Where is Standard of Living the Highest? Local Prices and the Geography of Consumption

Rebecca Diamond (Stanford and NBER)
Enrico Moretti (Berkeley and NBER)

November 2021

Abstract

Income differences across US cities are well documented, but little is known about the level of standard of living in each city—defined as the amount of market-based consumption that residents are able to afford. In this paper we provide estimates of the standard of living by commuting zone for households in a given income or education group, and we study how they relate to local cost of living. Using a novel dataset, we observe debit and credit card transactions, check and ACH payments, and cash withdrawals of 5% of US households in 2014 and use it to measure mean consumption expenditures by commuting zone and income group. To measure local prices, we build income-specific consumer price indices by commuting zone. We uncover vast geographical differences in material standard of living for a given income level. Low income residents in the most affordable commuting zone enjoy a level of consumption that is 74% higher than that of low income residents in the most expensive commuting zone.

We then endogenize income and estimate the standard of living that low-skill and high-skill households can expect in each US commuting zone, accounting for geographical variation in both costs of living and expected income. We find that for college graduates, there is essentially no relationship between consumption and cost of living, suggesting that college graduates living in cities with high costs of living —including the most expensive coastal cities—enjoy a standard of living on average similar to college graduates with the same observable characteristics living in cities with low cost of living—including the least expensive Rust Belt cities. By contrast, we find a significant negative relationship between consumption and cost of living for high school graduates and high school drop-outs, indicating that expensive cities offer lower standard of living than more affordable cities. The differences are quantitatively large: High school drop-outs moving from the most to the least affordable commuting zone would experience a 26.9% decline in consumption.

We thank Matthias Hoelzlein, Erik Hurst, Matt Turner, and seminar participants at Berkeley, Boston University, Cornell, Dallas Fed, Harris, Irvine, IZA, London Business School, Mannheim, MIT Sloan, NBER Summer Institute, Ohio State, OSUS, Penn, Stanford, Toronto, UCLA Anderson, UPF, Urbana, Wisconsin and Wisconsin Business School for useful comments. Abhisit Jiranaphawiboon, Anais Galdin, and Raymond Lee provided outstanding research assistance. Diamond acknowledges support from the Stanford Graduate School of Business and the National Science Foundation (CAREER Grant 1848036). Results using NielsenIQ data is based on researchers' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researchers and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

1 Introduction

Over the last three decades there have been increased differences in income among US communities. Economically vibrant cities, like New York, San Francisco, Boston, and Seattle, have experienced fast increases in mean household income. At the same time, less dynamic local labor markets have experienced more limited increases in income and, in some cases, even declines. What is less clear is how the actual standard of living of residents varies across communities. The standard of living of residents of a city—which we define as the amount of market-based consumption households are able to purchase—depends both on the income level that residents can expect there and the local cost of living. While we know that large, expensive cities tend to have jobs that offer higher nominal earnings, and small, affordable cities tend to have jobs that offer lower nominal earnings, we know little about where market-based consumption is the highest. Are residents of dynamic metro areas better or worse off in terms of consumption compared to residents of smaller, economically struggling communities? This lack of information is surprising, because the amount of market-based consumption is arguably a key component of utility and economic well-being. Despite the fundamental role of consumption for economic well-being, there is limited systematic empirical evidence on the differences in consumption across cities and how they relate to local cost of living.[1] The paucity of evidence likely reflects the lack of datasets that can measure consumption and are large enough to allow for a detailed geographical analysis.[2]

In this paper, we provide estimates of standard of living by commuting zone for households in a given income or education group, and we study how they relate to local cost of living. Our main data source is a 5% sample of US households' linked bank and credit card transaction data in 2014. We use it to measure the value of consumption expenditures as we observe essentially all debit and credit card transactions, check and ACH payments, and cash withdrawals conducted every day in 2014. For each commuting zone and income group, we create local price indexes and use it to deflate consumption expenditures and obtain estimates of consumption in real terms. This is our main measure of market-based standard of living enjoyed by residents with a given income level in each commuting zone. We quantify how consumption in expensive commuting zones compares with consumption in affordable commuting zones for a given income. We then endogenize income and compare consumption by high- and low-skill households in expensive commuting zones to consumption in affordable commuting zones once we account for geographical variation in both cost of living and expected income. Finally, we study the role played by geographic sorting of households in nationwide consumption inequality.

Relative to existing data sources on consumption, such as the CEX, our combined dataset has

[1] Prior work has mostly focused on groceries that can be tracked by scanner data (Handbury, 2019; Handbury and Weinstein, 2015) or focused exclusively on house price variation Moretti (2013); Ganong and Shoag (2017). A notable exception to this is Bertrand and Morse (2016) who use the CEX to study consumption of the low-income by state. There is a much larger literature on consumption inequality at the national level—for example see Aguiar and Bils (2015); Attanasio and Pistaferri (2016); Meyer and Sullivan (2017)—and the role of prices —see for example Broda et al. (2009); Jaravel (2019).

[2] Limited geographical detail and small samples make it difficult to measure consumption differences at the local level in the CEX or PSID.

important advantages. Our consumption data is comprehensive and include virtually all purchases conducted by individuals in our sample, and unlike other consumption data, it is not self-reported. It matches well both the mean household consumption in the National Accounts (NIPA) and NIPA's share of consumption for main consumption categories. Our data also matches estimates of consumption expenditures by mean of payment (cash, check, credit or debit card) reported by the Fed. Merchant-level expenditures in our data match sales reported by specific publicly traded merchants—Starbucks, Walmart, Home Depot, Macy's, etc.—in official SEC filings. Importantly, our data has detailed geographical information. This allows us to study consumption at the commuting zone level. Unlike the CEX, our sample is large enough that we have enough observations to cover most commuting zones, although larger commuting zones are over-represented.

Our data, however, have important limitations. The main one is that we miss all un-banked households, which account for 7% of the US population and are overwhelmingly low income (Federal Deposit Insurance Corporation, 2015). Second, not all accounts can be linked at the family level. Third, while we can identify the exact type of good and service purchased by credit card, debit card and ACH, we observe only the value but not the type of purchase when the purchase is paid for by cash or check.

To measure local prices, we build consumer price indexes that vary by commuting zone and income group. Our baseline price index is a Laspeyres index which mimics the index used by the BLS to estimate the official national CPI. It is a weighted average of the local prices of items consumed by the average household with income-specific weights reflecting the importance of each item in the bundle for consumers of a given income group. We also examine six alternative price indices based on alternative assumptions, including ones that correct for differences in variety and supply across cities (Handbury and Weinstein, 2015; Handbury, 2019). We augment our data with data on local prices of specific goods from the ACS, NielsenIQ and ACCRA.

The price indexes point to large differences in cost of living across commuting zones, especially for low-income households. The overall cost of living faced by low-income households (post-tax income <\$50,000) in the most expensive city—San Jose, CA—is 49% higher than in the median commuting zone, Cleveland, and 99% higher than the most affordable commuting zone—Natchez, MS. By contrast, we uncover significantly smaller geographical differences for high-income households (post-tax income >\$200,000). The spatial distribution of cost of living is not symmetric, but highly skewed to the right for all income groups. While the cost of living in most cities is between -20% and +20% of the median city, there are a handful of very expensive cities in the right tail, where cost of living is much higher than the median.

To investigate consumption differences across space, we begin by mapping consumption *expenditures* in each commuting zone for each income group. Of course, differences across areas in consumption expenditures reflect not just the quantity of goods consumed by the area residents, but also variation in local prices. We use our price indices to deflate expenditures in order to obtain measures of mean *quantity* of consumption by commuting zone and income group measured in real terms.

We find that geographical differences in material standard of living for a given nominal income level are economically large. The three commuting zones with the lowest consumption of low-income households are San Jose, CA; San Francisco, CA; and San Diego, CA, with consumption levels between 27% and 30% lower than the median commuting zone. At the other extreme of the spectrum, examples of commuting zones with high consumption of low-income households are Huntington, WV; Johnstown, PA; and Elizabeth City, NC, with consumption levels in real terms 22–23% higher than the median commuting zone. The range of consumption levels observed across U.S. communities is quite wide: Low-income families who live in the most affordable commuting zone enjoy a level of market-based consumption measured in real terms that is 74% higher that of families with the same income who live in the least affordable commuting zone.

We estimate that the elasticity of overall market consumption with respect to the local price index is -0.900 (0.009) and -1.016 (0.037) for low- and high-income households, respectively. The fact that low-income consumers cut consumption less than high-income consumers in response to higher local prices could reflect the fact that the former are closer to a minimum subsistence level, and small consumption cuts cost more in terms of utility.[3] Consistent with this possibility, we find that low-income households in expensive commuting zones have a higher incidence of financial distress than low-income households in affordable commuting zones: they are significantly more likely to have negative saving and pay more overdraft fees.

To validate our findings, we replicate the analysis using model free evidence based on direct and transparent measures of consumption quantities from NielsenIQ. The data contains information on the quantities purchased of specific grocery products measured in physical units, unlike our bank data that measure expenditures. For example, we measure the number of cans of beer, the number of light bulbs, or the number of pounds of nuts purchased in a year by each NielsenIQ consumer. There are 823,507 grocery products, divided in 116 product groups. Consistent with our findings on overall consumption, we find that for 79 out of 116 product groups, NielsenIQ consumers in expensive commuting zones buy fewer physical units, holding income and demographics fixed. The elasticity of consumption with respect to the local price index is lower than the elasticity estimated for overall consumption. This likely reflects the fact that grocery items are necessities. When faced with higher cost of living, households seem to cut consumption of necessities less than consumption of all other goods. In addition, groceries exhibit less geographic price variation, making them a relative bargain in expensive cities.

The analysis up to this point compares consumption of residents of expensive and affordable cities, holding their nominal income constant. But income levels are not necessarily the same across areas: for a given level of human capital, households in expensive cities tend to have higher incomes than households in affordable cities. Therefore, in the second part of the paper we turn to the question of how our results change if we endogenize income. Specifically, we measure the standard of living that low- and high-skill households can expect in each US commuting zone,

[3] Alternatively, it could reflect the possibility that more low-income households in expensive cities expect larger future income gains than high-income households; or that more low-income households in expensive cities expect to move to affordable cities in the future than high-income households.

accounting for geographical variation in both cost of living and expected income.

We focus on three skill groups, based on the schooling level of the household head: (i) four-year college or more; (ii) high school or some college; (iii) less than high school. We use the 2012–2016 ACS data to predict the income that a given household may expect in each commuting zone as a function of education and demographics under the assumption that location sorting across cities depends only on observables. We then map our previous estimates of consumption by income level into estimates of consumption by skill level.

We find that for all skill groups the spatial variation in consumption is much smaller than spatial variation in pre-tax income because high income cities tend to have high cost of living. For example, the standard deviation of consumption of low income households is less than half the standard deviation of pre-tax income.

For the high-skill group, San Francisco, New York, and Boston are among commuting zones with the highest pre-tax mean incomes. Accounting for taxation and cost of living reduces the purchasing power of households in these cities by over 40%. However, since pre-tax income is so high, even after accounting for cost of living and taxes, college graduates in San Francisco, New York, and Boston retain a high level of standard of living and remain in the top 30% of the distribution of market consumption across all US cities. Among large cities, Houston has the highest mean consumption, since it offers good expected income and moderate prices.

Overall, we find that for college graduates, there is no significant bi-variate relationship between expected consumption and cost of living. A regression of expected consumption on the local price index across all commuting zones yields a coefficient of -0.032 (0.047). This suggests that college graduates located in cities with high cost of living enjoy an expected standard of living similar to college graduates with the same observable characteristics located in cities with low cost of living. The reason is that for skilled households, expensive cities offer pre-tax incomes high enough to exactly offset the higher cost of living and personal taxes.

For less skilled households, the picture that emerges is markedly different. San Francisco, New York, and Boston do offer high pre-tax incomes to high school graduates but not high enough to offset cost of living and taxes. On net, standard of living of middle-skill households in these three cities are in the bottom third of the distribution. A regression of consumption by high school graduates on the local price index yields a coefficient of -0.237 (0.026), confirming that expensive cities offer standard of living that are systematically below that of affordable cities. The estimated coefficient implies that a middle-skill household moving from the median commuting zone (Cleveland) to the commuting zone with the highest price index (San Jose) would experience a 7.7% decline in their standard of living. Moving from the commuting zone with the lowest cost of living index (Natchez) to the commuting zone with the highest index would imply a decline in the standard of living by 12.7%.

The negative relationship between consumption and cost of living is significantly steeper for high school drop outs. The slope is -0.391 (0.032), suggesting that for this group standard of living in expensive commuting zones is quantitatively much lower than in cheaper commuting zones. For

households in this group, moving from Cleveland to San Jose implies a 15.6% decline in the standard of living. Moving from Natchez to San Jose implies a 26.9% decline in the standard of living.

Since consumption of college graduates is uncorrelated with local prices, while consumption of less skilled groups declines with local prices, consumption inequality within a commuting zone increases significantly with cost of living. In particular, we find that the difference in standard of living between high- and low-skill households living in the same commuting zone is much larger in expensive commuting zones than affordable commuting zones. This finding appears to validate the growing concerns in expensive cities about the declining standard of living of less skilled residents, who in recent decades have been exposed to increasingly affluent co-residents and higher local prices, raising questions about affordability and gentrification.

An important question for future work is how economically large differences in consumption can exist across communities within the US for less skilled households. The fact that consumption of high school graduates and high school drop-outs declines with local prices, while consumption of college graduates does not, may reflect higher mobility frictions faced by less educated households (credit constraints or lack of information) or stronger idiosyncratic preferences for certain expensive locations. It is difficult to draw strong conclusions on the exact reasons without analyzing local non-market amenities. However, we note that in order for amenities alone to explain the difference in our findings between high- and lower-skill households, it would need to be the case that high school graduates and high school drop-outs enjoy amenities in expensive, large, and well-educated cities more than college graduates—a possibility that we cannot rule out, but goes against prior research (Diamond, 2016).

We stress that the objective of our analysis is the measurement of consumption of *market goods*. We do not seek to quantify spatial differences in utility, which are a function of both market consumption and non-market local amenities, such as weather, crime, air quality, etc. There is a rich literature on amenity differences across cities, while less is known about geographical differences in consumption of market goods. We note, however, that while non-market amenities are a component of utility, market consumption is likely to be a very important component. Any future analysis of utility differences across locations would require estimates of market consumption by area as a key input.

We conclude the paper by studying the effect of geographic sorting of households into high- and low-cost commuting zones on income and consumption inequality in the US as a whole (Moretti, 2013). We find that in the absence of sorting by household characteristics and by city size, the nationwide mean difference between college graduates and high school graduates would be 16 percent smaller for pre-tax income and 10 percent smaller for consumption than the observed difference.

The remainder of the paper is organized as follows. Sections 2 and 3 describe the data and the cost of living indexes. Sections 4 and 5 present our estimates of consumption by income group and by skill level, respectively. Section 6 reports our findings on nationwide inequality. Section 7 concludes.

2 Data and External Validity

The source of our main data is a firm that provides financial software to banks. The data are in the form of transaction-level bank and linked credit and debit card data. In particular, for individuals who have an account in the banks served by the firm, we observe the amount and details of all transactions on the bank accounts and credit card accounts. For example, this includes the expenditure amount and merchant name for all debit and credit card purchases, expenditure and merchant name for all ACH credits and debits into and out of bank accounts, expenditure amount for all checks and cash deposits/withdrawals, and transfers between accounts (including transfers from/to accounts not observed in our data).

The sample includes 3,000,518 households observed in 2014. Selection into our sample is based on which banks the firm that provided the data works with. Our sample includes account holders in 78 banks, including the majority of the largest 10 US banks. For the banks in our sample, we have a random sample of active accounts. An advantage relative to data like Mint.com is the fact that selection into the sample does not depend on user sign-up.

An important limitation of our sample is that we miss unbanked households, which account for 7% of the US population (Federal Deposit Insurance Corporation, 2015) and are over-represented among low-income households. The unbanked will not be part of our analysis.

A second limitation has to do with multiple accounts. If a household has multiple bank accounts within the same bank, then these accounts are linked and we observe them as linked. On the other hand, if a household has accounts at other banks, we do not observe their transactions there. For these multi-banked households, we only have a partial view into their income and consumption patterns. The 2013 Survey of Consumer Finances (SCF) shows that 70% of all banked households maintain their checking accounts at a single bank. The 30% of households that are multi-banked maintain 74% of their checking account balances at the bank that services their "main" checking account. In an effort to focus on primary bank accounts, we restrict the analysis to active accounts. Our data provider uses a proprietary algorithm to identify accounts that are active and we drop all inactive accounts. Ganong and Noel (2019) deal with this problem using the same approach. In addition, we require accounts to have at least $10,000 of annual income and $1,000 of annual expenditures. If these restrictions leave us with households' main bank accounts, we expect to be missing only 7.8% of the average household's income and expenditures.[4]

Given these limitations, a crucial question is how representative our sample is for the population with income above $10,000. We compare our measures of income, consumption, and location to nationally representative established data sources.

2.1 Measuring and Validating Income

We estimate household income as the sum of all deposits into bank accounts excluding transfers between accounts, expense reimbursements, payment reversals, sales returns, and refunds. Since part of federal and income taxes are withheld from paychecks before arriving into a bank account,

[4]We are missing 0% of data for 70% of the sample and 26% of data for 30% of the sample.

for consistency we also exclude from our measure of income federal and state tax payments and refunds. Thus, our measure of income is after-taxes. More details are in Appendix A.

Our measure of income has two limitations, which are likely to be important for households with low levels of true income. First, we cannot observe income that is paid in cash and spent in cash, unless the cash is deposited into the bank account before being spent. Second, we cannot observe some government transfers. Our data include income from Social Security, Disability Insurance, and EITC—since these transfers are deposited into the household's bank account. But it misses Food Stamps and TANF—which in most states are paid through debit cards not linked to a bank account—and housing assistance.

Thus, our data are not great at tracking very low-income household income. Not only are these households more likely to be unbanked, they are also more likely to receive in-kind government transfers. Both these considerations further motivate our restriction to households with income above $10,000. We cannot observe very low income households in our data, so we do not attempt to study their consumption.

In practice, the omission of Food Stamps, TANF, and housing assistance does not appear to be an important source of bias in our context. In Section 5, we analyze the sensitivity of our estimates to including the imputed value of Food Stamps, TANF, and housing assistance and find that our main empirical results do not change.[5]

As a first step in assessing the representativeness of our sample for the population of households with income above $10,000, in Figure 1 we compare the income distribution in our data to the post-tax household income distribution in the 2012-2016 American Community Survey (ACS). To make the ACS data comparable to our data, we run income through TaxSim to calculate post-tax income for each household and drop incomes below $10,000. Our income distribution appears to trace the ACS distribution generally well. Low-income households are slightly underrepresented in our data and high-income households slightly overrepresented—likely reflecting unbanked individuals and the fact that ACS under-reports self-employment and business income (Rothbaum, 2015). The median household income in our data and in the ACS are $52,956 and $48,835, respectively. The difference is 8.4%. In the 2013 SCF, the median income of the banked population is 8.3% higher than the median income of the total population, suggesting the difference in the income distributions in the ACS and our data are mostly due to the missing un-banked households in our data.

