bureau

kids count A project of the Annie E. Casey Foundation

August 2006

Population Reference Bureau (PRB)

The Population Reference Bureau informs people around the world about population, health and the environment, and empowers them to use that information to advance the well-being of current and future generations. PRB's work is funded by private foundations, government agencies, and individual donors, and we frequently collaborate with other nonprofit organizations and universities. PRB brings broad expertise and innovative, cost-effective approaches to analysis, information sharing and capacity building to these partnerships.

The Annie E. Casey Foundation — KIDS COUNT

KIDS COUNT, a project of the Annie E. Casey Foundation, is a national and state-by-state effort to track the status of children in the United States. By providing policymakers and citizens with benchmarks of child well-being, KIDS COUNT seeks to enrich local, state, and national discussions concerning ways to secure better futures for all children.

At the national level, the principal activity of the initiative is the publication of the annual *KIDS COUNT Data Book*, which uses the best available data to measure the educational, social, economic, and physical well-being of children. The Foundation also funds a nationwide network of state-level KIDS COUNT projects that provide a more detailed community-by-community picture of the condition of children.

For more information or for a pdf version of this report, visit the Annie E. Casey Foundation's KIDS COUNT website at www.kidscount.org or PRB's website at www.prb.org.

This research was funded by the Annie E. Casey Foundation. We thank them for their support but acknowledge that the findings and conclusions presented in this report are those of the author alone, and do not necessarily reflect the opinions of the Foundation.

EXECUTIVE SUMMARY

In 2005, the U.S. Census Bureau launched its new American Community Survey (ACS) in 3 million households nationwide. The ACS will replace the traditional long form sent to a sample of homes during the decennial census. It will be the only source of small-area data on a wide range of important social and economic characteristics for all communities in the country.

Like the long form before it, the ACS will produce information that will guide policymakers in the public sector and corporate leaders in the private sector. ACS data will shed light on the demand for social services for the elderly, child care, English-as-a-second-language classes, job training, veterans' programs, and affordable housing assistance. Some ACS data, including information about poverty, will be used along with population counts to allocate federal funds among state and local governments or to establish eligibility for programs and grants. In the private sector, the ACS will help businesses decide where to locate new stores and plants, and how much to invest in capital improvements, marketing, and human resource development.

Because the ACS will replace the once-a-decade census long form, it is important to have confidence in the accuracy of the data. Anyone who follows the decennial census is familiar with longstanding concerns about accuracy. Most census debates revolve around the accuracy of the population count, especially in terms of basic characteristics such as age, race, and ethnicity. Policymakers and the public have focused less attention on the accuracy of the broader range of demographic and socio-economic data gathered on the long form. The quality of characteristics data—that is, how accurately the data reflect the "true" characteristics of our population and housing—will be a key factor in deciding whether the ACS can be used to guide policy decisions.

Census accuracy is best understood in terms of undercounts and overcounts. How many people were missed, and how many people were counted twice? Although evaluating ACS accuracy is not always as simple, there are several ways to measure the quality of both the survey and the results. These evaluations can help decision-makers assess the accuracy of ACS data and compare the results to the census long form.

After a brief overview of the American Community Survey, this report will examine:

- Primary measures of ACS data quality, including response and coverage rates, for key characteristics;
- Ways to improve the accuracy of the ACS; and
- Recommendations to improve the accuracy and quality of ACS data.

Two appendices offer more detailed information on sources of error in the ACS and scientific methods for minimizing the effects of potential errors.

THE AMERICAN COMMUNITY SURVEY IN BRIEF

The American Community Survey is a monthly survey designed to collect community-level information on key demographic, social, and economic characteristics of America's population and housing. The ACS samples about 250,000 housing units across the country each month, or 3 million each year.[1] Over a five-year period, the ACS will collect data on 15 million addresses, enough to produce estimates for areas as small as census tracts and block groups. By comparison, the 2000 Census long form, which also produced data for census tracts and block groups, was sent to 19.2 million addresses.

The Census Bureau has been developing the ACS for a decade. Testing of the survey's concept began on a small scale in 1996, in four counties. Over time, the Census Bureau added more test sites, eventually conducting the survey in 36 counties around the country. Starting in 2000, the Census Bureau also fielded a national sample of roughly 700,000 households—called the Census 2000 Supplementary Survey (C2SS)—in order to assess ACS operations on a larger scale and compare ACS estimates with information collected in the 2000 Census. The Supplementary Survey continued in the field in subsequent years (2001-04). Data from the 36 test sites and Supplementary Survey provide the basis for recent evaluations of the ACS. (References in this report to published ACS data generally mean information collected from the test sites or Supplementary Survey.)

Unlike the census, the ACS is not designed to provide total population counts. The census, which reaches every home and group facility in the United States, counts the number of people and produces official population numbers for all governmental units—the nation, states, counties and cities, townships—as well as for designated geographic units such as census tracts and blocks.[2] By contrast, the primary purpose of the ACS, which collects information from only a sample of the population, is to produce profiles of key population and housing characteristics for state and local communities. The decennial census (and population estimates issued annually during the decade) will continue to be the official sources of population counts for apportionment, legislative redistricting, and the allocation of federal program funds based solely on population. In contrast, the ACS will provide updated information on important characteristics of America's population, including income, employment and occupation, education, primary language, disabilities, marital status, commuting patterns, and housing (rent and mortgage payments, age and condition of homes, number of rooms, sources of energy). The ACS can also be used to track changes among key population subgroups, such as children, the elderly, racial and ethnic minorities, and immigrants.

MEASURING ACS QUALITY

The Census Bureau uses four primary "quality measures" to gauge the reliability and accuracy of ACS data: sample size, response rates, item allocation rates, and coverage rates.

Sample size

The Census Bureau determines sample size for different geographic areas in order to ensure that it is collecting information from enough homes to produce data that are statistically reliable.

Between 2000 and 2004, the ACS sampled more than 700,000 households each year, providing reliable annual estimates for states and geographic areas with at least 250,000 people. Beginning in 2005, the ACS was mailed to 3 million households nationwide and will provide reliable yearly estimates for smaller geographic areas with populations of 65,000 or more. Five-year averages are intended to provide reliable estimates for geographic areas down to the census tract and block group levels starting in 2010.

Response rates

Response and nonresponse rates are used to evaluate how successfully the ACS gathered information from homes in the sample. For the ACS, the overall response rate represents the proportion of homes from which the Census Bureau collected sufficient information by mail, telephone interview, and personal interview, compared to all homes from which information should have been collected.

