Simulated Versus Actual SNAP Unit Composition in Survey Households in Two States

John L. Czajka
Karen Cunnyngham
Randy Rosso

Mathematica Policy Research
1100 1st Street NE, 12th Floor
Washington, DC 20002-4221

Proceedings of the 2015 Federal Committee on Statistical Methodology (FCSM) Research Conference

I. Background and Overview[1]

Studies linking state administrative records for the Supplemental Nutrition Assistance Program (SNAP) to household survey data have focused almost exclusively on estimating the underreporting of SNAP participation and benefits in the surveys and on investigating the characteristics of households and individuals associated with the misreporting of participation (see, for example, Bollinger and David 1997; Meyer and Goerge 2011; Harris 2014). Such studies have been few in number because of the infrequency with which state agencies that administer SNAP have shared their data with external entities for research purposes.

A recent exception to the general focus of these studies used SNAP records from Texas linked to survey data from the American Community Survey (ACS) to estimate and analyze access to SNAP at both the state and county levels (Newman and Scherpf 2013). A key component of this study involved the simulation of eligibility for SNAP benefits in order to estimate SNAP access among households explicitly targeted by SNAP—that is, those individuals and families who appeared to be eligible for SNAP benefits based on the income and other characteristics reported in the survey.

The research described in this paper focuses on an aspect of eligibility simulation that has challenged researchers for decades—namely, the grouping of household members into the units whose eligibility is jointly determined. The research expands upon previous efforts to link SNAP administrative records to household survey data by including records from the administrative data that do not match directly to individuals in the survey data. With these additions, the data allow a fuller assessment of SNAP unit formation. Furthermore, our research generates findings that extend beyond SNAP eligibility simulation to the potential use of SNAP and other administrative records to improve the efficiency and accuracy of decennial census operations. Section A provides important background for this research. Section B gives an overview of the research goals and an outline of this paper.

A. Background

Unlike household surveys (but like other administrative data), SNAP administrative records are compiled and maintained at the program unit level—that is, the level at which SNAP eligibility is assessed and benefits are determined. A SNAP "household," or unit, is defined as "individuals who share a residential unit and customarily purchase and prepare food together" (U.S. Department of Agriculture 2014). Individuals who live together but do not purchase and prepare food together generally may apply as separate SNAP units. However, spouses must apply together and parents must apply with their children under age 22 who reside with them—even if those children have spouses or children of their own. Minor children may be eligible as "child-only" units if their parents do not meet the U.S. residency requirement or are otherwise categorically ineligible for SNAP. Elderly individuals who are unable to purchase and prepare food because of a substantial and permanent disability and their spouses may apply

[1] This work was conducted under a project cosponsored by the Census Bureau's Center for Administrative Records Research and Applications (CARRA) and the Economic Research Service (ERS) of the U.S. Department of Agriculture. The views expressed here are the authors' and not necessarily the sponsors.

as a separate unit if the gross monthly income of remaining household members is below 165 percent of the federal poverty guideline.
When a household survey is used to simulate eligibility for SNAP, a key step is to identify the SNAP unit(s) within each sample household. Researchers customarily rely on reported family and partner relationships for this. Countable income and resources, if the latter is available on the survey or imputed, are then aggregated across the members of each unit and compared to SNAP income and resource screens to determine if a unit is eligible for SNAP.

Typically, household surveys that collect data on SNAP participation do not identify the members of the SNAP unit. Consequently, comparisons of aggregate characteristics between SNAP units in administrative data and survey households that report receiving SNAP benefits show large differences. Even when the surveys identify the members of the SNAP unit, differences persist. Part of the reason is that the approach to identifying SNAP participants in most surveys does not lend itself very well to identifying child beneficiaries when their parents are not participants themselves. In the Survey of Income and Program Participation (SIPP), for example, a respondent cannot report that he or she receives SNAP benefits on behalf of a child and also report not being covered by those benefits. As a result, child-only units are rarely reported as such. Either no child participants are reported or they are reported with the parent incorrectly included as a member of the unit as well.

SNAP quality control (QC) data show that nearly half (46.7 percent) of the average monthly SNAP units in 2009 consisted of a single person; about a third contained two or three persons; and just 3.2 percent had six or more members (Czajka et al. 2012).[2] In the SIPP for that year, one-person units (based on reported unit composition and average monthly participation) represented 35.1 percent of the total; two- and three-person units were 39.4 percent, and units with six or more persons were 5.8 percent of the total. Substituting the SIPP household size for reported SNAP unit composition produces substantially greater discrepancies. Only 22.6 percent of the SIPP households reporting SNAP benefits were one-person households, 39.5 percent had two or three members, and 10.8 percent had six or more members. Households that reported receiving SNAP benefits in the Current Population Survey Annual Social and Economic Supplement (CPS ASEC) and the ACS resembled SNAP households in the SIPP. For example, 22.9 percent of average monthly SNAP households in the CPS ASEC and 21.6 percent of the ACS households with any SNAP participation in the prior 12 months had a single person. Similarly, 40.6 percent of SNAP households in the CPS ASEC and 39.7 percent of those in the ACS had two or three members. Likewise, 9.4 percent of the CPS ASEC households and 11.0 percent of ACS households with SNAP benefits had six or more members.

B. Overview

The objective of the larger project from which this paper was derived was to use linked New York State and Colorado administrative and survey data to conduct within-household comparisons of administrative SNAP units with full households and simulated SNAP units in the ACS, the CPS ASEC, and the SIPP. This paper describes the linkage process for all three surveys but focuses on the ACS in its presentation of findings from the comparative analysis. For the ACS analysis, we used the simulated SNAP units developed by Decision Demographics based on specifications provided by the Economic Research Service (ERS), which cosponsored this research with the U.S. Census Bureau.

In contrast to previous studies of linked survey and SNAP administrative records, this analysis takes account of unmatched SNAP recipients in the administrative units that are matched to survey households. The unmatched recipients are in SNAP units from which at least one person was matched to a survey household member. Including the unmatched SNAP recipients enables us to calculate the size and demographic composition of the matched administrative units and to compare the incidence of unmatched persons in both the survey and administrative data. Unmatched recipients may reflect match failures (when there are also unmatched survey household members), but they may also represent household members who were not enumerated with the survey household.

[2] SNAP QC data are monthly samples of SNAP case files drawn by each state to assess the accuracy of eligibility and benefit determination. Original and corrected values of selected fields are entered into a database and transmitted to the Food and Nutrition Service (FNS), which calculates and publishes error rates and produces an annual database of QC sample cases containing information on unit characteristics as well as benefit amounts.

