RESEARCH REPORT SERIES
(Statistics #2012-03)

Interpretation and Use of American Community Survey Multiyear Estimates

Michael Beaghen
Tucker McElroy
Lynn Weidman
Mark Asiala
Alfredo Navarro

Center for Statistical Research & Methodology
Research and Methodology Directorate
U.S. Census Bureau
Washington, D.C. 20233

Report Issued: April 18, 2012

Interpretation and Use of American Community Survey Multiyear Estimates

Michael Beaghen, Tucker McElroy, Lynn Weidman, Mark Asiala, and Alfredo Navarro
U.S. Census Bureau

Abstract

The American Community Survey (ACS) was designed to produce annually updated estimates for detailed demographic, socioeconomic, and housing topics that were formerly available only once-a-decade from the census long-form sample. The first one-year estimates from the full implementation ACS were released in August 2006 for areas with populations of 65,000 or more. The ACS released the first set of multiyear estimates (MYEs) in December 2008, consisting of three-year period estimates for all areas with populations of 20,000 or more, and in December 2010 the ACS released the first five-year period estimates for all standard tabulation areas. The introduction of MYEs has received a great deal of attention among statisticians and the general American public, and concerns over issues of statistical interpretation and usability, particularly the choice between one-, three-, and five-year estimates, have arisen. This paper addresses these concerns, summarizing recently published literature and internal Census Bureau documents.

Key Words: Estimation, Time Series, Trends, Usability.

Disclaimer This paper is released to inform interested parties of ongoing research and to encourage discussion of work in progress. Any views expressed are those of the authors and not necessarily those of the U.S. Census Bureau.

1 Introduction

The U.S. Census Bureau (USCB) began implementation of the American Community Survey (ACS) at its full sample size in 2005, releasing the 2005 one-year estimates in 2006. The ACS currently produces three types of period estimates annually: one-year estimates based on one year of collected data, three-year estimates based on three years of collected data, and five-year estimates based on five years of collected data. In 2008 the USCB published its first three-year estimates based on data collected from 2005 through 2007, and in 2010 it published its first five-year estimates based on data collected from 2005 through 2009. The ACS five-year estimates are calculated for subject and geographic detail comparable to the USCB's former decennial long form, but published annually. This paper summarizes some key ACS concepts, focusing on what will be novel to users of decennial sample data and ACS one-year data, namely, the choice between one-, three-, and five-year estimates, and the interpretation of estimates over time.

The ACS is an extremely important source of economic and demographic information about the American public. Intense interest – among external researchers, the U.S. Congress, and the American society at large – has focused upon the survey since its inception, due to a number of noteworthy factors: the ACS provides detailed information for small areas (e.g., school districts); the ACS provides information on economic, housing, demographic, and social characteristics with over 400 variables; the ACS has an annual publication schedule, thereby offering a more timely snapshot of the American society than previously available through the Census long form; and the ACS quantifies statistical uncertainty explicitly. So promising is the ACS, that a National Academies of Sciences panel was convened on small area estimation methodological issues (Committee on National Statistics, 2007). Already – due to the above-cited reasons – published ACS estimates are being utilized to assist other surveys and estimation programs, for example the Small Area Income and Poverty Estimates program – see Bell et al. (2007), and the discussion later in this paper.

The multiyear estimates (MYEs) defined below provide new opportunities and challenges for those who have in the past used decennial long form data or the ACS one-year estimates. The purpose of this paper is to help inform data users of some important issues of statistical interpretation of MYEs. Several papers and documents have already been written to address this subject. From within the USCB the main reference on MYE usability is Beaghen and Weidman (2008); there are also a series of user handbooks targeted to several data user communities, such as states, local governments, and rural areas. Thinking on this topic from outside the USCB is diverse and includes work by the Transportation Research Board (2008) and the New York City Department of Planning (Salvo and Lobo, 2009). For additional background on ACS data products and operational considerations see Torrieri (2007).

The paper is structured as follows. Section 2 provides a general background on the ACS, while Section 3 provides more detailed background on the ACS one-year and MYE methodology. Section 4 delves into the key statistical topics of MYE usability – the choice between using the one-, three-, and five-year estimates – and the analysis of estimates over time. The calculations are based on publicly available figures, computed using SAS. Section 5 puts some of these issues in context by giving an example based on real data of the challenges of using multiyear ACS data for decision making in the context of data changing over time. Section 6 presents a simple time series approach for comparing MYEs of different period lengths across geographies (the calculations were performed in Excel, and graphs were produced with R). These two illustrations are complemented in the final section by a more extensive discussion of analytical studies involving ACS data.

2 Background

The ACS and the Puerto Rico Community Survey (PRCS) produce estimates updated annually for detailed demographic, socioeconomic, and housing topics that were formerly available only once every ten years from the census long form sample. This is accomplished via data collection from annual samples of (a) about 3,000,000 housing unit (HU) addresses in the United States and 36,000 HU addresses in Puerto Rico, and (b) 2.5% of persons living in groups quarters (GQ) accommodations in both areas.

The annual HU sample is equally distributed across the 12 months of the year by assigning each HU address to a month in which it is mailed a questionnaire (a small percentage cannot be mailed to) which is to be completed by members of the HU and returned by mail during that or the two subsequent months. For the HUs that have not returned a questionnaire by the end of the first month and that have an available phone number, an attempt is made to interview them by computer-assisted telephone interviewing during the second month. A sub-sample of those who have not been interviewed by the end of the second month (including the unmailable cases) is selected for interview by computer-assisted personal interviewing in the third month. Each GQ sample person is assigned a month for data collection by personal interview only. In most GQ facilities six weeks are allowed for completion of these interviews.