The Census Bureau has published weighted response rates for the ACS test phase, from 2000 through 2004 (see Table 1).[3] The response rates are considered high for a household survey.[4]

Table 1
Response Rates (Weighted) in the ACS[1]

Year	Weighted response rates (percents)
2004	93.1[2]
2003	96.7
2002	97.7
2001	96.7
2000	95.1

1 Response rates are weighted to compensate for sampling.
2 Due to insufficient funding, the Census Bureau did not conduct follow-up interviews by telephone or in person for nonresponding households that received the ACS questionnaire in Jan. 2004.
Source: U.S. Census Bureau.

Conversely, nonresponse rates reflect the percent of homes from which the Census Bureau was unable to collect enough information for a successful interview. Some people simply do not answer the ACS—by mail, telephone, or when a survey taker knocks on their door. If the people who do not respond have the same characteristics or live in the same types of communities as those who do respond, their absence from the survey would not be as big a problem, because the data collected would still be representative of the population as a whole. But in the ACS, like the 2000 Census, data are less likely to be reported for racial and ethnic minorities than for non-Hispanic whites.

Evaluations of response in the ACS county test sites and in the 2000 Census show generally higher response (and, conversely, lower non-response) in the ACS than in the census. For example, an evaluation of housing unit response in the 36 ACS test sites from 1999 – 2001 showed an average nonresponse rate of 4.4 percent, compared to 9.7 percent for the Census 2000 long form. Nonresponse rates for occupied housing units were closer but still notably better for

the new survey: an average of 5.2 percent for the ACS counties and 8.7 percent for Census 2000.[5]

This finding is not necessarily surprising. Although a higher percentage of homes in the census than in the ACS return their questionnaires by mail, the ACS benefits from professional interviewers and stricter standards for collecting data in person, which improve the chances of gathering sufficient information in the field. In contrast, the census is focused on counting the population; collecting socio-economic data on the census long form has been an important but secondary goal. Census enumerators do not have as much time in the field as ACS interviewers to obtain information from unresponsive households; after several unsuccessful tries, census takers might obtain a population count for a household and move on.

Coverage

While response rates measure the extent to which the ACS successfully collects information from housing units in the sample, coverage evaluations focus on the data and the extent to which the survey results adequately capture different population groups. Coverage rates show how completely population groups are included in the ACS, compared to the official population estimates of these groups. The numbers presented in Tables 2 and 3 show ACS coverage rates before the ACS figures are controlled to census population estimates.[6] A coverage rate of 96.2 percent for non-Hispanic whites means that the ACS collected data for 96.2 percent of this population group when compared to the Census Bureau's official estimate of non-Hispanic whites. In Table 2, a coverage rate above 100 percent means that a population group is over-represented in the ACS compared to the independent population estimate. Over-representation can occur if individuals in the population group are erroneously counted twice or, as is more likely, the population estimate is lower than the true size of that group.

The Census Bureau calculates coverage rates for the ACS by gender and by race/ethnicity, but not for cross-tabulated subgroups of these characteristics (e.g. for Hispanic males) or for age cohorts, as it does for the census.[7] For groups with coverage rates below 100 percent, the Census Bureau must adjust the data upward to match independent population estimates. For groups with coverage rates above 100 percent, the data are adjusted downward.

ACS coverage rates from 2000 to 2004 show a pattern of undercoverage similar to the pattern of census undercount in many, but not all, respects (see Tables 2 and 3). Except for American Indians and Alaska Natives for all years except 2003, there is undercoverage of all race and ethnic groups in the ACS. Coverage of non-Hispanic whites is notably better than coverage of all other race and ethnic groups except American Indians and Alaska Natives, a result similar to the historic differential undercount of race and ethnic minority groups in the census. The "differential undercount" refers to the gap in undercount rates between population groups. Asians have the next highest coverage rates, while coverage rates for blacks and Native Hawaiians and Other Pacific Islanders are the lowest. With respect to gender, coverage of females in the ACS is higher overall than for males, mirroring a historical census pattern.

As Table 2 shows, coverage rates for American Indians and Alaska Natives suggest that these populations were over-represented in the ACS. Given that American Indians, especially those

living on reservations, have historically suffered from large undercounts in the census, it is likely that measurement of overcoverage for this proportionately small population group in the ACS stems from problems associated with sampling or statistical weighting, and not from an actual over-representation of American Indians or Alaska Natives in the survey.

Table 2
Coverage Rates in the ACS by Race and Ethnicity[1]

Year	Total population	White, non-Hispanic	Black, non-Hispanic	AIAN, non-Hispanic[2]	Asian, non-Hispanic	NHOPI, non-Hispanic[3]	Hispanic
2004	94.4	96.2	89.3	100.9	92.4	89.7	90.7
2003	94.1	95.7	89.2	97.9	92.0	86.2	91.3
2002	93.5	94.9	89.0	102.7	92.7	90.0	89.7
2001	94.1	95.2	90.1	102.6	92.8	89.7	91.8
2000	96.8	97.1	93.5	121.6	96.7	100.2	96.4
Average (2000-2004)	94.6	95.8	90.2	105.1	93.3	91.2	92.0

1 Coverage rates below 100 percent indicate under-coverage of the population subgroup, compared to independent population estimates. Coverage rates above 100 percent indicate over-coverage.
2 American Indian/Alaska Native
3 Native Hawaiian and Other Pacific Islander
Source: U.S. Census Bureau, "Using the Data: Quality Measures," (www.census.gov, accessed Jan. 5, 2006).

Table 3
Coverage Rates in the ACS by Sex

Year	Total Population	Male	Female
2004	94.4	93.2	95.6
2003	94.1	92.9	95.2
2002	93.5	92.3	94.6
2001	94.1	93.2	94.9
2000	96.8	96.1	97.6
Average (2000-2004)	94.6	93.6	95.6

Source: U.S. Census Bureau, "Using the Data: Quality Measures," (www.census.gov, accessed Jan. 5, 2006).

Item allocation rates

While some households fail to respond at all to the ACS, others might report some but not all of the requested information. *Item nonresponse* occurs when people do not answer all of the questions on the survey form or provide invalid answers. (See Appendix A, Sources of Error in the ACS, for a more detailed explanation of the types of nonresponse.)

