Using Administrative Records to Improve the Nonresponse Weighting Procedure in the Consumer Expenditure Survey: Follow Up Analysis

February 27, 2024

Barry Steinberg (U.S. Bureau of Labor Statistics), Stephen Ash (U.S. Bureau of Labor Statistics), and John Voorheis (U.S. Bureau of Census)

Using Administrative Records to Improve the Nonresponse Weighting Procedure in the Consumer Expenditure Survey: Follow Up Analysis

Barry Steinberg (U.S. Bureau of Labor Statistics), Stephen Ash (U.S. Bureau of Labor Statistics), and John Voorheis (U.S. Bureau of Census)

1. Introduction

The Consumer Expenditure Surveys (CE) are the most detailed source of spending data collected directly from U.S. consumers. The CE consist of two components: an Interview Survey, collected quarterly from a random sample of consumer units (CUs) across the U.S.; and a Diary Survey, collected daily for a two-week period from another random sample of CUs across the U.S.[1] [2] The CUs in the samples are representative of all CUs in the U.S., and they are assigned weights that add up to the total number of CUs in the U.S. Unfortunately, some of the CUs in the sample do not participate in the surveys because they are unable to be contacted, refuse to participate, etc. Noninterview adjustments are made to the weights of the responding CUs to mitigate the effects of nonresponse.

Both the Interview and the Diary surveys began using IRS's publicly available ZIP-code income data in their noninterview adjustments with 2014 data. In that process, the IRS's average reported income per ZIP-code was assigned to each CU in the sample and used as a proxy for the individual household's (HH) income. In this report, the effect of IRS's publicly available ZIP-code income versus its non-public Title 26 HH income on CE's noninterview adjustments is analyzed.

Although IRS income data is used in both CE surveys, only the Interview Survey's data is analyzed in this report. There are three reasons. First, the Interview Survey is the source of most of CE's expenditures. Over 80 percent of CE's expenditures come from the Interview Survey, which makes its impact on total expenditures greater than that of the Diary Survey. Second, the impact of ZIP-code income versus HH income is expected to be similar in both surveys, because the sample addresses for both surveys come from the same sampling frame. And third, this report is a follow-up study of two other studies, and those two other studies examined only the Interview Survey's data.

[1] Technically, the two CE samples are random samples of residential addresses, with every CU at the addresses being included in the surveys. Hence, the samples can be thought of as random samples of CUs as well, in a two-stage sampling process.

[2] Consumer units (CUs) are similar to households (HHs). A HH is a group of related family members and all the unrelated people, if any, living together in a housing unit. A person living alone, or a group of unrelated people sharing a housing unit, is also considered to be a HH. By contrast, a CU is a group of people living together in a housing unit who are related by blood, marriage, adoption, or some other legal arrangement; who are unrelated but pool their incomes to make joint expenditure decisions; or is a person living alone or sharing a housing unit with other people but who is financially independent of the other people. The key difference between CUs and HHs is the financial relationship between the people living in the housing unit. The people in a CU are financially interdependent on each other, while the financial relationship between the people living in a housing unit does not figure into the definition of a HH. Most HHs have only one CU, so the terms are often used interchangeably.

Administrative records are nonpublic data collected by other government agencies and surveys that are available at the U.S. Census Bureau for researchers to use after signing a strict confidentiality agreement. With the availability of the Census Bureau's own administrative records, Brummet et al. (2018) suggested that those administrative records "have the potential to improve the accuracy and quality of statistics produced from the CE." Steinberg et al. (2020) found that replacing the current ZIP-code income from the IRS with the HH Adjusted Gross Income (AGI) also from the IRS had little effect on the 2014 estimates of mean total expenditures for the Interview Survey.

The research presented in this paper is a follow-up to Steinberg et al 2020 (which used 2014 data) using 2015 and 2016 CE Interview Survey results. Using the 2015-2016 data and the same metrics, the results were very comparable to those in Steinberg et al. 2020, which had little effect on the estimates of the mean total expenditures for the Interview Survey.

