

Population Subgroup Price Indexes:

Evidence of Heterogeneity or Measurement Error?

by

Robert Cage, Joshua Klick, William Johnson

Office of Prices and Living Conditions

Bureau of Labor Statistics | U.S. Department of Labor

Traditionally, national consumer price index (CPI) statistics are designed and engineered to represent a single population target: the noninstitutionalized resident civilian population of the nation. To that end, aggregation weights typically represent the experience of an average consumer or representative household. In reality, household budget shares can vary widely depending on the socio-economic characteristics of the household members, and the inflation experience of a given consumer cohort could vary far from the average. This results in some CPI users criticizing the official statistic as flawed, especially if the user's perceived personal experience with inflation differs from the published number. To address this issue, statistical agencies can produce cohort indexes tailored toward the specific indexation needs of various users or groups. However, this approach is encumbered by methodological challenges and practical limitations including available funding.

Currently, the U.S. Bureau of Labor Statistics (BLS) produces official consumer price indexes for two population cohorts: all urban consumers (CPI-U), and urban wage-earners and clerical workers (CPI-W). In addition, an experimental index is produced for a third cohort: urban elderly consumers (CPI-E). A storied legislative history, examined in Section 1 below, explains how the three indexes came about. Of note, the CPI-E was created to potentially address shortcomings of the CPI-U and the CPI-W for one specific indexation need: the annual cost-of-living adjustment (COLA) made to Social Security benefits by the Social Security Administration (SSA). Debate among policymakers and analysts over which index would be most appropriate for the COLA started in the late 1970s and amplified when a fourth index was added to the fray in 2002. That year, BLS began publishing a second index for the all-urban consumer cohort, the Chained CPI-U, that some argue is the most appropriate index for the COLA.

This paper highlights the practical issues and methodological problems BLS faces in estimating indexes for population subgroups, using the Social Security COLA indexation need as a focus. Public arguments for and against the CPI-E are summarized in Section 2. Methodological issues impacting the estimation of cohort indexes are summarized in Section 3. A model-based approach to constructing cohort indexes, within the confines of producing a national index and constrained by current survey data limitations, is

proposed and outlined in Section 4 as a potential alternative to the current practice of varying expenditure weights only. Results are summarized in Section 5.

1. Historical Perspective

A review of seminal moments in BLS's history reveals why it currently produces CPIs for three different cohorts. The origins of the modern day U.S. CPI date back to World War I when rapid increases in prices created a demand for the statistic for use in employment contract negotiations.¹ To provide weighting patterns for the index, surveys of family expenditures were conducted in 1917 to 1919 exclusively in major industrial centers. The initial survey coverage was limited to married couples with at least one child, where the principal wage-earner in the family accounted for 75% or more of total family income. Periodic collection of prices started in 1919 and regular publication of a national index, commenced in 1921 with indexes estimated back to 1913. One of the first uses of the index was in cost-of-living escalator clauses built into the employment contracts of manual laborers, particularly shipbuilders working in industrial centers to support the war effort.

Over the next 50 plus years, the BLS produced the national CPI by tracking the buying habits of this specific, narrowly defined consumer cohort: families of manual laborers. During this time, however, the definition of the cohort and the geographic coverage expanded in scope.² By 1964, the target population had evolved to include all wage-earner and clerical worker households living in metropolitan areas of the country. To be considered a *wage-earner* household, at least one member of the household had to be employed full-time (i.e., 37 weeks out of the year) in an eligible occupation. Eligible occupation groups accounted for roughly 70% of the total U.S. workforce at that time, but excluded professionals, managers, and the self-employed. BLS's definition of *wage-earner and clerical worker* for the CPI has essentially remained unchanged since 1964.

1.1 Creation of the CPI-U

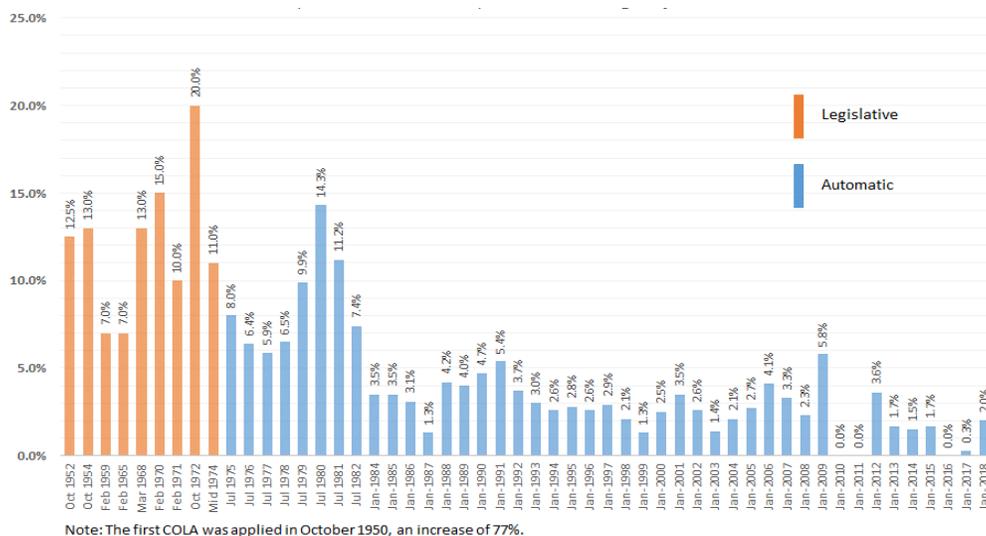
Then, in 1972, the passage of Public Law 92-336 changed the course of BLS's CPI history. This legislation, primarily aimed at increasing the national debt limit, contained an amendment to the Social Security Act of 1935. This 1972 amendment established an automatic annual cost-of-living adjustment (COLA) to Social Security benefits, to ensure the purchasing power of the otherwise fixed benefits would not be

¹ For a history of the first 100 years of the CPI, see Rippey, Darren. [The first hundred years of the Consumer Price Index: a methodological and political history](https://www.bls.gov/opub/mlr/2014/article/pdf/the-first-hundred-years-of-the-consumer-price-index.pdf). *Monthly Labor Review*, Bureau of Labor Statistics. April 2014 at <https://www.bls.gov/opub/mlr/2014/article/pdf/the-first-hundred-years-of-the-consumer-price-index.pdf>.

² For a complete history of how the CPI target population evolved over time, see Exhibit 1 in Chapter 17 of the *BLS Handbook of Methods* at <https://www.bls.gov/opub/hom/pdf/homch17.pdf>.

eroded by the forces of inflation.³ The 1972 statute mandated the Social Security Administration (SSA) to calculate the COLA using the CPI produced and published by the BLS. The law went into effect shortly thereafter, and the first automatic COLA was applied in 1975, an increase of 8%.⁴ At that time, the only CPI published by the BLS was the index representing the *wage-earner and clerical worker* cohort. Hence, it was the index adopted by the SSA.

Figure 1. History of Social Security cost-of-living-adjustments



Meanwhile, the BLS was in the midst of a comprehensive overhaul and revision of the CPI program, partly in response to recommendations made by a committee commissioned by Congress to examine federal price indexes, commonly referred to as the Stigler Committee.⁵ BLS planned to expand the population scope of the CPI to ‘nearly all consumers’ in 1978, engineering it to represent not just the existing urban wage-earner cohort, but all residents of metropolitan areas, including professionals, the self-employed, the poor, the unemployed, and retired persons.⁶

³ Prior to 1975, cost-of-living adjustments to Social Security benefits required legislative action. Congress enacted public laws to increase benefits 10 times prior to the 1972 amendment taking effect: 1950, 1952, 1954, 1959, 1965, 1968, 1970, 1972, and 1974.

⁴ For history of COLA increases see Whitaker, Julie. *Social Security: Cost-of-Living Adjustments*. Congressional Research Service, October 2017. <https://fas.org/sgp/crs/misc/94-803.pdf>

⁵ See Stigler, George et. al. *The Price Statistics of the Federal Government. Report of the Committee before the Subcommittee on Economic Statistics of the Joint Economic Committee*, Congress of the United States. Sec 5(a) of Public Law 304, 79th Congress. January, 1961. Six recommendations included: (1) periodic weight updates at least once every 10 years; (2) probability sampling, (3) prompt introduction of new products, (4) funding for research, (5) extend population target to include single persons as well as families, and cover rural nonfarm as well as urban workers. Rural farm was excluded because a separate index produced by Department of Agriculture represented prices received and paid by farmers.

⁶ Improvements made to the CPI in the 1978 Revision included: (1) updated weights using 1972-73 Consumer Expenditure Survey, (2) updated sample of items, (3) updated sample of outlets –introducing probability sampling, and (5) expanded population coverage. See *The Consumer Price Index Revision-1978*, Monthly labor Review. US Department of Labor, Bureau of Labor Statistics, 1978.

In 1974, BLS initially announced its intention to replace the urban wage-earner definition of the CPI with the broader all urban consumer definition, and publish just one index.⁷ However, this plan was criticized by some members of Congress who feared the new index would, ‘no longer be firmly grounded in the experience of low- and middle-income workers.’⁸ As a result, Congress increased BLS’s budget to allow both indexes to be calculated and published for at least three years. BLS introduced the all-urban index (CPI-U) in 1978, and relabeled the existing index as the CPI-W. From 1978 to 1980, separate outlet and item samples were selected to track price change for the separate CPI-U and CPI-W indexes. In 1981, Reagan era budget cuts to the BLS put the CPI-W on the chopping block once again. BLS could no longer afford to maintain two distinct outlet and item samples for the two indexes and could have terminated calculation of the CPI-W. However, because little difference between the CPI-U and CPI-W was observed during the 1978-1980 period, BLS decided to continue the CPI-W series using the CPI-U sample of areas, outlets, items, and prices. Hence, in 1981 the CPI-W became a reweighted version of the CPI-U, with the component weights used to calculate the aggregate index adjusted to reflect the spending habits of the urban wage-earner cohort. Since 1981, the CPI-U and CPI-W have differed only in the expenditure weights assigned to item categories and geographic areas.⁹

Had the BLS decided to terminate the wage-earner index in 1978, or 1981, the SSA would have had no option but to convert the calculation of the Social Security COLA to the lone index produced by the BLS: the CPI-U. Absent a legislative mandate identifying the specific index to use, the introduction of the CPI-U coupled with the continuation of the CPI-W presented the SSA with a rulemaking conundrum: which index was more appropriate for the stated purpose of the COLA?