Chapter II discusses the data sources used in this research, explains how we constructed the analysis files, and provides counts of survey households that were matched to one or more administrative records for each pairing of a survey with a state administrative file. Chapter III reviews our findings from the linked ACS and New York SNAP data, while Chapter IV discusses our findings from the linked ACS and Colorado SNAP data. Finally, Chapter V summarizes our key findings and discusses their implications.

II. Data and Methods

This chapter provides an overview of the three household surveys and the New York and Colorado administrative data to which they were linked. We then describe the construction of the analysis files, concluding with a summary of the matched households by survey.

A. Overview of the Three Surveys

The ACS is a continuous survey, with independent monthly samples of households. Sample households are first contacted by mail and asked to complete and mail back their survey forms. On average, roughly half do so. A sample of the households that do not respond by mail is selected for intensive follow-up, starting with telephone calls and progressing if necessary to in-person visits. The data collected by these follow-up methods are typically obtained two to three months after the sample month. Because the ACS reference period is defined relative to the interview month, the survey data are organized by the month in which the data were collected, then assembled into calendar year files for processing and analytical use. Although the survey month is not identified on public use files, it is identified on the Census Bureau's internal files. We used the survey month in linking the ACS data to the SNAP administrative data, as explained below.

The CPS is a monthly survey to which the Census Bureau appends a number of supplements each year to collect additional data on a variety of topics. The ASEC collects data on household demographic composition, income during the prior calendar year, health insurance coverage, and migration. The ASEC is conducted primarily in March of each year, with a relatively small number of additional interviews conducted in February and April.[3] As with the ACS, the Census Bureau does not release the interview month on the public use file—however, that information is available on the internal file. Because household composition is defined as of the interview month, we used the interview month in linking the CPS ASEC data to the SNAP administrative data, as explained below.

Unlike the ACS and CPS ASEC, the SIPP is a longitudinal survey. Members of households interviewed at the start of the survey are reinterviewed every four months over a period of a few years. At each interview, they are asked detailed questions about each of the four preceding months, yielding monthly data on household composition, income, participation in numerous programs, health insurance coverage, and other characteristics. Modules attached to most of the interviews collect additional information on specialized topics such as assets, retirement plans, and child care expenses, which vary by wave. Some of these topical modules are repeated at annual intervals; others are administered less often. Because the SIPP data are reported by reference month rather than interview month, we used the reference month in linking the SIPP data to the SNAP administrative data. Details are provided below.

The three surveys differ substantially in the number of households from which they collect data. In 2012, the ACS collected data from more than 2 million households. For the CPS ASEC, the interviewed sample size in 2012 was 74,000 households. For the SIPP, 2012 represented the end of the fourth year and the start of the fifth year of data collection for the 2008 panel. Although the 2008 panel started with about 40,000 interviewed households, the number of households responding to each of the three waves in 2012 was around 30,000.

Another way in which the surveys differ is in the amount of information they collect that is relevant to simulating SNAP unit composition. The ACS collects the most limited data. The questionnaire asks how each household member is related to the household reference person, or householder, but it does not ask how household members are related to each other. Some relationships can be inferred. For example, the Census Bureau reconstructs subfamilies among persons related to the householder. This effort is assisted by subsequent questions on marital

[3] The February and April interviews are part of a congressionally mandated sample expansion that was introduced in 2001 to improve the precision of state-specific estimates of uninsured children. The ASEC supplement became the new name for what was previously known as the March supplement.

status, whether a woman has given birth in the past 12 months, and whether an individual's grandchildren are in the household. Clearly, the limited relationship data has implications for the simulation of SNAP units, particularly among persons unrelated to the householder. The Census Bureau does not attempt to infer relationships among such persons. Users who wish to do so must base their inferences on age and gender.

The CPS ASEC collects sufficient relationship data to identify subfamilies directly; this extends to persons unrelated to the householder. However, subfamily identification among the latter is less complete than among persons related to the householder. For example, parent-child relationships among persons unrelated to the householder are established only for children under 18. No other relationships among adults are identified, besides spouses.

The SIPP collects full relationship information among all household members in one of its early topical modules, but that information is not maintained over time. Nevertheless, even the monthly relationship data collected in the SIPP are more complete than those collected in the CPS ASEC. In addition, a periodic module collects information on who shares food expenses, which is relevant to SNAP unit composition. Moreover, while all three surveys collect data on SNAP participation, only the SIPP determines which household members were covered by SNAP—and it does so for every month. The SIPP also collects information on which reported SNAP beneficiaries in each month formed SNAP units together. Although SNAP participation is underreported in the SIPP, the underreporting is far lower than in either of the two other surveys (Czajka et al. 2012). Given this, reported unit composition is available for a substantial majority of the actual SNAP participants who were included in the SIPP households.

B. Overview of New York and Colorado SNAP Administrative Data

SNAP administrative data are compiled and maintained at the state level. Although the systems that states use to administer SNAP vary, state administrative records generally (1) identify who received SNAP benefits in a given month and (2) contain a variety of information on the individual recipients and the unit as a whole. The household is not a relevant unit for SNAP administration—states prepare statistics on the numbers of units receiving SNAP benefits, not the number of households. Although address information, theoretically, may link units residing in the same household in a given month, the accuracy of addresses from month to month is not as important administratively as the accuracy of unit identifiers.

Under agreements between the state SNAP agencies and the Census Bureau, New York and Colorado supplied SNAP administrative data to the bureau for research purposes. New York provided data for multiple years. Colorado initially supplied data for just 2012 and subsequently added 2013. The New York and Colorado data were delivered in the form of person-months or pay-months, meaning that there is a record for every month that a person was a beneficiary.

As part of its standard preprocessing of administrative data, the Census Bureau adds a Master Address File Identifier and Protected Identification Key (PIK), where possible, and strips off direct personal identifiers. The Master Address File is a database of unique household addresses that the Census Bureau develops and maintains for household survey sampling and decennial census enumeration. The PIK is a unique personal identifier generated by the Census Bureau to facilitate the linking of records across databases while preserving the confidentiality of individual records. Each PIK corresponds to a social security number (SSN), but the SSN cannot be derived from the PIK.