Item nonresponse can lead to inaccuracies in the data because the sample for a specific piece of information, such as educational attainment, will be incomplete and may not accurately reflect the true characteristics of the population. The Census Bureau evaluates the potential effects of item nonresponse on data quality by looking at item allocation rates. The rates reflect the percent of data for a specific ACS question that were imputed using statistical procedures, instead of reported directly by the household. Imputation is a statistical procedure that uses data collected from responsive households to allocate information (fill in the blanks) for nearby housing units that did not answer all questions on the form. For example, if a family does not

report their income, they may be assigned a value based on the income level of their nearest neighbor. Table 4 shows item allocation rates for several population variables and compares them with data collected on the 2000 Census long form. Lower item allocation rates suggest more accurate data, because information reported directly by a respondent is generally of higher quality.

Table 4
ACS Item Allocation Rates[1]

	Percent Allocated					
Item	**2004 ACS**	**2003 ACS**	**2002 ACS**	**2001 ACS**	**2000 ACS**	**2000 Census Long Form**
Age	0.9	1.0	1.0	1.4	2.4	2.6
Wage/salary income	12.7	14.1	14.0	14.8	16.4	20.0
Educational attainment	3.3	3.8	3.6	3.7	4.8	7.2
English ability[2]	2.6	2.9	3.8	3.8	6.0	7.6

1 Item allocation refers to the imputation of missing questionnaire items using statistical procedures.
2 Household population ages five and older who speak a language other than English at home.
Sources: "Using the Data: Quality Measures," (www.census.gov, accessed Jan. 5, 2006); and C.F. Citro, D.L. Cork, and J.L. Norwood, eds., *The 2000 Census: Counting Under Adversity* (2004): table 7.4.

Item allocation rates for the ACS compare favorably with the 2000 Census long form, especially for the most recent evaluations available (2004). This finding suggests that skipped questions are not as great a problem in the ACS as in the census and that the effects of missing data on the accuracy of the results are probably smaller. It is also encouraging that the Census Bureau, with few exceptions, reduced the incidence of item allocation over time during ACS testing for all questions.

Comparisons with Census 2000

Comparisons of population data from the ACS and the 2000 Census also shed light on coverage of specific population subgroups in the ACS. The ACS does not produce population numbers for any official use and is not designed to produce numerically accurate population counts. However, comparing the proportions of population groups reflected in ACS data and in the census helps in evaluating the accuracy of the ACS.

C2SS estimates of race, Hispanic origin, age, gender, and household composition and relationships—items included on the Census 2000 short form sent to all homes—generally mirrored the 2000 Census population distributions at the national level. Comparisons of counties and selected census tracts in the test sites offer further insight into how well the ACS covers demographic subgroups. Below are the most important findings from comparisons between the ACS/C2SS and the 2000 Census for key population characteristics.

Age

Historically, research suggests, children are at higher risk of being undercounted in the census than the elderly. Analyses show a similar pattern in the ACS: lower proportions of children were included in the ACS than in the 2000 Census. Conversely, evaluations show that the ACS

included a higher proportion of persons 62 years or older than the 2000 Census. Overall, however, the age distributions of respondents in the ACS and the census were fairly similar.

Children under 10 were especially at risk of undercoverage in the ACS. As in the census, the undercoverage of children is greatest at the youngest ages (0 to 4), with progressively greater accuracy, compared to 2000 Census population counts, for subsequent age cohorts (ages 5 to 9, ages 10 to 14, and ages 15 to 19). Patterns of accuracy at the county level are consistent with those at the national level.[8]

Early evaluations suggest several possible causes for the undercoverage of children in the ACS. One involves the ACS "continuation forms," which are used to list additional household members in households with more than five persons. During early test phases, a processing error apparently led to a loss of data for people listed on ACS continuation rosters. Because respondents tend to list adults before children on a census form, the sixth and subsequent persons reported for a large household are far more likely to be children. The Census Bureau has corrected the processing glitch. Nevertheless, large households continue to present measurement challenges in both the ACS and census. The five-person limit for reporting complete data on the ACS mail questionnaire poses a coverage risk for children in large households and could affect the reliability of data on children's characteristics.

In addition, the ACS estimated a lower proportion of family households with children under 18 and married-couple households with children under 18 than the census.[9] The Census Bureau attributes the differences to several possible factors, including different residence rules, which could affect the reporting of college students or of children in joint custody situations, questionnaire wording disparities, and different editing procedures. The C2SS form, for example, did not offer more detailed options for reporting children, such as adopted or step children and son or daughter, that the census offered, which could have affected responses. The Census Bureau also believes that the census may have over-estimated some family and household types, especially married-couple families, compared to the ACS, because ACS editing relied on answers to the marital status question to verify answers to the relationship question, while the census could not use that additional source of information because it was collected on the long form only.[10]

Evaluations suggest that undercoverage of very young children has resulted in inaccurate ACS data on preschool and nursery school enrollment. This is important because enrollment figures are used to determine eligibility for various federal programs that benefit disadvantaged children. Accurate enrollment statistics are also needed to determine where new schools and teachers are most needed.

Race and ethnicity

There were notable differences in the ACS and 2000 Census estimates for four race categories: white alone, black alone, some other race alone, and "two or more races."[11] The ACS measured a higher proportion of whites and a smaller proportion of blacks than the 2000 Census, disparities that were reduced but not eliminated after applying population controls to the final

estimates. The percent of respondents reporting some other race and "two or more races" in the ACS also was lower than in the census.

The ACS and census measured identical proportions of the population (12.6 percent) of Hispanic origin, with no statistically significant differences in the estimates. There were, however, notable differences for Hispanic subgroups: the ACS measured a higher proportion of Mexicans, while the census showed a higher proportion of persons reporting other Hispanic or Latino. Because population controls are not broken down by Hispanic subgroup, the final weighting process did not change this outcome. Outcomes at the county level were consistent with differences seen at the national level, especially for counties such as Tulare, Calif., Pima, Ariz., and Bronx, N.Y.) with significant Hispanic populations.

The Census Bureau attributes much of the discrepancy in Hispanic subgroup data to a different ACS interview approach. Persons of Hispanic origin answering the ACS by telephone and in-person were offered subgroup examples such as Argentinian or Dominican if they described themselves as "other Hispanic or Latino" (that is, not Mexican, Puerto Rican, or Cuban, the three largest Hispanic subgroups), making them more likely to report a specific group. By contrast, the 2000 Census Hispanic origin question did not give subgroup examples (a change from the 1990 questionnaire), and census enumerators did not offer any examples during nonresponse follow-up.