In the late 1970s, Congress created a National Commission on Social Security to undertake ‘fundamental, long-term, comprehensive consideration for changes to the Social Security system.’¹⁰ The law establishing this commission expressly directed it to investigate the COLA calculation issue and which index would be most appropriate for said use. Notably, the directive called for an assessment of whether a special index for the *elderly* was most suitable. The report of the committee, issued in March 1981, recommended against the use of a CPI tailored toward any beneficiary group, such as the retired

⁷ See Reed, Stephen and Kenneth Stewart. *Why does BLS provide both the CPI-W and the CPI-U?* Beyond the Numbers. BLS. February 2014. <https://www.bls.gov/opub/btn/volume-3/why-does-bls-provide-both-the-cpi-w-and-cpi-u.htm>.

⁸ Joseph P. Goldberg and William T. Moyer, *First Hundred Years of the Bureau of Labor Statistics*, Bulletin 2235 (U.S. Bureau of Labor Statistics, September 1985).

⁹ Between 1983 and 1985, the indexes also differed in the treatment of the shelter component of the index. A rental equivalence approach was adopted in the CPI-U in 1983, and in 1985 for the CPI-W. See <https://www.bls.gov/opub/btn/volume-2/pdf/owners-equivalent-rent-and-the-consumer-price-index-30-years-and-counting.pdf> for more information on why this change was made.

¹⁰ See <https://www.ssa.gov/history/reports/80commission.html>.

elderly, primarily due to a lack of analytical data and research on how to properly construct such an index. The committee did, however, recommend use of the CPI-U over the CPI-W for the Social Security COLA, primarily because the CPI-U included retired persons in its population target while the CPI-W did not.

Neither Congress nor the SSA acted on the recommendation. Inaction was likely due, in part, to the financial crisis enveloping the Social Security program at the time. By the late 1970s, the Old-Age and Survivors Insurance Trust Fund (OASDI) was projected to become insolvent as early as 1983. A second National Commission on Social Security Reform was created by Congress in 1981, commonly referred to as the Greenspan Commission, to solve the financial problems facing Social Security.¹¹ The Greenspan Commission report, issued in 1983, *inter alia* recommended against the use of the CPI for automatic benefit increases, preferring instead to base them on changes in average wages, when the latter was lower.¹²

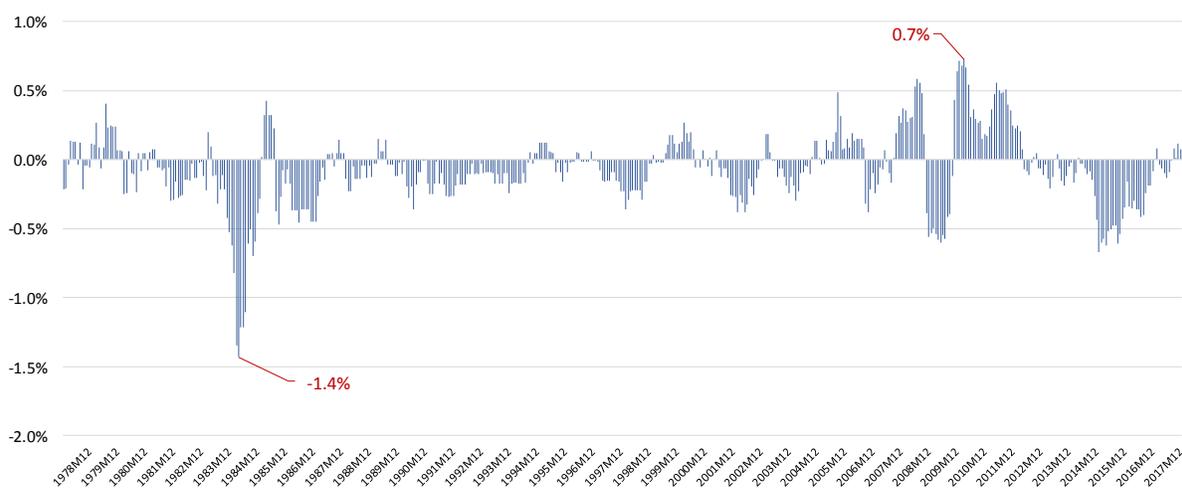
So within the span of two years, two separate national commissions had issued two conflicting opinions on the appropriate indexation method for Social Security COLAs. One, in effect, favored an index representing the inflation experience of a retired cohort, and the other favored an index representing the experience of workers.

Complicating matters was the double-digit levels of inflation experienced by the American economy at the beginning of the decade. In early 1983, the CPI-U was yielding higher estimates of inflation compared to the CPI-W, a gap which would turn out to be the largest difference in the history of the two series to date (see Figure 2). Perhaps fearing a switch from the CPI-W to the CPI-U would hasten the woes of the OASDI, both Congress and SSA remained idle. Automatic Social Security COLAs have been tied to the CPI-W ever since, while many other COLAs cited in federal legislation are tied to the CPI-U or Chained CPI-U.¹³ Public debate over the appropriate index for COLA indexation has not only endured, but has arguably amplified over the past 30 years.

¹¹ See <https://www.ssa.gov/history/reports/gspan.html>. Alan Greenspan was the chair of National Commission on Social Security Reform.

¹² See Greenspan Commission report, recommendation 15.

¹³ In 2017, legislation was enacted to index federal income tax brackets to changes in the Chained CPI-U. See Public Law 115-97, The Tax Cuts and Jobs Act of 2017.

Figure 2. Difference in annual inflation estimates, CPI-W minus CPI-U

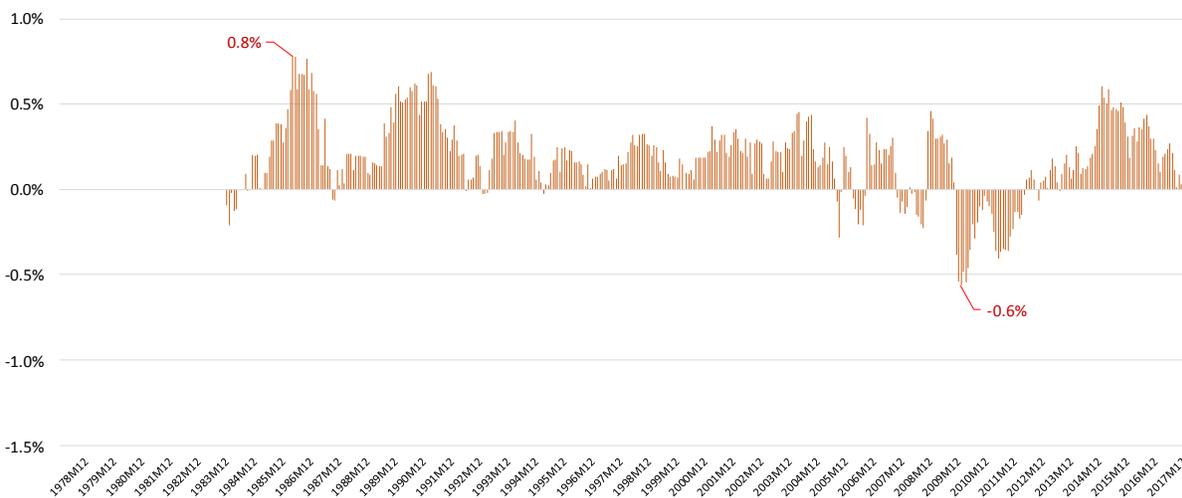
1.2 Creation of the CPI-E

With over 67 million beneficiaries, Social Security is by far the largest U.S. government outlay directly adjusted using a CPI. Hence, SSA's use of the CPI-W instead of the CPI-U has remained a topic of discussion since the initial debates of the early 1980s. The primary criticism of the CPI-W is its cohort definition: the index is calculated using expenditure weights of wage-earners. This population is employed and the spending habits are reflective of people who work. In contrast, most Social Security recipients are retired and not in the labor force.

Suitability of the CPI-W came into Congress's crosshairs again in 1987, when it directed the BLS to develop an index reflecting the spending patterns of seniors, so that it could be evaluated by SSA as a possible replacement for the CPI-W for Social Security COLAs. However, this Congressional mandate came without any additional funding. In turn, BLS released the experimental CPI for Americans 62 Years of Age and Older (CPI-E) in 1988, with estimates calculated back to 1982. The index was labeled experimental because, like the CPI-W, it differed from the CPI-U only in the expenditure weights used to aggregate component indexes together. The index was not constructed by selecting a dedicated sample of outlets and items based on the spending choices of the elderly, and it did not track the prices paid exclusively by the cohort, which would not have been possible without additional funding. For these reasons, and because the sample size used to estimate the weights for the elderly cohort were smaller than the wage-earner cohort, the BLS deemed the CPI-E less accurate than the CPI-W, reinforcing its experimental moniker.

BLS continues to calculate the experimental CPI-E and release it upon request. Historically, the index has yielded annual inflation estimates about 0.2 percent higher than the CPI-U, and 0.3 percent higher than the CPI-W (see Figure 3). Had the SSA used the CPI-E instead of the CPI-W for the most recent COLA, the average recipient would have received an additional \$1 in average monthly benefits in 2018. This seemingly *de minimis* difference would equate to an additional \$1.2 billion in outlays from the OASDI trust fund. Over the years, bills have been periodically proposed in Congress to mandate the use of the CPI-E or a similar index product targeting the elderly cohort for Social Security COLA indexation. In 2017 alone, as many as five different bills were sponsored but to date none have passed into law.¹⁴

Figure 3. Difference in annual inflation estimates, CPI-E minus CPI-U



1.3 Creation of the Chained CPI-U

In the early 1990s, concerns over potential biases in the CPI and subsequent impacts on the Social Security COLAs and financial solvency of the OASDI Trust Fund, led Congress to once again appoint a commission to study the CPI. The Advisory Commission to Study the Consumer Price Index, commonly referred to as the Boskin Commission, issued its report in 1996 which concluded the CPI overstated inflation by about 1.1 percentage points per year.¹⁵ The Boskin Commission summarily identified four specific sources of upward bias in the CPI aggregating to a plausible range of 0.80 to 1.60 percent per annum:

¹⁴ See <http://www.ncpssm.org/PublicPolicy/SocialSecurity/Documents/ArticleID/1159/The-CPI-E-%E2%80%93-A-Better-Option-for-Calculating-Social-Security-COLAs>.

¹⁵ See Boskin, Michael, et al., *Toward A More Accurate measure of the Cost of Living*, December 4, 1996. Michael Boskin was the chair of the Advisory Commission.