In order to compare mean income in each commuting zone obtained from our data to the one obtained from the ACS, in Appendix Figure A1 we plot mean log household income by commuting zone from our data on the y-axis against mean log household income from the 2012-2016 pooled tabulated ACS data. The figure shows a tight relationship between median income measured in the ACS and in our data, with a slope of 0.915 (0.101).[6]

[5]Food Stamps, TANF, and housing assistance are federal transfers with limited geographical variation and therefore limited effect on our estimated coefficients, which are identified by geographical variation.

[6]Measurement error stemming from imperfect geographical matching would lead to a slope of less than one, even if both data sources were representative. Commuting zones are not available in the ACS. In this figure, we assign households in the ACS to CZ's based on PUMAs. The geographical matching is not perfect and likely introduces some attenuation bias in the regression in the Figure.

2.2 Measuring and Validating Consumption Expenditure

We measure consumption expenditure as the total of all transactions flowing out of each household's bank accounts. This includes all checks, cash withdrawals, credit card bill payments, debit card transactions, and ACH (excluding transfers between own accounts and including external accounts).[7]

It is important to benchmark our measure of consumption expenditure against other known measures of consumption expenditure. The most accurate data come from the National Income and Product Accounts. Panel A in Figure 2 reports the average total expenditure per household as reported by NIPA, our bank data, and the other main data source on spending —the Consumer Expenditure Survey. Expenditure measured by NIPA is not exactly comparable to our spending data due to how NIPA treats spending on healthcare. The NIPA data contain health spending that includes both the out-of-pocket component and spending paid by insurers and the government. Our bank data only include out-of-pocket health spending. To make the two data more comparable, in Panel A we subtract out non-out-of-pocket health spending from the NIPA expenditure.[8]

Panel A shows that our bank data closely match NIPA. Our data estimate average household spending at \$74,631. NIPA estimates \$77,533. The CEX estimates \$53,495. The bank data mean is only 4% lower than NIPA's mean, while the CEX mean is 31% lower. The CEX is a survey-based dataset known to significantly under report spending (Sabelhaus and Groen, 2000; Aguiar and Bils, 2015; Sabelhaus et al., 2015).[9] We will return to healthcare in the next sub-section, where we discuss Panel B.

A second way to validate our expenditure data against NIPA and the CEX is to compare the type of consumption expenditure. In Figure 3 we plot expenditure shares by category in the three

[7]We exclude transfers not only between the linked accounts in the data, but also to external accounts using key words listed in the description of the transaction. We also exclude payments for credit card interest.

[8]In this Figure, we estimate non-out-of-pocket health spending in NIPA by taking the NIPA-reported spending on healthcare and net health insurance premiums and multiply it by 0.87, the 2014 share of health care costs that are out-of-pocket, as measured by the Centers for Medicare & Medical Services (CMS, 2021). There is an additional difference between the two datasets in the way NIPA measures housing expenditures. NIPA measures housing expenditure for homeowners by the imputed rental value of their home. Our bank data measure "actual" housing expenditure by homeowners, which includes spending on mortgages and interest, property taxes, and down payments. In this figure, we leave the NIPA data unadjusted as there is not a simple fix to the NIPA data. We will return to these housing expenditure differences in the following subsection. (It will turn out that both of these definitions lead to similar estimates of housing spending.) There are some small differences between the CEX housing and health spending measures and the measures in our bank data. The CEX health spending includes out-of-pocket spending, as well as payroll deductions towards health insurance premiums, but does not include any contributions towards health costs directly from employers or the government. This should lead to more health spending in the scope of the CEX than our bank data, but less than in the raw NIPA data. The CEX housing expenditure definition does not include spending on paying down mortgage principles or on down payments, both of which would be included in our bank data. This makes CEX housing spending slightly less than what our bank data would include.

[9]Our bank data have a slightly different sample than the NIPA data. Our data are restricted to households with bank accounts that earn at least \$10k of post-tax income, while NIPA includes all households. According to the SCF, our sample has 12% higher income than the average US households, which would lead our data to have a higher household spending level than NIPA. On the other hand, we miss spending out of unlinked bank accounts from other banks of multi-banked households. According to the SCF, un-linked bank accounts likely lead us to miss 7.8% of household spending. Combining these offsetting effects suggests we should overestimate spending by about 4%. We end up underestimating spending by 4%, but this bank-of-the-envelope adjustment suggests that we are in the right ballpark.

datasets. Our data classify each transaction in 21 high-level consumption categories based on the identity of the merchant. One limitation is that these categories only exist for goods purchased by credit card, debit card and ACH. When the purchase is paid for by cash or check, we observe the value but not the type of purchase. These transactions are in a category called "Unclassified". Thus, the shares that we report in this figure for our data are computed as the expenditure in each category as a share of total classified expenditures. We will return to these unclassified transactions in the next subsection.

In Panel A, we restrict the comparison to types of expenditure that are measured consistently in all three datasets.[10] It shows that our bank data line up closely with the NIPA expenditure shares. The correlation of the expenditure shares between our bank data and the NIPA data is 0.94. In contrast, the correlation between the bank data and the CEX is 0.52, and the correlation between the CEX and NIPA is 0.64. Bee et al. (2012) show that there is substantial variation in the underreporting rate of consumption across types of spending in the CEX creating poorly measured expenditure shares.

As a third way to probe the quality of our expenditure data, in Figure 4, we compare the fraction of consumption expenditure by mean of payment in our data with estimates from the Federal Reserve Report by Greene and Schuh (2016). On average, our data appear to closely match the corresponding fractions in the general population from the Fed report. The value of all credit card, debit card, and ACH transactions accounts for about 70% of all expenditure in our data, with cash and checks accounting for a smaller fraction. In the figure, we also break down the shares by income group. The Fed does not report these estimates by income.

Fourth, we compare our measures of expenditure for specific publicly traded merchants to corporate sales reported in SEC filings. SEC filings are precise measures of sales for merchants due to the penalties for misreporting, although corporate sales reported in SEC filings do not need to match exactly our measures of expenditure because they include oversea sales while our measures are only based in the US. Figure 5 shows this comparison for Abercrombie & Fitch, Chipotle, Costco, Dunkin' Donuts, Home Depot, Kroger, Macy's, McDonald's, Nordstrom, Starbucks, Walmart, Whole Foods. Overall, our expenditure data appear to track sales trends well.

As a fifth way to probe the quality of our expenditure data, we searched for cases of well-known sudden changes in merchants sales and compared them to our data. In general, our merchant-level expenditure tracks episodes of sudden changes in sales. Just as an example, Appendix Figure A2 shows the changes in expenditure at Chipotle following the *Salmonella*, *E. Coli* and Norovirus outbreaks in 2015. Our data detect an immediate drop in expenditure in the relevant locations.

Overall, we conclude that our measure of consumption expenditures appears to be generally consistent with other nationally representative data sources, both in terms of overall amount and composition.

[10]Since healthcare and housing in NIPA are not defined in the same way as in the other datasets, as previously discussed, we do not include these categories in Panel A, and calculate household average expenditure shares among the remaining categories. We also aggregate some categories with definitions that don't quite line up across datasets into "Other Goods" and "Other Services". See note to Figure for definition of "Other Goods" and "Other Services"

Appendix Figure A3 shows the relationship of log consumption expenditure and log income across the 3,000,518 households in our sample. Households with higher income have higher levels of consumption expenditure. The slope is 0.921 (0.002), indicating that low-income households tend to consume a higher fraction of their income than high-income households, as previously documented by Dynan et al. (2004).

2.3 Measuring and Validating Location

The geographical unit of observation in our analysis is a commuting zone. We do not observe residential address of account holders. However, we observe all transactions made by a consumer, the merchant's city and state and whether a transaction was in-person. In-person transactions include all purchases in physical retail establishments, ATM visits, etc. We assign account holders to commuting zones by taking the modal commuting zone across transactions that take place in-person.

Figure 6 plots the log size of our sample in each commuting zone against the log number of households from 2012-2016 American Community Survey (ACS). There appears to be a tight link between our sample size and the corresponding number of households in the ACS, with an R-squared of 0.81. The slope is 1.340 (0.028), indicating that we under-sample rural areas and over-sample larger cities. This likely reflects the geographical presence of the banks in our sample, which includes the majority of the ten largest banks in the US. These banks' locations are skewed to larger, urban areas. We use weights to adjust for sample representativeness, where weights are the ratio of the number of households in a given commuting zone in ACS data to the corresponding number in our data. However, this weighting does not impact our results.

2.4 Health and Housing Expenditures Adjustments

Two conceptual issues arise in measuring of consumption of health and housing services. First, out-of-pocket health expenditures do not necessarily equal the consumption of health services in any given year because most consumers pay out-of-pocket only a fraction of the actual value of the health services that they receive. Second, while for renters the amount paid on rent in a given year can generally be considered a good approximation of the value of housing services, for homeowners expenditures on housing do not necessarily equal the cost of purchasing one year of housing services. An homeowner who has paid off their mortgage, for example, does not have annual expenditures beside property taxes, but still enjoys housing consumption. Moreover, housing is an asset and its price likely reflects not just its user value but also expectations of future appreciation or depreciation.

These two issues are not specific to our paper but are common to all papers on consumption. The solution typically adopted by the literature is to adjust both expenditures and income (see, for example, Aguiar and Bils (2015)). In practice, this means (a) adding to an household's expenditures and income observed in the data an estimate of the value of their health expenditures that are not out-of-pocket; and (b) adjusting the housing expenditures and income of homeowners to reflect the

value of housing services. We follow the literature, and make these two adjustments.

Specifically, for health expenditures, we augment out of pocket expenditures to account for the value of medical expenditures that are not paid by a consumer directly but are paid by their insurance or the government. We use the Medical Expenditure Panel Survey to measure the relationship between total health care expenditure and out of pocket spending. We use this relationship to impute total health care expenditure for each household given their observed out of pocket spending. We add the estimated extra health care spending both to expenditure and income. Details are in Appendix B.

For housing, we face the additional limitation that the share of total expenditures that is spent on housing cannot be accurately quantified in our data because many consumers pay their rent with checks and mortgages with bill-pay transfers to banks. While the value of these transactions is included in our measure of total expenditure, our data puts them into a category called "Unclassified". To quantify the share of expenditures that is spent on housing and properly measure housing services for homeowners, we adopt the same methodology and same data that the BLS employs in measuring the CPI (Poole et al. (2005); Bureau of Labor Statistics (2007)). Namely, for renters, we estimate housing expenditures by commuting zone and income group using mean contract rent from the ACS. For homeowners, the BLS uses a measure of "rent equivalent" from the CEX, which is defined as the rental value of their home if they were to rent it out.[11] Details are in Appendix B.

After these two adjustments, we return to the comparison of our data to the NIPA and CEX. In Panel B of Figure 2 we compare our adjusted average total expenditure against NIPA and the CEX. (The NIPA data is now in its "raw" format, since we have adjusted our bank data to make the health and housing definitions consistent with those used by NIPA.) The adjusted average household expenditure is \$85,640 in the bank data, which is only 7.7% less than the raw NIPA estimate of \$92,779. Quantitatively, the main reason for why our adjusted average in Panel B is higher than the unadjusted average in Panel A is that the former includes health expenditures that are not out-of-pocket. The housing adjustment for homeowners is quantitatively less consequential on average. In the Figure, we adjust the CEX in a similar way to make it more comparable to NIPA by using imputed rents for homeowners and including healthcare spending paid by employers and the government. The adjusted CEX's mean is at \$66,907, still below NIPA.[12]

Figure 3b reports average spending by category in the CEX, NIPA, and bank data. The NIPA figures are visually close to our adjusted figure. Indeed, the correlation of spending across categories is the highest between NIPA and our data at 0.98.

[11] Homeowners in the CEX are asked the rental value of their home if they were to rent it out. Bee et al. (2012) show that the CEX Interview survey accurately tracks housing expenditure, when validated against NIPA. We take rent equivalent for each income group from the CEX, pooling the 2012–2016 data. We estimate average rental payments for renters by income group in the 2012–2016 pooled ACS. We then average these together, weighted by homeownership rates, to get total housing expenditure. To avoid double-counting, we subtract out the actual spending on housing from our unclassified spending and add back our estimated cost of a year of housing services.

[12] The baseline CEX average household spending is \$53,495. CEX reported homeowner costs are \$6,149 and estimate imputed rent is 10,896. NIPA estimates healthcare paid by employers at \$2,382 per household and \$6,284 paid by the government. Our adjusted CEX household expenditure is thus: 53,495-6,149+10,895+2,382+6,284=66,907.

Taken together, Figures 2 and 3 indicate that our data match well NIPA data, both in terms of overall mean household consumption and in terms of consumption by category. This appears to be the case for the raw data on the left and the adjusted data on the right.

2.5 Summary Statistics.

We classify households into three income groups (based on unadjusted income): low \$10,000-\$50,000; middle \$50,000-\$200,000; and high >\$200,000. In our empirical analysis, we only include commuting zones for which we have at least three low-income, three middle-income, and three high-income households. We end up with 443 commuting zones, accounting for 96.3% of US population. More details on the construction of the sample are in Appendix A.

Panel A in Appendix Table A1 shows summary statistics by income group. Our final sample includes 1,368,817 low-income; 1,449,978 middle-income; and 181,723 high-income households. Panel B is for adjusted expenditures and income —they are both higher due to the addition of health expenditures that are not out-of-pocket.

3 Local Cost of Living Indexes

The Bureau of Labor Statistics (BLS) releases an official Consumer Price Index (CPI-U) for the entire US. This index is not informative of price differences across space. We use the same methodology to create price indexes that vary across commuting zones and across income groups. This allows us to deflate the consumption expenditure of households in a given city and income group by the relevant price level. While the Bureau of Economic Analysis (BEA) produce annual estimates of local prices that cover some commuting zones, their indices do not vary across income groups. Since preferences vary across the income distribution, it seems important to estimate local price indexes that vary by income strata (Jaravel, 2019; Handbury, 2019).

3.1 Baseline Price Index

The BLS uses a Laspeyres index to calculate the CPI-U. This is defined as the average price change between period t and $t+1$ across a representative consumption bundle of goods, weighted by the average expenditure share of each good, measured in period t (Chapter 17 in Bureau of Labor Statistics, 2007). For our main analysis, we closely follow the methodology that the BLS uses to build its official CPI, but we generalize it to allow our index to vary across commuting zones and across income groups. Our baseline Laspeyres price index for commuting zone j and income group k is defined as:

$$P_{j,k}^{\text{Laspeyres}} = \sum_{i \in I} \frac{p_{i,j}}{\overline{p_i}} \cdot s_{i,k} \tag{1}$$

where $p_{i,j}$ is the price of good i in commuting zone j; $\overline{p_i}$ is the price of good i in the reference commuting zone: Cleveland, OH; $s_{i,k}$ is the nationwide average expenditure share of income group k on good i; and I is a set of consumption categories of goods and services. By allowing the

expenditure shares to vary by income group, we allow for preference heterogeneity across the income distribution. We choose Cleveland as the reference city because its monthly rent for a given vector of housing characteristics is roughly equal to the median rent across all commuting zones in our analysis sample. This normalization implies that the price index for Cleveland is by construction equal to 1 and that the indexes from other locations are to be interpreted as relative to Cleveland.

A desirable property of the Laspeyres index is that it is a first-order approximation of the true price index, but does not require us to specify the functional form of the utility function or estimate its structural parameters. A less desirable property of the Laspeyres index is that it is only a first-order approximation and does not capture higher order effects.[13] The Laspeyres index also does not easily allow for variation in variety and supply of goods and services across space (Handbury and Weinstein, 2015; Handbury, 2019). For these reasons, in the next sub-section we discuss several alternative indexes based on alternative assumptions.

To estimate Equation 1, we need data on local prices ($p_{i,j}$) and expenditure shares ($s_{i,k}$). Here we describe the general approach. We provide more details in Appendix C.

Measuring Prices of Consumption Items. To measure prices $p_{i,j}$, we combine data from NielsenIQ, ACCRA, and the ACS. First, we use price data from the 2014 NielsenIQ Retail Scanner data for six consumption categories: Grocery, General Merchandise, and Personal Care; and three additional categories for which we can find a one-to-one map to a product group in NielsenIQ: Baby Needs, Electronics, and Office Supplies. For these categories, we observe prices at the twelve-digit barcode level (UPC). To measure the local price for each of the 116 product groups in the data, we regress log prices on a UPC fixed effect and a dummy for each community zone. We run a separate regression for each product category. We use the commuting zone dummies as our estimate of local price for each product group. In this approach, quality is held constant since we are comparing the price that consumers in different commuting zones pay for the same twelve-digit barcode level product. Since NielsenIQ prices are reported before taxes, for goods categories that in a given state are subject to sale tax, we add the relevant sales tax.[14]

Second, we purchased data on prices from ACCRA, which is collected by the Council for Community and Economic Research. We use ACCRA prices for nine consumption categories: Automotive Expenses, Clothing/Shoes/Jewelries, Gasoline/Fuel, Healthcare/Medical, Hobbies/Entertainment,

[13]It is well understood that Laspeyres indexes are subject to substitution bias, where the true price index would account for the utility benefits of allowing consumers to substitute away from high price goods. By the envelope theorem, this substitution effect does not have a first-order welfare effect, but could matter for large price changes. Separately, we note that while our index is based on the Laspeyres price index used by the BLS to measure inflation over time, there are some conceptual differences in comparing prices across many geographic locations and across a pair of time periods. The standard Laspeyres index is defined for comparing a pair of time periods (or cities). However, it is ill defined for comparing a set of cities simultaneously, since the pairwise price differences between a pair of cities a and b multiplied by the price differences between cities b and c does not equal the Laspeyres price index between cities a and c. When comparing prices across many cities at once, there is no obvious "base city" to choose to use expenditures from. Instead we average the expenditures together across all cities and use this as the weights for the price differences across cities. This style index is sometimes called a Stone index. This is a well-known issue in the purchasing power parity literature that compares prices across countries (Deaton and Heston (2010)). We will draw on the methods developed in the PPP literature as robustness.

[14]We collect sales tax data from 2014 from Walczak and Cammenga (2021).

Home Maintenance/Improvement, Restaurants/Dining, Telecommunications, and Utilities. Within each category, ACCRA data report the price of specific products, with the quality of the product held constant across locations. For example, an item in the Automotive Expenses category is "Tire Balancing". ACCRA reports the price of "Tire Balancing" in each city for a specific type of tires.[15] For goods categories that in a given state are subject to sale tax, we add the relevant sale tax.

Third, we use the 2012–2016 ACS data (centered on 2014) to measure housing costs. In computing the CPI, the BLS uses rents to measure the cost of housing since they are arguably a better measure of the user cost than house prices, and we do the same. Houses are assets, and their prices reflect both the user cost as well as expectations of future appreciation. To account for different types of housing across locations, we estimate a household-level hedonic model where we regress the monthly contract rent excluding utilities on a vector of commuting zone identifiers; and a vector of housing characteristic, including the number of bedrooms, rooms, units; year the structure was built; and presence of kitchen and plumbing. We predict monthly rent at the commuting zone level using the commuting zone fixed effects.