An important property of ACS (from this point we refer to the combination of the ACS and the PRCS as the ACS) estimates is that they do not represent a single point in time but an average of the characteristics of a geography over a one-year, three-year, or five-year period, so they are referred to as period estimates. Data collected during the 60 months of five calendar years are combined together to produce estimates for the same levels of geography as did the Census 2000 long form. In addition, three-year and one-year estimates are produced for geographies containing populations of at least 20,000 and 65,000 people, respectively. These population thresholds are due to the smaller sample sizes available over these shorter periods and the resulting loss in precision of their period estimates compared to the five-year estimates (of course, a user willing to build a model with covariates may be able to obtain suitably precise single year estimates). Each year, new one-year estimates are produced, while the three- and five-year estimates produced in the previous year are updated by replacing their oldest 12 months of data with data collected in the most recent year.

The five-year ACS estimates differ from the long-form estimates with respect to sampling error and data quality. ACS estimates of counts and proportions have coefficients of variation that are roughly 1.5 to 2 times larger than their census sample counterparts because (a) the ACS total national HU sample size over five years is less than 2/3 of the census sample size and (b) the ACS sub-samples for the personal visit follow-up for HU nonresponse, whereas all such cases were followed up in the census. On the other hand, the quality of the ACS is improved through the use of permanent telephone and personal interviewing staffs who benefit from continual training and interviewing experience; in comparison, the census long from interviewers were employed on a temporary basis.

A major conceptual difference between the two surveys is that an ACS estimate is for an average over a year, three years, or five years, while a census estimate represented a period of time around April first of a single year. In order for ACS to capture an average of a characteristic that can change over time, the persons to be included as residents of a HU for data collection are generally only those staying there for longer than two months. This is referred to as the "current residence" rule. In contrast, the census identifies the residents of a housing unit based on the concept of "usual residence as of April 1" or where people stay "most of the time."

3 MYE Estimation

Both the one-year and the multiyear estimates are period estimates and their estimation processes are very similar. One can view the multiyear estimation as an extension of the one-year estimation, hence we present the one-year estimation first and then discuss the generalizations made for multiyear estimation. Additional references on the construction of MYEs – from which much of this material is drawn – include: Fay (2007), Starsinic and Tersine (2007), and Tersine and Asiala (2007).

The weighting for the ACS produces two sets of weights, HU and person weights, with person weighting for the HU and the GQ populations done separately. The housing unit weights are used to produce tabulations of housing unit, household, and family characteristics. The person weights are used to produce tabulations of person characteristics.

3.1 One-Year Estimation

The one-year estimation has three primary steps: calculation of the baseweights, adjusting for non-response, and the application of controls.

Baseweights: The baseweights are defined as the inverse of the sampling probabilities. For HUs these weights are the inverse of one of seven different rates used in the ACS sampling operation. For group quarters persons, these weights begin with the inverse of the first stage sampling probabilities (in most states equal to 40) and are then adjusted for the second stage field sampling probabilities that are calculated at the time of interviewing.

Non-response adjustment: In other surveys and censuses, characteristics that have been shown to be related to housing unit response include census tract, building type (single- versus multi-unit structure), and month of data collection (Weidman et al., 1995). This full cross-classification, however, produces too many cells (more than one million) for the sample to support so instead the non-response adjustment is conducted in two steps. The first step calculates a ratio estimate within the cross-classification of census tract and building type in each weighting area. The second step calculates a ratio estimate within the cross-classification of month of data collection and building type. In doing this two step process, information from all three characteristics is used in the adjustment in a manner that the sample can support. The GQ person noninterview adjustment is performed for county by GQ type, if possible, or after combing types and/or counties to meet collapsing criteria.

Application of controls: The nonresponse adjusted weights for HUs are controlled to a set of independent HU estimates produced by the U.S. Census Bureau's Population Estimates Program (PEP). The GQ person weights are controlled to independent GQ population estimates obtained from the PEP for state by major GQ type. The weighting for HU persons is done in one step where a three-dimensional raking is used to achieve certain data consistencies and to ratio-adjust the household person weights to the controls derived from independent total HU resident population estimates obtained from the PEP. As a final step in the HU weighting, the HU weights for all occupied HUs are set equal to the person weight of the householder. This corrects for differential coverage of households and produces consistent estimates of occupied housing units, households and householders from the survey.

3.2 Multiyear Estimation

The multiyear estimation methodology involves reweighting the data and making certain adjustments to geography and monetary values. To reweight the data, we pool all of the sample over the multiyear period and use the one-year weighting methodology with some changes. These weights are then used to produce the MYEs; the process is discussed in greater detail below.

Pooling of the data: All sample addresses over the multiyear period are pooled together into one file. The one-year base weights are adjusted by the reciprocal of the number of years in the period so that each year contributes its proportional share to the multiyear estimate. Further, for the non-response adjustments all responses in the same calendar month are pooled across the years in the multiyear period. For example, for the 2005-2007 MYE, the January 2005, January 2006, and January 2007 responses are pooled.

Geography: All sample addresses in the period are put into the common geography of the final year of the period. Thus all addresses that are considered to be inside the boundaries of a place in the final year of the period will be tabulated for that place regardless if they were considered to be inside the boundaries for that place at the time of interview.

Derivation of controls: Since the MYEs are an estimate for the period, their controls are the average of the annual independent population estimates over that period. Thus the MYE controls are not those for a particular year, e.g., the 2005-2007 controls are the average of the 2005, 2006, and 2007 independent population estimates. Also it is worth noting that the Population Estimates Program updates its independent estimates every year, and the MYE weighting uses the most current estimates.