The Census Bureau believes that differences in interviewer training may have contributed to a lower proportion of "some other race" responses in the ACS. Many ACS interviewers may have pressed for answers from one of the survey's other race categories, because they were used to working with other federal surveys that do not offer a "some other race" option.

Gender

Evaluations suggest that the census misses a higher percentage of males than females.[12] A similar pattern emerges from the ACS, with the survey including lower proportions of males than females at the national and county levels. Application of the independent population controls corrected the undercoverage of males in the ACS estimates.[13]

IMPROVING ACCURACY IN THE ACS

A number of operational and methodological factors could affect the accuracy of the ACS. (See Appendix A for reasons behind sampling and nonsampling error in the ACS and Appendix B for some of the procedures used by the Census Bureau to compensate for these errors.) None of these factors are unique to the ACS; they affect the census and many other household surveys. What is unique is the design of the ACS, which presents both advantages and disadvantages over the decennial census, as well as opportunities for stakeholders to help improve public participation, the quality of responses, and overall accuracy.

Accuracy of the Master Address File

The ACS sample is selected from the Census Bureau's Master Address File (MAF), a permanent list of all residential addresses that is used also in the census. Several evaluations of the 2000 Census MAF found significant problems with creation of the list. A National Academy of Sciences review of the 2000 Census concluded that building the MAF from multiple sources, such as local records, Postal Service address lists, and the 1990 Census MAF, was "appropriate in concept but not well executed."[14]

The Census Bureau estimated that a net of 0.4 million occupied housing units were not included in the 2000 Census. While this net undercount of occupied housing units appears to be small, the total number of housing units left out of the census, or included twice, was much higher. The Census Bureau estimated that as many as 2.7 million occupied housing units were excluded and 2.3 million were counted more than once in the 2000 Census, resulting in a gross error rate of 4 percent. These mistakes could affect the selection of the sample for the ACS.[15]

Census Bureau evaluations shed little light on whether people were missed in Census 2000 because an entire household was missed (whole-household misses) or because they were not included in a household that was otherwise counted (within-household misses). We also know little about how the accuracy of the MAF varies by population group and geographic area. We do know that more populous geographic areas were far more likely than smaller areas to participate in a pre-census program, called the Local Update of Census Addresses (LUCA) Program, which gave local governments a chance to review preliminary address lists and submit corrections. Local government participation in LUCA varied widely by geographic region, ranging from a low of 18.1 percent participation in New England to a high of 55.5 percent in the Pacific region. More thorough analysis of the relationship between missing housing units and missing persons, and of the demographic and geographic variability of mistakes in the address list, would help determine the extent to which problems with the address list contribute to inaccuracies in the ACS. If the Master Address File is not accurate, the universe from which the ACS sample is selected each month will be incomplete. Unless mistakes in the MAF affect all types of geographic areas and all demographic groups evenly—an unlikely occurrence—the sample will be unrepresentative of the nation as a whole.

The quality of survey data depends, in no small part, on the quality of the sample that is selected. Therefore, it is important that the Census Bureau, working in partnership with local governments and community-based stakeholders, improve the accuracy of its Master Address File on an ongoing basis. This can be accomplished in part through the ACS itself, as field staff correct mistakes discovered through the personal interviewing process. Improving address list development for the 2010 Census will also be a key factor in ensuring more complete coverage in the ACS in future years. Community service organizations, affordable housing advocates, and other stakeholders with access to information about the housing stock should work with the Census Bureau to identify newly-constructed units, especially in low-income neighborhoods and rural areas, as well as demolished units, which are not uncommon in urban areas undergoing significant regentrification.

Questionnaire format and question wording

A large body of research has shown that survey question order, question wording, and overall layout of a questionnaire can influence responses. The ACS questionnaire is substantially similar in content to the 2000 Census long form, but there are notable distinctions. Some of the differences relate to the different residence rules used in the census and the ACS. In some cases, the order of questions is different. For example, the ACS includes a question on marital status immediately following the question on relationship to householder, while the 2000 Census question on marital status followed the question on race, for respondents who received the long form.

The Census Bureau will begin final testing of questionnaire content for both the ACS and the 2010 Census in 2006. A mail-out, mail-back test will help refine the wording of questions to ensure the most accurate responses and test wording for any new questions. One especially noteworthy topic for further study is the wording of the race and Hispanic origin questions, as well as a possible new question on ancestry. Evaluations of the 2000 Census determined that the absence of examples in the Hispanic origin question likely contributed to an underreporting of some subgroups, such as Salvadorans, Guatemalans, and Dominicans.[16] While the C2SS and the 2000 Census used an identical Hispanic origin question, telephone and personal interviewers in the C2SS offered subgroup examples, which apparently encouraged respondents to provide more detailed answers about their national origin than were provided in the 2000 Census. Another well-documented phenomenon related to question wording involves the ancestry question, which was asked on the 2000 Census long form. The number of respondents reporting a particular ancestry often rises when that specific ancestry is among examples in the write-in question.

Effective implementation of the 2006 test and careful, timely evaluation of the results should help improve the accuracy and consistency of responses in the future. Data collected in the ACS and in the census are more likely to be comparable if the wording of questions in both surveys is similar, because interviewers and respondents are less likely to interpret the questions differently.

Population controls

Final ACS data are controlled, or adjusted, intercensal population estimates at the county level. Intercensal estimates are based on population counts from the most recent census and updated annually using records on births, deaths, immigration, migration, and other administrative sources. Flaws in the data used to compile the population estimates, including under- and overcounts in the census data, can influence the accuracy of ACS numbers. For example, final ACS estimates of males and young children, even after applying population controls, might still reflect an undercount because the census also misses greater proportions of males than females and of young children than older persons.

Intercensal estimates also are likely to be less accurate later in the decade, because they are farther in time from the last census. In addition, immigration, especially of undocumented persons and temporary residents, is notoriously difficult to quantify and is not measured well at the sub-national level; county to country migration also is difficult to track accurately. These

additional weaknesses in the numbers used as controls for the ACS can compromise ACS accuracy.

ISSUES FOR FURTHER STUDY

Not enough is known about coverage in the ACS to draw definitive conclusions about the extent of any "undercount" or "overcount" in the survey, as we understand those terms in the census context. Nationwide implementation of the ACS in 2005 and the addition of the group quarters population in 2006 offer much-needed opportunities to evaluate more fully the representation of race and ethnic groups, age cohorts, and males and females in the ACS, and whether nonresponse, mistakes in the address list, and other measures of quality affect some groups more than others.