The Census Bureau assigns PIKs to the records in each survey or administrative database by using either an exact or a probabilistic match between that database and an enhanced version of the Numident database, in which the Social Security Administration maintains data collected from applications for SSNs. Many administrative databases include SSNs, which are first validated against the Numident database; PIKs can then be assigned to records with validated SSNs by exact match. For probabilistic matches the match variables include full names, demographic information, and addresses. Once PIKs are assigned, files can be linked by an exact match on PIK—yet, the removal of direct identifiers prevents users with access to the data from identifying individual records. The fact that PIKs do not exist for persons without SSNs (or, recently, the alternate taxpayer identification numbers that the Internal Revenue Service assigns to individuals who do not qualify to receive SSNs) excludes a small percentage of the population. However, a bigger source of missing or erroneous PIK assignments is incomplete or incorrect identifying information. Widely shared surnames and varied spellings also contribute to PIK assignment failures. In household surveys, typically 9 to 10 percent of the sample cannot be assigned PIKs. Furthermore, because the match is

probabilistic, some of the PIK assignments may be incorrect, which in very large files can result in duplicate PIK assignments. Statistics on erroneous PIK assignments are not available. For the New York and Colorado SNAP files, the bureau was able to assign PIKs to 99.2 percent of the records, which meant that virtually all of the records were potentially eligible to be linked to household survey data from the two states (Tordella et al. 2014).

C. Construction of the Analysis Files

Each of the analysis files that we created combines one of the state SNAP administrative files with one of the three surveys. We constructed analysis files linking the New York SNAP administrative records to the ACS, CPS ASEC, and SIPP, and an analysis file linking the Colorado SNAP administrative records to the ACS. Linkages between the Colorado SNAP records and both the CPS ASEC and the SIPP would have produced too few records for useful analysis.

1. General Approach

Creation of each analysis file involved the following six steps:

1. Preprocessing the administrative data for the selected state (New York or Colorado)
2. Preprocessing the survey data (the ACS, CPS ASEC, or SIPP) for that state
3. Linking the administrative records to the survey records by PIK
4. Adding the unmatched members of each survey household with one or more matched members
5. Adding the unmatched members of each administrative SNAP unit with one or more matched members
6. Appending the ERS simulated unit ID to the survey records (for the ACS only)

Steps 3 through 5 constitute the main components to the construction of each analysis file. After we linked state SNAP records and survey person records by PIK, every survey household in which one or more members was matched to a state SNAP record was included in the analysis file. To these households, we added the survey records of any unmatched household members. This ensured that every survey household in the analysis file included all of its members as identified in the survey household roster. Next, we added state SNAP person records of unmatched SNAP participants who shared SNAP unit membership with a matched state SNAP person. This component, which was by far the most complex of the three steps, created an "enhanced survey household." We use this term to describe the analysis file households because some of the households contain more persons than were identified by the survey.[4]

Originally, we intended to include all unmatched persons in the state data who resided at the same address as the matched persons. This would have enabled us to add SNAP units with no members who matched to the survey records. Given that about 10 percent of the survey records lack PIKs, this could have been particularly useful for capturing small SNAP units—especially one-person units—that were excluded by the limitations of the matching process. These units might include survey household members who lacked PIKs and therefore could not be matched. However, we could not add units that had no unmatched members at the same address because the address information in both states was deficient.

The analysis files we constructed are at the person level and contain three types of records. Using the ACS–New York analysis file as an example, the file includes person records drawn from the ACS, the New York SNAP data, or both. For the ACS household members who were matched to New York SNAP participants, the person record contains data from both the ACS and the New York SNAP data. For the ACS household members who were not matched to New York SNAP participants, the person record contains only the ACS data. For members of New York

[4] As we will discuss later, when there are at least as many unmatched survey household members as unmatched state SNAP unit members, we cannot be sure that the unmatched SNAP members are truly additions to the survey household. Rather, it could be that all of the unmatched SNAP unit members would have been matched to the unmatched survey household members if PIKs had been assigned correctly to all members of the survey household and SNAP administrative unit.

SNAP units who were not matched to ACS records but whose units were matched, the person record contains only New York SNAP data. Therefore, for each ACS household, there is a record for every ACS household member, plus any additional unmatched New York SNAP participants whose SNAP units include at least one member who matched a member of the ACS household. This is the enhanced survey household.

2. Survey-Specific Procedures

The procedures that we followed in constructing an analysis file of linked survey and administrative data differed somewhat across the three surveys. The survey-specific procedures are outlined below.

a. The ACS

Although the reference period for reporting SNAP participation in the ACS is the past 12 months, we chose to restrict our matches to the survey month and the two preceding months. Specifically, for each New York ACS sample member in the 2012 annual sample, we searched the New York SNAP administrative file for a match on the PIK and the survey month. For records with no match, we searched the New York SNAP file for a second time for a match on the PIK but in the month prior to the interview month. Finally, for the remaining unmatched New York ACS records, we searched the New York SNAP file for a third time for a match on PIK two months prior to the interview month. As we will show in the next chapter, the number of matches dropped sharply between the survey month and each prior month. That SNAP enrollment was still rising in 2012 may have contributed to this outcome. Had our matches been performed in a period of declining enrollment, the falloff in matches between one month and the prior month may have been less precipitous. As it was, the pattern of declining matches provided strong support for our decision to focus on the survey month.

b. The CPS ASEC

Although the bulk of the CPS ASEC interviews are conducted in March, some are conducted in February and April. In linking the New York SNAP data to the CPS ASEC, we followed the same procedures used with the ACS. That is, we first linked SNAP records to New York records in the CPS ASEC by PIK and interview month. For unmatched New York CPS ASEC records, we then searched the New York SNAP file for matches in the month prior to the interview month. Finally, for the remaining unmatched New York CPS ASEC records, we searched the New York SNAP file for matches two months prior to the interview month.

c. The SIPP

Because the SIPP is longitudinal, with monthly data for every interviewed sample member, we employed a different linkage strategy than we did with the ACS and CPS ASEC. The SIPP interviews for a survey wave are conducted over four consecutive months. The sample is divided into four "rotation groups" that are equally sized. The members of each rotation group are interviewed in the same month and asked about their activities over the same previous four months. With this design, there is one month that appears in the reference period for all four rotation groups in a given wave. We used the SIPP data from waves 12, 13, and 14, which have as their common reference months April, August, and December 2012, respectively.