Adjustments to monetary values: All income and dollar value variables are inflation adjusted to the last year in the MYE period. For example, variables such as income, the value of housing units, and monthly owner costs are inflation adjusted.

4 Topics in MYE Usability

In preparation for the 2008 release of MYEs, the U.S. Census Bureau conducted the Multiyear Estimates Study (U.S. Census Bureau, 2007), using data collected in 34 test counties from 1999 to 2005. The study was a final evaluation of the procedure to be used to produce the MYEs and an analysis of their properties. The discussion in this section and the next is illustrated with data from the MYE Study, and is largely taken from the fuller treatment by Beaghen and Weidman (2008).

4.1 Relative Precision of the One-, Three-, and Five-Year Estimates for Totals

A three-year MYE is based on three times as much sample data as a one-year estimate, and a five-year MYE on five times as much. Since for uncontrolled estimates the standard error (SE) is a function of the sample size, there is an approximate relationship between the SEs of one-, three-, and five-year estimates of totals of persons, households, or housing units with certain characteristics. The SEs of the three-year estimates are about one over the square root of three, or about 58%, of the one-year estimates; and the five-year estimates are about one over the square root of five, or about 45%, of the one-year estimates. While this relationship holds approximately for estimates of totals, it does not hold up as well for estimates of proportions or means because they involve estimates in both the numerator and the denominator.

Published estimates of SEs are based upon simple sampling theory – see U.S. Census Bureau (2006, 2007). More sophisticated techniques for geographies with small populations or very small estimates, such as the use of covariates or hierarchical Bayesian modeling, are currently under study at the USCB.

Table 1: Percent Poverty by Family Type for Sevier County							
	2000-2004	2000-2004		2002-2004		2004	
	Families by Type	% in Poverty	SE	% in Poverty	SE	% in Poverty	SE
Fam	21,881	9.5	0.8	9.7	1.2	10.0	2.3
FamC	9,067	15.3	1.5	16.5	2.4	17.8	4.5
FamFemale	3,433	27.2	3.0	26.7	4.8	19.0	7.2
FamFemaleC	1,883	40.2	4.9	40.4	6.8	38.3	13.0

Table 1: Estimates of percentage in poverty, along with Standard Errors (SE) for Sevier County. Fam refers to all families; FamC refers to all familiar with children under 18 years of age; FamFemale refers to families with a female householder and no husband; FamFemaleC refers to families with a female householder and no husband, and with children under 18 years of age.

4.2 Precision of Estimates for Subpopulations

ACS data users need to be cautious when working with estimates for subpopulations. An estimate can be based on a larger area such as a county, yet if it applies to a smaller subpopulation it is the size of the subpopulation that determines how large the sample is for that estimate. For example, consider Sevier County, Tennessee, which had an estimated population of 77,270 in 2004 according to the PEP. This total is larger than the U.S. Census Bureau's 65,000 cutoff for publishing one-year estimates for geographic areas. However, some subpopulations will be much smaller than 65,000. In Table 1 we see that there are an estimated 21,881 families in Sevier County based on the 2000–2004 MYE; further, the number of families with a female householder, no husband, with related children less than 18 years, has an estimate of only 1,883. Not surprisingly, the SE for the 2004 one-year estimate of the poverty rate for this subpopulation is large (13%). In this example the five-year estimate has a SE of 4.9%, and the three-year estimate has a SE of 6.8%. Under the assumptions that the estimates of the SE are fairly accurate, and that one-year estimates and MYEs are estimating the same population quantities, it is apparent that for such small subpopulations, users obtain more precision using the MYEs. In short, our recommendation is that *MYEs are typically preferable to one-year estimates for examining estimates based on small subpopulations,* under the assumptions alluded to above.

4.3 Currency Versus Precision

When considering the choice among the one-year, three-year, and five-year estimates, or between the three-year and five-year estimates, the central statistical trade-off is between currency and precision. *MYEs yield smaller SEs but use less current data.* Conceptually, it is preferable to use a shorter period estimate as it uses data more relevant to what is happening currently. However, if that estimate is not precise enough to answer a data user's questions, this currency must be traded for the additional precision of a multiyear estimate.

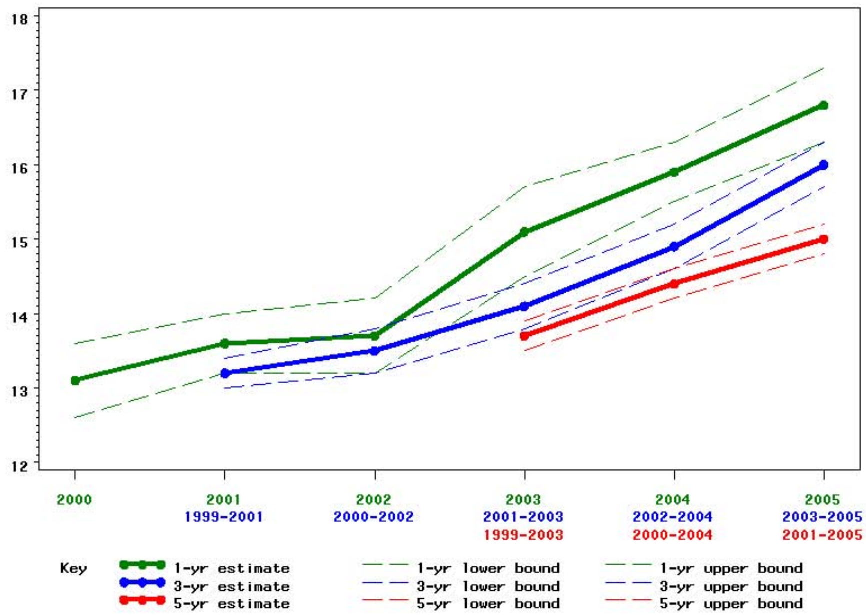

Figure 1: Percent Spanish speakers at home for population five years and older – Lake County, IL.