Coverage rates and sample completeness ratios point to disproportionate under-representation of race and ethnic minorities, males, and young children in the ACS (even after weights and population controls are applied). There are points in the ACS process that make the survey vulnerable to greater undercoverage of groups that are harder to count in the census.

1. Mistakes in the Census Bureau's Master Address File, which is the basis for the ACS sample, could contribute to undercoverage of housing units, especially in small, multiple-family structures in densely populated urban areas.

2. The five-person limit for reporting complete data on the ACS mail questionnaire poses a coverage risk for children in large households and could affect the reliability of characteristics data about children. The Census Bureau should continue to explore the possibility of requesting age and relationship information on the ACS continuation roster, to help improve the accuracy of data on children.[17]

3. The dependent relationship between the ACS and intercensal population estimates could make it more difficult to identify coverage problems in both programs. Final ACS data are adjusted based on the population estimates, and data collected in the ACS are used to help develop the estimates. The Census Bureau should continue to work on independent approaches to ACS coverage measurement, such as comparisons with administrative records, which will help it better detect patterns and causes of under- and overcoverage.

4. In a comprehensive 2004 report on the ACS, the U.S. Government Accountability Office (GAO), the audit arm of Congress, said that revised population and housing controls are needed to benchmark the full ACS starting in 2005. The controls used in the test phase generally are available only at the county level and for some subcounties. Controls used in the future must be consistently accurate—especially for key characteristics such as age, race, and gender—for increasingly smaller geographic areas, as the ACS continues to produce smaller-area data over time, down to the block group level. The GAO also suggested adjusting controls to more accurately reflect the "current residence" concept used in the ACS (see Box 1). Auditors were concerned about how slowly the Census Bureau was moving to develop new methodologies for population and housing controls. They also noted that it could be more difficult to compare ACS population and housing characteristics data with

data from other federal surveys and from the 2000 Census long form, once the controls are changed.

Box 1
Residence Concepts in the Census and ACS

The decennial census and American Community Survey use two different concepts of residence, which determine where people should be properly counted. The census counts people at their "usual residence," the place where they live or sleep most of the time as of Census Day (April 1). College students living away from home, for example, are counted at their college, not at their parents' residence.

The ACS uses a "current residence" rule: People are counted at a housing unit if they have lived there for at least two months before the survey or intend to stay for at least two months after the survey. So, for example, if a household with a college age child at home for the summer receives the questionnaire in July, the student would be counted at that residence, not at his university.

The Census Bureau did not agree with many of the GAO's recommendations regarding population and housing controls used to calibrate the ACS estimates. Nevertheless, the accuracy and appropriateness of the controls are an important factor in compensating for nonresponse and coverage deficiencies in the ACS.

5. During the ACS test phase, the Census Bureau assessed response and coverage by race, ethnicity, and gender. Similar evaluations for cross-tabulations of these characteristics (e.g. black females) and by age cohort would illuminate accuracy issues more fully. These additional evaluations would help the Census Bureau and data users better understand specific strengths and weaknesses in the data and the extent to which the ACS is susceptible to the same patterns of under- and overcoverage found in the census. (More focused research on the coverage of race and ethnic subgroups, such as Mexican Americans and Chinese Americans, also would be helpful in this regard.) They also would help build data users' confidence in the data.

6. There are two primary ways a census or survey fails to account for all persons: The counting process can overlook an entire housing unit (called "whole household misses") or it can overlook some of the people in a housing unit that is otherwise counted (called "within household misses). The former problem often stems from an incomplete housing roster which forms the basis for a census or from which a sample is drawn; the latter problem can signal a misunderstanding of rules governing whom should be reported as a household resident, unclear question wording or questionnaire format, or processing errors. Evaluations of how each type of "miss" contributes to overall inaccuracy in the ACS would help the Census Bureau correct methodological deficiencies and improve accuracy among historically vulnerable subgroups, especially children.

7. In 2010, ACS implementation will coincide with the decennial census. An extensive promotion and outreach campaign will accompany the short form-only census and, presumably, help boost participation in the constitutionally-mandated population count. While much attention will focus on the accuracy of the population figures, and while many organizations will promote census participation to their constituencies, the effects of conducting the census simultaneously with a full scale ACS are unknown. Households that receive the ACS in 2010 might ask why they are receiving a second questionnaire, especially if the census marketing campaign emphasizes the new, streamlined nature of the decennial count. People in ACS households may be confused, since the survey is promoted as part of the census. ACS response rates could drop, the personal interviewing process could temporarily become more challenging, and data quality could suffer. The Census Bureau should consider the potential consequences in advance and work with partner organizations to minimize confusion and maintain cooperation from households selected for the ACS.

   The Census Bureau and its partner organizations should continue to explore effective ways to promote ACS participation in the years between censuses. A comprehensive, national campaign during the census encourages cooperation but cannot be replicated in non-census years. Even so, the ACS is by far the government's largest ongoing survey, and effective means to promote participation, especially among population groups at risk for undercoverage, should be discussed and implemented.

8. Finally, the Administration and Congress should allocate sufficient funding, every year, for continued and expanded evaluations of ACS quality. Congress is a major census and ACS stakeholder. Lawmakers use census and ACS data extensively to identify areas of need, evaluate programs, allocate funds, and develop new initiatives. It is in their best interest to ensure the highest quality and most accurate data possible, to support effective policymaking.

APPENDIX A: SOURCES OF ACS ERROR

Like any survey, the ACS might not be accurate for several reasons. Because the ACS is sent to only a sample of homes, the information it gathers might not accurately represent the make-up of the population. Disproportionately high rates of nonresponse among minority households can also affect how well the survey accounts for different population subgroups (see Box 2).

Sampling Error

The American Community Survey is sent each month to a sample of roughly 250,000 housing units nationwide, representing all housing units in the United States. The addresses are selected from the Census Bureau's Master Address File, which also is the basis for the decennial census. The monthly sample includes a larger proportion of addresses in less populated counties, towns, and American Indian reservations.

Box 2
What Do We Mean by Accuracy?

Put simply, accuracy tells us how close the final data (e.g. population counts, characteristics data) are to the "truth." Two basic types of error can affect the accuracy of the American Community Survey: sampling error and non-sampling error. Sampling error affects any data collected from less than 100 percent of the universe (e.g. of people or households) being measured. Non-sampling error refers to all other mistakes in survey or census taking that can affect the quality of the final estimates, such as disproportionate nonresponse, interviewer mistakes, respondent confusion, and processing errors. Sampling error does not affect the accuracy of the decennial census population counts; non-sampling error can and does influence the reliability of census population numbers. While sampling error adds measurable uncertainty to any data derived from a sample, non-sampling error is more difficult to identify, measure, and quantify.