For the remaining six consumption categories—Charitable Giving, Education, Financial Fees, Insurance, Printing and Postage, and Travel—we have no data on geographical variation in prices. We assume that their prices do not vary geographically. This assumption may be violated in practice and the magnitude of any resulting bias is a function of how important these categories are. The sum of expenditure shares of these items for low-, middle-, and high-income households are 7.4%, 11.4%, and 18.5%, respectively.

Measuring Expenditure Shares. The expenditure shares for 22 high-level consumption categories by income group and are listed in Appendix Table A2. As we have shown in Section 2, these shares match well the NIPA shares. To get a sense of how spatial variation in prices contribute to spatial variation of our overall-income index, column 5 reports standard deviation of prices across commuting zones. Categories that exhibit large variation in prices include Hobbies and Entertainment (0.34); Healthcare/Medical (0.32); and Housing (0.30). By contrast, Gasoline/Fuel (0.06) and Electronics (0.03) have much lower geographical variation.

Three categories among the 22 in our data are very broad: Grocery, General Merchandise, and Personal Care. To improve precision, we use data from the 2014 NielsenIQ Consumer Panel Survey to obtain a more refined product definition nested within each of these three categories. For example, NielsenIQ identifies 17 subcategories within the Personal Care category: Cosmetics, Deodorant, Vitamins, etc. The shares for each subcategory are shown in Appendix Table A3.

3.2 Alternative Price Indexes

The choice of using a Laspeyres index as our baseline is motivated by the fact that it is the index used by the BLS to compute the official price index. The index in Equation 1 is a useful and transparent starting point. But it is not the only possible index we can use. For one, while it allows expenditure

[15] One limitation of the ACCRA index is that it is based only on a limited number of goods. Another limitation is that the sample size within each city—the number of observations per item—is smaller than the BLS sample size.

shares to vary across income groups, it does not account for the utility benefits of re-optimizing one's expenditure shares due to variation in local prices. By the envelope theorem, this cannot have a first order welfare effect, but it could matter when comparing cities with very different local prices. The baseline index also restricts all income groups living in the same commuting zone to face the same set of prices. In principle, it is conceptually correct to think of consumers in a city as facing similar prices. But in practice, housing segregation within a commuting zone combined with segregation in the clientele of merchants patronized by high- and low-income families may result in high- and low-income consumers facing different prices. In addition, the Laspeyres index rules out differences in the choice set across areas. Product variety differs across cities, as some goods exist in some commuting zones but not in others, and it been shown to be quantitatively important for measuring local prices (Handbury and Weinstein, 2015; Handbury, 2019).

Following Jaravel (2019); Handbury and Weinstein (2015); Deaton and Muellbauer (1980), and methods developed to measure purchasing price parities across countries (Deaton and Heston, 2010), we present additional estimates based on six alternative indices: the Törnqvist index, the price index implied by a CES utility function, the price index implied by a nested-CES utility function that accounts for variation in the variety of goods and services supplied in each CZ, the price index implied by an estimated EASI demand system, and two indices developed by the purchasing power parity literature: the Geary-Khamis price index and the GEKS-Fisher price index. We discuss the conceptual differences between these indices here, and the full details of their construction in Appendix C.

The Törnqvist index is a geometric average of expenditure shares between the nationwide average and the specific CZ. This has been shown to be a second-order approximation to the true price index. The CES index assumes that underlying preferences within an income group have constant elasticity of substitution across all product categories. That elasticity is implicitly inferred from a transformation of the expenditure shares within each CZ. The benefit of this index is that is an "exact" index, not an approximation. The downside is that it is only exact if the true underlying utility function is CES.

The Nested CES allows for more complex substitution patterns between products. There is an elasticity of substitution between the 22 high-level expenditure categories, and then expenditure category-specific substitution elasticities across product groups. Finally, within each product group, there is a product-group specific substitution elasticity between unique varieties of products. In addition, the nested-CES also accounts for differences in the choice set across commuting zones by allowing for differences in product variety. We follow Handbury and Weinstein (2015) and Broda and Weinstein (2010) in building the nested CES index and correcting for variety. To measure local variety we use the number of unique UPC codes sold in each CZ as observed in the NielsenIQ RMS (store sales) data. For product categories not covered by NielsenIQ, we use the number of unique merchants that we observed transacted at within each CZ in our bank data. We explore two choices of the elasticity parameter σ that have been found in the literature: 7 (Montgomery

and Rossi, 1999) and 11.5 (Broda and Weinstein, 2010).[16]

In addition, we estimate an approximate Exact Affine Stone Index (EASI) implicit Marshallian demand system as developed by Lewbel and Pendakur (2009) which generalizes the popular AIDS demand system (Deaton and Muellbauer, 1980). We follow Lewbel and Pendakur (2009) methods to estimate an approximate EASI demand model and derive the price index implied by the demand model estimates.

Finally, our last two alternative indexes are based on purchasing power parity (PPP) methods. The Geary-Khamis index is a Paasche index that compares the local prices in a given CZ to nationwide average prices. The weights on the relative prices differences between the CZ and the nationwide average are equal to the focal CZ's expenditure shares. This is the method used by the BEA to estimate local price indices. The second PPP index we estimate is the GEKS-Fisher index.[17]

All price indexes we have discussed so far assume that prices for all goods vary only across commuting zones ($p_{i,j}$). In addition to spatial price variation, we build another set of indexes that allows prices for some goods to vary across income groups within the same commuting zone ($p_{i,j,k}$). We recompute all these price indices discussed above with income group specific prices within each CZ.

3.3 Facts About Geographical Differences in Cost of Living by Income Group

Table 1 shows the 15 most expensive commuting zones, the 5 commuting zones around the median, and the 15 least expensive commuting zones based on our baseline indexes. Throughout the paper, we label each commuting zone using the name of its largest city, instead of the official commuting zone name. For low income families, the most expensive commuting zones are San Jose, CA; San Francisco, CA; and San Diego, CA where the low-income price index is 1.491, 1.477, and 1.410, respectively. This implies that prices faced by low-income residents of these cities are 41% to 49% higher than prices faced by low-income residents of Cleveland (which has index equal to 1 by construction). Other expensive commuting zones include Honolulu, HI; New York, NY; and Newark, NJ. The least expensive commuting zones for low-income residents are London, KY; Gallup NM; and Natchez, MS with price indexes equal to 0.758, 0.755, and 0.749, respectively.

The geographical price differences revealed by our price index differences are economically large, and this is particularly true for low income families. The overall cost of living in San Jose is estimated to be 99% and 49% higher than the cost in Natchez for low- and high-income households, respectively, suggesting that the range of prices that low income families are exposed to is much wider than the range of prices that high income families are exposed.

[16] Redding and Weinstein (2020) propose approach to measuring the cost of living for CES preferences that treats demand shocks as taste shocks that are equivalent to price shocks.

[17] A Fisher index is the geometric mean of a Laspeyres and a Paasche price index for a given pair of cities. The Fisher index is a second-order approximation for the true price index but it is only defined for pairs of cities, and it is not transitive. This means the Fisher index between cities A and B, multiplied by the Fisher index between cities B and C does not equal the Fisher index between cities A and C. The GEKS-Fisher index uses these pairwise Fisher indices to estimate price indices that impose transitivity.

Figure 7 displays the spatial dispersion of the our price indexes across all 443 commuting zones. Two features of this Figure are particularly interesting. First, the cost of living index for low-income households exhibits significantly higher spatial variation than the index for high-income households. The standard deviation equals 0.119 and 0.073 for the low- and high-income group, respectively. The 75-25 and 90-10 percentile differences for low-income households are nearly twice as large as the corresponding differences for high-income households. This finding likely reflects the fact that low-income households put higher weights on housing expenditure, which is the item in the consumption basket whose price varies the most across commuting zones.

Second, the figure shows that the distribution is far from symmetric, but highly skewed to the right for all three income groups. While the mass of the distribution is concentrated between 0.8 and 1.2—indicating that most cities have an index that is between -20% and +20% of the median—there are a handful of expensive cities in the right tail, where cost of living is much higher. For low-income families, there are 32 commuting zones with cost of living that is more than 20% above the median and 16 commuting zones with cost of living that is more than 30% above the median. Similar skewness is present for other income groups.

The consumption item that is most responsible for the spatial variation in the cost of living indexes is housing, since its share of consumption is the largest and its price varies over space more than the price of any other goods. By contrast, product categories with lower shares of consumption and smaller geographical variation in prices—Grocery or Electronics, for example—contribute much less to the spatial variation in the indexes. A regression of the log of the index on log rent yields coefficients of 0.269 (0.013) and 0.462 (0.011) for high- and low-income households, respectively (Appendix Table A4), while the share of housing in the indexes of high- and low-income households is 0.147 and 0.335, respectively. If the only source of geographical variation in prices of consumption items were housing costs, and all other items had the same price nationwide, we would find the coefficients equal to these shares. The fact that the coefficients are higher reflects the fact that the prices of non-housing nontradables tend to be higher in areas with more expensive land. In turn, this reflects the fact that it costs more to produce nontradable goods and services in areas where land is more expensive (Choi and Jo, 2020). For example, the cost of a haircut or a slice of pizza is higher in San Jose than in Cleveland, holding quality constant, because retail space and labor are more expensive in San Jose.[18] That said, our findings indicate that local price indices can be well-approximated by using data only on local housing costs, weighted appropriately, especially for the low-income group. The R-squared of the regression of our low- and high-income price indexes on housing rent are 0.97 and 0.89, respectively.

Alternative Indexes. These findings are generally similar if we use our alternative price indexes. The reason is that in practice all indexes are highly correlated with one another. This is shown in Appendix Table A5, where we report the correlation matrix of all the variants of all the indexes used in this paper. For parsimony, the table focuses on the indexes for all consumers. Based on

[18] Using NielsenIQ data, DellaVigna and Gentzkow (2019) find limited variation in grocery prices across locations, despite wide variation in consumer demographics and competition.

the 120 pairwise combinations formed by 16 different indexes (A1 to D), the associated correlation coefficients have a mean of 0.966 and a standard deviation of 0.027. The bottom row in the Table shows that our baseline price index has a correlation of 0.93 with the BEA price index, at least for the geographical areas that are covered by the BEA.

Appendix Table A6 quantifies the spatial dispersion of all our alternative price indexes and it compares it to our baseline index. For all indexes, the low-income index exhibits highest variation, followed by the middle-income index and the high-income index, respectively. Indexes that allow different prices across income groups within a location (Panel B) look generally similar to indices that hold prices fixed across income groups (Panel A).

Differences across commuting zones in the variety of products that are locally available are potentially important. Using data on on grocery products from NielsenIQ, Handbury (2019) and Handbury and Weinstein (2015) have shown that correcting for differences across cities in product variety has a large impact on measured prices. Handbury (2019) shows that the correlation of the variety-corrected price index and city income is negative — so that richer cities have lower effective prices — while the price indices that don't include the variety adjustment show a positive correlation of city income. We are able to replicate this finding with our variety-corrected nested CES index applied to Grocery, General Merchandise and Personal Care, as shown in Appendix Table A7. However, Appendix Table A7 also shows that this negative correlation does not hold for the overall price index. The reason is that the majority of other types of products—housing, gas, utilities, clothing, automotive, and healthcare—do not exhibit this inverse relationship when adjusting for variety.

The Nested CES index is highly correlated with our baseline index, as shown in Appendix Table A5. Below, we find that our empirical findings are not sensitive to using one or the other.

4 Geographical Differences in Consumption by Income Group

Consumption expenditures are unevenly distributed over space and this geographical variation is particularly pronounced for low-income households. The maps in Appendix Figure A5 show mean consumption expenditures in each commuting zone for which we have data. However, the maps are not informative of local consumption. Differences across areas in consumption expenditures reflect not just the quantity of goods consumed by the area residents, but also local prices. It is possible that areas with high expenditures enjoy a lower consumption than areas with low expenditures if local prices are high enough to more than offset the higher expenditures.

We use our price indexes to deflate expenditures and quantify "real" consumption for each commuting zone and income group. This allows us to study how real consumption varies across commuting zones as a function of local cost of living and compare the standard of living experienced by the residents of expensive and affordable commuting zones, *holding nominal income fixed* (Subsection 4.1). We then examine how the quantity consumed of specific grocery items, measured not in dollars but in physical units or weight vary as a function of local prices as "model free" evidence of consumption differences (Subsection 4.2). We also study how measures of financial distress for

low-income families vary as a function of cost of living (Subsection 4.3).

We reiterate that our goal is to measure *market consumption*. While quantifying geographical differences in market consumption is arguably an important step in ultimately understanding geographical differences in utility, in this paper we do not seek to quantify differences in utility. Estimating the value of non-market local amenities would require a separate analysis. However, we note that while non-market amenities are certainly a relevant component of utility, market consumption is likely to be a fundamental component. The methods needed to simultaneously estimate the value of non-market amenities and market consumption of beyond the scope of this paper and left to future research.

4.1 Overall Consumption

To estimate mean consumption in a given commuting zone and income group, we deflate consumption expenditure $C_{h,j,k}$ of household h in commuting zone j and income group k by dividing it by the relevant income-group-specific and commuting-zone-specific price index $P_{j,k}^{\text{Laspeyres}}$. We then estimate the following model:

$$\ln(C_{h,j,k}/P_{j,k}^{\text{Laspeyres}}) = \delta_{j,k} + \beta_k \ln Y_{h,j,k} + \varepsilon_{h,j,k} \tag{2}$$

where the vector of commuting zone-income group fixed effects $\delta_{j,k}$ represents our estimates of the conditional mean log consumption in commuting zone j of income group k; and $Y_{h,j,k}$ is household h adjusted post-tax income. We run this regression separately by income group and we condition on household income to control for possible income differences within income groups across cities. For example, low-income households in expensive cities may have a higher income than low-income households in affordable cities. We report estimates where consumption is evaluated at post-tax income equal to \$30,000, \$80,000, and \$285,000 for low-, middle-, and high-income consumers, respectively.

Table 2 shows the 15 commuting zones with the highest level of consumption, 5 commuting zones in the middle of the distribution, and the 15 commuting zones with the lowest level of consumption for low-income and high-income households.[19] The consumption levels are priced at the median cost city, Cleveland, OH and real consumption is measured by the expenditure a household would need to spend in Cleveland to achieve the same utility *from market consumption* as their actual bundle consumed in their city of residence.

The three commuting zones with the highest consumption for low-income households are Huntington, WV; Johnstown, PA; and Elizabeth City, NC. They have level of consumption measured in real terms equal to \$47,270; \$47,267; and \$47,068. Other examples of commuting zones that offer high consumption of low-income households are Mobile, AL and Charleston, WV. Three cities with the lowest consumption are San Diego, CA; San Francisco, CA; and San Jose, CA. The cor-

[19] In this table, to limit the role of sample error, we report empirical-Bayes shrunken estimates. In practice, we calculate $\hat{Y}_i^{shrunk} = \omega_i \cdot Mean(\hat{Y}_i) + (1-\omega_i) \cdot \hat{Y}_i$, where $\omega_i = SE^2_{\hat{Y}_i}/(Var(\hat{Y}_i) - Mean(SE^2_{\hat{Y}_i}) + SE^2_{\hat{Y}_i})$. We also restrict this list to CZ that have at least 20 households in our sample for each income group to further restrict the role of sample error.

responding values are $28,227; $27,304; and $27,139. Other examples of commuting zones with low consumption are Honolulu, HI ($28,598); New York, NY ($29,015); Washington, DC ($29,542); Seattle, WA ($30,527); and Los Angeles, CA ($29,468).

The geographical differences in standard of living are economically large. Low-income households who live in the top commuting zone in the top group enjoy a level of consumption that is 74% higher than households with the same income who live in the bottom commuting zone in the bottom group.

Note that despite the fact that we are holding income fixed, there is not a one-to-one correspondence between estimated consumption and price index. The reason is that households adjust the share of income that they devote to consumption vs. savings. The consumption share of low-income households in expensive areas is higher than in cheaper areas—a point that we will come back to below.

The right panel shows the corresponding estimates for high-income households. Three cities with the highest consumption for this group are Huntington, WV; Johnstown, PA; and Toledo, OH. They have level of consumption measured in real terms equal to $290,538; $290,219; and $290,030, respectively. Three cities with the lowest consumption for high-income households are San Diego, CA ($197,459); San Jose, CA ($196,000); and Honolulu, HI ($192,947). Other cities in this category include New York, NY ($207,563); Boston, MA ($211,947); Seattle, WA ($205,111); and San Francisco, CA ($205,274). The group of commuting zones with the lowest consumption for high-income families overlaps in part with the group of commuting zones with the lowest consumption for low-income families. Expensive places like Boston, Honolulu, New York, San Jose, and San Francisco are associated with lower consumption by both high- and low-income households.

To see more systematically the relationship between consumption and cost of living, Figure 8 plots log consumption expenditures (top panel) and log consumption (bottom panel) against log income-group-specific price index across all 443 commuting zones in our data. The top panel indicates that low-income families have higher consumption expenditures in expensive cities, whereas consumption expenditures for middle- and high-income families tend to be unrelated to cost of living. In the bottom panel, the relation is negative for all three groups, indicating that households in more expensive areas consume less than households with the same income in less expensive areas.

The elasticities of consumption with respect to income-group-specific local prices are economically large, and confirm large differences in the amount of consumption that households can afford in cheap and expensive communities. Specifically, the elasticities are -0.900 (0.009), -0.978 (0.019), and -1.016 (0.037) for low-, middle-, and high-income households, respectively. This means that a 10% higher cost of living index results in a 10.2% lower consumption for high income households, and 9.0% lower consumption for low income households. The effect of cost of living on consumption needs to be interpreted as an income effect, as opposed to a price effect (assuming that most consumers expect to be in their current city for a long time). If utility is locally homothetic, elasticity of consumption with respect to the price index should be -1 — meaning that a 10% higher cost of living index is equivalent to a 10% lower income, implying a 10% lower consumption. We cannot

reject that this holds for high-income consumers and for middle-income consumers, while we can reject it for low-income consumers (p-value = 0.0001).

Thus, the consumption of high- and middle-income consumers is more sensitive to local prices than the consumption of low-income consumers. There are a number of possible, not mutually exclusive explanations for this finding. First, low-income households might be closer to a minimum subsistence level, and small consumption cuts may cost more in terms of utility. This would also imply lower levels of savings by low-income households in expensive commuting zones compared to low-income households in affordable commuting zones—a fact that we confirm below. Second, it is possible that the share of low-income households expecting future income gains is larger in expensive cities than in affordable cities (compared to the relative shares of high-income households). Alternatively, the share of low-income households in expensive cities expecting to move to affordable cities is larger compared to the share of high-income households.[20]

This finding is generally robust to using alternative price indexes. In particular, Appendix Table A8 shows the elasticities of consumption with respect to income-group-specific local prices based on all our alternative price indexes. In every row, the elasticity for high income households remains significantly higher in absolute value than the elasticity for low income households, confirming that irrespective of the specific index used, the consumption of high-income consumers is more sensitive to local prices than the consumption of low-income consumers.