The following example illustrates the trade-off between currency and precision. Figure 1 shows the one-year, three-year, and five-year estimates of percent Spanish spoken at home in Lake County, Illinois, from 2000 through 2005, based on the MYE Study. Lake County had a population of 702,682 in 2005 according to the PEP. On the horizontal axis is the estimation period. The three period estimates are aligned by the last year of the estimation period so as to compare those estimates that are released at the same time and thus represent the choice of estimates data users will be presented with. For example, the 2005, the 2003–2005, and the 2001–2005 estimates are placed in the same horizontal position. The solid lines connecting the estimates are to help visualize the change over time, while the dashed lines are upper and lower 90% confidence interval bounds.

In this example we see how the lack of currency in MYEs can be quite apparent when there is a strong linear trend over time. As measured by one-year estimates, the percent Spanish speakers increases from 13.1% in 2000 to 16.8% in 2005. In Figure 1 we see the lag clearly for both the three- and five-year estimates, though the five-year estimates' lag is greater. For example, consider the three estimates whose last estimation year is 2005; the one-year estimate is at 16.8%, while the three-year estimate lags about a year behind at 15.9%, and the five-year about two years behind at 15.1%. For these data we suggest the one-year data for its greater currency. For situations with weaker linear trends relative to the standard error the MYEs would often be preferable.

4.4 Interpreting Change Over Time with Multiyear Estimates

Because the ACS releases one-, three-, and five-year estimates annually, ACS data users have opportunities to analyze estimates over time that they didn't have with decennial census estimates. When considering change over time, an important distinction is whether the time periods being compared overlap or not. By overlapping estimates we mean that the data collection periods of the two estimates share years in common and consequently also share the sample housing unit and person data collected in those years. Analyzing change by examining estimates of non-overlapping time periods is both mathematically simpler and more straightforward to interpret, and thus recommended to users. Since estimates for non-overlapping years are nearly independent, a very good approximation to the SE for a difference between non-overlapping estimates is just the square root of the sum of the variances of the two estimates. (Because of the way ACS selects its sample, there is a small correlation between estimates of non-overlapping years). That is, if x_1 and x_2 are two non-overlapping MYEs, or any two one-year estimates, then a close approximation to the standard error of the difference between x_1 and x_2 is given by the square root of the sum of the variances.

Data users may prefer overlapping estimates because they smooth differences over time. However, the estimates of differences between overlapping MYEs suffer from difficulty in interpretation, as *the difference between overlapping MYEs is driven by the difference between the non-overlapping years.* To illustrate this point, we can approximate a MYE by assuming that it is equal to the average of the one-year estimates in its period. For example, consider two overlapping five-year estimates $x_{5,5} = (y_1+y_2+y_3+y_4+y_5)/5$, and $x_{5,6} = (y_2+y_3+y_4+y_5+y_6)/5$, where y_t is the estimate for the t-th year; then the difference between these two estimates is $x_{5,5} - x_{5,6} = (y_1 - y_6)/5$. Note that the SE of the difference of overlapping MYEs could be better estimated with a knowledge of the correlations in sampling errors across years, but unfortunately this information is not available (i.e., it is not easily estimable).

It would be easy for a naïve data user to come to incorrect conclusions when directly comparing overlapping estimates. They might interpret the difference between consecutive MYEs as the difference between the most recent years in the estimation periods or between the middle years. For example, they might see the difference between the 2000–2004 and 2001–2005 estimates as being indicative of the difference between 2004 and 2005, or between the difference between 2002 and 2003, while it actually is based on the difference between 2000 and 2005.

5 Currency versus Precision: a Funding Allocation Example

This section presents an example, condensed from material in Beaghen and Weidman (2008), of how ACS estimates could be used to distribute state funding among counties. Its purpose is to demonstrate how a decision on whether to use one-, three-, or five-year estimates can affect results

Table 2: Persons 5 and older of Limited English Proficiency, by Size of County (2000-2004 MYE Study)			
Shortest MYE Period/ County Size	Number of Counties	Number of Persons of Limited English Proficiency	Proportion of Persons of Limited English Proficiency
Five-year	4	4,213	0.4%
Three-year	11	34,952	3.0%
One-year	19	1,197,032	96.6%
Total	34	1,236,197	100.0%

Table 2: This table provides a break-down of counties by their size, also providing data on English proficiency.

when characteristics are changing over time.

5.1 A State for Analysis

The example takes the 34 counties in the MYE Study and creates a "New State" of 9,813,462 million residents. Let's suppose that New State would like to strengthen English as a Second Language (ESL) programs. Table 2 shows the breakdown of these 34 counties into three groups defined by the shortest period estimates that are released for each, which is determined by population size. Table 2 also shows the number and proportion of persons of limited English proficiency as determined by the 2000-2004 ACS estimates, where we will define persons of limited English proficiency as persons 5 and older who speak a language other than English at home and speak English "less than very well."