Figure 4. *Estimated biases in the CPI, Boskin Commission, 1996*

| | |
|-------------------------------|------|
| Upper Level Substitution Bias | 0.15 |
| Lower level Substitution Bias | 0.25 |
| New Products/Quality Change | 0.60 |
| New Outlets | 0.10 |

Among the CPI market basket categories, the Commission concluded the quality change bias was most problematic for medical care, estimating an upward bias in prescription drug, hospital service, and physician service indexes as high as 3 percent per annum.¹⁶

Meanwhile, BLS was already considering measures to improve the handling of some of these biases. To mitigate lower level substitution bias, BLS adopted a hybrid approach to estimating lower-level or component indexes in 1999, whereby the Geometric Mean index formula was adopted for the majority of market basket strata.¹⁷ This index assumes consumers respond to relative price change by holding the budget shares of the specific items purchased constant over time (i.e., a unitary elasticity of substitution). The Laspeyres index, which assumes consumers maximize utility by holding the quantity mix of specific items purchased constant over time (i.e., zero elasticity of substitution) was retained for a few categories.¹⁸ The CPI-U, CPI-W, and CPI-E each use this hybrid combination of Geometric Mean and Laspeyres indexes at the component level.

To address the upper level substitution bias, BLS began publication of the Chained CPI-U (C-CPI-U) in August 2002 with data going back to 2000. The C-CPI_U utilizes a Tornqvist index to aggregate component indexes. This index uses observed monthly expenditure changes made by consumers, ostensibly capturing marketplace substitutions brought about by changing relative prices.¹⁹ The formula was not – and has not been -- adopted for the flagship CPI-U index, because the requisite monthly expenditure data needed for index calculation become available with roughly a one year lag. Failing to meet BLS’s timeliness standard of a 2-week lag for the headline national inflation index, BLS decided to

¹⁶ See Boskin Commission, Table 2, pp. 39-40.

¹⁷ See Dalton, Kenneth, John Greenlees and Kenneth Stewart. *Incorporating a geometric mean formula into the CPI*. *Monthly Labor Review*, October 1998. <https://www.bls.gov/mlr/1998/10/art1full.pdf>.

¹⁸ Elementary items using the Laspeyres formula include rent, owner’s equivalent rent, electricity, natural gas, housing at school, water and sewerage maintenance, most, medical care services, and vehicle registration and licensing fees.

¹⁹ See Cage, Robert, John Greenlees and Patrick Jackman. *Introducing the Chained Consumer Price Index*. Meeting of the International Working Group on Price Indexes, Paris France, May 2003. <https://www.bls.gov/cpi/additional-resources/chained-cpi-introduction.pdf>.

release the Chained CPI as a new, separate index product. Preliminary versions are published on the same schedule as the CPI-U, but are subject to revision.²⁰

Over its history, the Chained CPI has yielded annual inflation estimates that are on average approximately 0.25 percentage points less than the CPI-U (see Figure 5). As a consequence, some observers, notably fiscal conservatives, argue it is most suitable for the Social Security COLA.²¹ As part of the 2012 and 2013 “fiscal cliff” negotiations, President Obama proposed switching the COLA to the Chained CPI-U as a way to address budgetary shortfalls.²² Congress did not adopt his proposal. But five years later, with the passage of the Tax Cuts and Jobs Act of 2017, Congress changed the indexation of federal income tax brackets from the CPI-U to the Chained CPI-U.²³ The new legislation, however, did not mandate this change for the Social Security COLAs which remain indexed to the CPI-W.

Figure 5. *Difference in annual inflation estimates, Chained CPI-U minus CPI-U*



2 Criticisms of population subgroup indexes and the current COLA debate

2.1 Population Targeted

Proponents of the CPI-E argue it is the more appropriate index for Social Security COLAs precisely because the cohort definition, elderly consumers, more properly aligns with the targeted group

²⁰ See Klick, Joshua. Improving initial estimates of the Chained Consumer Price index. Monthly Labor Review, February 2018. <https://www.bls.gov/opub/mlr/2018/article/improving-initial-estimates-of-the-chained-consumer-price-index.htm>.

²¹ See, e.g., Rosenberg, Adam and Marc Goldwein. Measuring Up: The Case for the Chained CPI. Moment of Truth Project, December 2012. <http://www.momentoftruthproject.org/sites/default/files/Measuring%20Up.pdf>.

²² The term fiscal cliff was popularized by Ben Bernanke former chairman of the US Federal Reserve, describing the point when Bush era tax cuts would expire and spending cuts would be triggered in 2013.

²³ See Public Law 115-97, The Tax Cuts and Jobs Act of 2017.

addressed by the adjustment. Research has shown that spending patterns differ between the elderly and the general population, especially in the medical care category. Seniors 65 and older spend more than twice on health care, and those 75 and older spend nearly three times more on health care than younger consumers.²⁴ Moreover, price change for medical care items have increased at a higher rate than other items in the CPI market basket. Critics of using the CPI-W argue its primary flaw is its failure to take these differences for the elderly population into account, and is therefore biased downward for the COLA purpose.²⁵

Figure 6. Relative importance for select CPI categories, 2013-2014

| | Category | CPI-U | CPI-W | CPI-E |
|----|----------------------------------|-------|-------|-------|
| A1 | Adult clothing | 1.7% | 1.6% | 1.3% |
| A2 | Children and infant clothing | 0.4% | 0.5% | 0.1% |
| E0 | Telephone and electronics | 6.0% | 6.9% | 5.5% |
| E1 | Tuition | 2.7% | 2.0% | 1.2% |
| F1 | Food at home | 10.9% | 12.6% | 9.9% |
| F2 | Food away from home | 5.1% | 5.1% | 4.4% |
| F3 | Alcohol and tobacco | 1.4% | 1.6% | 1.1% |
| H0 | Shelter | 32.1% | 30.6% | 35.5% |
| H1 | Household utilities | 5.5% | 6.0% | 5.7% |
| H2 | Housefurnishings and operations | 3.7% | 2.6% | 4.1% |
| M0 | Medical Care | 8.3% | 6.5% | 11.1% |
| R0 | Entertainment and recreation | 4.5% | 3.6% | 4.3% |
| T0 | Vehicles | 7.0% | 7.9% | 6.4% |
| T1 | Gasoline and vehicle maintenance | 9.5% | 11.4% | 8.0% |
| T2 | Public transportation | 1.3% | 0.9% | 1.3% |

Note: Estimates from CE Interview Survey only, research sample data.

Opponents of the CPI-E are just as numerous. Arguments against its use for indexing public transfer payments generally fall into the following categories:

1. Medical Care. Out-of-pocket medical care expenses account for the majority of the growth rate differences between the CPI-E and the other indexes. But, as the 1996 Boskin Commission hypothesized, failure to adequately capture quality change in the medical sector could bias these component indexes upward as much as 3 percent per year. Switching the COLA to CPI-E, where medical care has a higher weight, would exacerbate this perceived bias.
2. Substitution Bias. BLS does not currently produce a chained version of the CPI-E utilizing the Tornqvist version to mitigate upper level substitution bias. For this reason, some commentators

²⁴ See <http://www.ncpssm.org/PublicPolicy/SocialSecurity/Documents/ArticleID/1159/The-CPI-E-%E2%80%93-A-Better-Option-for-Calculating-Social-Security-COLAs>.

²⁵ *Ibid.*

argue the CPI-E is flawed due to its upper-level aggregation formula. Accordingly, they maintain the Chained CPI-U is a more accurate measure of overall inflation, making it a better candidate for the COLA.²⁶

3. Accuracy of Aggregation Weights. CPI-E aggregation weights are derived from smaller CE survey sample sizes, leading to larger sampling error in the weights and a less accurate index compared to the CPI-U and Chained CPI-U. Households classified as CPI-E comprised about 30% of the CE sample in 2015.
4. Outlet and Item Samples. The CPI-E fails to adequately incorporate outlets patronized by the elderly, nor track the prices paid by the elderly, and therefore the index falls short of accuracy standards. Due to known third degree price discrimination for some items, e.g., senior citizen discounts for food and cinema admissions, the CPI-E cohort faces different prices for the same goods and services to some extent. To the extent that inflationary change for the set of actual prices paid by the cohort for the subset of specific goods and services purchased at the subset of retail establishments differs from price change experienced in the overall economy, then the CPI-E is of unknown accuracy.
5. Cohort Definition. Some argue if the dedicated goal of the COLA is to protect the purchasing power of Social Security income, as opposed to other sources of income, then the adjustment should be narrowly defined to represent the spending of this income by Social Security recipients. This viewpoint essentially argues the current cohort definition of the CPI-E is ill-fitting. In 2007, Burdick and Fisher of the SSA argued the most practical objection to using the CPI-E for Social Security COLAs was its cohort definition.²⁷ They posited one-fifth of OASDI beneficiaries were under the age of 62, and therefore excluded from the CPI-E. Likewise, one-fifth of persons age 62 or older were not beneficiaries but were included in the index.²⁸ In effect, Burdick and Fisher suggested the targeted cohort of the COLA should be all OASDI recipients, regardless of age. Some argue the intent of the COLA is to guard against the erosive effects of *overall* inflation, and hence a national index measuring the macro economy is best suited, such as the CPI-U or Chained CPI-U. These commentators argue a switch to a cohort-

²⁶ See Rosenberg and Goldwein, *supra*.

²⁷ See Social Security Bulletin, Vol. 67, No. 3, 2007 at <https://www.ssa.gov/policy/docs/ssb/v67n3/v67n3p73.html>.

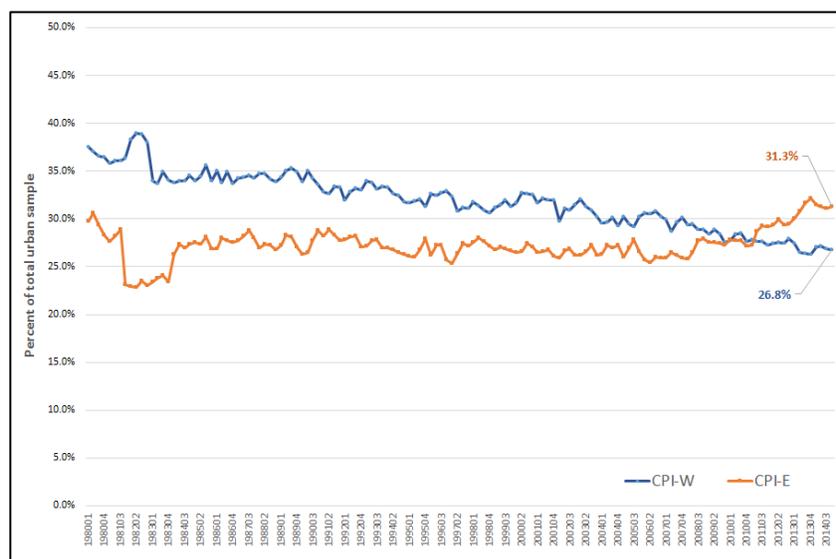
²⁸ In 2013-2014, approximately 7% of non CPI-E households in the Consumer Expenditure Interview Survey had the majority of their income sourced from Social Security, and 20% of CPI-E households had the majority of total household income from non-Social Security sources.

specific index targeting a population subgroup would constitute a public policy change, rather than a change necessitated to improve the accuracy of the COLA.²⁹

2.2 Accuracy of the CPI-W

Although largely overlooked by analysts, each of these potential criticisms levied at the CPI-E can also attach to the index actually used for the COLA: the CPI-W. Of particular concern is the declining sample size of the CPI-W in the CE household expenditure survey. Once constituting 100% of the sample prior to the expansion to urban population coverage in 1978, the CPI-W sample has been in decline. In fact, it has been smaller than the CPI-E sample since mid-2011. The CPI-W accounted for a little over a quarter of all survey respondents by the end of 2015 (see Figure 7). This phenomenon is partly explained by higher survey non-response rates among worker households compared to elderly households, but also indicative of a labor force transitioning into the modern digital economy. The out-of-date delineation of occupations eligible for the CPI-W cohort is shrinking as a percentage of the total workforce.