For each of the three months, we searched the New York SNAP data for matches to New York SIPP records in that month. We then added unmatched SIPP household members and unmatched SNAP unit members to create a set of enhanced SIPP households in each month. Because of the panel nature of the survey and the continuity of SNAP participation, a given household might appear in all three of the match months. However, its membership might not be the same in all three months and its SNAP participation might be different. Furthermore, a given household in April might be two households in December if one or more members left to form a new household. Limiting households to one appearance in the data set is not only challenging—requiring us to determine which of two or three versions of a household is best—it also reduces the representativeness of the final data set. In other words, although multiple appearances of the same household overstate the effective sample size of the data set based on raw sample counts, they reduce bias relative to a data set that limits each household to one appearance. Given these considerations, we opted to pool the April, August, and December data sets without removing any SNAP households. This effectively tripled the size of the SIPP analysis file.

3. The Possibility of Missed Matches

Primarily because of missing PIKs in the survey data, it is theoretically possible that there were SNAP administrative units that we were unable to add to the survey households because they contained no matched members. If we missed some units that should have been included, they were most likely to have contained only one or two members, as the likelihood that a SNAP administrative unit will contain at least one matching member rises with the unit size. Similarly, it is highly likely that we missed some SNAP households entirely. If the sole member of a one-person SNAP household in the survey data lacked a PIK, then a match was not possible. With a 90 percent PIK rate for the ACS and CPS ASEC (the SIPP is somewhat lower), 10 percent of the one-person SNAP households would be missed if successful PIK assignment were independent of SNAP participation. Because SNAP participants are generally required to have SSNs, their PIK rate should be higher than the average survey respondent. Furthermore, some of the findings we present below suggest that survey PIK assignment is more successful for householders than for other household members. For these reasons, the fraction of one-person SNAP households missed by our linkages is probably well below 10 percent. Nevertheless, small SNAP households are likely to be underrepresented in our analysis files.

Unidentified SNAP units in households with one or more matched units will lead us to underestimate the average number of SNAP units among our matched households. At the same time, if unidentified SNAP households are disproportionately one-unit households, their exclusion from our analysis files will contribute a bias in the opposite direction to our estimates of the average number of SNAP units per household. We cannot predict the net effect of these opposing biases.

4. Summary of Match Results by Survey

Actual and projected match results for the six survey and administrative data linkages for 2012 are reported in Table II.1. Linking the largest of the three surveys, the ACS, with administrative data from New York yielded nearly 20,000 households with matched data. Linking the same state administrative data to the CPS ASEC produced a small fraction of matched households—just over 600—while the SIPP match produced an average of about 290 matched households over the three months. Combining the matches from the three months gave us 860 matched households in our analysis file.

Table II.1. Number of Enhanced Survey Households in Each Analysis File, Actual and Projected

Analysis File	Actual Size	Projected Size
ACS–New York SNAP	19,856	
CPS ASEC–New York SNAP	604	
SIPP–New York SNAP[a]	1,274	
ACS–Colorado SNAP	2,899	
CPS ASEC–Colorado SNAP		90
SIPP–Colorado SNAP[a]		190

Source: Mathematica Policy Research, from linked 2012 survey and SNAP administrative data.

[a] File combines households from three waves containing many of the same households interviewed in different months.

In terms of total participants, the Colorado SNAP program is one-sixth the size of New York's, so we would expect the matching process for each of the three surveys to yield roughly one-sixth the number of matched sample households as New York. For the ACS, we obtained close to 3,000 matched Colorado households—a little under one-sixth the number we obtained in New York but almost five times as many as in the CPS ASEC–New York SNAP file. However, linking the Colorado SNAP data with the CPS ASEC and the SIPP would appear likely to yield too few matched households to support very useful analysis in isolation from other data. The projected counts of 90 for the CPS ASEC and 130 for the SIPP would severely limit the types of analysis that could be conducted on the matched households.

III. Findings from the ACS: New York

We present findings from our analysis of the ACS data matched to New York and Colorado SNAP records, beginning with the former. In discussing the findings from our analysis of the ACS data linked to New York SNAP administrative data, we begin by examining the results of the linkage at the person level. Following that, we review key findings at the level of the enhanced ACS household, then we turn to an examination of the comparative characteristics of ACS households, simulated SNAP units, and actual SNAP units.

A. Linkage Results at the Person Level

We examine the linkage results at the person level beginning with the ACS and then turning to the New York SNAP data. In all, the 19,586 ACS households to which at least one New York SNAP record was matched contain 57,442 individuals (Table III.1). Of these, 39,371 or 68.5 percent were matched to New York SNAP records. Of those who were not matched, 15,260 or 26.6 percent of the total had a PIK, implying that they would have been matched if they had participated in SNAP during the survey month or one of the two preceding months. The remaining 2,811 or 4.9 percent of the total had no PIK, so they could not have been matched to the New York SNAP data even if they had participated in SNAP in New York during the three-month reference period. The fraction lacking a PIK is less than half the proportion of all 2012 ACS person records that could not be assigned a PIK (Tordella et al. 2014). New York could be above the ACS average with respect to PIK success, but we suspect that most of the difference is due to a higher than average PIK rate among households in which at least one member could be linked to a SNAP participant record.

Table III.1. Distribution of ACS Household Members by Match Result to New York SNAP Administrative Data

Match Result	Number	Percentage
All ACS household members	57,442	100.0
Matched a New York SNAP administrative record	39,371	68.5
Did not match a New York SNAP administrative record	18,071	31.5
Has a PIK	15,260	26.6
Does not have a PIK	2,811	4.9

Source: Mathematica Policy Research, from linked 2012 ACS and New York SNAP administrative data.

It is significant that most of the unmatched ACS records have PIKs. These records could not be matched to New York SNAP administrative records, on which PIKs were absent less than 1 percent of the time. It is highly likely, therefore, that these individuals did not participate in SNAP during the three-month reference period. We discuss the implications at the household level below. In addition to the 39,371 New York SNAP administrative records that were matched to the ACS records, another 3,766 New York SNAP records were added to the enhanced ACS households as unmatched members of the New York SNAP units of the matched persons (Table III.2). All but 19 of these records had PIKs. That a record with a PIK failed to match anyone in the ACS household to which another member of the SNAP unit was successfully matched could have one of several explanations:

- The New York SNAP record's ACS counterpart did not have a PIK.
- One of the two records had an incorrect PIK.
- The person represented by the New York SNAP record did not live in the ACS household during the survey month.
- The person represented by the New York SNAP record lived in the ACS household during the survey month but was not reported.