2. Background and Literature Review

Within the calculation of the Interview Survey sample weights, CE adjusts the weights of the respondents to account for the nonrespondents. This noninterview adjustment uses the traditional cell adjustment method where all the consumer units in the sample are partitioned into 192 cells using variables based on a few demographic characteristics[3]. The weights of the respondents in a cell are then increased to account for the nonrespondents by multiplying them by an adjustment factor equal to the inverse of the cell's weighted response rate.

The variables used by a nonresponse adjustment are critical to its effectiveness. Having good nonresponse variables can reduce nonresponse bias, which can be large when the respondents and nonrespondents have different characteristics. A successful nonresponse adjustment depends on having variables that are available for both respondents and nonrespondents and, as suggested by Vartivarian and Little (2002a, 2002b), are correlated with the propensity to respond, the estimate of interest, or both. Finding good variables is difficult due to limited information on nonrespondents[4].

Prior to 2014, the nonresponse adjustment for the CE Interview Survey used the following variables: region of country (Northeast/South/Midwest/West), Consumer Unit size (1, 2, 3-4, and 5+ persons), housing tenure (owner/renter), race (Black/Non-Black), and rotation group (1-4) which identifies the quarter of the year where the CU is first introduced. The variables are summarized in Table 1.

[3] Steinberg, B. (2016). U.S. Bureau of Labor Statistics requirement "Phase 2 Interview Production Requirement – Weighting, 2017 Release, version 1.0," dated October 22, 2016.

[4] While characteristics such as age and education are not available for nonrespondents, others, such as region of country are available, as the interviewer has information about the address of the nonresponding consumer unit(s) within the CU at that address, but not information about the persons residing therein.

Table 1: Nonresponse Adjustment Variables Used in the CE Interview Survey

	1986 – 2013	2014 – present
1	Region of country	Region of country
2	Consumer Unit size	Consumer Unit size
3	Housing tenure	Number of contact attempts
4	Race	IRS Adjusted Gross Income (AGI) by ZIP-Code
5	Rotation Group	

Sabelhaus et al. (2013) examined CE's response rates by income. Taken from their paper, Figure 1 shows their main result.

CE had approximately 6,300 zip codes in the 2006-2010 sample and each CU in their zip code had an IRS publicly available adjusted gross income (AGI) associated with it. Each CU with its AGI was sorted from lowest AGI to highest AGI and the percentiles were created after sorting. Response rates associated with each of these percentiles were then calculated and plotted on Figure 1 and appear as 100 dots. The vertical axis refers to the response rate and the horizontal axis the income percentile from 1 to 100.

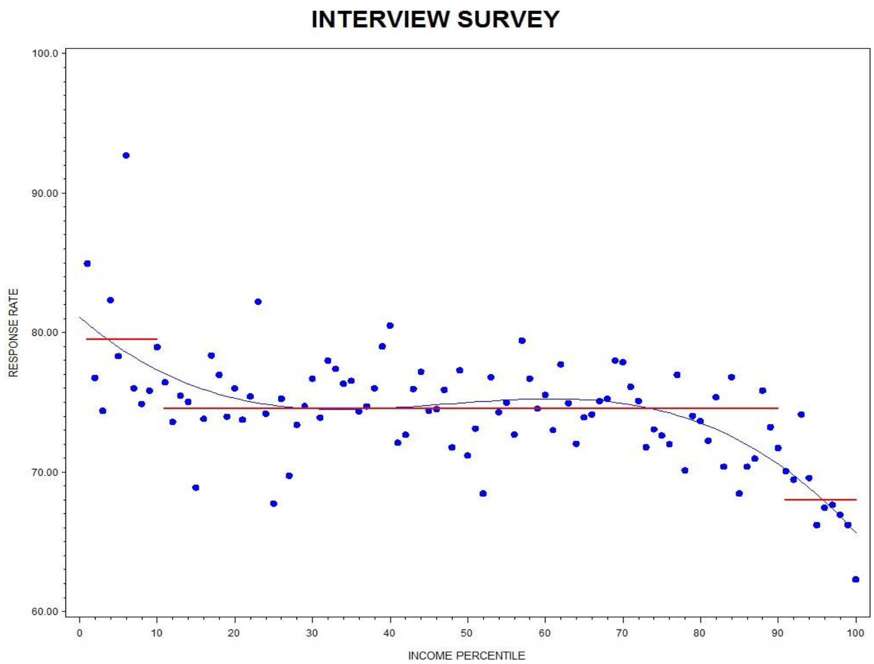

Figure 1. Response Rates for CE Interview Survey by Income Percentile: 2006-2010. From Sabelhaus, Johnson, Ash, Swanson, Garner, and Henderson. (2013). "*Is the Consumer Expenditure Survey Representative by Income?*," (NBER Working Paper No. 19589, October 2013).