Figure 7. CPI-U and CPI-W sample size, as proportion of total sample Consumer Expenditure Interview Survey



Furthermore, BLS and SSA currently utilize different customized definitions of the *wage-earner* cohort. Historically, the term has been used pejoratively to describe the poor, underprivileged, and propertyless proletariat.³⁰ It has also been used to describe a class of workers, usually manual laborers, distinguished by method of recompense: paid by the hour earning a fixed hourly wage rate, rather than an annual

²⁹ See Rosenberg and Goldwein, *supra*.

³⁰ See George Bernard Shaw, *What Socialism Is*. The Fabian Society, 1890.

salary.³¹ Neither of these connotations, low-income and earning method, are technical criterion used in the current BLS nor SSA wage-earner definition.

The CPI-W cohort is currently defined as urban households with a majority of income sourced from earnings in a limited list of specific occupation types.³² This cohort represented roughly 29% of all U.S. households in 2015-2016, and only 40% of full-time labor force participant urban households. In contrast, the SSA defines wage-earner as an individual who earns Social Security credits while working for wages or for self-employment income.³³ Initially excluding the self-employed, federal and state government employees, and other classes of workers, Social Security eligibility has expanded since 1935 and over 90% of the U.S. labor force is now covered under the program.³⁴ Hence, there is a large gap between the CPI and SSA wage-earner definitions. If the targeted population of the index used for the Social Security COLA is the cohort of Social Security tax payers in the labor force, then the CPI-W represents only a fraction of this group.

Moreover, CPI-W households have lower incomes, on average, than non CPI-W households with earnings. Stores and retail establishments patronized by lower-income households, and the specific variety and brands of products purchased within CPI component categories, may be different from the all-urban population at large – the same concern associated with the CPI-E. Any bias associated with using the all-urban component indexes to construct the CPI-W likely increases as the cohort shrinks in size. Nonetheless, these aforementioned shortcomings of the CPI-W have yet to gain much attention in the public debate over the COLA.

3 Methodological Issues in Cohort Index Construction

3.1 The Schultze Advisory Panel

The difficulty of linking the actual prices observed on a monthly basis in the CPI, in the sampled set of retail establishments, to the characteristics of the consumers paying the prices, hampers the ability of statistical agencies to construct accurate indexes for population subgroups. This was a major finding of a National Research Council panel of the National Academy of Sciences sponsored by BLS to investigate conceptual, measurement, and other statistical issues in the development of cost-of-living indexes,

³¹ In most English dictionaries, wage-earner is defined as ‘a person who works for wages, especially as distinguished from one paid a salary.’ American English dictionaries typically omit the wage-salary distinction, defining wage-earner as ‘a person who works for wages or salary.’

³² Delineated occupations included in the CPI-W include construction workers, mechanics, forestry, fishing, grounds keeping, machine or transportation operators, protective service, retail sales workers, administrative support including clerical, and other service workers such as janitors and food industry workers.

³³ See <https://www.ssa.gov/agency/glossary/>.

³⁴ See <https://www.ssa.gov/policy/docs/ssb/v48n4/v48n4p33.pdf>.

commonly referred to as the Schultze Panel.³⁵ The panel’s report, issued in 2002, contained a comprehensive examination of population group indexes, from conceptual basis to practical estimation methods³⁶

Among the notable findings and recommendations of the Schultze Panel:

- Group indexes should account for two sources of heterogeneity: across stratum product categories and within stratum product categories; within stratum differences arise from idiosyncratic tastes and preferences, but also from differences associated with age, income, family composition, and geographic location. Most experimental group indexes control for the former, but there is a lack of evidence on the latter to enable adequate exploration.³⁷
- In the absence of an index that can capture differences in price or qualities of good purchased by a subgroup, there is no rationale for switching to an index (e.g., the CPI-E) for purposes of Social Security COLA indexation. Differences between the CPI-U and CPI-W are not large enough to warrant a change.³⁸
- If inflation rates differ among population cohorts, then use of an overall CPI to index Social Security COLAs may overcompensate some groups and undercompensate others, in a way most people would deem unfair or unjust.³⁹
- BLS should explore collecting prices in a way that allows them to be associated with household characteristics, and Congress should fund it.⁴⁰

The panel seemed to suggest the preeminent weakness in cohort indexes, such as the CPI-E and CPI-W, is their failure to incorporate cohort-specific outlet and item samples, and to track cohort-specific transaction prices. They went as far to proclaim, ‘the current BLS data collection system cannot produce group indexes.’⁴¹ Most consumer population subgroup indexes produced by national statistical agencies suffer from this problem.

Since the issuance of the Schultze and Mackie advisory report in 2002, research has begun to emerge using scanner and large transaction data, seeking evidence of within-category differences across cohorts. Diamond, Watanabe, Watanabe (2016) used individual level purchase and demographic household scanner data, linked to 33 million transactions, to decompose group indexes into weight and

³⁵ See Schultze, Charles and Christopher Mackie. *At What Price? Conceptualizing and Measuring Cost-of-Living and Price Indexes*. National Research Council, 2002. <https://www.nap.edu/read/10131/chapter/1>. Professor Schultze was the chair of the panel.

³⁶ Schultze, Chapter 8.

³⁷ Schultze, p. 224.

³⁸ Schultze, p. 198.

³⁹ Schultze, p. 226.

⁴⁰ Schultze, Finding 8.1.

⁴¹ Schultze, p. 226.

price effects in Japan— essentially controlling for heterogeneity across and within category groupings.⁴² They concluded price and weight effects do not vary across age group, except at the in the oldest group, but warn measures of variation are highly dependent on how the groups are defined.

Beyond the within product category heterogeneity issue, two additional conceptual problems plague the estimation of cohort indexes: treatment of owner-occupied housing and aggregation method. However, a consensus approach among researchers and practitioners appears to be emerging for these.

3.2 Treatment of Owner Occupied Housing

The definition of the weight for owner occupied housing is controversial. BLS switched from a user cost to rental equivalence approach for this component in 1983.⁴³ Most elderly consumers in the Consumer Expenditure Survey sample are homeowners, and the majority of these own their home without a mortgage, having disposed of the debt obligation earlier in life.⁴⁴ Hence, the current expenditure needed to maintain the home, or user cost, is limited to property taxes, insurance, utilities, and maintenance. Insurance, utilities, and household maintenance are explicitly captured in the CPI in different component categories. In effect, the rental equivalence approach discards the mortgage payment and property tax expense, in favor of an estimate of the market rental value of the home. In particular, this approach imputes a hypothetical expenditure typically larger than the carrying cost of the home, especially for homeowners not holding any mortgage debt. For these consumers, this lowers the relative importance of other CPI items – items actually purchased in the current period. When shelter inflation outpaces inflation for other market basket items, the rental equivalence approach may overstate the actual inflation of the elderly cohort.

While many economists agree the rental equivalence approach is appropriate for an index needed to gauge overall national inflation, many argue it may not be appropriate for indexation needs of targeted population subgroups. For example, Hobijn and Lagakos questioned its use for indexing a cash benefit program like OASDI because rental equivalence measures an *opportunity cost* rather than actual out-of-

⁴² See Diamond, Jess, Watanabe, Kota and Tsutomu Watanabe, Estimating Consumer Prices Inflation by Household at <http://www.price.e.u-tokyo.ac.jp/img/event/2015SWET-Jess.pdf>.

⁴³ See Gillingham, Robert and Walter Lane. Changing the treatment of shelter costs for homeowners in the CPI. *Monthly Labor Review*. June 1982.

⁴⁴ In 2013-2014, 77% of CPI-E households were classified as homeowners, and 67% of these owned without a mortgage.

pocket expenses for many elderly.⁴⁵ Boskin argues the rental equivalency method, in effect, compensates homeowners for capital gains on their homes.⁴⁶

3.3 Aggregation across Households

The choice of proper aggregation method consistently arises in the price index literature. To aggregate household specific indexes into aggregate measures for a group, one can weight each household by their total expenditures on the CPI market basket items, a plutocratic weighting approach, or assign each household an equal weight, referred to as a democratic approach. The plutocratic approach tends to skew the aggregate index toward the behavior of high-income households, who get a larger relative weight. The CPI-U, CPI-W, CPI-E, and Chained CPI-U are all plutocratically weighted. Deaton (1998) suggested the representative household suggested by the plutocratic CPI-U was around the 75th percentile of income. Schultze argues, ‘it is hard to imagine anyone would deliberately make decisions about public pensions by tracking households at the 75th percentile.’⁴⁷ The Schultze panel suggested the optimal aggregation choice for a national index may differ from the choice for subgroups, essentially recommending the plutocratic approach for the index used as a barometer for the macroeconomy, and a democratic approach for subgroup indexes. The U.K. Royal Statistical Society concurred with this recommendation in 2015.⁴⁸

4 Improving population subgroup index estimates

To understand the dynamics generating current differences in CPI-U, Chained CPI-U, CPI-W, and CPI-E inflation estimates, and to seek evidence of current inflation heterogeneity across cohort groups in the United States, we scrutinized the sample unit level data used to compile the U.S. official indexes from 2014 to 2017. Then, we utilized a model-based machine learning clustering algorithm to control for heterogeneity across households, in an experimental cohort aggregation approach.

⁴⁵ See Hobihn, Bart and David Lagakos. *Inflation Inequality in the United States*. *Review of Income and Wealth*, December 2005.

⁴⁶ See Boskin, Michael et al. *Consumer Prices, the consumer price index, and the cost of living*. *Journal of Economic Perspectives*. Vol.12, No. 1. Winter 1998. <https://pubs.aeaweb.org/doi/pdf/10.1257%2Fjep.12.1.3>.

⁴⁷ Schultze, p 240.

⁴⁸ See Bentley, Alan. *Household-group inflation: methods to combine expenditure patterns*. Meeting of the Group of Experts on Consumer Price Indices. Geneva, Switzerland. May 2016.