The last possibility implies a deficiency in the identification of household members in the ACS. Such cases have direct implications for the decennial census and how administrative records from SNAP might be helpful in improving the quality of the census count. Our analysis did not attempt to explain the missed matches, however.

Table III.2. Distribution of New York SNAP Administrative Records by Match Result to the ACS and by Match Month

Match Status	Number	Percentage
New York SNAP unit members in enhanced ACS households	43,137	100.0
Matched an ACS record	39,371	91.3
Did not match an ACS record	3,766	8.7
Has a PIK	3,747	8.7
Does not have a PIK	19	0.0
Records matched to ACS records	39,371	100.0
Matched in ACS interview month	37,707	95.8
Matched one month prior to ACS interview month	1,023	2.6
Matched two months prior to ACS interview month	641	1.6

Source: Mathematica Policy Research, from linked 2012 ACS and New York SNAP administrative data.

The lower portion of Table III.2 shows how often the New York SNAP administrative records that matched ACS records were drawn from either the survey interview month, one month earlier, or two months earlier. The vast majority of matches—95.8 percent—occurred in the survey interview month. Another 2.6 percent occurred one month earlier; the remaining 1.6 percent occurred two months earlier.[5] Adding matches from the two previous months does not expand the survey household; rather, it expands the SNAP unit to include more members of the survey household. The small numbers added from the two months preceding the interview month support our strategy of limiting the search for matches to just three months.

B. Linkage Results at the Household Level

In turning from the person level to the household, we examine first the prevalence of SNAP units—both simulated and actual—within households participating in SNAP. We then look at the combined incidence of unmatched individuals from the survey and administrative data and the alignment of simulated and actual units within the same household.

1. SNAP Units in Households

Table III.3 directly addresses the important issue of the prevalence of multiple SNAP units within the same household. The table provides a cross-tabulation of the number of administrative SNAP units by the number of simulated SNAP units in New York ACS households with at least one member linked to a SNAP administrative record. The unit simulations presented in this table do not make use of the ERS eligibility simulation or a participation imputation, so we are comparing participating administrative units with simulated units that could be (1) ineligible or eligible and nonparticipating as well as (2) eligible and participating. Of course, all of the households have at least one member participating in SNAP, which may have implications for the eligibility and participant status of the other members.

Of the 19,586 ACS households, only 7.2 percent had more than one SNAP administrative unit, while 33.9 percent had more than one simulated unit. Thus, the SNAP unit simulation creates multiple SNAP units more frequently than participating units exist in the New York data. The numbers of units of each type are related, however. Nearly two-thirds (64.8 percent) of the households had only one of each type of unit. Another 3.0 percent had two of each type of unit and another 0.6 percent had three or more of each type of unit. Furthermore, if there was only one simulated unit, there was nearly always only one administrative unit. On the other hand, if there was only one administrative unit, then 30 percent of the time there was more than one simulated unit. And if there were exactly

[5] Not shown in the table: all 3,766 nonmatching NY SNAP records, which were added to the analysis file because they shared SNAP unit membership with one or more matched records, were drawn from the survey interview month. Limiting these additions to the survey month increases the likelihood that the persons in these units were in the household at the time of the survey.

two administrative units, then more than half the time there was a different number of simulated units, with one simulated unit being as likely as three simulated units.

Table III.3. Prevalence of Multiple SNAP Units in New York ACS Households: All Simulated Units Included

Administrative SNAP Units with Matches to the Same ACS Household	Total ACS Households	Simulated SNAP Units Within the Same ACS Household				
		One	Two	Three	Four	Five or More
Number	19,586	12,946	4,282	1,555	533	270
One	18,178	12,688	3,668	1,246	413	161
Two	1,267	253	596	267	94	59
Three or more	141	5	18	42	26	50
Percentage of Total ACS Households	100.0	66.1	21.9	7.9	2.7	1.4
One	92.8	64.8	18.7	6.4	2.1	0.8
Two	6.5	1.3	3.0	1.4	0.5	0.3
Three or more	0.7	0.0	0.1	0.2	0.1	0.3

Source: Mathematica Policy Research, from linked 2012 ACS and New York SNAP administrative data.

When we restrict the simulated units to those that the ERS program simulates to be eligible financially, we find that 23.0 percent of the ACS households have no simulated units, while only 13.7 percent have more than one simulated unit (Table III.4). This latter figure contrasts with 33.9 percent when all simulated units are included. Below, we compare the overall number of simulated SNAP units with and without the eligibility screen and contrast these totals with the number of actual SNAP units in these households.

Table III.4. Prevalence of Multiple SNAP Units in New York ACS Households: Only Eligible Simulated Units Included

Administrative SNAP Units with Matches to the Same ACS Household	Total ACS Households	Simulated SNAP Units Within the Same ACS Household					
		None	One	Two	Three	Four	Five or More
Number	19,586	4,510	12,393	2,150	413	85	35
One	18,178	4,316	11,947	1,617	249	41	6
Two	1,267	179	429	497	124	28	12
Three or more	141	15	17	36	40	16	17
Percentage of Total ACS Households	100.0	23.0	63.3	11.0	2.1	0.4	0.2
One	92.8	22.0	61.0	8.3	1.3	0.2	0.0
Two	6.5	0.9	2.2	2.5	0.6	0.1	0.1
Three or more	0.7	0.1	0.1	0.2	0.2	0.1	0.1

Source: Mathematica Policy Research, from linked 2012 ACS and New York SNAP administrative data.

2. Match Status Within Households

Table III.5 breaks down the 19,586 enhanced ACS households by the match status of their members and those of the linked SNAP units. In 55.6 percent of the households, all members of the ACS household were matched to New York SNAP administrative unit members. For most of these households, or 52.2 percent of all households with one or more matches, there were no unmatched SNAP administrative unit members. Of the 10,892 households in which all ACS members were matched, there were 660 households (3.4 percent of the total) in which one or more administrative unit members were not matched (that is, we added an unmatched administrative unit member to the ACS household.)

Table III.5. Match Status of ACS Household Members and New York SNAP Administrative Unit Members

Match Status Summary	Number	Percentage
ACS households containing matched person	19,586	100
All ACS household members matched	10,892	55.6
All administrative unit members matched	10,232	52.2
Some administrative unit members not matched	660	3.4
Some ACS household members not matched	8,694	44.4
All administrative unit members matched	7,284	37.2
Some administrative unit members not matched	1,410	7.2

Source: Mathematica Policy Research, from linked 2012 ACS and New York SNAP administrative data.