- In Figure 1, we see that response rates differ by income. HHs in higher income ZIP-codes had lower response rates, and HHs in lower income ZIP-codes had higher response rates. In general, Sabelhaus et al. (2013) drew three main conclusions. First, high-income HHs were under-represented and low-income households were over-represented in the CE Survey. Second, high-income HHs under-reported both their incomes and their expenditures. Third, CE's noninterview adjustment should have accounted for the different response propensities by income, but they did not.

As a result of preliminary work eventually published as Sabelhaus et al. (2013), Dumbacher et al. (2012) conducted research on the variables used in the noninterview adjustment.[5] Their research suggested keeping some of the old variables but replacing others, including income. Table 1 above summarizes the variables before and after the change in 2014.

With the traditional cell adjustment method, all of the variables used in the noninterview adjustment need to be categorical. Therefore, continuous variables, like income, must be made categorical in some way. For the IRS ZIP-code AGI variable, based on the distinct horizontal lines representing response rates in Figure 1 from Sabelhaus et al. (2013), CE decided to classify three groups for the IRS Zip-code AGI variable: top 10%, middle 80%, and bottom 10%. The three categories are based on the horizontal lines in Figure 1 from Sabelhaus et al. (2013). For brevity, we refer to the categorical variable defined by IRS AGI by ZIP-code as simply *ZIP-code income* for the remainder of the paper.

The Census Bureau's administrative records are a valuable source of information for the CE because it has IRS HH-level income. Brummet et al. (2018) examined the use of these administrative records and their relationship to response rates. One of their conclusions conveys evidence of systematic survey non-response across the income distribution by linking responding and non-responding households to administrative records by address. Non-responding sample units were much more likely to be richer than responding households, measured either by wage and salary income or by broader AGI. Furthermore, nonresponse rates were higher at the top of the income distribution, and high income non-responding sample units were substantially richer even than high income responding sample units. This indicates that very high-income individuals were systematically less likely to respond to the CE. Given their conclusions, Brummet et al. (2018, p. 45) suggested the following:

> These results suggest potential ways forward to improve income statistics produced from the CE. In particular, both the PIK and MAFID based analysis suggest areas where current imputation and weighting procedures may be producing less than optimal outcomes. Incorporating the linkage of administrative records into these production processes has the potential to improve the accuracy and quality of statistics produced from the CE.

Krieger et al. (2019) also compared the set of nonresponse variables from 2013 and prior to those for 2014 and forward and showed the newer variables (which included ZIP-code income) made some improvement in stratifying the sample HHs by response propensity, but they did not noticeably affect estimates of mean total expenditures.

Steinberg et al. (2020) defined a new variable based on IRS HH income from the Census Bureau's administrative records. This new variable was constructed with the same three categories that define ZIP-code income but uses IRS HH income in place of ZIP-code income. For brevity, we refer to the categorical variable defined by the nonpublic HH-level IRS AGI as *HH income*. Their research examined the effect of replacing ZIP-code income with HH income on the noninterview adjustment, on the final weights, and on the estimates

[5] Sabelhaus et al. presented this work in December, 2011 at the CRIW/NBER Conference before the work was formally published. The paper, Is the Consumer Expenditure Survey Representative by Income? was formally presented at the National Bureau of Economic Research in October 2013.

of mean total expenditures for the 2014 Interview Survey. They found that "the increase in expenditure estimates were not significantly higher."

This research is a follow-up to Steinberg et al. (2020) in which a new variable was defined based on the IRS HH income from the Census Bureau's administrative records. This new variable was constructed as done in Steinberg et al. (2020). After deriving HH income and incorporating it into the noninterview adjustment, a similar analysis as in Steinberg et al. (2020), was then conducted which included:

- 1) How much information is lost by using ZIP-code income rather than HH income? This is discussed in section 5.