4.2 Consumption of Specific Goods Measured in Physical Units

We have found that households in more affordable commuting zones enjoy a lower level of market consumption than households with similar incomes in more expensive commuting zones. We now use NielsenIQ data to replicate the analysis focusing on the consumption of specific goods, where consumption is measured in number of physical units or weight. For example, we measure the number of cans of beer, the number of light bulbs, or the number of pounds of nuts purchased in a year by a NielsenIQ consumer. Unlike the previous sub-section, we don't need any deflation, because we observe physical quantity of consumption directly from the raw data.

The sample includes 59,755 households in the 2014 NielsenIQ Consumer Panel data with income above $10,000. A product is defined as a twelve-digit UPC. There are 823,507 UPC codes in the data, in 116 product groups. As UPCs may come in different units within a product group, we convert all UPCs within each product group to have the same unit as the most prevalent or "modal unit" within that product group following Allcott et al. (2019).[21] We assign 0 to households that did not purchase that product in 2014. To allow a comparison of the coefficients across products, we divide the household consumption of each UPC by its nationwide income-specific-group mean, and we use this mean-adjusted quantity as the dependent variable. This allows for an elasticity

[20] The existence of government assistance may allow low-income households to cut savings, since future consumption is at least partially insured.

[21] Whenever direct conversion is possible, e.g., from pound to kilogram or from liter to kilogram we do this directly. When direct conversion is not possible (e.g., from kilogram to count), we assume that the log of quantity has the same underlying distribution across different units within the same product group, equate their z-scores across these different units, and then convert all non-modal units to the modal unit. See Appendix D for details

interpretation. We regress the mean-adjusted quantity of a good purchased by a household in a year on the log of the price index controlling for household income; presence of children; type of residence; household size; head's age, gender, race, marital status, education, and employment status.[22]

We run one regression for each product. To summarize the results, we compute the average of the estimated elasticities for each of the 116 product groups in our data. Table 3 shows some examples. The first row reports results for the consumption of carbonated beverages measured in kilograms per year. Entries in this row are the average elasticity across all types of carbonated beverages. They indicate that the elasticity of consumption of carbonated beverages with respect to the cost of living index is -0.790 (0.080) and -1.277 (0.170) for low- and high-income households, respectively. For many products in the table, we find that households cut their consumption as local prices increase and that the magnitude of the elasticity is increasing with income.

Of course, it is difficult to draw strong conclusions based on selected examples. Figure 9 plots the distribution of all the 116 estimated elasticities—one for each product group—weighted by average household expenditure on these product groups. Three features of the figure are noteworthy. First, the majority of the coefficients are negative, confirming a lower level of consumption in more expensive commuting zones. Of the 116 coefficients, 64%, 66%, and 75% are negative for low-, middle-, and high-income households.[23] Second, the effects appear to be more negative for high income households than low income households. This indicates that low income households in expensive cities cut consumption of grocery goods less than high income households in expensive cities, consistent with what we observed for overall consumption. The median values for the low-, middle-, and high-income groups are -0.067, -0.113, and -0.289. Third, the elaticities for all three groups are much smaller than the elasticities estimated for overall consumption above. This likely reflects the fact that most of the consumption items in the NielsenIQ data are grocery items and many grocery items are necessities. When faced with higher cost of living, households seem to cut consumption of necessities less than consumption of all other goods. In addition, groceries exhibit less geographic price variation, making them a relative bargain in expensive cities.[24]

NielsenIQ data have excellent product detail, but cover only a subset of the consumption bundle. We replicate the analysis using information on consumption of some non-grocery items from our bank data. Recall that our data is at the transaction level. For some types of goods, we can measure quantity purchased by counting the number of transactions. For example, we can measure the number of times a consumer buys gas, the number of online subscription service payments like Netflix or Hulu, and the number of times they go to the movies, as long as they pay with credit or debit card. By contrast, a swipe at stores like Costco or Walmart is not be informative in this

[22] In the NielsenIQ data income is top-coded at $100,000. For this part of the analysis, we define middle- and high-income households as having income $50,000-$100,000; and above $100,000, respectively.

[23] Of course, some of the variation in our estimates of the coefficients reflects small sample noise.

[24] We have run similar regressions relating the price of each goods to the cost of living index, controlling for the same set of household characteristic indicators. We find that goods are more costly in expensive areas than in cheap areas. This is despite the fact that virtually all grocery goods can be considered traded. The median estimated price coefficients on the index for the low-, middle-, and high-income groups are 0.072, 0.109, and 0.188.

respect, as we do not know what goods are purchased within the transaction.

In Table 4, we report the results for 7 groups of goods. The dependent variable is the number of a specific type of transactions each household makes in 2014 divided by the relevant income group mean, which is also reported in the table. Like for NielsenIQ, we assign 0 to households that did not purchase a given product group in 2014. We control for log household income and indicators for number of unique card accounts within the household. For most goods, we find that quantities consumed are lower in commuting zones with higher local prices. The estimated elasticities tend to be larger than the ones found for groceries. The largest elasticities are the ones for gasoline and fuel; movies; and streaming services.

Overall, from Figure 9 and Table 4 we conclude that households in more expensive commuting zones tend to have significantly lower levels of consumption of grocery items and some non-grocery items than households with the same income level in less expensive commuting zones. The evidence in this section measures consumption in physical units and does not depend on deflating expenditures by a price index. This provides some "model free" evidence of dramatic consumption differences across space and generally confirms the evidence on overall consumption in the previous subsection.

4.3 Negative Savings and Measures of Financial Distress

We have seen that households in expensive areas can afford a lower level of consumption compared to households with the same income in cheaper areas. It is possible that as a result, some low-income households in expensive areas have trouble in making ends meet, or more precisely they experience a level of consumption expenditures that in some years exceeds their available income. We use our data to estimate the fraction of low-income households in each commuting zone who have zero or negative savings—defined as having yearly consumption expenditures equal to or larger than yearly annual income—and ask whether this fraction tends to be higher in more expensive areas.[25]

Panel A of Figure 10 shows that low-income households in expensive commuting zones have a higher probability of negative savings than low-income households in cheap commuting zones. The slope is 0.072 (0.011).[26]

Of course, having negative savings in a given year does not necessarily imply financial distress since the household may be borrowing from expected future income to smooth consumption across years. A cross-section is poorly suited to draw strong conclusions on the dynamics of consumption. But Panels B and C in the Figure show a positive relationship between the price index and the share of income spent by low income households on overdraft fees (where overdraft fees are identified from entries in bank account statements); and the existence of bankruptcy fees as a proxy for bankruptcy

[25]For this analysis we use "raw" expenditure and income, meaning don't use imputed rents for housing or add in extra healthcare spending. We want to measure savings out of market income.

[26]Like we did in Equation 2, we regress each household indicator for zero or negative savings on CZ indicators controlling for log household income. The Figure plots the coefficients on CZ indicators against the low income price index.

(where bankruptcy fees are identified as transactions that contain the words "bankrupt").[27] Overall, we conclude that for low-income households, the probability of financial distress appears to be higher in more expensive cities, although the magnitude of these estimates is admittedly hard to interpret.

5 Where is Standard of Living the Highest? Expected Consumption by Skill-Level

The analysis in the previous section identifies average consumption by city *for a given income level*. The analysis is useful because it is informative of the differences in standard of living across cities experienced by current high- and low-income residents as a function of differences in the local prices and spending.

But of course the income level that a specific household can attain varies significantly across space: for a given level of human capital, some cities offer high labor earnings (and therefore high income), while others offer low labor earnings (and therefore low income). Ultimately, a household's standard of living in a given city is determined by the relation between the income level that it can achieve there and the local cost of living. In equilibrium, local earnings and cost of living are jointly determined and typically, cities that offer higher labor earnings tend to have higher cost of living, while cities that offer lower labor earnings tend to have lower cost of living.

In this section we seek to measure the standard of living that low-, middle-, and high-skill households can expect in each US commuting zone, once we account both for geographical variation in cost of living (as we did in the previous section) and also for geographical variation in expected income. The analysis in this section complements the analysis in the previous section because it allows us to answer the question of where in the US a household with a specific level of human capital can expect the highest standard of living. In turn, this allows us to understand how standard of living of each skill group varies across space as a function of local costs of living. We seek to answer the following two questions: (a) Is standard of living higher or lower in cities where income and prices are high, compared to cities where income and prices are low? (b) Is the relationship between standard of living and local cost of living the same for high- and low-skill households? We focus on three skill groups, based on the schooling level of the household head: (i) 4-year college or more; (ii) high school or some college; (iii) less than high school.

The analysis in the previous section is descriptive in nature and does not require any assumptions on how income is generated or how it may vary across cities, since it takes income of residents as observed in the data. By contrast, the analysis in this section inevitably requires an assumption on how income of households that in the data are observed in a given commuting zone may vary if they were to move to different commuting zone. We use 2012–2016 ACS data in combination with estimates from the previous Section to predict the income and consumption that a given household may expect in each commuting zone as a function of education and demographics, under the assumption of selection on observables.

To do so, we first use our previous estimates of consumption by income level to assign to each

[27] This is consistent with Keys et al. (2020), who find a sizable effect of location on personal bankruptcy.

household in the ACS an estimate of their expected consumption. Specifically, we bin our bank data into 20 income ventiles by commuting zone. Similarly, we assign households in the ACS to income ventiles by commuting zone using the same income bounds. For each household in the ACS, we then take a random draw of expenditures-to-income ratios from our bank data given income ventile and commuting zone and multiply the drawn ratio by household post-tax income to obtain consumption expenditures (and therefore consumption).

Next, for each commuting zone and skill level, we predict mean pre-tax income, post-tax income and consumption holding constant the other observable characteristics of the household. To do so, we define household types as the combination of characteristics of the household head and spouse (if present): mean age of head and spouse; gender; race; Hispanic origin; education; marital status; and number of children. We end up with 664 household types. We run household-level regressions of pre-tax income, post-tax income or consumption on commuting zone indicators and 664 indicators for household types. For each skill group, we then predict mean pre-tax income, post-tax income or consumption evaluated at nationwide weighted-average fractions across types. We provide more details in Appendix E.

This approach assumes that there are no systematic geographical differences in unobserved determinants of household income across cities, conditional on household observable characteristics, or if there are, they are uncorrelated with local prices. While the assumption of sorting on observables is widely used in the literature, we caution that this is a strong assumption and that sorting on worker effects has been shown to explain some of the geographical differences in earnings across US cities (Card et al., 2021). A violation of this assumption would occur, for example, if households located in more expensive cities tend to have better unobservable determinants of household income than households with the same combination of education, age, gender, race, Hispanic origin status, marital status, and number of children who are located in less expensive cities. In this case, our imputation would overestimate the income that household of a particular type can expect to obtain in expensive cities and consequently it would also overstate the expected household consumption in expensive cities. The ultimate effect would be that our estimates of the differences in standard of living between expensive and affordable cities would overstate the true differences.

5.1 Standard of Living in the Largest 50 Commuting Zones

The maps in Appendix Figure A6 show the geographical distribution of consumption by skill level. Since the maps are not easy to read, Tables 5, 6, and 7 present our main findings for the largest fifty commuting zones, ordered by pre-tax nominal income for each skill group. For each commuting zone, we report estimates of average household adjusted pre-tax income, post-tax income, and consumption. These estimates hold constant the combination of household characteristics that define a type (education, age, gender, race, Hispanic origin, education, marital status, and children). For each variable, we also report its corresponding percentile among all 443 commuting zones in our data (not just the 50 shown in the table). The last two columns report the difference between pre-tax income and consumption, as a percent difference and as a percentile difference.

Table 5 is for households where the head has a college degree or more. The first three rows show that San Jose, CA; San Francisco, CA; and Washington, DC are the three CZs where high-skill households have the highest expected adjusted pre-tax income: $144,255; $139,677; and $138,924, respectively (column 1). White Plains, NY—a suburb of New York—and New York, NY follow closely. Column 3 reports the corresponding after-tax income obtained by subtracting personal federal and state taxes from column 1. Column 5 shows our estimates of the levels of expected consumption. It quantifies the standard of living that a family with this level of schooling can expect in each commuting zone. For San Jose, San Francisco, Washington DC, and New York, the corresponding values are $70,692, $71,580, $73,634, and $70,416. The entries in column 5 are substantially lower than column 1 because high-skill residents face a particularly high local cost of living, and, to a lesser degree, because they face high state taxes. But in terms of consumption percentile, the decline for these four cities is modest. In terms of adjusted pre-tax income, these four cities are at the 99th or 100th percentile (column 2), while in terms of consumption San Jose, San Francisco, Washington, and New York drop to 78th, 84th, 91st, and 76th percentile respectively (column 6).

Thus, despite some of the highest costs of living in the US, Washington, DC, San Francisco, and New York remain in the top quartile of the distribution of all US commuting zones in terms of the standard of living distribution for college graduates. Given the general perception of the Bay Area and New York as regions that are unaffordable even for high-skill workers, this finding may come as a surprise. While these cities are indeed incredibly expensive, they offer a before taxes nominal income level so high that even after local prices and taxes are taken into account, standard of living of the highly educated remain higher than in most other US cities. (Combining this finding with the fact that local amenities in the Bay Area and New York tend to be considered desirable by many college graduates may explain why these regions have attracted so many college graduates over the past three decades.)

Boston consumption level is at the 77th percentile, while the one of Chicago is at the 69th percentile. Los Angeles and San Diego experience larger drops in relative standings. In terms of pre-tax income, these two cities are at the 97th percentile, while in terms of consumption percentiles they drop by 60 and 88 to the 37th and 9th percentiles, respectively. Other examples of cities with large negative percentile changes are Portland (-78), Seattle (-74), Minneapolis (-73), Sacramento (-66), Denver (-67), and Providence (-72). By contrast, Cincinnati (+8), Cleveland (+7), Pittsburgh (+12), and Buffalo (+15) improve their relative rankings as we go from pre-tax income to consumption.

Among the largest 50 commuting zones in the table, the one that offers the highest standard of living for college graduates is Houston because local income is relatively high, cost of living is moderate, and there are no state taxes. Cincinnati follows closely.[28]

A feature of the table worth noting is that spatial variation in consumption is much smaller

[28]The three commuting zones that offer the highest standard of living among all 443 commuting zones in the sample are Gallup, NM (82,598); Greenville, MS (79,559); and Huntington, WV (78,375).

than spatial variation in pre-tax income. For example, despite different levels of pre-tax income, consumption in San Jose is found to be similar to that in Nashville, TN and San Antonio, TX. Similarly, consumption in Washington, DC is found to be similar to Dallas, TX and Cleveland, OH. Overall, across all 443 commuting zones, the standard deviation in adjusted pre-tax income is 10,827, while the standard deviation in consumption is only 4,714—or less than half. This is to be expected if households are at least in part mobile and have a tendency to move toward areas that offer high standard of living.

Table 6 is for households where the head has a high school degree or some college. The picture that emerges is different, in the sense that the most expensive cities appear to offer significant lower consumption than affordable cities. For example, while in terms of adjusted pre-tax income, San Jose, CA and San Francisco, CA remain at the top, in terms of consumption they drop to the bottom quartile of the distribution. New York, Boston and Chicago are at the 12th, 29th, and 19th percentile, respectively. Los Angeles, San Diego, and Seattle are in the bottom 10 percent of the consumption distribution. Overall, it seems that while nominal pre-tax incomes are higher in expensive cities for this group, they are not high enough to offset the high costs of living.

Among the largest 50 commuting zones in the table, the one that offers the highest standard of living for high school graduates is Buffalo—due to its low cost of living—followed by Houston and Pittsburgh.[29] Across all 443 commuting zones, the standard deviation in adjusted pre-tax income is 6,218, while the standard deviation in consumption is only 3,182, confirming that spatial variation in consumption is smaller than spatial variation in pre-tax income.

Table 7 shows that local prices in high costs commuting zones take an even larger toll on the consumption of households where the head has less than high school. For example, the consumption of high school drop-outs in San Francisco, New York, Denver and Los Angeles is in the bottom 10 percent of the distribution. In these cities adjusted pre-tax nominal salaries are higher than in most other commuting zones, but cost of living is so high that low skill residents' standard of living is among the lowest in the nation. Boston and Philadelphia fare slightly better, although their consumption remains in the bottom quartile of the distribution.

Among the largest 50 commuting zones, the ones that offer the highest standard of living to high school drop outs are Buffalo—as it was the case for high school graduates—and Harrisburg, thanks to low prices.[30] Across all 443 commuting zones, the standard deviation in pre-tax income is 4,533, while the standard deviation in consumption is only 3,364.

5.2 Correlation of Standard of Living with Local Price Indexes

To understand more systematically how standard of living of each skill group varies as a function of local prices, Figure 11 plots household adjusted pre-tax income, post-tax income and consumption as a function of the local cost of living index. The figure includes all 443 commuting zones. Since

[29]The three commuting zones that offer the highest standard of living for this group among all commuting zones in the sample are Ketchikan, AK (61,665); Gallup, NM (60,418); and Summersville, WV (58,040).

[30]The three commuting zones that offer the highest standard of living for this group among all commuting zones in the sample are Summersville, WV (55,209); Galesburg, IL (53,712); and Paris, TX (52,042).

income, consumption, and local prices are all simultaneously determined, these relationships do not have a causal interpretation. Rather, they need to be interpreted as describing the cross-sectional equilibrium relationship between income, consumption, and local prices.

For all three skill groups, there is a positive correlation between pre-tax income and local cost of living. This is hardly surprising, as more expensive cities have long been known to offer higher equilibrium earnings. Crucially, the elasticity is above 1 for the high-skill group, and below 1 for the middle- and low-skill groups. In particular, the slope is 1.036 (0.053), 0.865 (0.035), and 0.704 (0.042) for high-, middle-, and low-skill households, respectively. This implies that a 10% increase in cost of living is associated with a more than proportional increase in expected pre-tax income for the high skill group —specifically: an increase of 10.4%—and a less than proportional increase in expected pre-tax income for the low skill group—specifically: an increase of only 7.0%.

We also observe a positive relationship for post-tax income. For all three groups, the slopes for post-tax incomes are smaller than the slopes for pre-tax incomes—since expensive cities tend to be located in states with higher income taxes—with the difference in slope larger for the high-skill group than for the low-skill group—due to tax progressivity. The intercept for post-tax income is lower than the one for pre-tax income—reflecting mean tax burden in the least expensive commuting zones—and the drop in intercept is the largest for high-skill households and minimal for low-skill households—again reflecting tax progressivity.

The findings for consumption are remarkable. For high-skill households, there is essentially no relationship between consumption and cost of living. The coefficient is -0.032 (0.047)—close to zero and not statistically significant at conventional levels. This suggests that college graduates living in cities with high costs of living enjoy a standard of living that is similar to that enjoyed by college graduates with the same observable characteristics living in cities with low cost of living. This appears to be true for the entire range of values observed for the cost of living index, including at the very top of the cost of living distribution. Compared with affordable cities, expensive cities appear to offer incomes high enough to exactly offset the difference in cost of living and personal taxes.