Among the 8,694 households (44.4 percent of the total) in which there was at least one unmatched member of the ACS household, about three-quarters (7,284 or 37.2 percent of the total) had no unmatched administrative SNAP unit members. A small subset of the matched households with one or more unmatched ACS members (1,410 households or 7.2 percent of the total) also had unmatched SNAP administrative unit members.

C. Characteristics of Simulated Versus Administrative SNAP Units

To better understand what the SNAP unit simulation does, we compare the characteristics of entire ACS households with (1) those of SNAP units simulated using the ERS algorithm, (2) the subset of these that were also simulated to be eligible, and (3) the actual SNAP units matched to the ACS households from the New York administrative data.

1. Demographic Characteristics

Table III.6 compares the full households and both simulated and actual SNAP units with respect to size (number of persons) and other demographic characteristics. The ERS simulation algorithm yielded 29,852 simulated SNAP units across the 19,586 ACS households, or about 1.52 units per household. Of these, 18,463 simulated units or fewer than 1 unit per household, were judged to be eligible to receive SNAP benefits. The number of New York SNAP administrative units that linked to the ACS households through one or more members numbered 21,159, or about 1.08 units per household. The excessive simulated units when eligibility is not considered could be the result of the simulation creating too many small units. As described by Newman and Scherpf (2013), the simulation assumes that, subject to certain mandatory rules of unit formation, "households will form the smallest units possible." This would lead us to expect a greater proportion of small units among the simulated units than among the actual SNAP administrative units.[6]

FMWWith the focus on creating small units, one-person units are more common among the simulated units—both the total units and eligible units—than among the actual units, accounting for 58.1 percent of both sets of simulated units versus 52.9 percent among SNAP administrative units.[7] Both two- and three-person units are more common

[6] Simulations of SNAP-eligible units have tended to produce too few one-person units compared to what is observed in actual caseload statistics. Designing a simulation to produce the smallest units consistent with the mandatory rules would seem to be an appropriate strategy to address this chronic limitation of past simulations.

[7] In Chapter II, we explained that if a one-person SNAP household had a missing or incorrect PIK in the ACS, then that household would not have been matched and, therefore, would not be included in our analysis file. Larger households could have been excluded as well, but the number of missing SNAP households is probably greatest for one-person households, given that a single missing PIK would prevent a match. If the number of one-person households is understated in the analysis file, this would imply that the numbers of simulated and actual one-person units are also understated. In Table III.1, we saw that the percentage of ACS household members who were missing PIKs (in households with one or more matched members) was 4.9 percent. We suggest that this is a reasonable rate of missing PIKs to assume for one-person households participating in SNAP, as their missing PIKs would be due to problems in assigning a PIK rather than their lack of an SSN. If we apply this rate of "missingness" to the number of matched one-person households, we find that we would be missing fewer than 300 households. This would raise the

among SNAP administrative units than among simulated units, offsetting the difference in one-person units. Units of size four or larger are about equally common among simulated and actual SNAP units.

Table III.6. All Matched New York ACS Households Compared to Simulated SNAP Units and Administrative SNAP Units

Household/Unit Characteristic	ACS Households Matched with New York Data	All SNAP Units Simulated with ERS Algorithm	Eligible SNAP Units Simulated with ERS Algorithm	SNAP Units from New York Administrative Data
Total	19,586	29,852	18,463	21,159
Household/SNAP unit size				
1	28.4	58.1	58.1	52.9
2	21.7	18.1	16.8	19.7
3	17.2	9.7	9.9	12.6
4	13.8	7.4	7.6	7.6
5	9.1	3.9	4.1	3.7
6 or more	9.8	2.9	3.5	3.4
Household/unit contains:				
Children (birth to 17)	43.9	29.7	32.0	38.3
Age: birth to 4	19.8	13.1	14.8	17.8
Age: 5 to 17	36.3	24.3	26.3	30.8
Children only	0.0	0.1	0.1	4.4
Non-elderly adults (age 18 to 59)	77.5	72.5	69.6	65.8
Non-elderly disabled members (birth to 59)	23.8	17.1	21.4	19.7
Elderly adults (age 60 or older)	41.7	29.9	32.4	32.9
Living alone	17.1	22.2	25.6	25.1

Source: Mathematica Policy Research, from linked 2012 ACS and New York SNAP administrative data.

Demographic composition is of particular interest given that simulated units have tended to differ from SNAP quality control units in being more likely to have children and more likely to have elderly members (60 or older), but less likely to consist of an elderly person living alone. With respect to children, we find the opposite, with administrative units being more likely to have both preschool and school-age children than simulated units—although, less likely than the ACS household. Consistent with prior findings, however, there are virtually no child-only units among the simulated SNAP units and effectively no child-only households—yet, child-only units represent 4.4 percent of all SNAP administrative units in New York State.

IV. Additional Findings from the ACS: Colorado

In reviewing results from the ACS Colorado SNAP match, we were particularly interested in how they differ, if at all, from those obtained with the ACS New York SNAP match. One difference that we observed immediately is that matched ACS household members made up a smaller share of all ACS household members in the households with matches than they did in New York. Specifically, 61.5 percent of the ACS household members in Colorado were matched (Table IV.1). In New York, this figure was 68.5 percent. The fraction of household members without a PIK, 5.3 percent, was essentially the same as in New York, while the overall PIK rates for New York and Colorado SNAP administrative records were about the same (Tordella et al. 2014). Therefore, the lower fraction of matches cannot be attributed to a lower PIK rate.

The proportion of SNAP records added through direct matches to ACS records rather than through a SNAP case affiliation—89.0 percent—was also lower than in New York, but by only 2 percentage points (Table IV.2). We also

percentage share of one-person households to 29.5 percent while raising the percentage shares of simulated and actual SNAP units by somewhat smaller fractions, given their larger totals.

found that a slightly smaller fraction of administrative records matched in the interview month versus the two prior months: 93.2 percent in Colorado and 95.8 percent in New York. Compared to the direct matches, a smaller fraction—89.4 percent—of the SNAP records added through case affiliations came from the interview month.

Table IV.1. Distribution of ACS Household Members by Match Result to Colorado SNAP Administrative Data

Match Result	Number	Percentage
All ACS household members	9,214	100.0
Matched a Colorado SNAP administrative record	5,666	61.5
Did not match a Colorado SNAP administrative record	3,548	38.5
Has a PIK	3,064	33.3
Does not have a PIK	3	5.3

Source: Mathematica Policy Research, from linked 2012 ACS and Colorado SNAP administrative data.