5.2 Promoting English as a Second Language in New State

New State passes legislation providing $10,000,000 in funding to help persons of limited English proficiency by subsidizing English as a Second Language (ESL) classes. The law stipulates that New State should allocate the $10,000,000 by county proportional to the number of persons of limited English proficiency, as determined by ACS data. Table 3 shows the number of persons of limited English proficiency in those New State counties for which the ACS produces one-year estimates. These data, which show an increase over time in persons of limited English proficiency, motivated the state to pass a law subsidizing ESL classes.

5.3 Question and Challenge

The question facing New State is which ACS estimates to use to allocate funds among counties: one-year, three-year, or five-year estimates? New State would like to use the most current data, but the ACS doesn't provide one-year data for all counties. New State sees two obvious approaches:

Table 3: Number of Persons of Limited English Proficiency by Year in the 19 Counties with One-Year Estimates					
Year	2000	2001	2002	2003	2004
Number	1,027,959	1,059,002	1,111,733	1,153,066	1,197,032

Table 3: This provides a time series for limited English proficiency for the 19 larger counties.

Table 4: Distribution of Funds			
County Size	Current	Five-Year	% Difference
Under 20,000	\$34,080	\$36,883	7.60%
20,000-64,999	\$282,738	\$297,543	4.98%
65,000+	\$9,683,182	\$9,665,574	-0.18%

Table 4: The distribution of funds between groups of counties divided according to population. County Size refers to the population of the county, and the column marked Current refers to the most current estimates available for each county. The column marked Five-Year refers to the five-year MYE for the period 2001-2005.

- use the most current estimate available for each county;
- use five-year data for each county.

A drawback of the first approach is that the allocations to those counties not among the largest 19 are lower due to the inclusion of older data and the resulting lag effect (assuming the number of persons of limited English proficiency is increasing across years in the same way as in the largest 19). We see this in Table 4 – if we used the most current data for each county, then the allocation for the smallest counties would be \$34,080, instead of \$36,883 when only using five-year data, a difference of almost 8%. A drawback of the second approach is that it uses older data to allocate funds among the larger counties, which have the bulk of the persons of limited English proficiency.

5.4 A Hybrid Approach to Funding Allocation

An alternative approach to the two discussed is the following hybrid method of allocation.

- Form three groups of counties based on the most recent data available for each county; i.e., one-year counties, three-year counties and five-year counties.
- Divide up the funds among these three groups according to five-year estimates.
- Within each group allocate funds based on the most recent shortest period estimates available.

This approach would be fair to smaller counties with only five-year data, avoiding the losses pointed out in Table 4. However, it also uses the most recent data available for larger counties to distribute funds among themselves. To see this consider Table 5, which shows that using the

Table 5: Allocation of Funds among Large Counties via Two Schemes (Parentheses Indicate Negative Differences)				
County	By Most Recent One-Year (Hybrid)	By Five-Year MYE	Difference	Percent Difference
Pima	$668,643	$641,704	$26,939	4.03%
Jefferson	$1,946	$4,351	($2,405)	-123.59%
San Fran.	$1,413,122	$1,527,932	($114,810)	-8.12%
Tulare	$610,724	$633,159	($22,436)	-3.67%
Broward	$1,763,165	$1,706,211	$56,954	3.23%
Lake, IL	$594,760	$554,202	$40,559	6.82%
Miami	$30,038	$31,158	($1,120)	-3.73%
Calvert	$5,765	$6,470	($704)	-12.22%
Hampden	$311,188	$343,207	($32,019)	-10.29%
Madison	$9,455	$8,107	$1,349	14.26%
Flathead	$6,670	$4,220	$2,450	36.73%
Douglas	$198,603	$189,616	$8,987	4.53%
Bronx	$2,607,810	$2,532,690	$75,121	2.88%
Rockland	$300,239	$297,674	$2,564	0.85%
Franklin	$390,957	$376,947	$14,010	3.58%
Multnomah	$460,479	$497,979	($37,500)	-8.14%
Schuylkill	$9,205	$12,685	($3,480)	-37.81%
Sevier	$14,236	$7,459	$6,777	47.60%
Yakima	$268,570	$289,804	($21,234)	-7.91%
Total	$9,665,574	$9,665,574	$0	

most recent one-year data would allocate about 7% more to Lake County, IL, and 48% more to Sevier County than using the five-year data. To summarize, this example shows that which ACS estimates are used will affect the decisions that are made, especially in the context of data that are changing over time.

6 Interpretation of Trends

Unlike the decennial long-form that was conducted only once a decade, the ACS produces estimates annually, opening the door to time series analysis. The previous sections of this paper have outlined some of the challenges with interpreting MYEs. In this section a simple time series approach is outlined that serves as an example of the possibilities for new approaches that the ACS offers. A fuller description of the methodology is given in McElroy (2009).

6.1 Comparing Trends Across Geographies

For geographies with smaller populations, the one-year (and possibly the three-year) estimates will be unavailable. How then can comparisons be made to one-year estimates for other geographies? Since a five-year MYE is lagged, it is not valid to compare it directly to the one-year estimate for another geography, as spurious differences in the trends could be indicated. For example, if the five-year MYE for a small geography is smaller than the one-year estimate for a large geography, this discrepancy could simply be due to the fact that the five-year MYE is lagged two years. Perhaps the (unavailable) one-year estimate for the former geography is actually larger than the latter one-year estimate; in this case one might be misled into thinking that the levels are higher for the large geography, whereas the opposite is the case. Documentation of this phenomenon can be found in McElroy, Titova, and Nagaraja (2011).