For less skilled households, the picture that emerges is markedly different. For high school graduates, we find a negative relationship between household consumption and cost of living, indicating that expensive cities offer standard of living that are not as good as more affordable cities. The negative slope reflects the fact that pre-tax income for this group is higher in expensive cities than more affordable cities, but not high enough to offset cost of living and taxes.

The elasticity of consumption with respect to cost of living is -0.237 (0.026), implying that a middle-skill household moving from the median commuting zone (Cleveland) to the most expensive commuting zone (San Jose) would experience a decline in standard of living by 7.7%. Moving from the commuting zone with the lowest cost of living index (Natchez) to the commuting zone with the highest index (San Jose) would imply a decline in the standard of living by 12.7%.

The negative relationship between consumption and cost of living is significantly steeper for high school dropouts, suggesting that for this group the standard of living in expensive commuting

zones is much lower than in cheaper commuting zones. The slope is -0.391 (0.032), implying vast geographical differences in consumption. Moving from Cleveland to San Jose implies a 15.6% decline in the standard of living. Moving from Natchez to San Jose implies a 26.9% decline in the standard of living. The finding that the elasticity of consumption with respect to cost of living for this group is the most negative of the three groups reflects the fact that the correlation between pre-tax income and cost of living is the lowest.

Overall, our findings are consistent with the growing concern that high cost cities are becoming unaffordable to the middle class and low-income households. The concern appears particularly serious for the low-skill households, who are increasingly exposed high costs of living and are found to be significantly worse off in terms of market consumption compared to similar households in more affordable areas.

Our findings also have important implications for within-commuting zone inequality. Since consumption of college graduates was found to be uncorrelated with local prices while consumption of less skilled groups was found to decline with local prices, consumption inequality within a commuting zone should increase with local prices. In our data, consumption of college graduates in San Francisco, San Jose, and New York is 1.94, 1.87, and 1.91 times higher than consumption of high school drop-outs. The corresponding ratios in the three cheapest commuting zones, London, KY; Gallup, NM; and Natchez, MS are 1.60, 1.61, and 1.56. The top panel of Figure 12 shows more systematically how the difference in mean consumption between high- and middle-skill households who live in the same commuting zones varies as a function of local cost of living across all commuting zones in the sample. The bottom panel shows the difference in mean consumption between high- and low-skill households. The slopes are 0.205 (0.028) and 0.359 (0.032), respectively, confirming that within-commuting-zone consumption inequality increases significantly with cost of living. This is particularly true for the difference in mean consumption between high- and low-skill households.[31]

This finding further underscores the growing concern raised by many residents of expensive cities about declining standard of living of less skilled residents, who in recent decades have been exposed to increasingly affluent co-residents and skyrocketing local prices, raising questions about affordability and gentrification.

In this respect, one possible concern is that our income data misses Food Stamps, TANF, and housing assistance. If low income residents in expensive commuting zones tend to receive more generous transfers than low income residents in affordable commuting zones, this could induce bias in the elasticity of consumption with respect to local prices estimated for less skill households. To assess the magnitude of the problem, we analyze the sensitivity of our estimates of the elasticity of consumption with respect to local prices to including the imputed value of Food Stamps, TANF and housing assistance. For the imputation, we use data from the CPS on the average receipt of Food Stamps, TANF and housing assistance by income level, martial status, number of children

[31] Bertrand and Morse (2016) report that poor households consume a larger share of their income when exposed to a higher number of rich residents in a state. See also Charles et al. (2009)

and state. Since the CPS has been shown to under-report government transfers, we inflate the reported amounts based on Table 2 in Meyer and Mittag (2019). See Appendix F for more details. Panel B in Appendix Table A9 shows that our main empirical results are not particularly sensitive. The regression coefficients of log consumption (inclusive of government transfers) and log price index change from -0.032 to -0.040 for high skill, from -0.237 to -0.242 for middle skill, and from -0.391 to -0.389 for low skill. The main reason for the robustness of our estimates is that Food Stamps, TANF and housing assistance are federal transfers with limited geographical variation and therefore limited effect on our estimated coefficients, which are identified by geographical variation.

Panel C in the same table probes the robustness of our estimates to the adjustments that we made to consumption expenditures. Recall that our expenditure measure includes an adjustment for health expenditures that are not out of pocket; and one for housing costs paid by homeowners. Entries in Panel C show that these two adjustments make the three elasticities more negative. The elasticity for the high skill group is now significantly different from zero, although it remains much smaller (in absolute value) than the elasticity for the other groups.

5.3 Correlation with City Size and College Share and Robustness to Alternative Indexes

We conclude this section by investigating how standard of living varies across cities as a function of two other city characteristics that have been prominent in the literature on spatial wage differences: city size and share of residents with a college degree. We also investigate how robust our findings are to alternative cost of living indexes.

The top panel in Figure 13 presents the results for size, measured by commuting zone population. For all three groups, there is a positive correlation between pre-tax income and population. This is unsurprising, and has been documented by a large literature on the wage premium offered by large cities over small cities and rural areas. What is more interesting is the relationship between consumption and city size. For high-skill households, there is no significant relationship between consumption and city size, indicating that large and small cities offer similar consumption. For middle- and low-skill households, however, there is a significant negative relationship between consumption and city size. Therefore, the standard of living of lower skill households is on average lower in large cities compared to small cities.

The bottom panel focuses on the share of residents with a college degree or more. The positive correlation between pre-tax income and college share is consistent with previous work (Moretti, 2004, 2013). More novel is the relationship between consumption and college share. For high-skill households, there is no significant relationship between consumption and college share. For middle- and low-skill households, there is a significant negative significant relationship between consumption and college share indicating that residents in cities with many college graduates enjoy standard of living that is lower than residents in cities with fewer college graduates.

Since cities' prices, population, and college shares are all positively correlated, in Table 8 we investigate a multivariate regression to see which of these characteristics consumption is most

related to including the same set of controls used above. We allow the coefficients on each of these three regressors to vary by skill group. In interpreting this table, it is important to keep in mind that local prices, city size, and college share are all simultaneously determined and the Table's entries do not reflect causal estimates, but rather equilibrium relationships. Entries indicate that conditional on city size and college share, the correlation of consumption and the price index is negative for all three groups, with similar elasticities. Interestingly, conditional on prices and college share, the correlation of consumption and city size is positive for the high-skill group, close to zero for the middle-skill group, and negative for the low-skill group. The correlation of consumption and college share appears close to zero for all three groups, but the effects are a bit noisy.

This indicates that large, lower price commuting zones offer best consumption to college educated households. By contrast, low and middle skill households maximize consumption in small, lower price commuting zones. It also indicates that part of the overall elasticity of consumption with the respect to the price index uncovered in the previous sub-section reflects the correlation between consumption and city size combined with the fact that larger cities tend to be more expensive.

Finally, we turn to robustness. In Section 3.3 we found that the alternative price indexes are highly correlated with the baseline index. We now investigate how sensitive are our estimates to the use of the alternative price indexes. In Appendix Table **??**, we report the estimates corresponding to the Table 8 based on all the alternative price indexes. Quantitatively, the coefficients vary somewhat in magnitude, as one might expect. But in most cases, the qualitative picture that emerges is similar to the one in Table 8. In general, we find a negative correlation with the price indexes that is similar across skill groups and a correlation with city size that is positive for the high skill and negative for the less skilled. The coefficients on the price index for the high-skill group range from -0.341 to 0.032, the differential effects for the middle-skill group from 0.039 to 0.158 and differential effects for the low-skill group from -0.044 to 0.118. The coefficients on city size for the high-skill group range from 0.023 to 0.032, differential effects for the middle-skill group from -0.026 to -0.024 and differential effects for the low-skill group from -0.040 to -0.036.

6 Implications for Aggregate Consumption Inequality

Over the last four decades, high income and high-skill individuals have increasingly sorted into in expensive communities. Evidence of sorting is shown in Figure 14, which plots mean of log price index by nationwide household income percentile. The confidence band indicates that 95% of households whose income percentile is between 0 and 70 experience a roughly similar cost of living—namely an index between 0.98 and 1. By contrast, richer households–those with income percentile above 70—are exposed to an index that is exponentially increasing in income.

In this section, we study the implications of geographic sorting into high- and low-cost commuting zones for nationwide inequality (Moretti, 2013). Figure 15 shows the income and consumption differences between high- and middle-skill households (top panel), high- and low-skill households (middle panel), and middle- and low-skill households (bottom panel). The first set of bars—labeled "Baseline"—indicate that the difference in consumption is smaller than the corresponding differ-

ence in pre-tax income.[32] In particular, the pre-tax income difference between college graduates and high school graduates is 0.440, while the corresponding difference in consumption is only 0.303, or 31 percent smaller. This reflects both the progressivity of taxation, savings behavior, and the fact that college graduates are more likely to live in expensive cities than high school graduates. The corresponding differences between high and low skill households are 0.667 and 0.459. Namely, the consumption gap is 31 percent smaller than the income gap.

To isolate the role of sorting and local prices, the next set of bars shows estimates where we re-weight households so that the distribution of observable household types within each commuting zone equals the nationwide distribution. The graph indicates that if all commuting zones were to have the same composition of household types, the pre-tax income gap between high and middle skill and high and low skill would decline to 0.40 and 0.62, respectively. However, the consumption gaps would remain basically the same at 0.31 and 0.46, respectively. Interestingly, geographic sorting of households across CZs seems to have no impact on nationwide consumption inequality between skill groups.[33]

To investigate the role of city size, we further re-weight the data to equate population across CZs, in addition to equalizing household types across CZs. The third set of bars show that this narrows the pre-tax income gaps to 0.36 and 0.57 and the consumption gaps to 0.26 and 0.40.[34] Equalizing population lowers high-middle skill and high-low skill consumption in equality by 14% and 12%, respectively. This is driven by large cities offering higher consumption to the high skill, but lower consumption to middle and low skill households.

Overall, population and demographic sorting across CZs can explain 19% of high-middle skill income gap, 15% of the high-low skill income gap, and 7% of the middle-low skill income gap. In terms of consumption inequality, these forces explain 13% of the high-middle skill consumption gap, 12% of the high-low skill consumption gap and 10% of the middle-low skill consumption gap.

7 Conclusion

We draw two main conclusions. First, we uncover vast geographical differences in material standard of living for a given level of income. Low income residents in the most affordable commuting zone enjoy a level of consumption that is 74% higher that of low income residents in the most expensive commuting zone. When we replicate the analysis focusing on the consumption of specific goods measured in physical units we also find significantly lower consumption in expensive areas.

Second, when we estimate the standard of living that low- and high-skill households can expect

[32] The height of each bar is the conditional differences observed in the data from a household-level regression of log income or log consumption on skill-group indicator, controlling for household type indicators.

[33] This may appear contradictory with our previous evidence high college share CZs disproportionately benefit high skill households. However, this exercise not only equation college shares across CZs, but all other demographic characteristics too, including household size, age, marital status, and presence of kids.

[34] Specifically, the adjustment weight for household h of type t living in commuting zone j is defined as $\widetilde{\text{hhwt}}_{h,j(h),t(h)} = \text{hhwt}_{h,j(h),t(h)} \times \frac{\text{share}_t}{\text{share}_{j,t}} \times \frac{\sum_{j \in J} \text{population}_j}{|J| \times \text{population}_j}$, where $\text{hhwt}_{h,j(h),t(h)}$ denotes the original household weight in ACS; share_t is the nationwide share of type t; $\text{share}_{j,t}$ is the commuting-zone-j share of type t; population_j is the number of households in commuting zone j.

in each US commuting zone once we account both for geographical variation in cost of living and also in expected income, we find marked differences between low- and high-skill households. For high-skill households, we find no relationship between expected consumption and cost of living, suggesting that college graduates living in cities with high costs of living enjoy a standard of living generally similar to college graduates living in cities with low cost of living. For high school graduates and high school drop-outs, we find a significant negative relationship between consumption and cost of living, indicating that expensive cities offer lower standard of living than more affordable cities. The differences are quantitatively large. A high school drop-out household moving from the most affordable commuting zone to the most expensive one would experience a 26.9% decline in market consumption.

Establishing the relationship between internal migratory flows and the geography of standard of living in the US, and the precise reasons for the persistence of large difference in standard of living for less educated households should be two primary objectives of future research in this area. Future work should also explore the appropriate model of spatial equilibrium that is consistent with our findings. A simplistic version of the Rosen-Roback framework where amenities perfectly offset differences in market consumption across space poorly fits the consumption differences across space that we have uncovered. Through the lens of Rosen-Roback, our finding that the lowest skill households have the largest consumption differences between expensive and cheap cities would indicate that the lowest skilled households have the highest willingness to pay for the amenities available in the most expensive cities. While we have not included amenities in any of our calculations, this possibility appears to be inconsistent with prior work (Diamond, 2016). A richer model with preference heterogeneity within and across these skill groups is likely needed to understand these equilibrium relationships.

References

Aguiar, M. and M. Bils (2015): "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review*, 105, 2725–2756.

Allcott, H., R. Diamond, J.-P. Dubé, J. Handbury, I. Rahkovsky, and M. Schnell (2019): "Food deserts and the causes of nutritional inequality," *The Quarterly Journal of Economics*, 134, 1793–1844.

Argente, D. O., C.-T. Hsieh, and M. Lee (2020): "Measuring the Cost of Living in Mexico and the US," *NBER Working Paper No. 27806*.

Attanasio, O. P., E. Hurst, and L. Pistaferri (2015): "The Evolution of Income, Consumption, and Leisure Inequality in the US, 1980–2010," in *Improving the Measurement of Consumer Expenditures*, ed. by C. D. Carroll, T. F. Crossley, and J. Sabelhaus, University of Chicago Press, chap. 4.

Attanasio, O. P. and L. Pistaferri (2016): "Consumption Inequality," *Journal of Economic Perspectives*, 30, 3–28.

Bee, A., B. Meyer, and J. Sullivan (2012): "The Validity of Consumption Data: Are the Consumer Expenditure Interview and Diary Surveys Informative?" *NBER Working Paper No. 18308*.

Bernard, D., C. Cowan, T. Selden, L. Cai, A. Catlin, and S. Heffler (2012): "Reconciling medical expenditure estimates from the MEPS and NHEA, 2007," *Medicare & medicaid research review*, 2.

Bertrand, M. and A. Morse (2016): "Trickle-Down Consumption," *Review of Economics and Statistics*, 98, 863–879.

Broda, C., E. Leibtag, and D. E. Weinstein (2009): "The Role of Prices in Measuring the Poor's Living Standards," *Journal of Economic Perspectives*, 23, 77–97.

Broda, C. and D. E. Weinstein (2010): "Product Creation and Destruction: Evidence and Price Implications," *American Economic Review*, 100, 691–723.

Bureau of Labor Statistics (2007): "Handbook of methods," Tech. rep., US Department of Labor, Washington, DC.

Card, D., , J. Rothstein, and M. Yi (2021): "Location, location, location," *UC Berkeley*.

Charles, K. K., E. Hurst, and N. Roussanov (2009): "Conspicuous Consumption and Race," *The Quarterly Journal of Economics*, 124, 425–467.

Choi, C. Y. and S. Jo (2020): "How Do Housing Markets Affect Local Consumer Prices? – Evidence from U.S. Cities," *Globalization and Monetary Policy Institute Working Paper No. 398*.

CMS (2021): "Centers for Medicare Medicaid Services– National Health Expenditure Accounts (NHEA)," *https://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/NationalHealthExpendData/NationalHealthAccountsHistorical*.

Couture, V., C. Gaubert, J. Handbury, and E. Hurst (2019): "Income Growth and the Distributional Effects of Urban Spatial Sorting," *NBER Working Paper No. 26142*.

Crossley, T. F. and J. K. Winter (2015): "Asking Households about Expenditures: What Have We Learned?" in *Improving the Measurement of Consumer Expenditures*, ed. by C. D. Carroll, T. F. Crossley, and J. Sabelhaus, University of Chicago Press, chap. 4.

Cutler, D. and L. Katz (1992): "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980s," *NBER Working Paper No. 3964*.

Deaton, A. and A. Heston (2010): "Understanding PPPs and PPP-based national accounts," *American Economic Journal: Macroeconomics*, 2, 1–35.

Deaton, A. and J. Muellbauer (1980): "An Almost Ideal Demand System," *American Economic Review*, 70, 312–326.

DellaVigna, S. and M. Gentzkow (2019): "Uniform Pricing in U.S. Retail Chains," *The Quarterly Journal of Economics*, 134, 2011–2084.

Diamond, R. (2016): "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000," *American Economic Review*, 106, 479–524.

Diamond, R., M. J. Dickstein, T. McQuade, and P. Persson (2018): "Insurance without Commitment: Evidence from the ACA Marketplaces," Tech. rep., National Bureau of Economic Research.

Dynan, K. E., J. Skinner, and S. P. Zeldes (2004): "Do the Rich Save More?" *Journal of Political Economy*, 112, 397–444.

Federal Deposit Insurance Corporation (2015): "FDIC National Survey of Unbanked and Underbanked Households, Executive Summary," .

Ganong, P. and P. Noel (2019): "Consumer spending during unemployment: Positive and normative implications," *American economic review*, 109, 2383–2424.

Ganong, P. and D. Shoag (2017): "Why has regional income convergence in the US declined?" *Journal of Urban Economics*, 102, 76–90.

Greene, C. and S. Schuh (2016): "The Diary of Consumer Payment Choice," *Federal Reserve Bank of Boston*.

Handbury, J. (2019): "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities," *NBER Working Paper No. 26574*.

Handbury, J. and D. E. Weinstein (2015): "Goods Prices and Availability in Cities," *Review of Economic Studies*, 82, 258–296.

Jaravel, X. (2019): "The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector," *Quarterly Journal of Economics*, 134, 715–783.

Keys, B. J., N. Mahoney, and H. Yang (2020): "What Determines Consumer Financial Distress? Place- and Person-Based Factors," *NBER Working Paper No. 26808*.

Lewbel, A. and K. Pendakur (2009): "Tricks with Hicks: The EASI demand system," *American Economic Review*, 99, 827–63.

Meyer, B. and N. Mittag (2019): "Using Linked Survey and Administrative Data to Better Measure Income: Implications for Poverty, Program Effectiveness, and Holes in the Safety Net," *American Economic Journal: Applied Economics*, 11, 176–204.

Meyer, B. and J. Sullivan (2017): "Consumption and Income Inequality in the U.S. since the 1960s," *NBER Working Paper No. 23655*.

Montgomery, A. L. and P. E. Rossi (1999): "Estimating Price Elasticities with Theory-Based Priors," *Journal of Marketing Research*, 36, 413–423.

Moretti, E. (2004): "Estimating the Social Return to Higher Education: Evidence From Longitudinal and Repeated Cross-Sectional Data," *Journal of Econometrics*, 121.

——— (2013): "Real Wage Inequality," *American Economic Journal: Applied Economics*, 5, 65–103.

Poole, R., F. Ptacek, and R. Verbrugge (2005): "Treatment of Owner-Occupied Housing in the CPI," *Paper presented to the Federal Economic Statistics Advisory Committee (FESAC)*.