4.1 Data set

The experimental research data set started with all households (consumer units) participating in the 2013 and 2014 CE Quarterly Interview Survey, the expenditure reference period of the CPI-U, CPI-W, and CPI-E in 2016 and 2017 (n=24,080). This data set was then restricted to the subset of respondents participating in four quarterly interviews during this period, to directly obtain annual expenditure estimates for each household (n=6,606).⁴⁹ Due to the reduced sample size, expenditure data were summed across the 243 component CPI item categories to form household (j) shares (w) at a broader, expenditure class (k) level (n=70) for the arbitrarily selected base-period of $\theta = 2013-2014$.⁵⁰ See Equation 1 and 2.

$$(1) \quad E_{k,j,\theta} = \sum_{i \in k,j}^i (PQ)_{i,j,\theta} \quad (2) \quad w_{k,j,\theta} = \frac{E_{k,j,\theta}}{\sum_{k \in j}^k E_{k,j,\theta}}$$

The CE Interview Survey, a 3-month recall survey, is not designed to capture detailed expenses for all of these categories. Detailed food expenses are only collected in the CE Diary Survey. Official CPI aggregation methods integrate data from the Diary and Interview survey components to estimate expenditures at the component cell level (n=9,234) by calendar month. Hence, the official method neither requires the collection nor calculation of household specific budget shares for each component item. This research jettisoned the Diary data, and used expenditures collected in the Interview for the categories officially sourced from the Diary, e.g., food at home and some apparel items. This reduced the set of component cells to 52 item categories.⁵¹ Then, an initial Laspeyres and Geometric Mean index for each household was computed using corresponding component (k) national average price change from December 2014 to December 2017. See Equations 3 and 4.⁵²

$$(3) \quad I^L_{j;\theta \rightarrow t} = \sum_{k \in j}^k w_{k,j,\theta} \left(\frac{P_{k,t}}{P_{k,\theta}} \right) \quad (4) \quad I^G_{j;\theta \rightarrow t} = \prod_{k \in j}^k \left(\frac{P_{k,t}}{P_{k,\theta}} \right)^{w_{k,j,\theta}}$$

⁴⁹ Using households with partial year reported data would bias household budget shares toward items with strong seasonal spending behavior. Use of 4-panel households eliminated any seasonal bias in the household specific shares, at the expense of jettisoning 75% of the data used in the constructing of the official BLS indexes.

⁵⁰ This results in a loss of precision in measuring heterogeneity across household and groups of households, notably variation in spending behavior that might exist across component items within expenditure class.

⁵¹ Garter, et al. demonstrate a model-based approach to link Diary households to Interview, in order to estimate item-specific weights, sourced from the Diary, to Interview households. Adoption of this method would allow for a complete matrix of spending for each Interview household at the CPI component stratum level. See <https://www.bls.gov/ore/pdf/ec020060.pdf>.

⁵² Using national level indexes sacrifices precision and variation attributable to geographic variation. Future enhancements of this research will control for geographic variation.

Descriptive statistics for the experimental research sample are shown in Figure 8. The restricted research sample has slightly more elderly and retired households, and fewer wage-earning households, compared to the full CPI-U sample.

Figure 8. Descriptive statistics for research sample

| Descriptive Statistics 2013 and 2014 Consumer Expenditure Interview Survey | 4-Interview CPI-U Households | | All CPI-U Households | |
|---|---------------------------------|-------|-------------------------|-------|
| | n | % | n | % |
| Sample size | 6,606 | | 24,080 | |
| Sample with: | | | | |
| Income > 0 | 6,596 | 99.8% | 22,488 | 93.4% |
| One or more members age 16 or under | 2,571 | 39.0% | 8,772 | 39.0% |
| One or more owned motor vehicles | 5,756 | 87.3% | 18,782 | 83.5% |
| One or more members age 64 or over | 2,134 | 32.4% | 5,821 | 25.9% |
| One or more members earning income | 4,832 | 73.3% | 17,308 | 77.0% |
| Only one member (single consumer) | 1,894 | 28.7% | 7,235 | 32.2% |
| Living in center city of metropolitan area | 1,906 | 28.9% | 7,574 | 33.7% |
| 50% or more of income sourced from Social Security | 1,424 | 21.6% | 3,726 | 16.6% |
| 50% or more of income sourced from wage earnings | 3,671 | 55.7% | 14,224 | 63.3% |
| Total household annual income before taxes | | | | |
| First quintile | \$ 26,889 | | \$ 21,720 | |
| Third quintile | \$ 89,520 | | \$ 85,015 | |
| Average | \$ 68,058 | | \$ 65,223 | |

4.2 Evidence of price change dispersion across categories

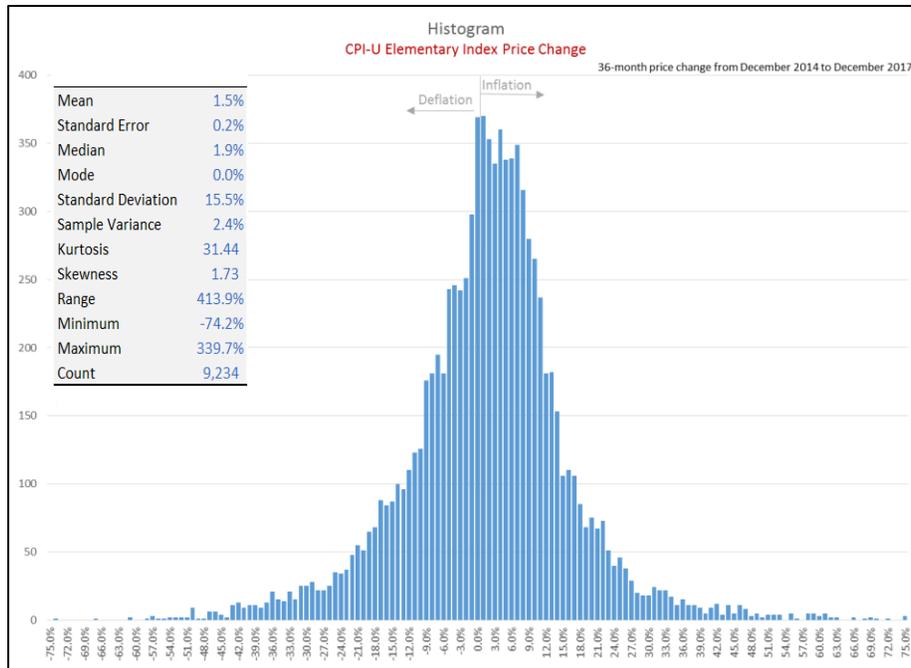
Overall inflation estimates from the study period, December 2014 to December 2017, are shown below. The research example exhibits the same ordinal ranking of estimates as the official indexes, CPI-E > CPI-U > CPI-W for this time period.

| | Official BLS Index | RESEARCH SAMPLE | | |
|---------------|--------------------|-----------------|------------|-------------|
| | | Sample Mean | Democratic | Plutocratic |
| CPI-U | 4.988 | 4.912 | 4.845 | 4.796 |
| Chained CPI-U | 4.220 | na | na | na |
| CPI-W | 4.618 | 4.504 | 4.440 | 4.333 |
| CPI-E | 5.611 | 5.530 | 5.498 | 5.385 |

If there is minimal variation in price change (1) within the component items in the CPI market basket, and (2) across the component items, combined with (3) homogeneous budget share patterns across subgroups, then the four CPI products will tend to produce similar estimates of inflation. Component price change for CPI elementary items exhibited significant variation during the study period, ranging from a minimum of -74.2% (car rentals in St. Louis) to a maximum of 339.7% (moving expenses in Houston). See Figure 9. The distribution of price change at the broader expenditure class (k), nationwide level is narrower, ranging from a low of -19.1% for recreational goods to a high of +22% for vehicle insurance. Categories exhibiting above average inflation during this period include medical care,

shelter, food away from home, and college tuition. Households with disproportionately large budget shares for these items will have higher than average household specific inflation. Categories exhibiting the most deflation during this period include apparel (except jewelry and watches), household goods, and telephone and internet services. Households with large budget shares for these categories will have lower household-specific inflation experience.

Figure 9. *Distribution of CPI component item stratum and area price change*



4.3 Evidence of inflation heterogeneity across households

The distribution of the household specific Laspeyres index estimates for the 2013-2014 experimental research sample is shown in Figure 10. The distribution exhibits less variation than the component indexes, with a min of -2%, a max of 11.9%, and a coefficient of variation of 33%. The distribution is slightly skewed toward the higher inflation tail, with 4% of the sample having inflation rates beyond two standard deviations from the mean, compared to 2% in the corresponding lower tail. Scatterplots of price change by age indicate a significant upward trend, although age does not explain much of the household inflation variation. See Figure 11. Similarly, a scatterplot of price change by income shows a downward linear trend, with a heteroskedastic range at lower incomes. See Figure 12. Note the households with the highest inflation values are the lower income elderly households. Hence, this

evidence is consistent with prior BLS research suggesting the poor tend to face higher inflation rates compared to the rich, and the elderly face higher rates than the young.⁵³