One solution is to only compare MYEs of the same period length, e.g., compare a five-year MYE with another five-year MYE. However, these may not be the most timely estimates available. Another approach is to allow comparisons of MYEs of different period lengths, provided that the data (viewed chronologically as a time series) are first weighted appropriately, such that the linear aspects of the data are preserved. For example, a sequence of five-year MYEs in a straight line would be shifted *forward in time* by two years. This method is discussed in more detail in the next section.

6.2 Producing Comparable Trends

Trends can be computed from the various sequences of MYEs by taking weighted averages, such that these trends of different period lengths can be compared. Mathematically, the process is very simple and requires only a little notation. Suppose we have a sequence of MYEs of different period lengths for two different geographies, and we view these as two time series $x_{i,t}$ and $y_{j,t}$, where i, j index the period length ($i, j = 1, 3, 5$). Here the time index t corresponds to the final year in the rolling sample. For example, if we are considering a one-year estimate and a three-year MYE for the respective geographies, then $x_{1,2005}$ is the one-year estimate for year 2005 for the first geography, and $y_{3,2005}$ is the three-year MYE for the 2003–2005 period for the second geography. The procedure involves taking a temporal weighted average of each time series, where the weights only depend on whether it is a one-year, three-year, or five-year estimate. The formulas are given as follows, respectively for one-year ($z_{1,t}$), three-year ($z_{3,t}$), and five-year ($z_{5,t}$) estimates:

$$\begin{aligned}
z_{1,t} &= \frac{4}{15}x_{1,t} + \frac{1}{3}x_{1,t-1} + \frac{2}{5}x_{1,t-2} + \frac{1}{5}x_{1,t-3} \\
&+ \frac{1}{5}x_{1,t-4} - \frac{1}{15}x_{1,t-5} - \frac{2}{15}x_{1,t-6} - \frac{1}{5}x_{1,t-7} \\
z_{3,t} &= \frac{4}{5}x_{3,t} + \frac{1}{5}x_{3,t-1} + \frac{1}{5}x_{3,t-2} + \frac{1}{5}x_{3,t-3} \\
&+ \frac{1}{5}x_{3,t-4} - \frac{3}{5}x_{3,t-5} \\
z_{5,t} &= \frac{4}{3}x_{5,t} + \frac{1}{3}x_{5,t-1} + \frac{1}{3}x_{5,t-2} - x_{5,t-3}.
\end{aligned}$$

Similarly, one computes weighted averages for the $y_{j,t}$ series as well. Note that the one-year estimate weighted average requires seven past values, whereas the three-year and five-year MYE averages require five and three past values, respectively. The full derivations of these weights are given in McElroy (2009), but we give a very brief heuristic justification here. These weights have the property that if the underlying time series is a perfectly straight line, then the one-year estimate weighted average exactly replicates it, while the three-year MYE weighted average forecasts it ahead one year, and the five-year MYE weighted average forecasts ahead by two years. Now actual data will not be given as a perfectly straight line, but an underlying linear trend effect in economic or demographic data will be an important aspect to capture in any trend analysis. But even when the data are not close to linear, the weighted average device will make the trends comparable so long as the rolling sample that is used to generate MYEs is closely approximated by a simple moving average – the extent to which this hypothesis is true or false is quantified in McElroy, Titova, and Nagaraja (2011). In summary, the resulting MYE weighted averages will be more comparable in general, and will suffer only non-linear distortions since the underlying linear dynamics are appropriately handled.

This procedure is illustrated on the estimates of "Percent speaking Spanish at home for population five years and older" for Lake County, Illinois and Otero County, New Mexico. For Otero County – in the MYE Study – there are only three-year and five-year estimates available, so it is impossible to compare one-year estimates. Since only six, five, and three years of one-year, three-year, and five-year MYEs were available for Lake County, the data were augmented to length eight, six, and four respectively via simple forecast and backcast extension (see McElroy (2009) for details) – for the sake of creating an example (with similar extensions for Otero County)[1]. Application of the weights then yields the values 17.3%, 17.1%, and 17.0% for the one-year, three-year, and five-year MYEs respectively for Lake County, and 25.5% and 24.2% for the three-year and five-year MYEs for Otero County. The fact that these values are different (within each county) reflects the fact that MYEs are not an exact moving average of the same underlying variable. So

[1] Although one could augment the data from the MYE Study with more current published ACS MYEs, it is unclear whether this forms a meaningful time series for counties given that the sample designs are somewhat different.

valid comparisons would be: 17.1% to 25.5% (trend values for the three-year MYEs) and 17.0% to 24.2% (trend values for the five-year MYEs). Figure 2 displays the extended MYE data (plotted by final year in the rolling sample), with the weighted averages (trends) indicated by the dots in year 2006; these values essentially forecast the three-year and five-year MYEs one and two years ahead, respectively. Note that although a trend typically consists of a sequence of values, in this case we are only able to produce trends at one time point due to the shortness of the time series.

7 Additional Applications of ACS Data

Even before the ACS began its publication schedule, there was much interest in the anticipated data release among demographers and statisticians throughout the country. Meetings and conferences provided early forums for discussion between applied researchers and the ACS staff at USCB. We review a sample of the published literature, with the aim of illuminating the manifold applications of MYEs.

We first make a philosophical point. Deming (1953, 1975) made the distinction between enumerative and analytical uses of data, which in this context refers to simple enumeration (i.e., publication) of demographic and economic information, versus an implementable use of the data for public policy decisions. Whereas the ACS is enumerative in its nature, it can indeed be used to make decisions, such as the allocation of federal funding or the optimal flow of resources. Below we consider several applications, which tend to fall in the analytical category of data use.