Sampling error occurs when information is collected from a subset of people or addresses being measured, instead of from all units (i.e., the short form portion of the census). A larger sample increases the likelihood that an address or person will be included, and lowers the uncertainty in the final data, because there is greater confidence that homes and people in the sample will be representatives of all homes and people. [18] A smaller sample results in greater sampling error and increased uncertainty, often called "variability," in the final data.

Like the census long form it is replacing, the ACS relies on scientific sampling principles to produce information that experts consider accurate enough to use for important governmental purposes such as allocating funds to communities in need. The sample of homes that receives the ACS must be representative of all homes, in order for data users to have confidence in the accuracy of data on characteristics such as educational attainment, employment status, and language spoken at home. The more housing units that are included in the survey, the greater the likelihood that they are indeed representative of all housing units, giving data users more confidence in the information produced. Statisticians describe this level of confidence as "sampling error." The larger the number of homes in the sample, the smaller the level of sampling error and the greater confidence there is in the results.

The ACS design includes sampling at several points. Like the census long form, the ACS is sent to only some homes in the U.S.—about three million each year starting in 2005, representing three percent of all residential addresses nationwide.[19] The ACS sample is not spread evenly across geographic areas. It includes a greater proportion of housing units in sparsely populated rural communities and a lower proportion in densely populated areas. The sample size is big enough to produce annual information for the 50 states and the District of Columbia, as well as cities, counties, metropolitan areas, and population groups with 65,000 or more people.

Not all of the homes that receive an ACS questionnaire send it back. The Census Bureau tries to reach unresponsive households first by telephone through Computer Assisted Telephone Interviewing, or CATI. Next, it selects only a sample of homes that still haven't responded to the survey, and sends interviewers to knock on those doors and gather the information in person through Computer Assisted Personal Interviewing, or CAPI.

Originally, survey takers were sent to knock on doors of one in every three homes that did not mail back a questionnaire or respond by telephone. The Census Bureau staff realized, however, that this uniform sampling rate did not produce the level of representation it needed to have confidence in the results, because members of minority race and ethnic groups are less likely than whites to mail back their questionnaires or to respond by telephone (see Table 5). Interviewers visited one of every three homes that did not respond by mail or telephone, even if the number of unresponsive homes in a given county was much higher than in others. Consequently, in counties with low mail and telephone response, interviewers collected information from a smaller number of unresponsive homes, making it less likely that the sampled homes were representative of the entire county. Because minority groups are less likely than whites to respond to the ACS by mail, some data users questioned whether the original ACS plan produced less accurate data for these groups and for smaller national origin populations such as Arab Americans. In response to this concern, the Census Bureau modified its plan so that survey takers personally visited a larger

Box 3
Selecting a Sample for In-Person Interviews

For each monthly sample of 250,000 addresses, the Census Bureau calculates the combined response by mail and telephone for each census tract. For tracts with response rates below 35 percent, survey takers are sent to visit half of all households in the tract (a 50 percent sample) that received the ACS form to collect information directly from residents. For tracts with response rates between 35 and 50 percent, two in five ACS households (a 40 percent sample) will receive a personal visit. For tracts with response rates of 50 percent or higher, survey takers visit one in three households (a 33 percent sample).

In census tracts that historically produce response rates of 60 percent or higher, the Census Bureau reduces the initial ACS sample size, so that slightly fewer addresses will receive the questionnaire and be included in the monthly sample overall. Housing units with "unmailable addresses," such as addresses with incomplete information, do not receive a questionnaire by mail. Survey takers are sent to gather data from two-thirds of these homes.

proportion of unresponsive homes in areas with lower mail response rates. (See Box 3 for a description of the new sampling plan.[20])

By varying the proportion of households that are interviewed in-person, based on the proportion that respond initially by mail or telephone, the Census Bureau hopes to make the accuracy of all data more consistent from place to place. The modified sampling plan is an important step towards ensuring that data for all population subgroups, especially those that suffer from disproportionate undercounts in the census, are of similar quality.

Table 5
Percent of Respondents for Each Response Mode, by Race and Ethnicity, in the Census 2000 Supplementary Survey[1]

Predominant[2] race or ethnicity in census tract	Percent mail	Percent telephone	Percent in-person
White	60.5	7.3	28.1
African American or Black	34.9	8.9	48.6
American Indian or Alaska Native	16.6	2.6	69.9
Asian	58.6	4.1	32.5
Hispanic	34.2	8.3	53.3

1 The rows for each race and ethnic category do not add to 100 percent because not all homes responded to the ACS.
2 Defined as 75 percent or more of the population reporting the same race or ethnicity.
Source: U.S. Census Bureau, "Meeting 21st Century Demographic Data Needs: Implementing the American Community Survey, Report 2: Demonstrating Survey Quality," (www.census.gov, accessed Jan. 5, 2006).

Non-Sampling Error

Non-sampling error can be divided into two basic categories: nonresponse, and other mistakes in gathering or processing the data.

Nonresponse is the most well-known source of non-sampling error. There are several reasons why the Census Bureau might be unable to collect information from some homes in the ACS. In ACS tests, refusal to respond to the survey accounted for the largest proportion of non-interviews, followed by failure to find anyone at home during the entire interview period.[21] High nonresponse can contribute to inaccuracy if the remaining homes and people are no longer representative of the larger population. The Census Bureau must compensate for different rates of response between race and ethnic groups, so that ACS data more accurately reflect the nation's true characteristics.

Nonresponse, and the related problem of people who respond more than once, contribute to *coverage* problems in the ACS. *Coverage* describes the extent to which the survey accurately reflects the size and characteristics of the population being measured. If the ACS misses people or housing units, the result is *undercoverage.* If the survey includes some people or housing units twice, the result is *overcoverage.* If the survey misses or double counts a larger proportion of some demographic groups such as African Americans and Hispanics, than others such as whites, it results in *differential coverage.*

Nonresponse is one significant reason for undercoverage in the ACS. Nonresponse can affect either an entire housing unit selected to be part of the ACS sample (unit nonresponse) or specific

questions on the survey form (item nonresponse). (See Box 4 for explanation). The ACS like the census also can miss people because the person filling out the questionnaire fails to include someone who should have been counted as a member of that household, or because an entire housing unit that should have been in the sample is somehow overlooked, for example, because of an incorrect address. The Census Bureau has evaluated coverage in the ACS by comparing results from the Supplementary Survey and 36 test sites with 2000 Census data.