Figure 10. Distribution of household specific price change

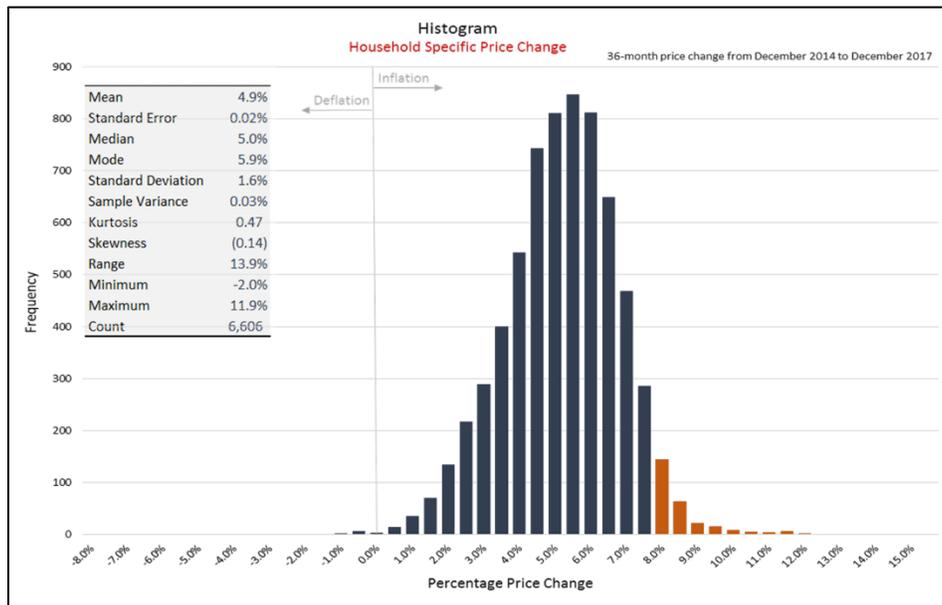
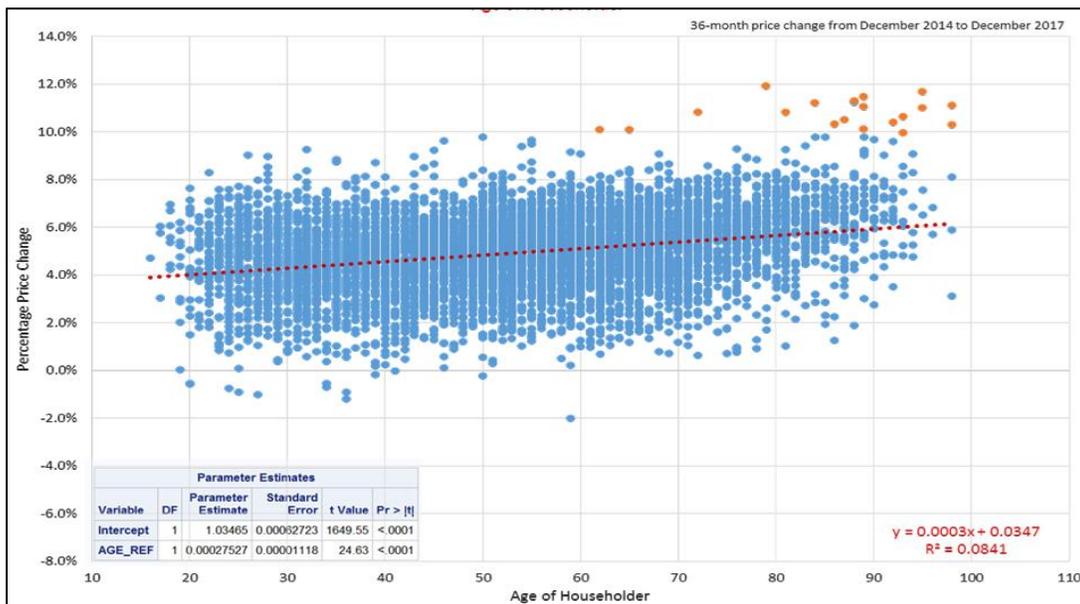
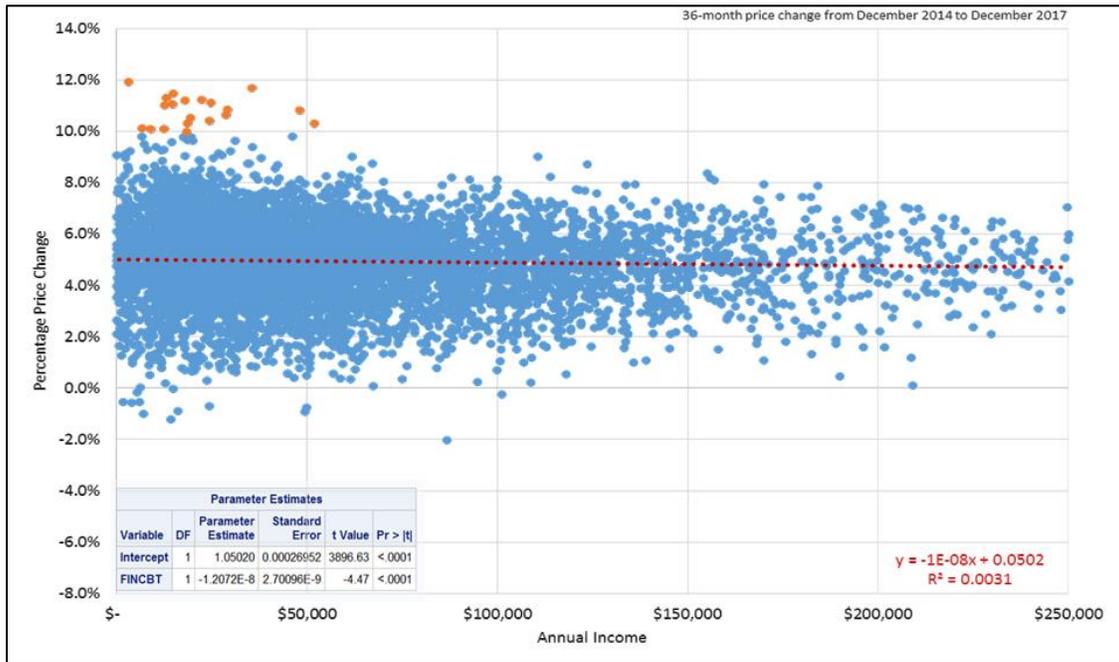


Figure 11. Inflation scatterplot by age of householder



⁵³ See Garner, Thesia, et al. An experimental consumer price index for the poor. Monthly Labor Review. September 1996.

Figure 12. Inflation scatterplot by household income



4.4 Evidence of variation within categories

To address the failure of the cohort indexes to account for within stratum variation, we stratified the CPI sample of rents and goods and services into three price level groups based on the price of each item in the sample as of December 2014: lowest quintile prices, highest quintile prices, and middle 60%. The stratification was executed at the CPI elementary item level. Then, the stratified samples were pooled by expenditure class (*k*) and component indexes were recalculated over the 36-month study period. Due to sample rotation, the sample of usable rents and price quotes with available price data in both periods was roughly 50% of the full CPI sample (n=67,268). Descriptive statistics for this restricted sample are shown in Figure 13. Price change in the study period for the low price items, on average, is higher compared to high price items.

Assuming lower income households have income to purchase only the lower price items within each category, and higher income households purchased the higher priced items, the stratified component indexes were matched to households by income quintile. This experiment is a heuristic way to capture potential variation existing in the marketplace, lacking the actual prices paid for specific brand and

variety of items purchased by each household. In effect, the price level of the items proxy for the hypothesized limited vector of all marketplace prices paid by income cohort groupings.

Figure 13. Descriptive statistics for CPI sample data, by price level strata

| | ALL | LOWEST | MIDDLE | HIGHEST |
|---------------------------|---------|--------|--------|---------|
| Mean | 1.07102 | 1.0824 | 1.0709 | 1.0600 |
| Standard Error | 0.00149 | 0.0036 | 0.0019 | 0.0031 |
| Mode | 1 | 1 | 1 | 1 |
| Standard Deviation | 0.38661 | 0.4150 | 0.3854 | 0.3595 |
| Sample Variance | 0.14947 | 0.1722 | 0.1485 | 0.1293 |
| Kurtosis | 347.8 | 270.3 | 422.2 | 175.0 |
| Skewness | 11.7 | 10.9 | 12.9 | 8.5 |
| Range | 19.95 | 17.34 | 19.95 | 12.17 |
| Minimum | 0.05 | 0.05 | 0.05 | 0.05 |
| p10 | 0.3983 | 0.4555 | 0.3901 | 0.3762 |
| Q1 | 0.9542 | 0.9479 | 0.9574 | 0.9515 |
| Median | 1.0308 | 1.0213 | 1.0337 | 1.0313 |
| Q3 | 1.1393 | 1.1438 | 1.1410 | 1.1298 |
| p90 | 1.3017 | 1.3333 | 1.3011 | 1.2800 |
| Maximum | 20 | 17.39 | 20.00 | 12.22 |
| Count | 67268 | 13540 | 40196 | 13532 |

4.5 A model-based clustering aggregation approach

Lacking data that connects the prices and quantities of marketplace transactions to the household characteristics of the consumers making the purchases due to lack of funding, how can statistical agencies improve subgroup index estimation methods using the survey data available to them? Using a subset of the household, rent, and price sample data, we developed an experimental approach to constructing a hypothetical subgroup index targeting Social Security, within the confines of estimating a national index. The experimental approach follows the guidance of the Schulze advisory panel, and differs from the official CPI methods in four ways: (1) treatment of owner-occupied housing, (2) cohort definition, (3) component prices, and (4) aggregation method.

Step 1. Stratify the starting or base period sample of households into a limited number of mutually exclusive but exhaustive cohorts of most interest to users. This could be income quintiles or CPI market basket expenditure quintiles, to control for hypothesized inflation variation across the income distribution. In the U.S., the prominent use of the CPI for Social Security COLA might suggest stratifying the sample by Social Security recipient status. We stratified the research sample into three income

(lowest quintile, medium quintiles, highest quintile) and two Social Security (recipients, non-recipients) groups.⁵⁴

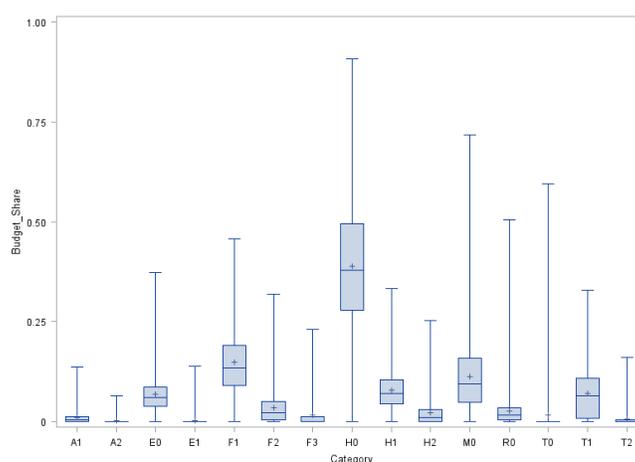
Figure 14. Initial population strata sample sizes

| | | n | % |
|------|---------------------------------|------|-----|
| Code | Social Security households: | 1424 | 22% |
| SL | Low income | 648 | 10% |
| SM | Middle Income | 776 | 12% |
| | Non Social Security households: | 5171 | 78% |
| NL | Low income | 671 | 10% |
| NM | Middle income | 3181 | 48% |
| NH | High income | 1319 | 20% |

Note: No households classified as Social Security recipients were in the high income quintile group, resulting in five initial population strata instead of six.

In order to apply the economic approach to index number theory, it is necessary to assume that each household (j) within each of the five initial strata (s) defined in Step 1, behave as a single household.⁵⁵ That is, the member households in each (s) should be homogeneous with respect to their CPI budget shares such that the assumption of homothetic preferences applies. However, there remains a great deal of variation in budget shares within these initial strata (see Figure 15). To strengthen the assumption of homothetic preferences, we cluster the households within each strata (s) into cohorts (c) based on their revealed spending shares.

Figure 15. Boxplot of household budget shares for select CPI categories, Social Security low-income strata



⁵⁴ To be classified as a Social Security household, 50% or more of total household income had to be sourced from Social Security benefits.

⁵⁵ See International Labour Office. Consumer Price Index Manual, Theory and Practice. Chapter 18, Section 18.1 p. 337.