A fairly extensive internal application of ACS data is its utilization in the Small Area Income and Poverty Estimates (SAIPE) program, described at length in Bell et al. (2007). The previous survey data source for obtaining these important estimates, which are mandated by Congress, was the Annual Social and Economic Supplement (ASEC) of the Current Population Survey (CPS). The ACS provides increased annual sample sizes over the ASEC, about 3 million compared to 100,000 housing units annually, which decreases the occurrence of zero estimates in the published data. Moreover, ACS single-year estimates from 3,141 counties are available for use, versus 1,100 counties from the ASEC, providing a more granular viewpoint of the country. While the ACS sample is spread evenly across all twelve months of each year, the ASEC is conducted in the month of March, which can impact certain variables. Because the ASEC involves phone calls and a personal visit – whereas the ACS utilizes mail, followed by phone, followed by a physical visit – a greater non-response is typically induced. Although there are some drawbacks to the use of the ACS (only 8 of the survey's questions are pertinent for the purposes of SAIPE, whereas all 50 questions of the CPS are relevant), it has been adopted as the main data source for this important federal program.

Another internal application involves the Department of Justice and fair-play in the electoral process. Section 203 of the Voting Rights Act requires that determinations of limited English proficiency and limited education be made, within specified small domains (race and ethnicity groups) and small areas (i.e., county or minor civil division). These determinations are made in order to determine which of these small domains must provide language assistance to voters. USCB research has provided a small area model-based estimate derived from ACS five-year MYE data – as well as 2010 Census data. The currency of the ACS and the availability of five-year MYEs for small areas is an important feature for this application.

Much of the USCB literature on the ACS is concerned with issues of data quality and usability: Bennett and Griffin (2002), Griffin (2002), Beaghen and Weidman (2008), and Schwartz (2009) are examples. Some of these papers are concerned with non-response among minorities, whereas others compare results from the ACS demonstration phase (U.S. Census Bureau, 2006) to the Census 2000 Long Form. Gage (2004) also makes this comparison for two counties in California.

Several studies by the New York City Department of Planning involving ACS data have been published on the topic of demographic shifts in New York City. Salvo and Lobo (2002, 2003a, 2003b) examine issues of non-response and data quality in the ACS at the Bronx test site, whereas Salvo and Lobo (2009) provide an overview of immigration patterns based on the ACS. This latter work provides an application of MYEs that we anticipate to become widespread in the future.

The five-year ACS estimates are the only MYEs available for many geographies, so a basic comprehension of the annual five-year MYEs is essential for data users. But many users will want to explore and understand the possibilities offered by the three-year and one-year estimates. While this paper aims at aiding both types of data user, it is more helpful to the latter type. A brief summary of when one would use and would not use MYEs is appropriate at this point: use one-year estimates for larger geographies and populations, and when currency is important; use MYEs for estimates of tracts and other smaller geographies, for estimates of smaller subpopulations of larger geographies, and for obtaining estimates with lower standard errors that are smoother over time.

We expect that as users become more familiar with the ACS data their practices and needs will grow, perhaps in ways that the authors cannot anticipate. We expect more statistical developments will suggest themselves as the needs of data users approach the limits of current practices and methodologies.

References

Beaghen, M., and Weidman, L. (2008), "Statistical Issues of Interpretation of the American Community Survey's One-, Three-, and Five-Year Period Estimates", U.S. Census Bureau, 2008 American Community Survey Research Memorandum Series #ACS08-R-4.

Bell, W., Basel, W., Cruse, C., Dalzell, L., Maples, J., O'Hara, B., and Powers, D. (2007), "Use of ACS Data to Produce SAIPE Model-Based Estimates of Poverty for Counties."

http://www.census.gov/did/www/saipe/publications/methods.html

Bennett, C. and Griffin, D. (2002) "Race and Hispanic origin data: a comparison of results from the Census 2000 Supplementary Survey and Census 2000." U.S. Census Bureau report.

Committee on National Statistics (2007), "Chapter 6: Weighting and Interpreting ACS Multiyear Estimates," in *Using the American Community Survey: Benefits and Challenges,* National Academies Press.

Deming, W. (1953) "On the distinction between enumerative and analytic surveys." *Journal of the American Statistical Association* **48**, 244–255.

Deming, W. (1975) "On probability as a basis for action." *The American Statistician* **29**, 146–152.

Fay, R. (2007) "Imbedding model-assisted estimation into ACS estimation." *Proceedings of the 2007 Joint Statistical Meetings on CD-ROM*, American Statistical Association, 2946–2953.

Gage, L. (2004) "Comparison of Census 2000 and American Community Survey 1999–2001 Estimates, San Francisco and Tulare Counties, California." U.S. Census Bureau report.

Griffin, D. (2002) "Measuring survey nonresponse by race and ethnicity." U.S. Census Bureau report.

McElroy, T. (2009) "Incompatibility of Trends in Multi-Year Estimates from the American Community Survey." *The Annals of Applied Statistics* **3**, 1493–1504.

McElroy, T., Titova, N., and Nagaraja, C. (2011) "Quantifying Discrepancies Between Diverse Multi-Year Estimates." U.S. Census Bureau report.

Salvo, J. and Lobo, A. (2002) "The American Community Survey: Quality of response by mode of data collection in the Bronx test site." *Proceedings of the 2002 Joint Statistical Meetings on CD-ROM*, American Statistical Association, 3033–3042.

Salvo, J. and Lobo, A. (2003a) "Evaluating continuous measurement: data quality in the Bronx test site of the American Community Survey." *Journal of Economic and Social Measurement* **28**, 263–277.