Redding, S. J. and D. E. Weinstein (2020): "Measuring Aggregate Price Indexes with Taste Shocks: Theory and Evidence for CES Preferences," *The Quarterly Journal of Economics*, 135, 503–560.

Rothbaum, J. L. (2015): "Comparing Income Aggregates: How do the CPS and ACS Match the National Income and Product Accounts, 2007-2012," *US Census Bureau, SEHSD Working Paper*, 1.

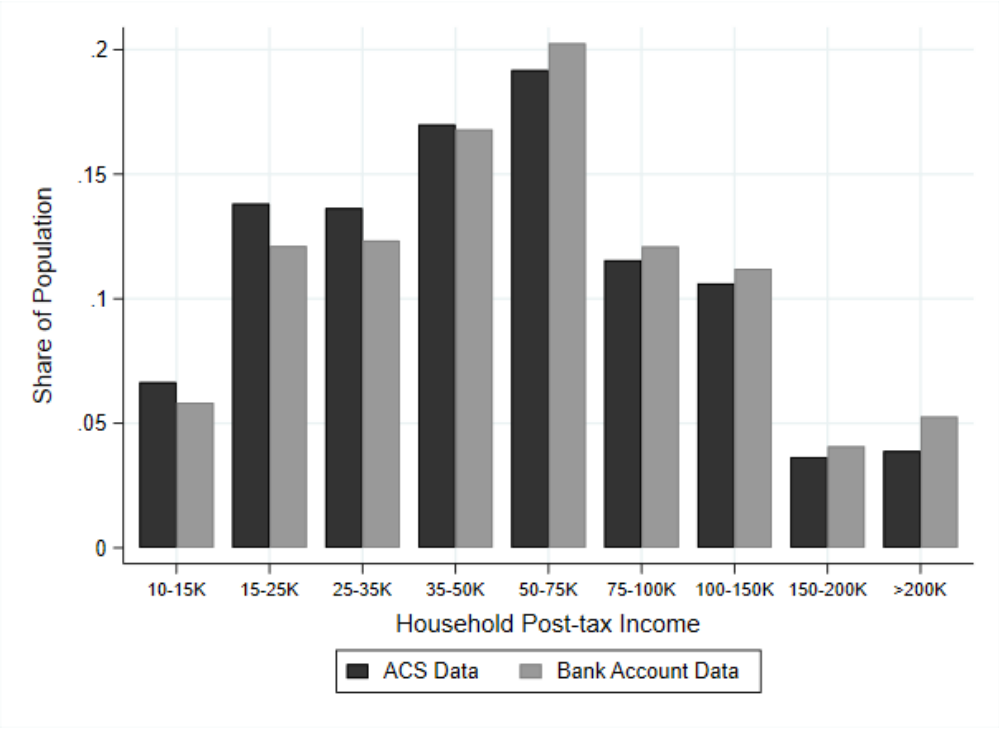
Sabelhaus, J. and J. A. Groen (2000): "Can Permanent-Income Theory Explain Cross-Sectional Consumption Patterns?" *Review of Economics and Statistics*, 82, 431–438.

Sabelhaus, J., D. Johnson, S. Ash, D. Swanson, T. I. Garner, J. Greenlees, and S. Henderson (2015): "Is the Consumer Expenditure Survey Representative by Income?" in *Improving the Measurement of Consumer Expenditures*, ed. by C. D. Carroll, T. F. Crossley, and J. Sabelhaus, University of Chicago Press, chap. 8.

Slesnick, D. T. (1991): "The Standard of Living in the United States," *Review of Income and Wealth*, 37, 363–386.

Walczak, J. and J. Cammenga (2021): "State Business Tax Climate Index," Tax Foundation.

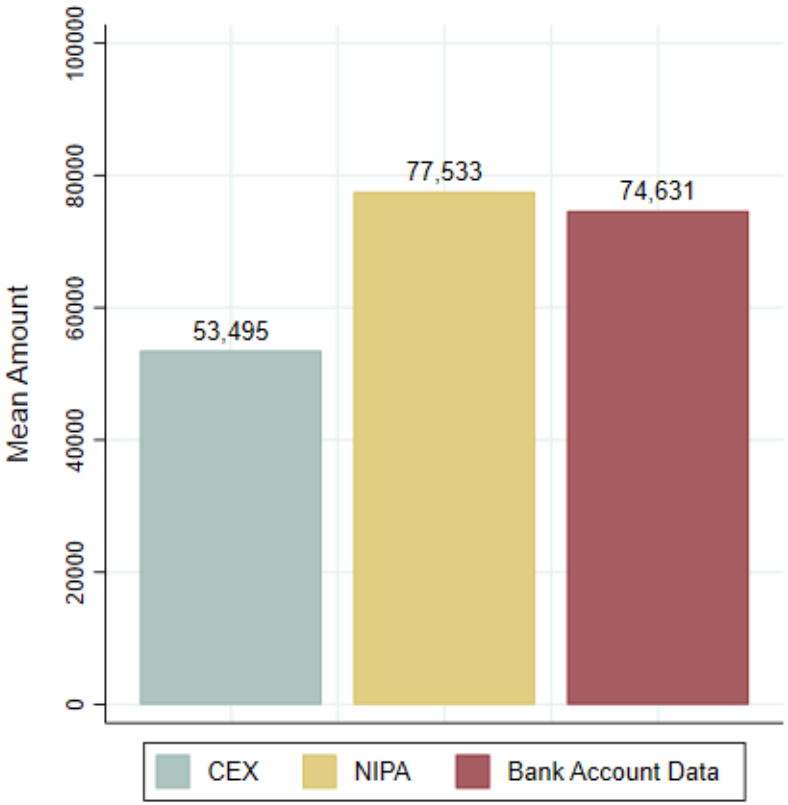
Figure 1: **Income Distribution: Bank Account Data and American Community Survey**



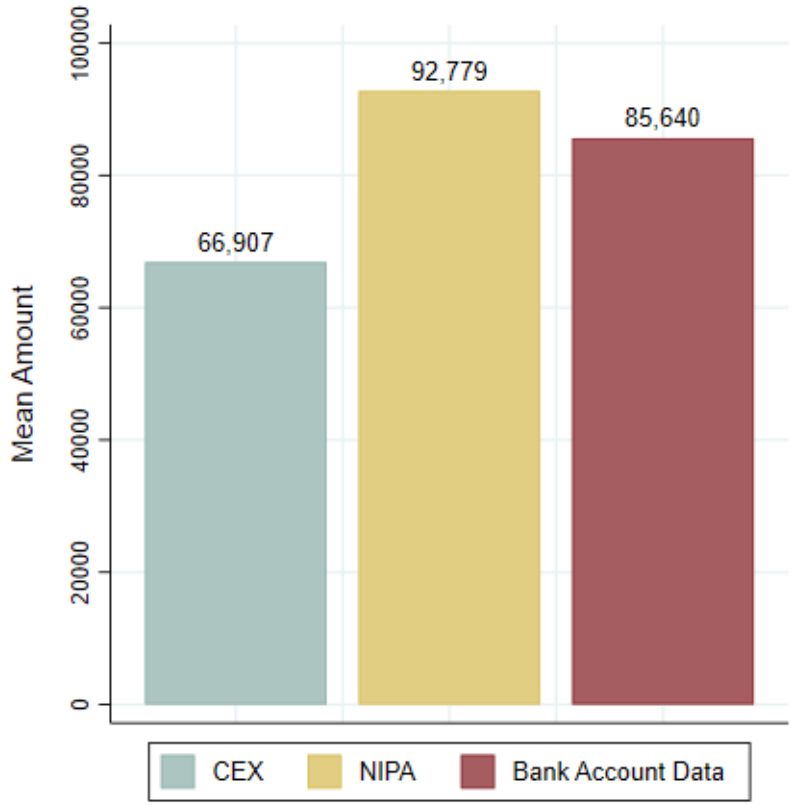
Notes: This figure compares the distribution of households post-tax income in our data and the 2012-2016 ACS. In the ACS, we use NBER TAXSIM to calculate income taxes and then subtract it from household pre-tax income, yielding post-tax income. The median (mean) in our data and in the ACS data are $52,956 ($81,011) and $48,835 ($67,628), respectively.

Figure 2: **Consumption Expenditure: Bank Data, NIPA and CEX**

(a) Raw Average Consumption Expenditure



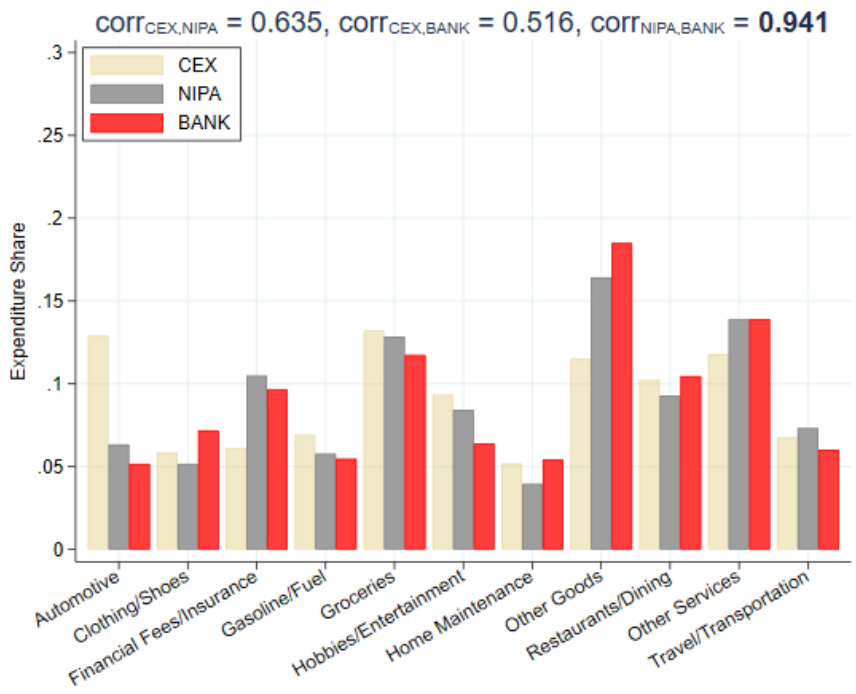
(b) Adjusted Consumption Expenditure



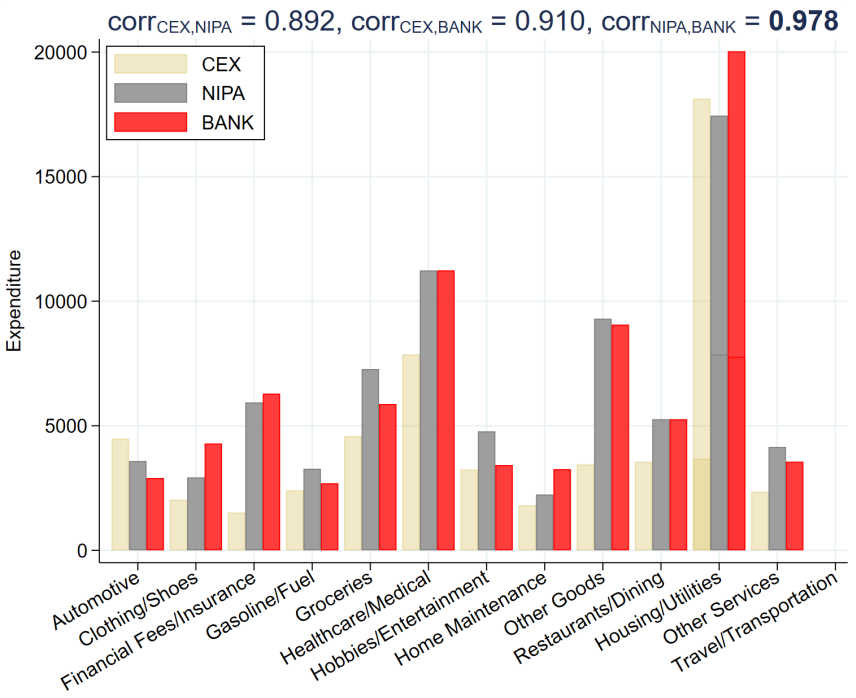
Notes: This figure compares average household annual expenditure in 2014 in the CEX, NIPA, and our data. Panel (a) shows average household expenditure from the raw data. The only adjustment done in panel (a) is for NIPA: we subtract healthcare spending paid by insurance companies and the government from the NIPA average total household expenditure. To do this, we take the NIPA total healthcare spending and spending on net health insurance premiums and multiply by it by 0.87, the 2014 share of health care costs that are not out-of-pocket (CMS, 2021). In panel (b), we adjust the data to make health expenditures and housing expenditures more directly comparable. For our bank data, we add healthcare costs paid by the government and insurers; and we adjust housing costs for homeowners. For NIPA, we add health insurance premiums paid by employers towards employee health insurance. For the CEX: we adjust housing costs for homeowners. See text for details.

Figure 3: **Categories of Expenditure: Bank Data, NIPA and CEX**

(a) Expenditure Shares (raw)

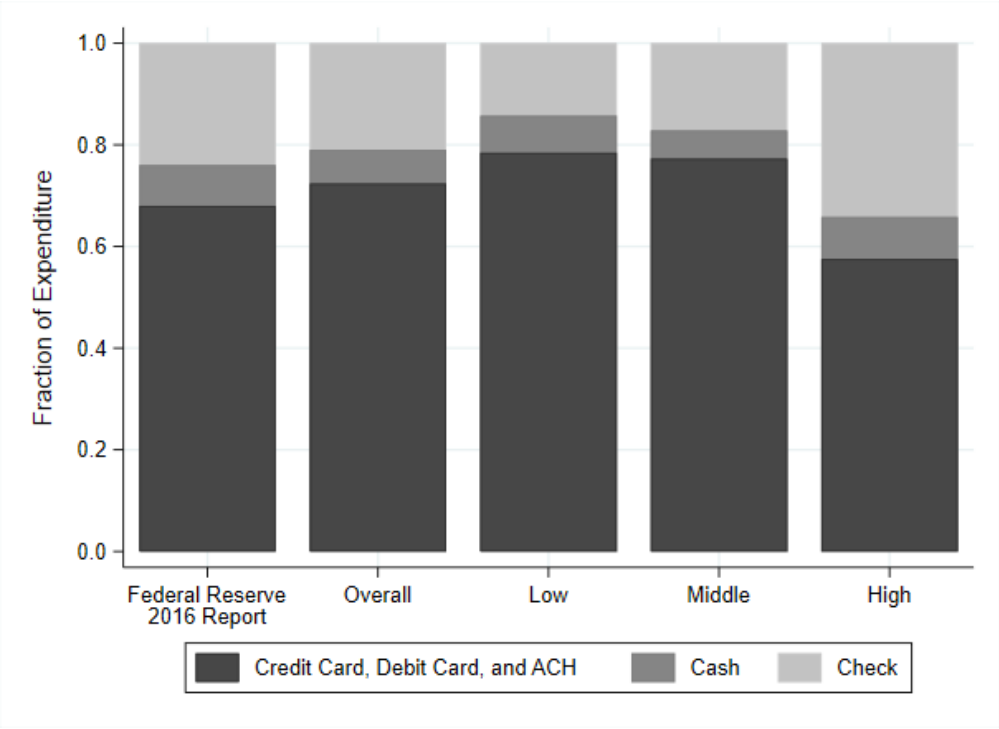


(b) Expenditure Levels (adjusted)



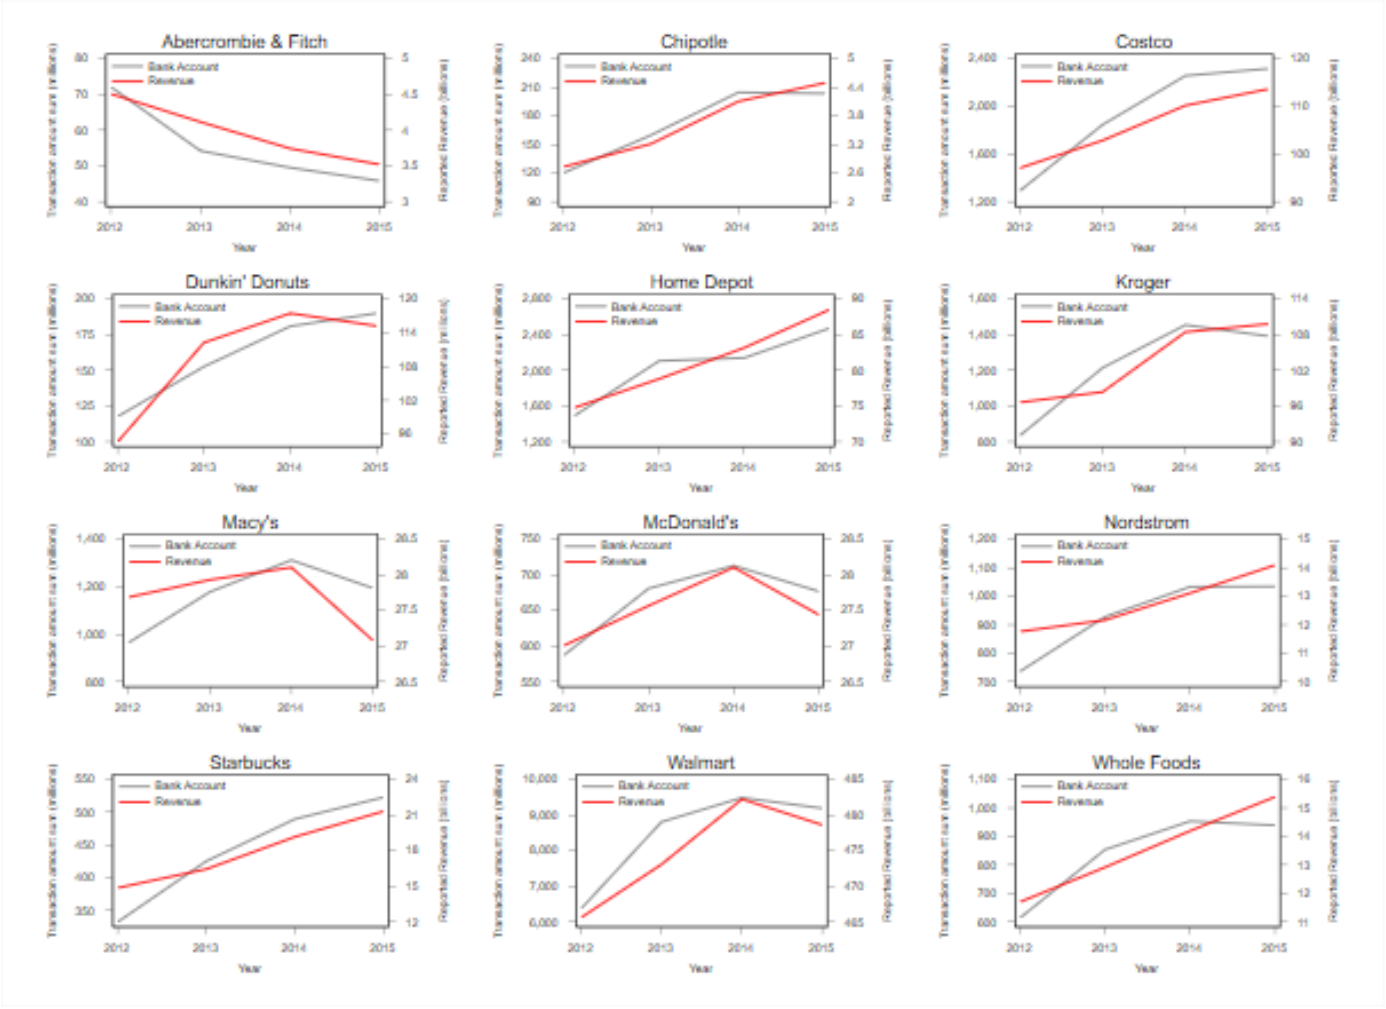
Notes: For CEX, we pool 2012-2016 Interview Survey data to measure annual spending. For NIPA, we use aggregate nationwide personal consumption expenditure in 2014. Panel (a) compares expenditure shares among comparable spending categories across the three datasets. Panel (b) compares average household expenditure levels across spending categories, where we adjust health and housing expenditures. As spending categories do not align perfectly across the three datasets, we restrict to types of expenditure that are defined consistently in each and we aggregate some categories with definitions that don't quite line up across datasets into "other goods" and "other services". The "Other Goods" category includes (i) communication equipment, household supplies, personal/personal care items, reading materials, and tobacco in NIPA; (ii) laundry and cleaning supplies, other household products, stationery, tobacco, and miscellaneous items in CEX; and (iii) electronics, general merchandise, and office and school supplies in our data. The "Other Services" category includes (i) communication, education, and personal/social/religious services in NIPA; (ii) child-related, education, personal care, postage, and telephone services in CEX; and (iii) charitable giving, child-related, education, personal care, printing and postage, and telecommunication services in our data. See text for details.