Box 4
Types of Nonresponse in the ACS and Census

Unit nonresponse occurs when no information is collected about a housing unit in the sample. Interviewers might be unable to locate a housing unit. Residents might refuse to answer the survey, or they might be unable to participate because of insurmountable language barriers. Sometimes interviewers fail to find anyone at home, or they determine that the residents are temporarily absent, during the entire nonresponse follow-up period. Unique circumstances such as impassable roads after a hurricane also affect the Census Bureau's ability to obtain responses from an address and contribute to nonresponse. Failure to collect enough information from an occupied housing unit is also considered nonresponse.

Item nonresponse occurs when people do not answer all of the questions on the survey form or provide invalid answers (for example, if a two-year old is also described as a high school graduate). Respondents might not understand a question or might think a question does not apply to them. Some people refuse to answer questions because they believe the government does not need to know that information.

When all attempts to gather missing information or correct invalid answers from the respondents fail, the Census Bureau imputes data. Imputation is a statistical procedure that involves using data collected from responsive households to fill in the blanks for nearby housing units that do not respond or do not answer all questions on the form. Imputation rates for the ACS are substantially lower for all basic population items and many other socio-economic characteristics than comparable rates for the 2000 Census because of more extensive follow-up procedures and the superior training of ACS field staff.

Other primary sources of non-sampling error include measurement error and processing error. Measurement errors generally relate to questionnaire design or data collection methods. For example, the wording or order of questions can affect responses.[22] Processing errors occur during the data capture phase of a census or survey. Mistakes in processing the data sometimes are not identified until after the results are published. For example, in the 2000 Census, data on specific Hispanic origin subgroups were later shown to be inaccurate, in part, because the data capture machines misinterpreted stray marks on the questionnaires.[23]

Responses also can vary depending on how interviewers probe for information. This is especially problematic when race and ethnic data are collected during a personal interview, instead of self-reported on a questionnaire returned by mail.

As discussed earlier, minority groups have higher nonresponse rates than whites during the mail phase of the ACS. Problems associated with lower mail response rates in the ACS are offset, to some degree, by extensive follow-up procedures for interviewing unresponsive households. ACS interviewers are more experienced than census enumerators. Census takers are temporary workers who receive only a few days of training and who might work for only several weeks during the nonresponse follow-up phase of the census. ACS interviewers, on the other hand, are full-time, permanent employees who receive extensive training in survey techniques and often work on other Census Bureau surveys. The greater skill of ACS interviewers has helped to contain the overall level of nonresponse in the survey and, ultimately, to reduce undercounting caused by failure to respond.

Reasons for Undercoverage

While *nonresponse* describes the failure of households in the ACS sample to respond to the survey, *undercoverage* refers to the survey's failure to account for all housing units and people in the sample. Some causes of undercoverage, such as difficulty enumerating everyone in large households, also are known to affect coverage in the census (see Box 5). Other potential causes, such as language barriers, might affect the ACS more than the census because of design differences. For example, the ACS questionnaire is available only in English and Spanish, while the 2000 Census questionnaire was also available in Chinese, Korean, Tagalog, and Vietnamese, and questionnaire assistance guides were available in 36 languages. Households with language barriers might fail to answer the ACS entirely or might not include everyone who should be reported.

Box 5
Counting Large Households in the ACS

The ACS questionnaire provides space to report information for up to five people in a household. Households can report the names of additional residents (up to 12) on a continuation roster. By comparison, the 2000 census questionnaire provided space to report complete information for up to six people in a household. ACS staff contact all households that list six or more persons by telephone, to collect detailed data for the additional residents. In the 2003 ACS, the Census Bureau successfully collected detailed data for roughly three-quarters of persons listed on continuation rosters.

As in the decennial census, large households appear to be a potential source of undercoverage of children in the ACS. A recent analysis of persons listed on continuation rosters in the ACS found that three-quarters were children. However, the instruments used to conduct computer-assisted telephone and personal interviews do not have a five-person limit, allowing for the full range of data to be collected on all household residents. Because a greater proportion of ACS households, compared to the census, are interviewed by telephone and in person the risk of undercoverage of children in large households is reduced.

APPENDIX B: ACCOUNTING FOR MISSING PEOPLE AND HOUSEHOLDS

As with any survey, non-sampling errors diminish the quality of ACS data, although there are no direct measures of how these mistakes affect the final data. Nonresponse and undercoverage, in particular, will result in ACS data that are incomplete. Missing responses to specific questions (item nonresponse) are filled in by imputation procedures. But the Census Bureau must also apply a complex set of statistical weights to compensate for other types of errors before final estimates are derived.[24] Without weighting to account for various factors that lead to nonresponse and undercoverage, the ACS estimates published every year would not accurately reflect the characteristics of the nation's overall population.

ACS weighting is designed to compensate for two broad sets of circumstances that might contribute to errors and reduce the accuracy of the data. The first set of weights, applied through a stratification procedure, compensates for the collection of information from only a sample of households (e.g. sampling error). If data are collected from four out of ten households, then those four households must represent the remaining six in the final analysis, and weights are applied to ensure proportional representation of key demographic and housing characteristics, as determined by the census. The second set of weights relates to conditions surrounding the data collection. They include the methodology used to gather the information or failure of the survey to collect data from all housing units in the sample and individuals within households (e.g. non-sampling error).

Factors considered in the weighting plan, some of which are applied to housing unit counts and some to individual counts, include:

- Oversampling of small governmental units, such as small towns and counties;
- Sampling of unresponsive households during Computer Assisted Personal Interviewing (CAPI) phase;
- Monthly variations in the percentage of the population responding to the survey through different modes (e.g. mail, telephone, in-person interview), and bias associated with response by different modes;
- Noninterviews;
- Differences in nonresponse between individuals and households; and
- Undercoverage of households and persons.[25]

The survey's ability to accurately account for all individuals and housing units in the ACS sample also bears directly on the quality of the demographic and socio-economic information the survey produces. If the ACS fails to capture an adequate sample of the African American population, for instance, then the characteristics about African Americans will also be inaccurate.

The estimation procedures for the ACS are technically complex and are described in detail in a number of Census Bureau reports.[26] All characteristics data related to persons in the ACS, such as language spoken at home, educational attainment, income, and disability, receive the same weights as the persons associated with those data. All characteristics related to housing units, such as type of fuel used, number of rooms, and year built, receive the same weights as the

housing units associated with those data. The average data user should simply understand that an appropriate weighting plan is an important factor in producing accurate and reliable ACS data.