Table IV.2. Distribution of Colorado SNAP Administrative Records by Match Result to the ACS and by Match Month

Match Status	Number	Percentage
Colorado SNAP unit members in enhanced ACS households	6,367	100.0
Matched an ACS record	5,666	89.0
Did not match an ACS record	701	11.0
Has a PIK	696	10.9
Does not have a PIK	5	0.1
Records matched to ACS records	5,666	100.0
Matched in ACS interview month	5,282	93.2
Matched one month prior to ACS interview month	247	4.4
Matched two months prior to ACS interview month	137	2.4
Records not matched to ACS records, in same Colorado SNAP unit	701	11.0
Same SNAP unit in ACS interview month	627	89.4
Same SNAP unit one month prior to ACS interview month	45	6.4
Same SNAP unit two months prior to ACS interview month	29	4.1

Source: Mathematica Policy Research, from linked 2012 ACS and Colorado SNAP administrative data.

Distributions of simulated and actual SNAP units in Colorado households were similar to what we found in New York. Households with a single SNAP administrative unit and a single simulated SNAP unit were slightly more common in Colorado than they were in New York: 94.5 percent of households in Colorado had a single SNAP administrative unit and 68.3 percent had a single simulated unit (Table IV.3). In New York, these two figures were 92.8 percent and 66.1 percent, respectively (Table III.3). When we restricted the simulated units to those that were eligible, the percentage of households with no simulated units was about 9 percentage points higher than New York: 31.9 percent in Colorado versus 23.0 percent in New York (Table IV.4). Likewise, the fraction of households with one unit was 6 percentage points lower than in New York, and the fraction with two or more simulated units was 3 percentage points lower.

The distribution of matched sample households by the match results for their members reveals some notable differences from New York. In Colorado, 42.7 percent of the matched households had all of their ACS members matched, while 40.1 percent of matched households also had all of their SNAP administrative unit members matched (Table IV.5). The comparable figures for New York were 55.6 percent and 52.2 percent, respectively—substantially higher. Among the Colorado households with at least one unmatched ACS member (57.3 percent of the total), four-fifths (or 46.3 percent of the total) had no unmatched administrative SNAP unit members. The corresponding figures for New York were 44.4 percent and 37.2 percent, respectively. Overall, 13.7 percent of the SNAP unit members in Colorado were unmatched compared to 10.6 percent in New York.

Table IV.3. Prevalence of Multiple SNAP Units in Colorado ACS Households: All Simulated Units Included

Administrative SNAP Units with Matches to the Same ACS Household	Total ACS Households	Simulated SNAP Units Within the Same ACS Household				
		One	Two	Three	Four	Five or More
Number	2,899	1,979	642	200	56	22
One	2,741	1,944	572	167	43	15
Two or more	158	35	70	33	13	7
Percentage of total ACS households	100.0	68.3	22.1	6.9	1.9	0.8
One	94.5	67.1	19.7	5.8	1.5	0.5
Two or more	5.5	1.2	2.4	1.1	0.4	0.2

Source: Mathematica Policy Research, from linked 2012 ACS and Colorado SNAP administrative data.

Table IV.4. Prevalence of Multiple SNAP Units in Colorado ACS Households: Only Eligible Simulated Units Included

Administrative SNAP Units with Matches to the Same ACS Household	Total ACS Households	Simulated SNAP Units Within the Same ACS Household					
		None	One	Two	Three	Four	Five or More
Number	2,899	926	1,660	262	37	14	2,899
One	2,741	877	1,621	211	22	10	2,741
Two or more	158	49	39	51	15	4	158
Percentage of total ACS households	100.0	31.9	57.3	9.0	1.3	0.5	100.0
One	94.5	30.3	55.9	7.3	0.8	0.3	94.5
Two or more	5.5	1.7	1.3	1.8	0.5	0.1	5.5

Source: Mathematica Policy Research, from linked 2012 ACS and Colorado SNAP administrative data.

Table IV.5. Match Status of ACS Household Members and Colorado SNAP Administrative Unit Members

Match Status Summary	Number	Percentage
ACS households containing matched person	2,889	100
All ACS household members matched	1,237	42.7
All administrative unit members matched	1,163	40.1
Some administrative unit members not matched	74	2.6
Some ACS household members not matched	1,662	57.3
All administrative unit members matched	1,342	46.3
Some administrative unit members not matched	320	11.0

Source: Mathematica Policy Research, from linked 2012 ACS and Colorado SNAP administrative data.

Compared to New York, SNAP households and simulated units in Colorado are larger. Also, SNAP administrative units have a slightly higher proportion of larger units than in New York, but the proportion of units that consist of just one person is slightly larger than in New York as well, at 55.7 percent versus 52.9 percent (Table IV.6). Also, unlike New York, one-person units are more common among SNAP administrative units than among simulated units in Colorado, where they represent 50.9 percent of all simulated units and 52.9 percent of simulated eligible units.

Because of Colorado's generally younger population, there are notable differences in the age distribution of SNAP beneficiaries and the members of SNAP households. Differences between SNAP households, simulated units, and SNAP administrative units are similar to what we saw in New York, however. Compared to New York, SNAP households and units are more likely to include children and much less likely to include elderly adults. For example, 54.0 percent of Colorado ACS households and 44.7 percent of SNAP administrative units have children under 18. In New York, these fractions are 43.9 percent and 38.3 percent, respectively. Similarly, child-only units represent 9.3

percent of all SNAP administrative units in Colorado, compared to 4.4 percent in New York.[8] Units with elderly adults make up only 18.2 percent of the administrative units in Colorado, compared to 32.9 percent in New York. Elderly adults living alone are 15.6 percent of the SNAP administrative units in Colorado versus 25.1 percent in New York. Within Colorado as in New York, the SNAP administrative units are more likely to contain school-age children and, very slightly, younger children than the simulated units. However, the SNAP administrative units are less likely to contain school-age children than SNAP households. In Colorado, elderly adults living alone are most common among simulated eligible units at 17.8 percent, compared to 15.6 percent among actual SNAP units.