Figure 4: **Consumption Expenditure by Mean of Payment: Bank Data and Fed Data**



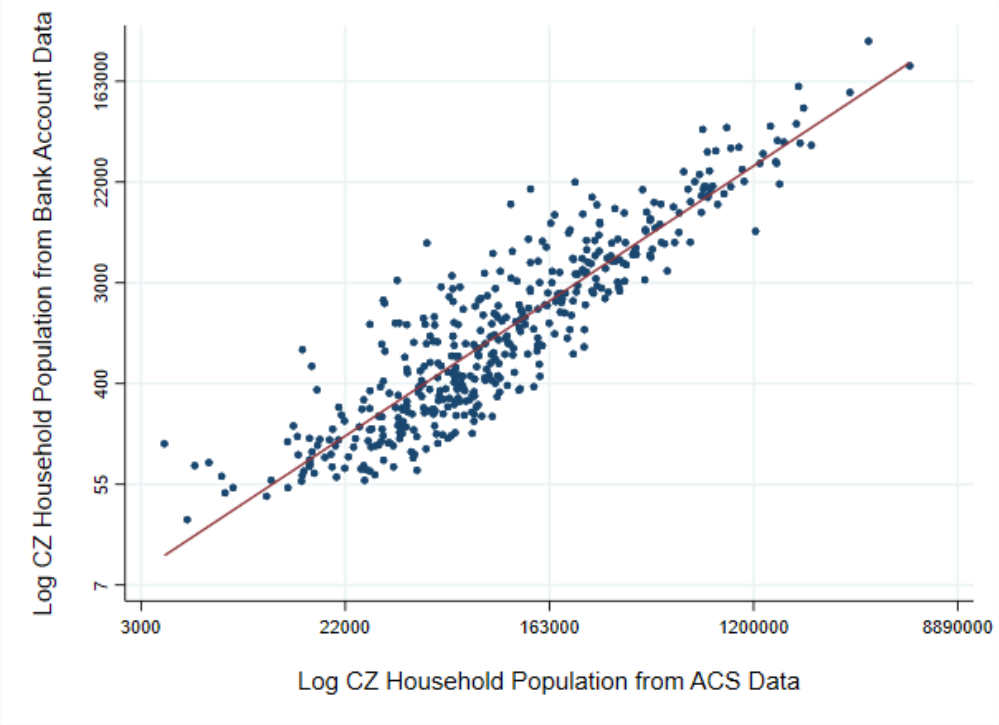
Notes: The bar on the left shows shares of consumption expenditure by mean of payment from the Federal Reserve Report by Greene and Schuh (2016). The Fed shares are based on consumers of all income levels. The Fed does not report shares by income. The four bars on the right are from our data.

Figure 5: **Transaction Dollar Amount from Our Data vs. Reported Revenue from SEC Filings**



Notes: The gray line is the total value of transactions at each merchant in our bank data, from 2012 to 2015. The red line is the reported revenues from SEC 10K filings from 2012 to 2015.

Figure 6: **Number of Households by Commuting Zone: Bank Data vs. American Community Survey**



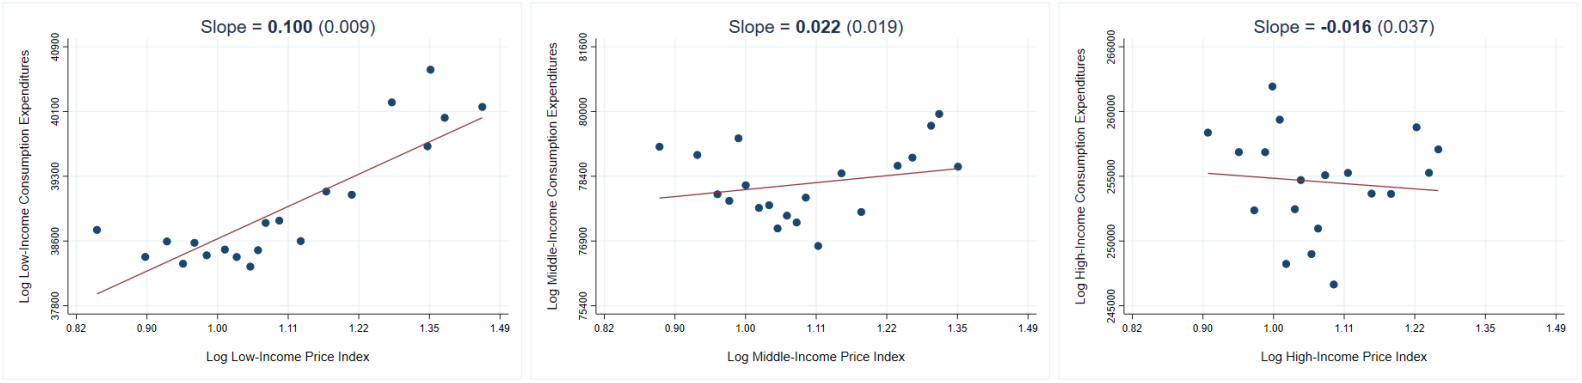
Notes: This figure plots log number of households from our data against log number of households from 2012-2016 ACS data. Each dot is a commuting zone. To make ACS data consistent with our data, we drop households in the ACS whose income is less than $10,000. Values on both x-axis and y-axis are measured in log scale but we label actual values for easier interpretation. The estimated slope is 1.340 (0.028). R^2=0.8147

Figure 7: **Spatial Distribution of Price Indexes**



Notes: The level of observation is a commuting zone. $N = 443$.

Figure 8: **Expenditure or Consumption vs. Price Index**

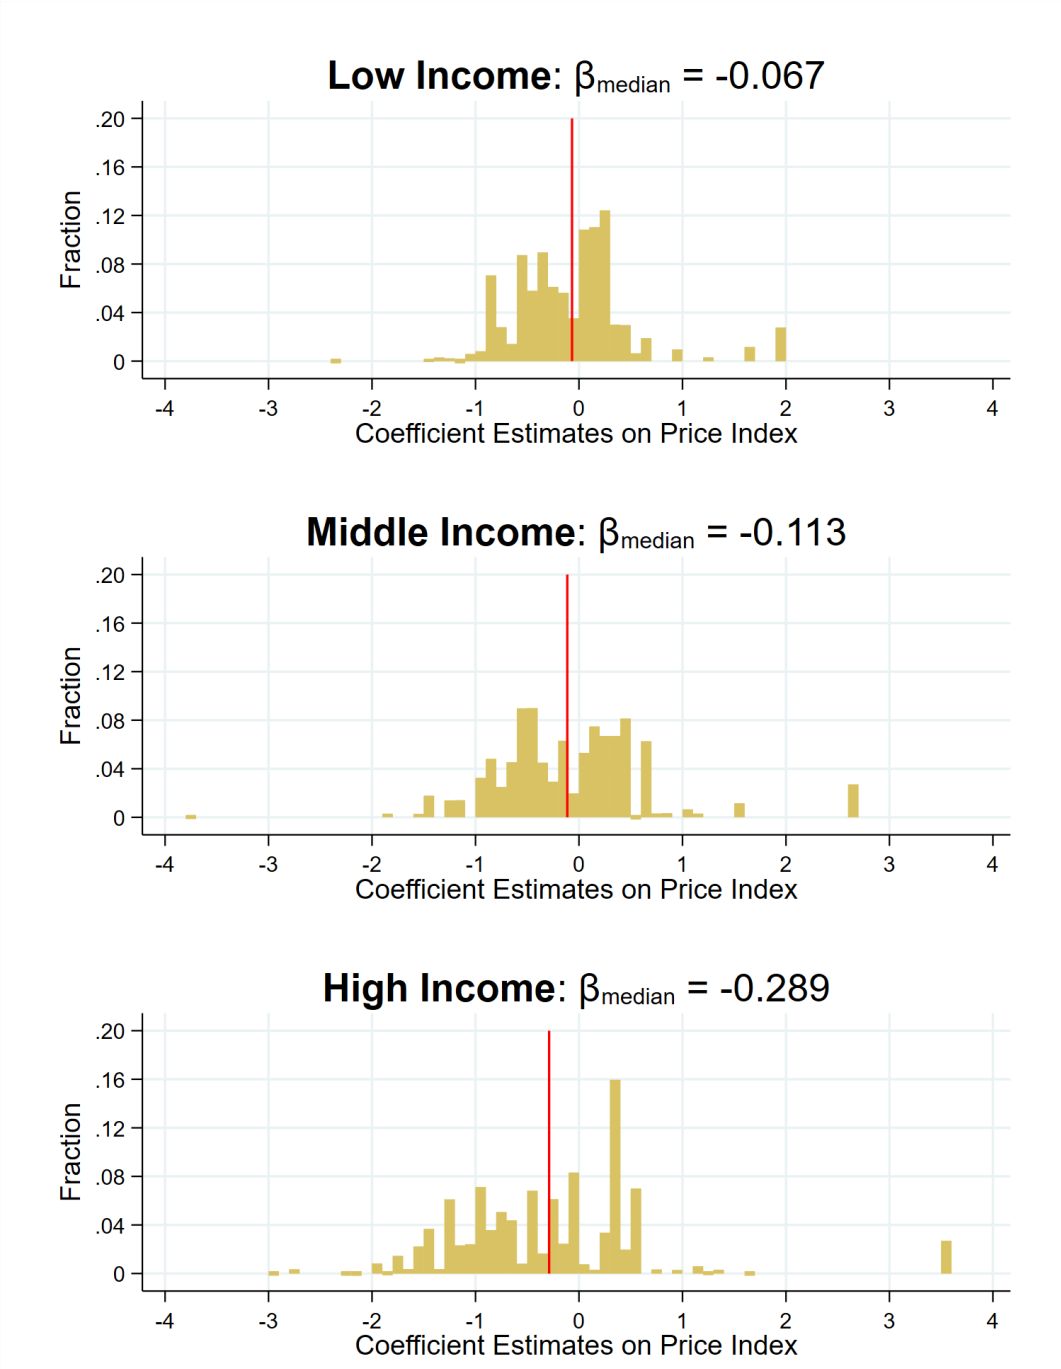


(a) Expenditure, Low Income (b) Expenditure, Middle Income (c) Expenditure, High Income

(d) Consumption, Low Income (e) Consumption, Middle Income (f) Consumption, High Income

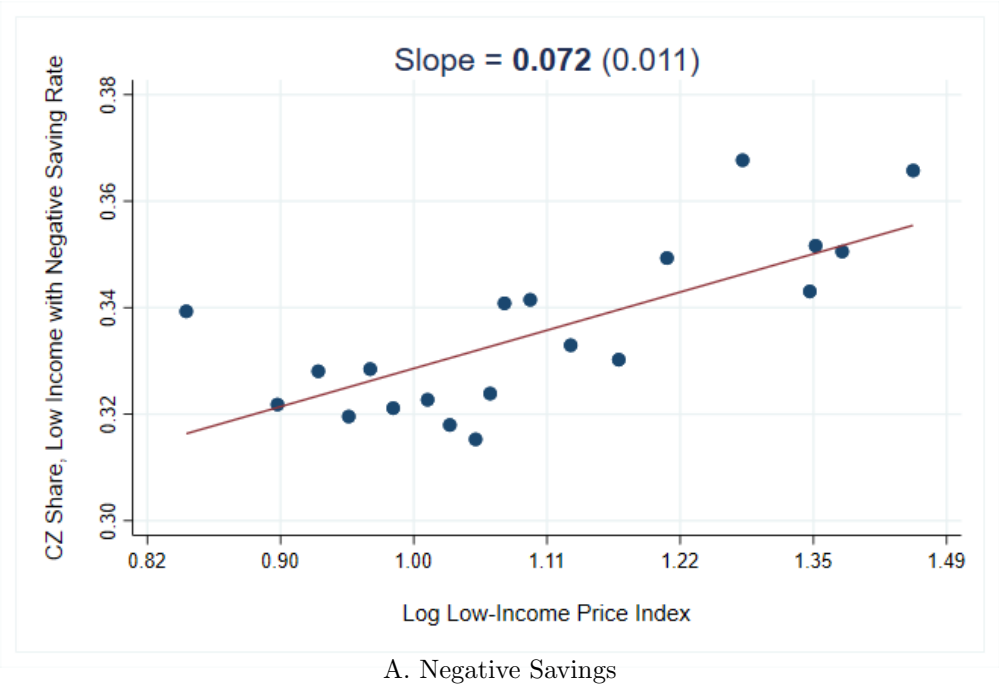
Notes: Values on both x-axis and y-axis are measured in log scale, but we label actual values for easier interpretation. $N = 443$.

Figure 9: **Distribution of Estimated Elasticities — NielsenIQ Data**



Notes: Each panel plots the distribution of estimated elasticities by product group. We weight by average household expenditure on each product group. Vertical lines denote the median. Elasticities are from regressions of mean-adjusted quantity of consumption on the local price index controlling for household characteristics: household income; household size; age and presence of children; type of residence; household composition; household head's characteristics including age, gender, race, marital status, education, employment status, and education. We average elasticities by product group. There are 116 product groups.

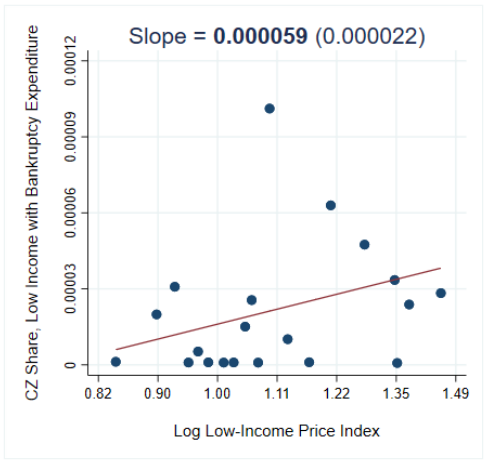
Figure 10: **Negative Savings, Overdraft, and Bankruptcy**



A. Negative Savings



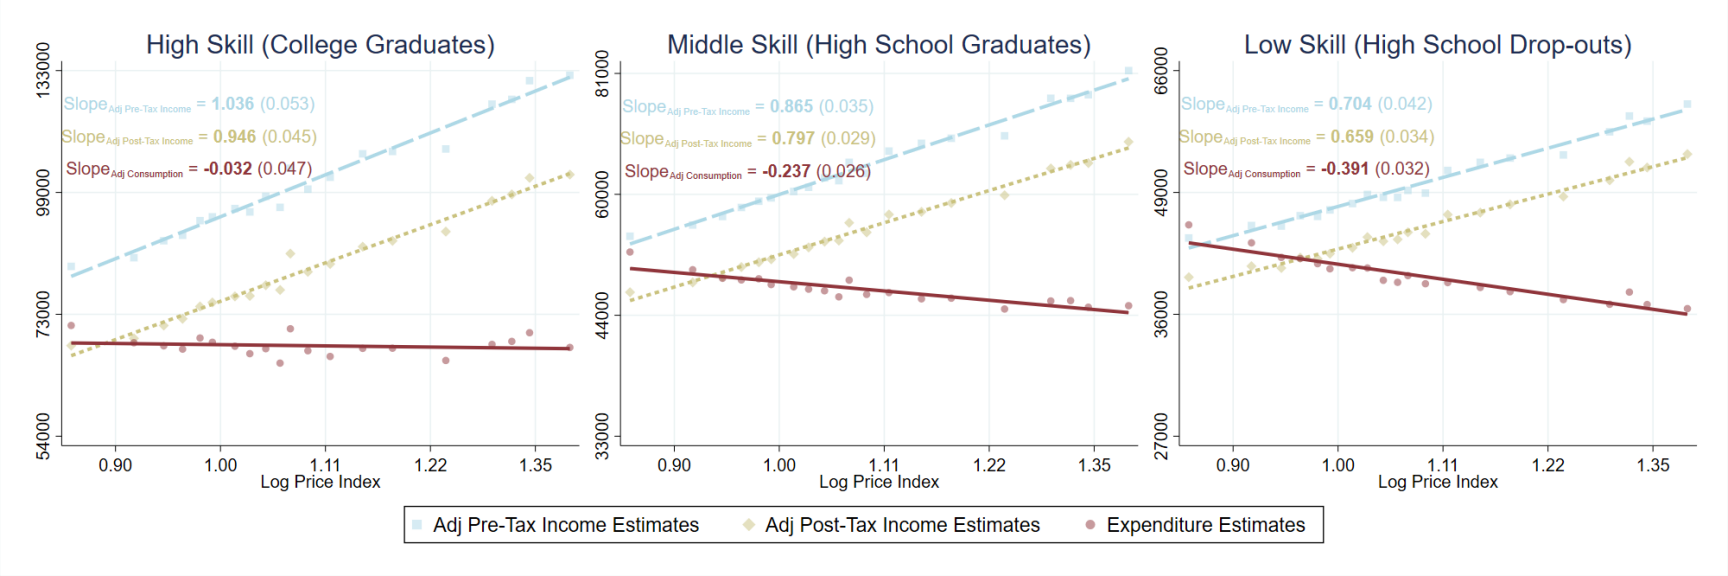
B. Overdraft Fees



C. Bankruptcy

Notes: Panel A plots the share of low-income households with zero or negative savings as a function of the low income price index. Panels B plots the share of income paid by low-income households on overdraft fees as a function of the low income price index. Panels C plots the share of low-income households who pay bankruptcy fees as a function of the low income price index. Overdraft fees and bankruptcy fees are identified from entries in bank account and credit card statements. See text for details.