In addition to applying statistical weights to account for sampling and non-sampling error, the Census Bureau uses population and housing unit controls to account for coverage error (undercoverage and overcoverage). In effect, final ACS population estimates are adjusted to match, proportionately, the sex, age, race, and Hispanic origin characteristics reflected in intercensal population estimates; final ACS housing unit counts are controlled to independent estimates of total housing.

Endnotes

[1] Housing units include both occupied units, which are also described as households, and unoccupied, or vacant, units.

[2] A census tract is a geographic unit created by the Census Bureau that includes roughly 1,700 housing units or 4,000 people. A census block is the smallest unit of geography for which the Census Bureau reports the number of people. Census blocks are put together to form census tracts.

[3] The published response rates are "weighted," meaning they are adjusted to compensate for the fact that information is collected from a sample of households.

[4] Daniel L. Cork, Michael L. Cohen, and Benjamin F. King, eds., "Reengineering the 2010 Census: Risks and Challenges," (Washington, DC: National Academies Press, 2004): 115.

[5] "Meeting 21st Century Demographic Data Needs – Implementing the American Community Survey, Report 7: Comparing Quality Measures: The American Community Survey's Three-Year Averages and Census 2000's Long Form Sample Estimates" (Washington, DC: U.S. Census Bureau, June 2004): 20.

[6] Controlling the ACS estimates to intercensal population numbers helps reduce, but does not eliminate, the effects of coverage problems on the accuracy of the final ACS data, because the estimates themselves might not be accurate.

[7] The ACS coverage rates are weighted (similar to the response rates) to account for sampling, but do not reflect final adjustments to independent population controls.

[8] Undercoverage of children was reduced but not eliminated after applying population controls.

[9] Family households refer to all households with two or more related persons. Married-couple households are a subset of family households.

[10] U.S. Census Bureau, "Meeting 21st Century Demographic Data Needs- Report 7," 21-27.

[11] All differences noted in this report are statistically significant.

[12] "Technical Summary of A.C.E. Revision II for the Committee on National Statistics" (Washington, DC: U.S. Census Bureau, March 2003).

[13] "Meeting 21st Century Demographic Data Needs – Implementing the American Community Survey, Report 4: Comparing General Demographic and Housing Characteristics With Census 2000" (Washington, DC: U.S. Census Bureau, May 2004): 17-18.

[14] Constance F. Citro, Daniel L. Cork, and Janet L. Norwood, eds., *The 2000 Census: Counting Under Adversity*, (Washington, DC: National Academies Press, 2004): 146.

[15] Citro, Cork, and Norwood, *The 2000 Census*, 139-43.

[16] Elizabeth Martin, "The Effects of Questionnaire Design on Reporting of Detailed Hispanic Origin in Census 2000 Mail Questionnaires," *Public Opinion Quarterly*, 66, no. 4 (2002):582-593.

[17] The ACS questionnaire can be used to provide information for up to five household members. If there are more, the respondent is instructed to list the names of additional people living in the household on a "continuation roster" so that ACS staff can collect information about these individuals through follow-up interviews.

[18] The error is expressed by showing not only the derived estimate, but also a range of estimates within which it is 90 percent likely that the "true" estimate falls. The range is known as a confidence interval.

[19] The ACS will also survey 2.5 percent of the group quarters population and about 36,000 addresses in Puerto Rico every year.

[20] Although the CAPI sample is drawn for each county, response rates on which the CAPI sampling rate is based are determined by census tract.

[21] U.S. Census, "Using the Data: Quality Measures," accessed online at www.census.gov, on April 5, 2006.

[22] Research has shown that responses to the race and Hispanic origin questions are particularly susceptible to variations in wording. In the 1980 and 1990 censuses, people were asked to check one of five races. The next question asked if they were Hispanic, because the federal government considers Hispanic an ethnicity, not a race. But many Hispanics failed to answer the race question, presumably because they did not identify with any of the categories. The order of the race and Hispanic origin questions was reversed in the 2000 census, resulting in a higher rate of response for both questions.

[23] U.S. Census Bureau, "Meeting 21st Century Demographic Data Needs - Implementing the American Community Survey, Report 4: Comparing General Demographic and Housing Characteristics with Census 2000" (Washington, DC: U.S. Census Bureau, May 2004): 30.

[24] Committee on National Statistics, National Research Council "The American Community Survey: Summary of a Workshop" (Washington, DC: National Academy Press, 2001).

[25] Committee on National Statistics, "The American Community Survey," 25.

[26] U.S. Census Bureau, "Accuracy of the Data (2004)," accessed online at www.census.gov, on April5, 2006.

SUGGESTED RESOURCES

Cork, Daniel L., Michael L. Cohen, and Benjamin F. King, eds., "Planning the 2010 Census: Second Interim Report," report of the Panel on Research on Future Census Methods, National Research Council (Washington, DC: National Academies Press, 2003).

Lowenthal, Terri Ann, "Counting Kids in Census 2000: Results and Challenges" (Baltimore, MD: The Annie E. Casey Foundation, 2005).

U.S. Census Bureau, "American Community Survey Operations Plan, Release 1" (Washington, DC, March 2003).

U.S. Census Bureau, "Comparing Characteristics from the American Community Survey and Census 2000 – Methodology," presented to the Census Advisory Committee of Professional Association Meetings, April 10-11, 2003.

U.S. Census Bureau, "Meeting 21st Century Demographic Data Needs – Implementing the American Community Survey, Report 2: Demonstrating Survey Quality" (Washington, DC, May 2002).

U.S. Census Bureau, "Meeting 21st Century Demographic Data Needs – Implementing the American Community Survey, Report 8: Comparison of the American Community Survey Three-Year Averages and the Census Sample for a Sample of Counties and Tracts" (Washington, DC, June 2004).

U.S. Government Accountability Office, "American Community Survey: Key Unresolved Issues," GAO-05-82 (Washington, DC: Government Printing Office, October 2004).

ACKNOWLEGEMENTS

The author thanks Constance Citro, Mark Mather, William O'Hare, Edward Spar, and ACS staff at the U.S. Census Bureau for reviewing this report.

Population Reference Bureau
1875 Connecticut Avenue, NW, Suite 520
Washington, DC 20009
www.prb.org

The Annie E. Casey Foundation
701 Saint Paul Street
Baltimore, MD 21202
www.aecf.org