- 2) How much do the noninterview adjustment factors and the final weights change by replacing ZIP-code income with HH income? This is discussed in section 6.; and

- 3) How much do the estimates of mean total expenditures change by replacing ZIP-code income with HH income? This was our main objective and is discussed in section 7.

Lastly, we mention that other authors, including Garner (2009) and Sabelhaus et al. (2013), suggest that CE's estimates of mean total expenditures are under-reported, so results that show an overall increase in mean expenditures could be considered an improvement in data quality.

3. Data Description

This paper uses CE Interview Survey data collected from February 2015 through December 2016. (January 2015 was the transition month between data collected under the 2000 and 2010 sample designs, while the next months described were all collected under the 2010 design.) The data include both respondents (households that completed interviews) and nonrespondents (households that could not be contacted or refused to give interviews). The data exclude residential addresses that are not occupied and nonresidential addresses. For 2015, the data include information from 36,692 eligible cases, which consist of 23,574 respondents and 13,118 nonrespondents. For 2016, the data include information from 40,375 eligible cases, which consist of 25,441 respondents and 14,934 nonrespondents. CE's sample of addresses was originally drawn from the Census Bureau's Master Address File (MAF), which is the Census Bureau's official list of all residential addresses in the U.S. Therefore, every address in CE's sample has a MAF identification number (a MAFID) linking it back to the MAF. Then the addresses on the MAF are linked to the IRS's administrative records using the same process that Brummett et al. (2018) used, a process that is described in the next section.

4. Data Linkage

As is shown by Brummett et al. (2018; pp. 4-8), record linkage between CE's data and IRS's data was also done in a two-step process. The first step linked CE's households to the Master Address File by their MAFIDs that linked CE's households to their addresses. Then, the second step linked those addresses to the IRS's dataset by the addresses on the two files.

The process of linking addresses on the two files was accomplished by first running them through two different software packages designed to clean up the addresses and put them in a standard format that matched addresses on the U.S. Postal Service's Delivery Sequence File. One was a commercially available software package called "SAS® Dataflux," and the other was a software package developed by the Census Bureau's Geographic Division to accomplish the same task of cleaning addresses. Both software packages corrected minor misspellings; standardized words and abbreviations like "Road" and "Rd." or "Street" and "St."; appended the ZIP-code's 4-digit extension if needed; and so forth. Then after the addresses on the two files were cleaned up and put in the standard format of the U.S. Postal Service's Delivery Sequence File, the two files were merged. The resulting data showed nearly 70 percent of CE households were successfully linked to at least one tax return with an adjusted gross income (AGI).

After linking the data, the software addressed other problems to ready the data for use. This included fixing problems like duplicate tax years for an address and other data anomalies, before deeming it to be ready for use. For example, sometimes an address had more than one tax return. When that happened, the AGIs from all the tax returns were summed together (up to a maximum of six) to represent the address's total AGI. Addresses with more than six tax returns were treated as non-linkages since they were suspected of being apartment buildings or addresses with data linkage problems, and not individual households.

Of course, a linkage rate of 70 percent means 30 percent of the households still needed to have an income value imputed for them. The imputation was done by assigning them the average AGI for their ZIP-code using IRS's publicly available data, as is done in CE's production weighting process.

5. Comparing ZIP-Code Income with HH Income

The first objective of the research was to determine how much information is lost by using ZIP-code income rather than HH income. Tables 2 and 3 compare the household counts by the three income groups of HH and ZIP-code income for 2015 and 2016, respectively.