Table IV.6. All Matched Colorado ACS Households Compared to Simulated SNAP Units and Administrative SNAP Units

Household/Unit Characteristic	ACS Households Matched with Colorado Data	All SNAP Units Simulated with ERS Algorithm	Eligible SNAP Units Simulated with ERS Algorithm	SNAP Units from Colorado Administrative Data
Total	2,899	4,221	2,359	3,075
Household/SNAP unit size				
1	23.2	50.9	52.9	55.7
2	20.0	17.6	15.7	14.3
3	18.0	11.5	11.4	12.3
4	15.7	9.5	8.8	8.7
5	11.1	6.2	6.8	5.5
6 or more	12.0	4.3	4.5	3.4
Household/unit contains:				
Children (birth to 17)	54.0	38.1	38.9	44.7
Age: birth to 4	26.6	18.4	20.0	21.3
Age: 5 to 17	43.4	30.4	30.8	35.6
Children only	0.0	0.1	0.0	9.3
Non-elderly adults (age 18 to 59)	86.5	81.5	79.4	73.8
Non-elderly disabled members (birth to 59)	28.7	21.9	26.9	22.5
Elderly adults (age 60 or older)	28.0	20.7	22.3	18.2
Living alone	10.5	14.8	17.8	15.6

Source: Mathematica Policy Research, from linked 2012 ACS and Colorado SNAP administrative data.

V. Summary and Implications

In this final chapter, we summarize our principal findings and discuss their implications for SNAP unit simulation and the potential use of SNAP administrative records in decennial census enumeration.

A. Summary of Findings

Our principal findings from the analysis of linked survey and administrative SNAP data are the following:

- A substantial fraction of households receiving SNAP benefits—between 44 percent and 47 percent in New York State and 57 percent in Colorado—included persons who were not recipients.
- More than 1 in 10 households receiving SNAP benefits had participants identified through administrative records who were not identified as members of the survey household. Most of these households had unmatched survey household members as well, suggesting that incomplete matches may have contributed to these first two findings.

[8] This striking difference is similar to what was recorded in the 2013 SNAP QC data, where child-only units were 8.2 percent of the total in Colorado and 5.4 percent of the total in New York.

- A very substantial majority of households receiving SNAP benefits in New York—93 percent—contained only one SNAP unit, while nearly all of the remaining SNAP households included only two SNAP units.
- Simulation of SNAP units in households receiving SNAP benefits produced a greater incidence of multiple SNAP units—nearly a third of all SNAP households—than was observed with linked administrative data.
- In nearly a quarter—23 percent—of the households receiving SNAP benefits in New York and nearly a third—32 percent—in Colorado, no one was simulated to be eligible for SNAP.
- Simulated SNAP units were very similar to actual SNAP units in their distribution by size. In New York, however, one-person units were more common among simulated than actual SNAP units, while the reverse was true in Colorado.
- SNAP units containing elderly adults, including those living alone, were about equally common among administrative and simulated eligible units in New York but somewhat more common among simulated eligible units in Colorado.
- Using the entire household as a proxy for the SNAP unit yields too few units overall and far too few one-person units.

These findings carry implications for the simulation of SNAP units and also for ways in which SNAP administrative data might be useful in improving decennial census enumeration. These implications are discussed below.

B. Implications

This research has focused on SNAP unit composition and the correspondence between actual and simulated SNAP units, so the primary implications involve ways to improve the simulation of SNAP units. However, findings regarding SNAP members who do not appear to have been counted in the survey household have implications for household rostering in both surveys and the decennial census. We consider such implications after discussing how our findings bear on SNAP simulation.

1. SNAP Simulation

A widespread perception among those modeling SNAP eligibility has been that SNAP households often contain multiple units. This appeared the most plausible way to reconcile the numbers of one-person SNAP units with the numbers of eligible one-person households. In linking SNAP units to survey households, however, we found that households with a single SNAP unit accounted for the vast majority of households receiving SNAP but that SNAP households often have members who are not receiving SNAP benefits. These two findings suggest a change in tactics toward finding single eligible subunits in more households and away from searching for multiple eligible subunits within each household.

2. Census Enumeration

Most of the implications that we see for census enumeration bear further research. These implications depend upon researchers first establishing whether or not the New York and Colorado SNAP data correctly identify household residents missed by major household surveys conducted by the Census Bureau. If omissions from survey household rosters can be confirmed, then whatever can be learned about the causes may suggest ways to reduce such omissions in the future. Because the SIPP follows sample members over time and collects data on all of the persons with whom they share households, further research with the linked SIPP and SNAP administrative data will be useful in determining what role timing may play in the correspondence between survey and administrative households.

References

Bollinger, Christopher R., and Martin H. David. "Modeling Discrete Choice with Response Error: Food Stamp Participation." *Journal of the American Statistical Association,* vol. 92, no. 439, 1997, pp. 827–835.

Czajka, John L., Anne Peterson, Brittany McGill, Betsy Thorn, and Catharine Warner-Griffin. "The Extent and Nature of Underreporting of SNAP Participation in Federal Surveys." Final report submitted to the U.S. Department of Agriculture, Food and Nutrition Service. Arlington, VA: Insight Policy Research, July 2012.

Harris, Benjamin Cerf. "Within and Across County Variation in SNAP Misreporting: Evidence from Linked ACS and Administrative Records." CARRA Working Paper 2014-05. Washington, DC: U.S. Census Bureau, Center for Administrative Records Research and Applications, July 2014.

Meyer, Bruce, and Robert Goerge. "Errors in Survey Reporting and Imputation and Their Effects on Estimates of Food Stamp Program Participation." U.S. Census Bureau Center for Economic Studies Working Paper Series No. CES-WP-11-14. Washington, DC: U.S. Census Bureau, 2011.

Newman, Constance, and Erik Scherpf. "Supplemental Nutrition Assistance Program (SNAP) Access at the State and County Levels: Evidence from Texas SNAP Administrative Records and the American Community Survey." Economic Research Report No. 156. Washington, DC: U.S. Department of Agriculture, Economic Research Service, September 2013.

Tordella, Stephen, Nancy Wemmerus, Denae Johnson, John Czajka, Karen Cunnyngham, and Randy Rosso. "Analyzing the Data Quality of SNAP Administrative Records. SNAP Data Quality Project: Tasks 2 and 3, New York and Colorado." Preliminary Analysis Memo. Alexandria, VA: Sabre Systems, September 2014.

U.S. Department of Agriculture, Food and Nutrition Service, Office of Policy Support, *Characteristics of Supplemental Nutrition Assistance Program Households: Fiscal Year 2013*, by Kelsey Farson Gray. Project Officer, Jenny Genser. Alexandria, VA, 2014.