Salvo, J. and Lobo, A. (2003b) "The effect of neighborhood characteristics on nonresponse in the Bronx test site of the American Community Survey." *Proceedings of the 2003 Joint Statistical Meetings on CD-ROM*, American Statistical Association, 3624–3635.

Salvo, J. and Lobo, A. (2009) "New Immigrants to New York." *The Hon. Jerome and Helene Berg Public Policy Papers on Government Reform*, Issue No. 6.

Schwartz, M. (2007) "2007 American Community Survey: a comparison to selected housing and financial characteristics for the United States from the 2007 American Housing Survey." U.S. Census Bureau report.

Starsinic, M. and Tersine, A. (2007) "Analysis of variance estimates from American Community Survey multiyear estimates." *Proceedings of the 2007 Joint Statistical Meetings on CD-ROM*, American Statistical Association, 3011–3017.

Taeuber, C. (2006) "Chapter 9: Multi-year estimates," in *American Community Survey for Community Planning,* Trafford Publishing.

Tersine, A. and Asiala, M. (2007) "Methodology for the production of American Community Survey multiyear estimates." *Proceedings of the 2007 Joint Statistical Meetings on CD-ROM*, American Statistical Association, pp. 3018–3023.

Torrieri, N. (2007) "America is Changing, and So is the Census: The American Community Survey," *The American Statistician*, 61, 16–21.

Transportation Research Board (2008), "A Guidebook for Using American Community Survey Data for Transportation Planning," National Cooperative Highway Research Program Report #558.
http://onlinepubs.trb.org/onlinepubs/nchrp/nchrp_ rpt_ 588.pdf

U.S. Census Bureau (2006), "Technical Paper 67," *Design and Methodology, American Community Survey.* http://www.census.gov/acs/www/methodology/methodology_main

U.S. Census Bureau (2007), Multi-Year Estimates Study, "Accuracy of the Data."

Weidman, L., Alexander, C., Diffendal, G., and Love, S. (1995), "Estimation Issues for the American Community Survey," *Proceedings of the Survey Methods Research Section*, ASA, 596–601.

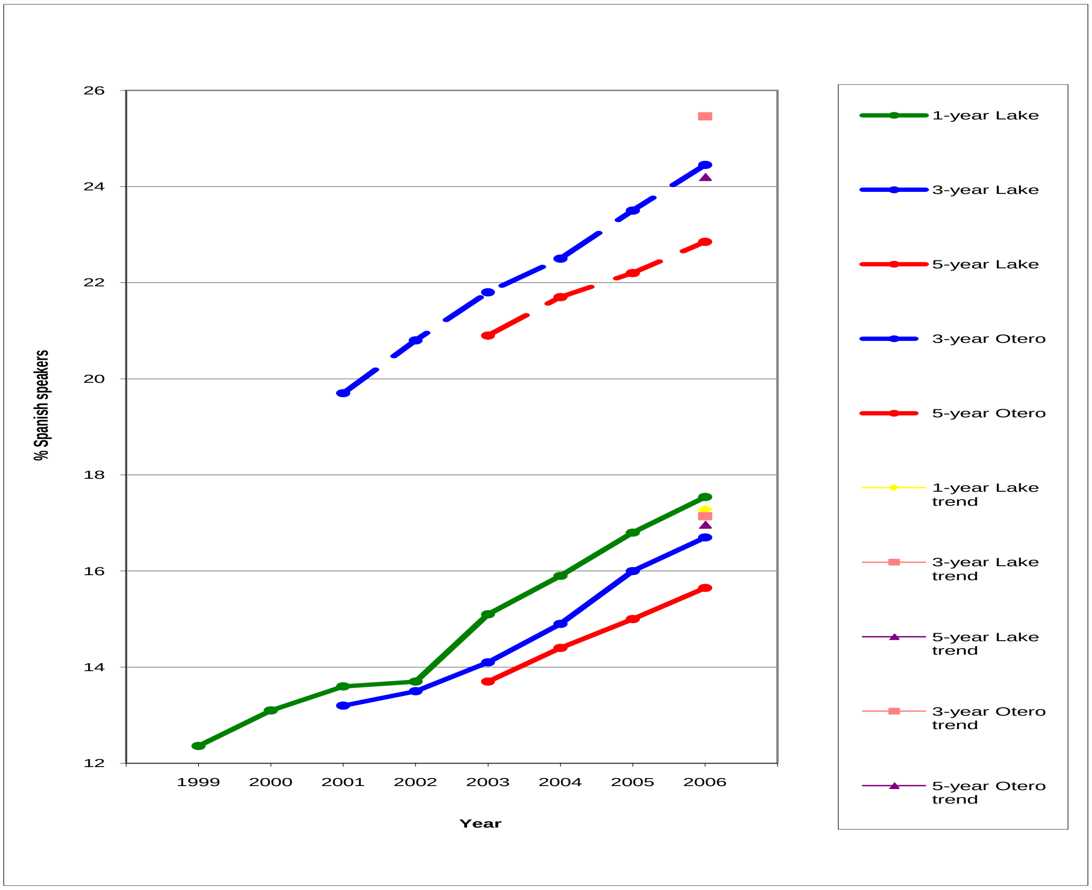

Figure 2: Various MYEs for "Percent speaking Spanish at home for population five years and older" in Lake County, Illinois (solid lines) and Otero County, New Mexico (dashed lines). The MYEs are plotted by final year in the rolling sample, and are forecast and backcast extended. The points at the right hand indicate trend values for each MYE, obtained by taking weighted averages – circles for Lake County and squares for Otero County.