Table 2. Comparison of 2015 Household and ZIP-Code Income

		Household Income			Total
		Bottom 10%	Middle 80%	Top 10%	
ZIP-Code Income	Bottom 10%	1,450	2,250	40	3,740
	Middle 80%	2,080	25,300	1,900	29,280
	Top 10%	150	1,840	1,750	3,740
	Total	3,680	29,390	3,690	36,760

Table 3. Comparison of 2016 Household and ZIP-Code Income

		Household Income			Total
		Bottom 10%	Middle 80%	Top 10%	
ZIP-Code Income	Bottom 10%	1,550	2,480	40	4,070
	Middle 80%	2,300	27,850	2,200	32,350
	Top 10%	200	2,130	1,700	4,030
	Total	4,050	32,460	3,940	40,450

Tables 2 and 3 show the ZIP-code income and the HH income agreed for 77.5 percent (= (1,450 + 25,300 + 1,750)/36,760) of the HHs in 2015. Similarly, they agreed for 76.9 percent (= (1,550 + 27,850 + 1,700)/40,450) of the HHs in 2016. We suggest that this is a fairly high percentage of agreement, which means that using ZIP-code income does not result in the loss of too much information, since over 75 percent of the sample units would have been put into the same nonresponse cell using either HH or ZIP-code income.

In line with standard procedure, the counts in Tables 2 and 3 are rounded in accordance with Census Bureau Disclosure Review Board guidelines.[6] Because of rounding, for 2015 data, there are 36,760 households in this table as opposed to 36,692 stated earlier, and for 2016 data, there are 40,450 households in this table as opposed to 40,375 stated earlier.

6. Impact on Weights

The second objective of the research was to determine how much the nonresponse adjustment factors and final weights would change at the CU level. If the nonresponse factors and final weights do not change from the production values, then we would expect the estimates of expenditures not to change substantially either.

After replacing the ZIP-code income with the HH income, the effect on the nonresponse adjustment factors and the final weights was substantial. In 2015, the average nonresponse adjustment factor was 1.57, which corresponded to a response rate of 63.7 percent (1.57=1/0.637). After replacing the ZIP-code income with the HH income, the average nonresponse adjustment factor remained at 1.57. However, the values for the individual CUs in the sample changed substantially. Half of their values increased, and half of their values decreased, with the average absolute change in their values being 0.13. Any change over 5 percent of the average nonresponse adjustment factor (i.e., anything over 0.08=1.57x0.05) was considered substantial by the authors.

Likewise, in 2016, the average nonresponse adjustment factor was 1.59, which corresponded to a response rate of 62.9 percent (1.59=1/0.629). Again, after replacing the ZIP-code income with the HH income, the average nonresponse adjustment factor remained at 1.59. However, the values for the individual CUs in the sample changed substantially. Half of their values increased, and half of their values decreased, with the average absolute change in their values being 0.13. Again, that change was considered substantial by the authors.

CE's "final weights" are defined as the survey's base weights, multiplied by a nonresponse adjustment factor, and then a calibration adjustment factor. That is, CE "final weights" (FW) are defined as: FW = BW x NAF x CAF, where BW are the survey base weights; NAF is a nonadjustment factor; and CAF is a calibration adjustment factor. Calibration adjustment factors are applied to make the estimates from CE consistent with population totals for certain demographic estimates as estimated from the Current Population Survey.

For our research, the final weights were computed two different ways. The first approach used the base weights provided by the Census Bureau and the nonresponse adjustment factors that were used by BLS in production (those based on ZIP-code level income categories). The second approach used the base weights provided by the Census Bureau

[6] The Census Bureau Disclosure Review Board (DRB) guidelines were not provided but they have a system of rounding guidelines that performed these calculations.

and the nonresponse adjustment factors that were customized for our research (those based on HH level income categories). Then both sets of nonresponse-adjusted base weights were sent through the same calibration algorithm described above to produce two sets of final weights.

After producing the final weights with the HH income, the average final weight in 2015 was 19,955. That means each respondent household in the survey represented 19,955 households in the U.S. population, itself plus 19,954 other households that were not selected for the survey or were selected for the survey but did not participate in it. The average absolute value of the change in the final weights, between those produced with ZIP-code level income categories and those produced with HH level income categories, was 3,181, which is 15.9 percent of the average final weight. It was more than a difference of simple rounding. For 2016, the average final weight was 20,354. The absolute value of the change was 1,634, which is 8.0 percent of the average final weight. These values were less than those in 2015, but more than simple rounding.