Step 2. Calculate household specific budget shares. See equations (1) and (2) in section 4.1. Calculate household specific indexes. See equations (3) and (4) in section 4.1. We calculated two sets: one using a rental equivalence approach for owner-occupied housing, and one using a payment approach as outlined in section 4.5. Household expenditures were summed to a broader 15-category level (used for clustering in subsequent steps) as shown in Figure 16.

Figure 16. Budget shares for select CPI categories, by experimental cohort

| RENTAL EQUIVALENCE | | | | | | PAYMENTS APPROACH | | | | | |
|-------------------------------------|-------|-------|-------|-------|-------|-------------------------------------|-------|-------|-------|-------|-------|
| Category | NL | NM | NH | SL | SM | Category | NL | NM | NH | SL | SM |
| A1 Adult clothing | 1.2% | 1.6% | 2.1% | 1.1% | 1.0% | A1 Adult clothing | 1.4% | 1.7% | 2.3% | 1.4% | 1.3% |
| A2 Children and infant clothing | 0.6% | 0.4% | 0.4% | 0.1% | 0.1% | A2 Children and infant clothing | 0.7% | 0.5% | 0.5% | 0.2% | 0.2% |
| E0 Telephone and electronics | 6.8% | 6.6% | 5.3% | 6.2% | 5.5% | E0 Telephone and electronics | 7.5% | 7.2% | 5.7% | 7.8% | 7.1% |
| E1 Tuition | 2.6% | 1.9% | 4.9% | 0.3% | 0.5% | E1 Tuition | 2.8% | 2.0% | 5.2% | 0.4% | 0.6% |
| F1 Food at home | 15.7% | 11.9% | 8.8% | 12.3% | 10.2% | F1 Food at home | 17.2% | 13.0% | 9.5% | 15.6% | 13.1% |
| F2 Food away from home | 4.2% | 5.1% | 5.8% | 3.6% | 3.9% | F2 Food away from home | 4.6% | 5.6% | 6.3% | 4.6% | 5.0% |
| F3 Alcohol and tobacco | 1.8% | 1.5% | 1.3% | 1.3% | 1.0% | F3 Alcohol and tobacco | 2.0% | 1.7% | 1.4% | 1.7% | 1.2% |
| H0 Shelter | 35.2% | 31.7% | 30.2% | 39.4% | 34.5% | H0 Shelter | 28.8% | 25.6% | 25.1% | 23.1% | 16.1% |
| H1 Household utilities | 7.2% | 5.8% | 4.3% | 7.1% | 6.3% | H1 Household utilities | 7.9% | 6.4% | 4.7% | 9.0% | 8.1% |
| H2 Housefurnishings and operations | 2.2% | 3.0% | 4.9% | 2.6% | 3.4% | H2 Housefurnishings and operations | 2.4% | 3.3% | 5.3% | 3.3% | 4.3% |
| M0 Medical Care | 5.7% | 7.6% | 7.3% | 12.0% | 13.9% | M0 Medical Care | 6.3% | 8.3% | 7.8% | 15.2% | 17.8% |
| R0 Entertainment and recreation | 3.3% | 3.9% | 5.9% | 2.8% | 3.7% | R0 Entertainment and recreation | 3.6% | 4.3% | 6.3% | 3.5% | 4.7% |
| T0 Vehicles | 3.8% | 7.0% | 8.0% | 3.2% | 6.9% | T0 Vehicles | 4.2% | 7.6% | 8.6% | 4.1% | 8.9% |
| T1 Gasoline and vehicle maintenance | 8.9% | 10.7% | 8.8% | 7.2% | 8.3% | T1 Gasoline and vehicle maintenance | 9.8% | 11.7% | 9.5% | 9.1% | 10.6% |
| T2 Public transportation | 0.9% | 1.0% | 1.8% | 0.7% | 0.9% | T2 Public transportation | 1.0% | 1.1% | 2.0% | 0.9% | 1.1% |

Figure 17. Difference in rental equivalence and payment approach shares, by experimental cohort

| Category | NL | NM | NH | SL | SM |
|-------------------------------------|-------|-------|-------|--------|--------|
| A1 Adult clothing | 0.1% | 0.1% | 0.2% | 0.3% | 0.3% |
| A2 Children and infant clothing | 0.1% | 0.0% | 0.0% | 0.0% | 0.0% |
| E0 Telephone and electronics | 0.7% | 0.6% | 0.4% | 1.7% | 1.6% |
| E1 Tuition | 0.3% | 0.2% | 0.4% | 0.1% | 0.1% |
| F1 Food at home | 1.5% | 1.1% | 0.6% | 3.3% | 2.9% |
| F2 Food away from home | 0.4% | 0.5% | 0.4% | 1.0% | 1.1% |
| F3 Alcohol and tobacco | 0.2% | 0.1% | 0.1% | 0.3% | 0.3% |
| H0 Shelter | -6.3% | -6.1% | -5.0% | -16.2% | -18.4% |
| H1 Household utilities | 0.7% | 0.5% | 0.3% | 1.9% | 1.8% |
| H2 Housefurnishings and operations | 0.2% | 0.3% | 0.4% | 0.7% | 1.0% |
| M0 Medical Care | 0.6% | 0.7% | 0.5% | 3.2% | 3.9% |
| R0 Entertainment and recreation | 0.3% | 0.4% | 0.4% | 0.7% | 1.0% |
| T0 Vehicles | 0.4% | 0.6% | 0.6% | 0.9% | 1.9% |
| T1 Gasoline and vehicle maintenance | 0.9% | 1.0% | 0.6% | 1.9% | 2.3% |
| T2 Public transportation | 0.1% | 0.1% | 0.1% | 0.2% | 0.2% |

Step 3. Using the sample of all households, identify budget shares of most significance in explaining variation in the household specific inflation estimates. The objective of the clustering problem is to group households into like-kind, homogenous groups based on budget shares. In the U.S. CPI, with 243 component strata, this becomes a high dimensional data problem, with vector n number of budget shares equal to the number of component strata in the CPI classification scheme. To reduce the number

of dimensions, we summed the shares into 15 groupings.⁵⁶ We then executed the following LASSO model:

$$(5) \quad I^L_{j;\theta \rightarrow t} = v(w_{k_i})$$

to determine the optimal set of variables for clustering. LASSO modifies OLS regression by adding a penalty term to the error. Mathematically the problem is to find $(\beta_1, \beta_2, \dots, \beta_{15})$ which minimizes the error of:

$$(6) \quad I^G_{j;\theta \rightarrow t} = \sum_{i=1}^{15} \beta_i w_{k_i} \quad \text{where the error is } \sum_{\text{observations}} (I^G_{j;\theta \rightarrow t} - \sum_{i=1}^{15} \beta_i w_{k_i})^2 + \lambda \sum_{i=1}^{15} |\beta_i|$$

The penalty term increases the error if there are large coefficients. This has a tendency to reduce or drive to zero the coefficients of shares which are not strongly predictive of the response. The value of λ is determined by cross validation.⁵⁷

We selected the following shares as most explanatory: f1- food at home, f2-food away from home, h0-rent or owner's equivalent rent, m0-medical area, and t1-gasoline and vehicle maintenance.

Step 4. In order to reduce the variation in spending behavior *within* the income-Social Security strata, we executed a hierarchical clustering algorithm to form clusters of homogeneous households based on the revealed budget shares.⁵⁸ This approach, in effect, superimposes homothetic utility preferences inside the cluster as the vector of budget shares on CPI items will be approximately equal across all households within the cluster. If the clustering algorithm performs well, then a democratic average and plutocratic average index will converge for the cluster, eliminating the difficult choice of aggregation method.

⁵⁶ A1=Adult apparel, A2=Infant and children apparel, E0=Telephone, internet, audio, video, electronics, E1=Tuition, F1=Food at home, F2=Food away from home, F3-Alcohol and tobacco, H0=Shelter, H1=Household utilities, H2=Housefurnishings and operations, M0=Medical care, R0=Entertainment and recreation, T0=Vehicles, T1=Gasoline and vehicle maintenance, T2=Public transportation.

⁵⁷ Typically 10-fold cross validation is used. This means the data are divided into 10 blocks, the betas are fit using 9 blocks of data and used to predict the 10th block. The combined error of predicting the 10 blocks in turn is used to select the best value of λ and the coefficients kept are the ones corresponding to this value of λ . If a coefficient is exactly zero then the variable should be dropped from the model. If the coefficients are close to zero relative to other coefficients then it is up to the modeler's judgement whether or not to keep these variables.

⁵⁸ Clustering is an unsupervised machine learning method to partition a set of data into mutually exclusive clusters where data points within a cluster are similar. There are many ways to cluster data but here we use hierarchical clustering.

Hierarchical clustering is a two-step process. First a tree is formed and then the tree is cut at a level chosen by the researcher to provide a partition into the desired number of clusters. It starts with every one of N data points being its own cluster. Then it combines the two nearest points into a cluster resulting in $N-1$ clusters at the next higher level of the tree. Again the two nearest clusters are combined so there are $N-2$ clusters at the next highest level of the tree. This is repeated until all points are in a single cluster. The main choice in hierarchical clustering is the choice of a distance function for two clusters. We chose Ward's minimum variance method, which minimizes the increase in total within cluster variance from combining two clusters.⁵⁹

The formula for the distance between two clusters K and L is $D(K, L) = \frac{\|\bar{X}_K - \bar{X}_L\|^2}{\frac{1}{N_K} + \frac{1}{N_L}}$ where the numerator is the squared distance between the centroids of clusters K and L and N_K is the number of points in cluster K (likewise for N_L).

Once the tree is constructed then one has to determine how many clusters to divide the data into. There is no method that is guaranteed to select the optimal number of clusters so one is left to rely on diagnostic criteria.⁶⁰ After deciding on the number of clusters, Proc Tree reads in the hierarchical tree and cuts it at the appropriate stage to produce the desired number of clusters.

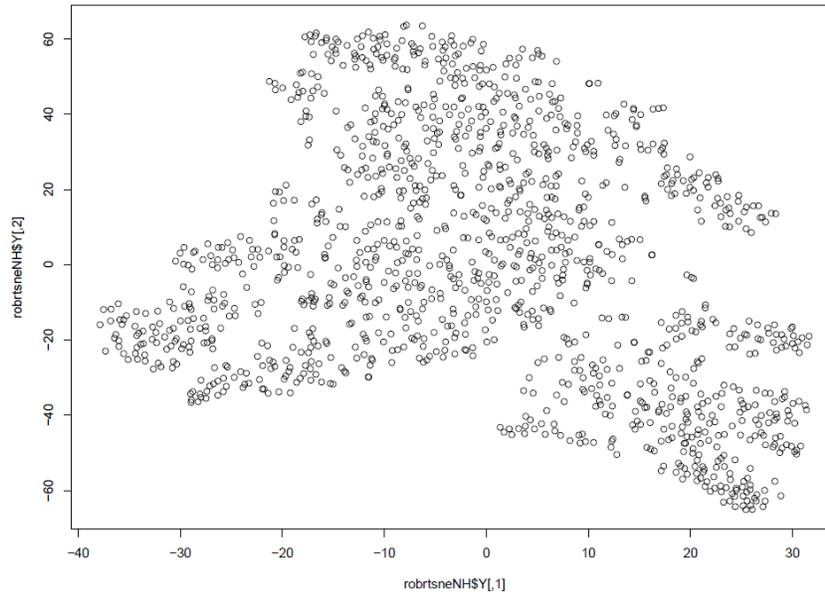
To visualize the clustering problem, the t-distributed Stochastic Neighbor Embedding method was used to visualize the structure of the high dimensional budget share data in two dimensions.⁶¹ An example is shown in Figure 18 below.