7. Impact on Estimates of Expenditures

The third and most important objective of our research was to examine the effect of IRS HH-level data on expenditure estimates. In general, the expenditure estimates using IRS ZIP-code income and IRS HH income were not much different. For example, as Table 4 below shows, for the "All Items" expenditure category the Interview Survey estimates were only 0.48 percent larger for 2015 and 0.19 percent larger for 2016. Table 4 also compares some of the larger summary variables for the 2015 and 2016 estimates of mean expenditures derived from ZIP-code income and HH income.

Table 4. Comparison of 2015 and 2016 National Estimates of Mean Expenditures

Year	Expenditure Category	The Nonresponse Adjustment Used…				Difference	Percentage Difference (%)
		ZIP-Code Income		HH Income			
		Estimate	Standard Error	Estimate	Standard Error		
2015	All items	$52,560	543	$52,810	$549	$250	0.48
	Housing	17,360	194	17,410	199	50	0.29
	Transportation	9,332	211	9,327	214	-5	-0.05
	Food	7,851	48	7,845	47	-6	-0.08
	Insurance	5,590	116	5,631	116	41	0.73
	Health	4,124	54	4,132	52	8	0.19
	Apparel	1,183	33	1,228	49	45	3.80
	Education	1,153	78	1,209	84	56	4.86
	Income before taxes	57,630	912	57,880	885	250	0.43
2016	All items	53,150	587	53,250	599	100	0.19
	Housing	17,640	167	17,660	167	20	0.11
	Transportation	8,802	196	8,834	198	32	0.36
	Food	7,969	55	7,972	55	3	0.04
	Insurance	5,974	156	5,984	157	10	0.17
	Health	4,367	76	4,358	75	-9	-0.21
	Apparel	1,239	41	1,246	45	7	0.56
	Education	1,099	67	1,104	68	5	0.45
	Income before taxes	62,780	1,547	62,810	1,547	30	0.05

We first note that the differences for the estimated mean expenditures are small. The differences are all less than one percent of the expenditure estimates that were computed with ZIP-code incomes. For example, for the “All Items” category, the percentage differences were 0.48 percent and 0.19 percent for 2015 and 2016, respectively. The percentage differences for the other expenditure categories are also small. They are all less than 1%, and most of them are less than 0.50 percent. However, there were a couple of smaller summary variables that showed slightly higher changes in expenditures, such as “Apparel” and “Education” in 2015. It showed an increase of 3.80 percent and 4.86 percent respectively due to using IRS HH income.

Additionally, the differences between the two methods were all smaller than the standard errors of the estimates. Because the samples are not independent, it is not appropriate to conduct statistical tests of the differences between their means. However, we suggest that appreciable differences should be larger than the standard errors, and therefore we consider the differences in Table 4 to be small.

Lastly, the estimate of the mean expenditure for the “All Items” category increased slightly for both 2015 and 2016 when weighting was changed from ZIP-code income to HH income. However, we suggest this increase was unimportant and negligible because it was small relative to the size of the estimate of expenditure. The mean expenditure increased slightly for a few other categories as well, but again we suggest those increases were unimportant and negligible because they were also small relative to the size of their estimates of expenditures.

At the regional level, there were similar patterns for the average annual expenditures for the “All Items” category that were seen at the national level. As shown in Table 5, all four regions saw a small increase or decrease, but the differences were all less than 1%.

Table 5. Comparison of 2015 and 2016 Estimates of Mean Expenditures for the All Items Category by Census Region

Year	Census Region	The Nonresponse Adjustment Used…				Difference	Percentage Difference (%)
		ZIP-Code Income		HH Income			
		Estimate	Standard Error	Estimate	Standard Error		
2015	Northeast	$56,940	1,508	$57,440	1,595	$500	0.88
	Midwest	50,810	953	51,130	966	320	0.63
	South	49,080	894	49,300	908	220	0.45
	West	56,680	1,360	56,670	1,282	-10	-0.02
2016	Northeast	56,190	1,118	56,480	1,235	290	0.52
	Midwest	50,590	917	50,700	893	110	0.22
	South	49,070	1,384	49,190	1,399	120	0.24
	West	60,350	791	60,240	813	-110	-0.18

8. Conclusions

Overall, HH income did not make an appreciable change to the estimates of total expenditures compared to ZIP-code income. It made an appreciable change to the nonresponse adjustment factors and the final weights, but not to the estimates of total expenditures. Comparing the results between ZIP-code income and HH income, the nonresponse adjustment factors changed on average 8.3% in 2015, and 8.0% in 2016; and the final weights changed on average 15.9% in 2015, and 8.0% in 2016. However, the estimates of total expenditure changed only 0.48% in 2015, and only 0.19% in 2016. These results are consistent with Steinberg et al. (2020), who found that in 2014 the nonresponse adjustment factors changed on average 8.67%, the final weights changed on average 8.60%, but the estimate of total expenditure changed only 0.24%. These results show there was an appreciable change to both the nonresponse adjustment factors and the final weights, but not to the estimates of total expenditure.

Also consistent with the research of Steinberg et. al. (2020) was the agreement rate being greater than 70 percent in weighting group classifications between HH income and ZIP-code income for all three years (70.6 percent for 2014, 77.5 percent for 2015, and 76.9 percent for 2016). These high agreement rates exhibit uniformity between the two sources but are still a considerable amount below 100 percent since there is heterogeneity in household incomes within most zip-codes. Even though there was a high agreement rate, it was still enough to change the nonresponse factors and final weights substantially as mentioned above.

Our conclusions are consistent with Sabelhaus et al. (2013) and Brummet et al. (2018) who found an association between HH income and response. We agree that there is an association between HH income and response, but the relationship does not appear to be as strong as originally thought.

References

Brummet, Q., Flanagan-Doyle, D., Mitchel, J., Voorheis, J., Erhard, L., and McBride, B. (2018). Investigating the Use of Administrative Records in the Consumer Expenditures Survey, CARRA Working Paper 2018-01.
https://www.census.gov/content/dam/Census/library/working-papers/2018/adrm/carra-wp-2018.pdf

Dumbacher, B., Swanson, D., Reyes-Morales, S., Steinberg, B., Ash, S., and Castelo, D. (2012). U.S. Census Bureau memorandum entitled "Consumer Expenditure Surveys 2010 Redesign Research: Noninterview and Calibration Weighting Adjustments for the Consumer Expenditure Surveys," from Ruth Ann Killion to Jay Ryan, dated September 26, 2012.

Garner, T.I., McClelland, R., and Passero, W. (2009). Strengths and Weaknesses of the Consumer Expenditure Survey from a BLS Perspective. Presentation at the National Bureau of Economic Research, Summer Institute, Conference on Research on Income and Wealth. Bureau of Labor Statistics, U.S. Department of Labor. Washington, DC.
https://www.bls.gov/cex/research_papers/pdf/strengths-and-weaknesses-of-the-consumer-expenditure-survey-from-a-bls-perspective.pdf

Krieger. S., McBride, B., and Wilson, T. (2019) Bureau of Labor Statistics memorandum entitled "Analysis Item #2 – Measuring the Effectiveness of the New Variables Used in Nonresponse Weighting," from Sharon L. Krieger, Brett McBride, and Taylor J. Wilson to Laura Erhard, dated September 23, 2019.

Sabelhaus J., Johnson, D., Ash, S., Swanson, D., Garner, T., and Henderson, S. (2013). "*Is the Consumer Expenditure Survey Representative by Income?*" (NBER Working Paper No. 19589, October 2013).
https://www.nber.org/system/files/working_papers/w19589/w19589.pdf

Steinberg, B., Voorheis, J., Reyes-Morales, S., and McBride, B. (2014). "Using Administrative Data to Improve Nonresponse Weighting Procedures in the Consumer Expenditure Survey," JSM, Section on Government Statistics.
https://www.bls.gov/osmr/research-papers/2020/st200050.htm

Vartivarian, S. and Little, R.J.A. (2002a). On the formation of weighting adjustment cells for unit nonresponse. Proceedings of the Survey Research Methods Section, American Statistical Association.
http://www.asasrms.org/Proceedings/y2002/Files/JSM2002-000178.pdf

Vartivarian, S. and Little, R.J.A. (2002b). Does Weighting for Nonresponse Increase the Variance of Survey Means? Survey Methodology, 31, 161-168.
https://www150.statcan.gc.ca/n1/pub/12-001-x/2005002/article/9046-eng.pdf