⁵⁹ See <https://arxiv.org/pdf/1111.6285.pdf>.

⁶⁰ We used the CCC statistic output from the PROC CLUSTER procedure in SAS. SAS Technical Report A-108 describes the formulas and empirical testing behind the CCC statistic. SAS Proc Cluster provides a plot of the CCC for each number of clusters. One should look for a peak in the CCC curve with a value greater than 2.

⁶¹ Laurens van der Maaten and Geoffrey Hinton, Visualizing Data using T-SNE, *Journal of Machine Learning Research* volume 9(Nov):2549—2578, 2008. To define the structure of the 15-dimensional data, a measure of similarity is constructed. For a point P_i use a Gaussian distribution centered at P_i . Then using the distance between P_i and P_j one can calculate the probability that P_j is the neighbor of P_i . This measure is not symmetric for P_i and P_j so the similarity of P_i and P_j is the average of the similarity of P_j to P_i and the similarity of P_i to P_j . This defines a distribution of similarities in the original 15 dimensional space. The goal of t-SNE is to find a set of points in 2-dimensional space such that the distribution of similarities among points is as similar as possible to the distribution of similarities in the original space. The construction of similarities in 2-dimensional space uses a Student's t distribution with one degree of freedom instead of a Gaussian distribution. The t-distribution has thicker tails than the Gaussian distribution which results in points separated by a moderate distance in the 15-dimensional space being somewhat closer together in the 2-dimensional space. The purpose of this is to hold together any cluster structure in the original 15-dimensional space and to improve the separation of clusters.

Figure 18. *t-SNE plot of stratum NH = Non Social Security, High Income Households*



The final clustering algorithm optimized on 40 clusters for each strata, except the NM (Non Social Security, Middle Income) strata which required 75, for a total of 235 strata-cluster combinations over the entire sample. The average household sample size per strata-cluster was roughly 30, with a min=2, q1=13, q3=37, max=126. The strata-cluster groups form the component building blocks for subsequent index aggregation.

Step 5. Calculate strata-cluster indexes using the Laspeyres formula (to support CPI-U calculation), and a superlative formula (Tornqvist) to mitigate substitution bias, by aggregating across clusters within stratum. Lacking current-period weights for the research sample, we estimated a preliminary version of the superlative using the base-period weights and base-period household clustering, by applying the CES formula and an estimate of substitution elasticity (here $\sigma=0.6$) across all clusters.⁶² This step can be performed using democratic or plutocratic weights, as the two schemes should converge. We found an average difference of 0.015% between the two weighting approaches across all 235 clusters.

$$(7) \quad I^L_{s,c;\theta \rightarrow t} = \sum_{k,j \in s,c} w_{k,j,\theta} \left(\frac{P_{k,t}}{P_{k,\theta}} \right) \quad (8) \quad I^C_{s,c;\theta \rightarrow t} = \prod_{k \in j} \left(\frac{P_{k,t}}{P_{k,\theta}} \right)^{w_{k,j,\theta}}$$

⁶² Ideally, the value of σ should be optimized for each stratum-cluster and could vary. This will be a future enhancement of this research.

Step 6. Democratically average the strata indexes into cohort specific indexes needed by specific users for tailored indexation needs. Here, we calculated a democratic average of the low-income and middle-income Social Security recipient strata to estimate an index for all Social Security recipient households.

Step 7. Plutocratically average the strata indexes into the national index.

Step 8. In order to calculate Tornqvist indexes at the strata-cluster level, a mechanism to classify the sample of current period CE survey households into the base-period within stratum clusters is needed. Since the longitudinal period of the CE survey is only four quarters, households in the survey constantly rotate and we lack current-period spending for the exact set of household as in the base period.

We plan to use the PROC DISCRIM procedure available in SAS on demographic variables to identify current period households with the same characteristics of base-period households, in a future enhancement of this research.⁶³

By using classification for the new current-period households, each household can be assigned to base-period strata-cluster based on different variables than were used to perform the original clustering (i.e., the budget shares) as the classification models make no assumptions about how the values of the response in the training dataset were produced. Ostensibly, this will allow the vector of budget shares of the current-period members of the cluster to vary from the vector of the original members, where the variance can be assumed to be response to changing relative prices. This could be an improvement to the current method used to construct Tornqvist indexes, as the clustering should ostensibly control for

⁶³ Classification is a supervised machine learning method. There is a response variable, which is the cluster number based of the clustering using the 15 budget shares described in step 4. Supervised machine learning means that there is a training dataset with the values of the response filled in. The classification model fitted to the training data (here the clustering in step 4) will predict which cluster a new point should belong to. There are a great many classification models. Most classification models in SAS are part of the SAS Enterprise Miner, such as neural networks, support vector machines, and decision trees. However SAS Proc Discrim is capable of classifying new households based on demographics and the clustering of the initial dataset. Proc Discrim by default assumes a multivariate normal distribution of the demographic variables for points in the training data (the households which were clustered). Since this assumption is unlikely to hold, we use a non-parametric classification method. This method is called K-nearest neighbors. The K-nearest neighbors classification algorithm works as follows: Choose the value of K, the number of neighbors to be used in classification. For a new point P, find the K closest points in the training dataset to the point P. The cluster number containing the most points from the K-nearest neighbors is assigned to the point P. Choosing an odd value for K reduces the likelihood of ties. One can use cross validation to find a very good value for K. In cross validation, the training data is broken into say 10 blocks. Then the K-NN procedure is used to classify the 10th block on the basis of the other 9. Each of the 10 blocks is classified in turn and the error rate or AUC (area under curve) are used to select the value of K which most accurately reproduces the original clusters.

exogenous factors that also influence spending behavior, such as changes in household composition and demographics.

5 Results

We computed aggregate indexes for each stratum using (1) a rental equivalence and (2) payments approach for owner occupied housing. Laspeyres and CES indexes for each stratum were computed using three different sets of component indexes: (1) official CPI-U expenditure class (k) national indexes; (2) a restricted sample of CPI-U (k) indexes, and (3) a restricted paired price sample of (k) indexes, pairing low-price items to the low income quintile strata, low and medium prices to middle income strata, and all prices to high incomes strata. Results are in Figure 19.

Figure 19. Index results, rental equivalence approach

| Formula | Owner Occupied Housing | Component Indexes | Low Income Quintile | High Income Quintile | Social Security Cohort | All Urban Consumers |
|------------------|------------------------|-------------------------|---------------------|----------------------|------------------------|---------------------|
| Laspeyres | Payment | Full sample, urban | 4.24% | 4.50% | 4.33% | 4.19% |
| Laspeyres | Payment | Restricted price vector | 6.12% | 4.44% | 5.47% | 4.71% |
| Laspeyres | OER | Full sample, urban | 5.06% | 4.86% | 5.56% | 4.80% |
| Laspeyres | OER | Restricted price vector | 7.05% | 4.75% | 6.81% | 5.34% |
| CES $\sigma=0.6$ | Payment | Restricted price vector | 5.98% | 4.32% | 5.31% | |
| CES $\sigma=0.6$ | OER | Restricted price vector | | | | 5.21% |

Conclusion

In this paper, we examined micro-level CPI source data from 2014 through 2017 seeking evidence of inflation heterogeneity across households, and experimented with an alternative technique to calculate cohort specific indexes that controls for the heterogeneity. To control for the presence of heterogeneity in budget shares across households, we clustered an experimental research sample into homogeneous groups defined by the similarity of spending patterns on select CPI items. Using these homogenous clusters as building blocks, we constructed a national all-urban consumer index and a hypothetical population subgroup index for Social Security recipient households, with alternative treatments of owner-occupied housing. To incorporate hypothesized within category heterogeneity for low-income

groups, we calculated separate component indexes for the lowest-priced products and services in the CPI sample, and imputed those indexes to the low-income households.

Due to the underlying sample size of the research data set, we do not draw statistical inferences from the exercises performed herein or sanction the accuracy of the index estimates. However, we note the following observations:

- The experimental indexes calculated herein are within the 95% confidence interval of the official CPI-U for the corresponding time period, with two exceptions. Low income quintile and Social Security cohort indexes using the restricted, low-price vector and rental equivalence definition for owner occupied housing were larger.
- The treatment of owner-occupied housing and within-item inflation variation have a large impact. The payment approach for shelter yields a lower inflation estimates during the study period, notably for the Social Security cohort. Use of restricted samples for component indexes (i.e., the low price level quintile) for low income cohorts yields higher overall inflation estimates.
- Stratified clustering of households based upon budget shares may improve the accuracy of superlative index estimation

Statistical agencies are responsible for producing accurate, reliable, timely statistics and should resolve to meet the needs of all index users. In the case of the index needed by the Social Security Administration to index COLAs, if the index is biased too high, there is a heightened risk of hastening the insolvency of the OASDI trust fund absent a commensurate increase in taxes levied on workers to populate the fund. If it is biased too low, it undermines the intent of the law and the purchasing power of recipients, especially the poorer subgroup that depends on the income for livelihood. Due to the sheer number of beneficiaries in the program, a difference of one tenth of one percent in the CPI translates into more than a billion dollars of outlays. Accuracy is paramount.

When COLAs are provided in public transfer payments, such as Social Security, an overarching question is whether the adjustment should compensate the recipients for changes in the overall cost-of-living for the nation as a whole, or should it account for any significant inflation difference experienced by the intended recipients? This is a question for policy makers to resolve. If the latter, the ability of statistical agencies to produce accurate subgroup indexes is hampered by a lack of sufficient household spending data, both with respect to survey sample size and consumption detail.

The BLS is currently redesigning its household consumption survey. BLS plans to add point of purchase questions to the Consumer Expenditure Survey and outlets will then be linked to households. This will enable BLS, potentially, to add population groups to its stratification of items and areas. This could, partially at least, enable BLS to control for within-category heterogeneity currently not accounted for in the CPI-E and CPI-W index estimates. Clustering households into homogeneous groups, and scientifically selecting samples independently for each group, could improve cohort index estimation.