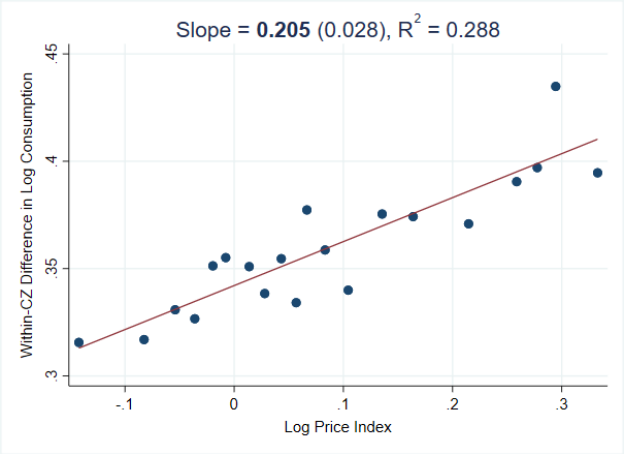
Figure 11: **Pre-Tax Income, Post-Tax Income and Consumption Against Price Index, by Skill Group**



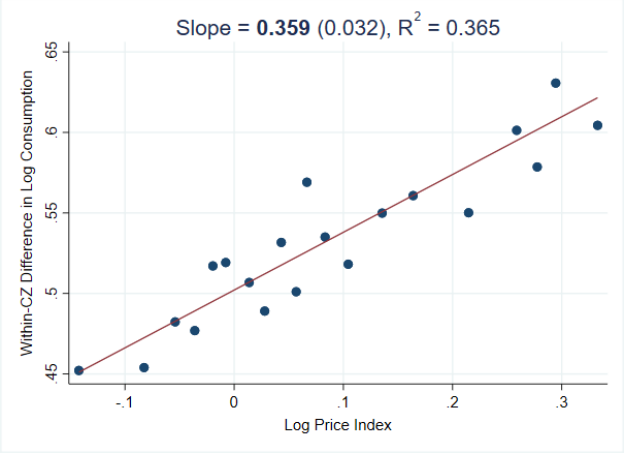
Notes: We plot expected adjusted pre-tax income (light blue), adjusted post-tax income (yellow), and consumption (red) on the y-axis against the relevant price index on the x-axis, across 443 commuting zones.

Figure 12: **Inequality in Consumption Within a Commuting Zone**

(a) College Grads vs. High School Grads



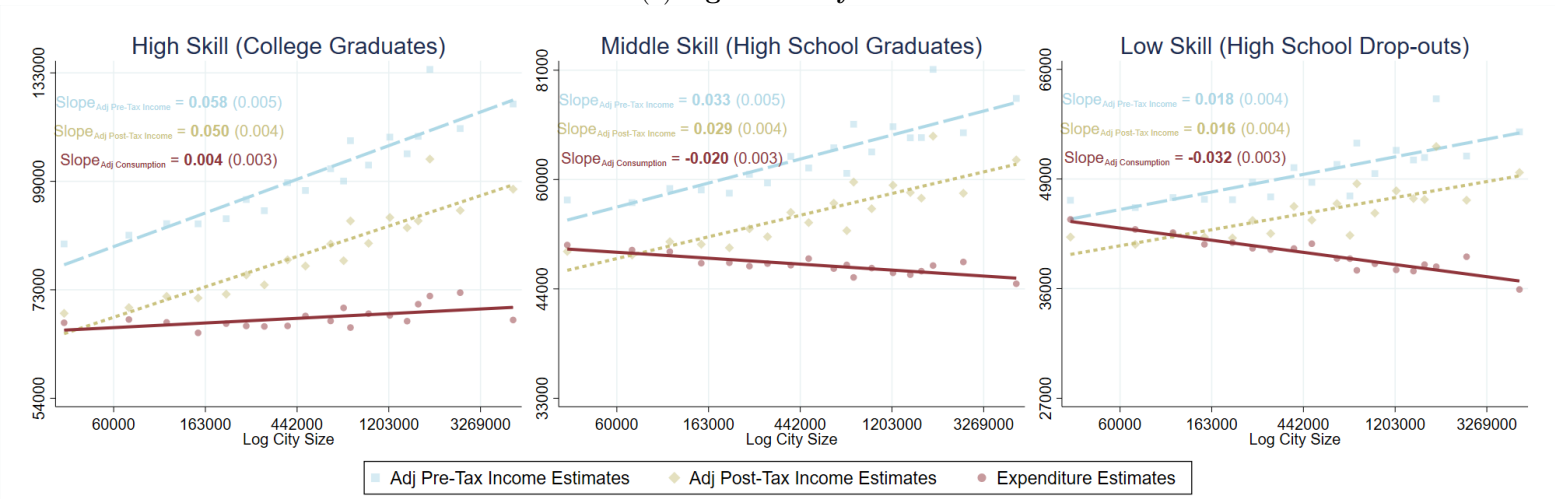
(b) College Grads vs. High School Dropouts



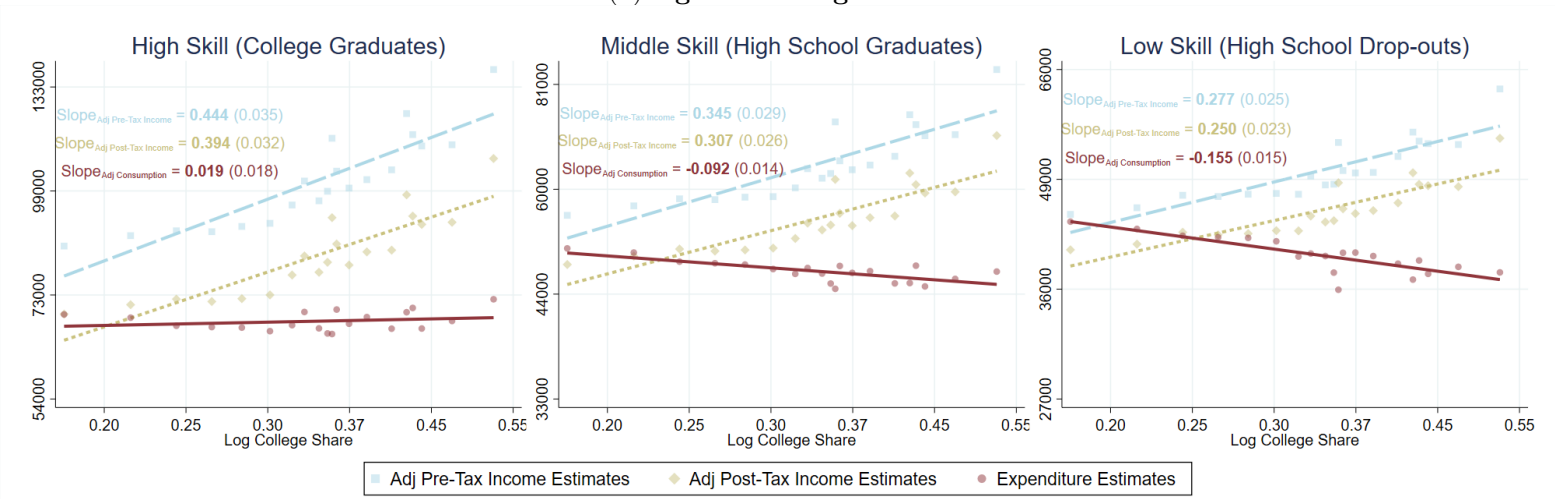
Notes: The top panel plots the difference in mean consumption between high- and middle-skill households who live in the same commuting zones as a function of the price index across all commuting zones in the sample. The bottom panel plots the difference in mean consumption between high- and low-skill households as a function of te price index across all commuting zones in the sample.

Figure 13: **Pre-Tax Income, Post-Tax Income and Consumption Against City Size or College Share, by Skill Group**

(a) **Against City Size**

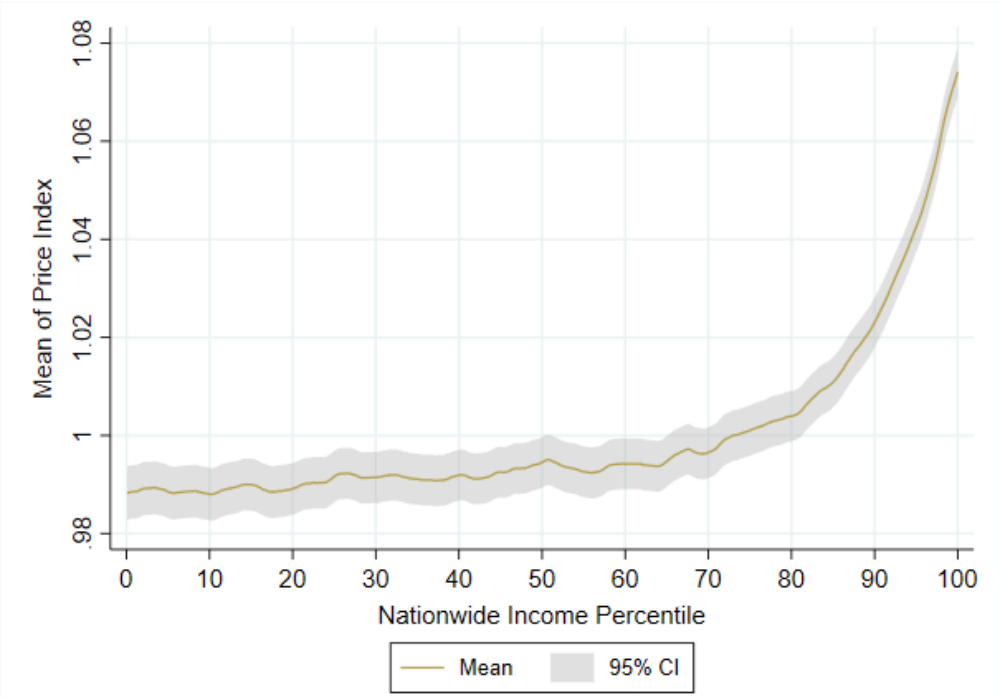


(b) **Against College Share**



Notes: Panel A: We plot expected adjusted pre-tax income (light blue), adjusted post-tax income (yellow), and consumption (red) against city population, across 443 commuting zones. Panel B: same, but college share is on the x-axis.

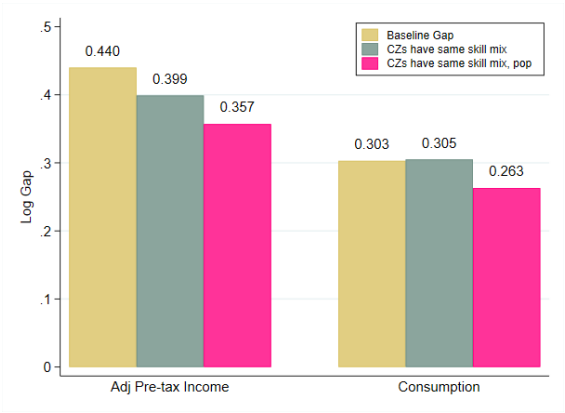
Figure 14: **Mean Price Index by Income Percentile**



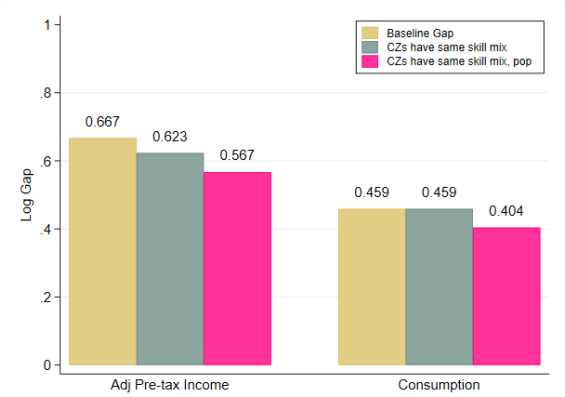
Notes: Mean price index of household's CZ vs household's percentile rank in the nationwide income distribution. Confidence band is in grey.

Figure 15: **Adjusted Pre-tax Income and Consumption Inequality at the National Level**

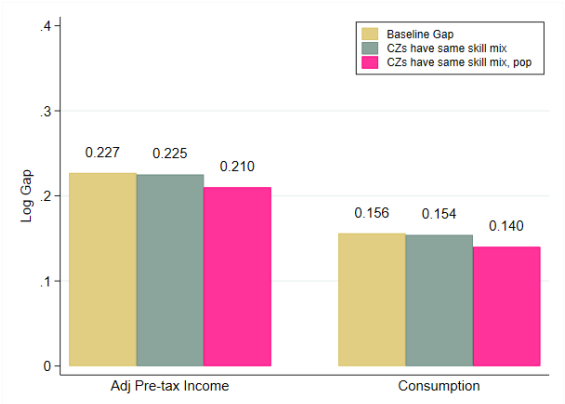
(a) College Grads vs. High School Grads



(b) College Grads vs. High School Dropouts



(c) High School Grads vs. High School Dropouts



Notes: The Figure shows the adjusted income and consumption differences between high- and middle-skill households (top Panel), high- and low-skill households (Middle Panel), and middle- and low-skill households (bottom panel). The first set of bars refers to the baseline differences. The second next set of bars shows the differences when we re-weight households so that the distribution of observable household types within each commuting zone equals the nationwide distribution. The third set of bars show the differences when we further re-weight the data to equate population across CZs, in addition to equalizing household types across CZs.

Table 1: **Commuting Zones by Price Index**

City	Price Index Low Income	Price Index Middle Income	Price Index High Income
Most Expensive			
San Jose, CA	1.491	1.385	1.283
San Francisco, CA	1.477	1.355	1.246
San Diego, CA	1.410	1.329	1.231
Honolulu, HI	1.405	1.322	1.240
New York, NY	1.388	1.317	1.248
Newark, NJ	1.364	1.313	1.257
Santa Barbara, CA	1.360	1.289	1.216
White Plains, NY	1.352	1.309	1.254
Los Angeles, CA	1.349	1.269	1.186
Washington, DC	1.340	1.259	1.172
Edison, NJ	1.335	1.292	1.240
Boston, MA	1.305	1.258	1.231
Seattle, WA	1.276	1.222	1.170
Miami, FL	1.271	1.228	1.170
Hartford, CT	1.265	1.251	1.223
Median			
Jacksonville, NC	1.004	1.017	1.019
Corpus Christi, TX	1.002	1.028	1.040
Cleveland, OH	1.000	1.000	1.000
Coeur d'Alene, ID	1.000	1.014	1.017
Stephenville, TX	0.999	1.035	1.050
Least Expensive			
Union City, TN	0.814	0.864	0.902
Bluefield, WV	0.801	0.854	0.894
Somerset, KY	0.801	0.853	0.894
Summersville, WV	0.801	0.853	0.894
Batesville, AR	0.791	0.845	0.887
Pikeville, KY	0.788	0.842	0.885
North Platte, NE	0.786	0.840	0.884
Marquette, MI	0.785	0.840	0.883
Great Falls, MT	0.775	0.831	0.877
Greenville, MS	0.774	0.831	0.876
Presque Isle, ME	0.773	0.830	0.875
Elkins, WV	0.768	0.826	0.872
London, KY	0.758	0.817	0.865
Gallup, NM	0.755	0.815	0.864
Natchez, MS	0.749	0.810	0.859

Notes: The price indexes for Cleveland are by construction equal to 1. The indexes from other locations are to be interpreted as relative to Cleveland.

Table 2: **Mean Household Consumption by Commuting Zone**

Low-Income Households			High-Income Households		
City Name (1)	Consumption (2)	Price Index (3)	City Name (4)	Consumption (5)	Price Index (6)
Highest Consumption			**Highest Consumption**		
Huntington, WV	47,270	0.832	Huntington, WV	290,538	0.914
Johnstown, PA	47,267	0.819	Johnstown, PA	290,219	0.905
Elizabeth City, NC	47,068	0.988	Toledo, OH	290,030	0.979
Mobile, AL	45,757	0.927	Kalamazoo, MI	288,913	0.960
Florence, SC	45,111	0.863	Erie, PA	288,067	0.958
Traverse City, MI	45,021	0.934	South Bend, IN	286,738	0.945
Bangor, ME	45,002	0.882	Warsaw, IN	285,052	0.949
Charleston, WV	44,914	0.854	Canton, OH	282,762	0.953
Youngstown, OH	44,826	0.863	Cincinnati, OH	276,565	0.987
McAllen, TX	44,774	0.860	Pittsburgh, PA	275,698	1.013
State College, PA	44,052	0.924	Sandusky, OH	275,165	0.991
Beaumont, TX	43,927	0.948	Bloomington, IN	274,972	0.945
Sunbury, PA	43,907	0.871	Fort Wayne, IN	274,955	0.956
Cadillac, MI	43,840	0.876	McAllen, TX	273,986	0.899
Kingsport, TN	43,792	0.883	Cleveland, OH	273,976	1.000
Median Consumption			**Median Consumption**		
Baton Rouge, LA	38,610	0.995	Hagerstown, MD	244,038	1.030
Gainesville, GA	38,568	0.980	Killeen, TX	243,675	0.983
Cleveland, OH	38,553	1.000	Atlanta, GA	243,405	1.017
Prescott, AZ	38,339	1.034	Utica, NY	243,393	1.033
New Orleans, LA	38,268	1.063	Kansas City, MO	243,353	1.050
Lowest Consumption			**Lowest Consumption**		
Hartford, CT	31,496	1.265	Providence, RI	212,144	1.197
Anchorage, AK	31,249	1.246	Boston, MA	211,947	1.231
Boston, MA	31,193	1.305	Olympia, WA	211,209	1.095
Seattle, WA	30,527	1.276	White Plains, NY	210,988	1.254
White Plains, NY	30,076	1.352	Edison, NJ	210,153	1.240
Edison, NJ	29,762	1.335	Virginia Beach, VA	209,649	1.126
Washington, DC	29,542	1.340	Yuma, AZ	207,989	1.041
Los Angeles, CA	29,468	1.349	New York, NY	207,563	1.248
Santa Barbara, CA	29,150	1.360	Newark, NJ	206,557	1.257
Newark, NJ	29,091	1.364	San Francisco, CA	205,274	1.246
New York, NY	29,015	1.388	Seattle, WA	205,111	1.170
Honolulu, HI	28,598	1.405	Medford, OR	198,858	1.067
San Diego, CA	28,227	1.410	San Diego, CA	197,459	1.231
San Francisco, CA	27,304	1.477	San Jose, CA	196,000	1.283
San Jose, CA	27,139	1.491	Honolulu, HI	192,947	1.240

Notes: The consumption levels are priced at the median cost city, Cleveland, OH and real consumption is measured by the expenditure a household would need to spend in Cleveland to achieve the same utility from market consumption as their actual bundle consumed in their city of residence. Only commuting zones with at least 20 individuals in each income group are reported in this table.

Table 3: **Elasticity of Consumption wrt Price Index — NielsenIQ Data**

Product	Unit	Low Income		Middle Income		High Income	
		$\hat{\beta}_{index}$	$\bar{Y}$	$\hat{\beta}_{index}$	$\bar{Y}$	$\hat{\beta}_{index}$	$\bar{Y}$
Beer	KG	0.404	21.2	0.097	22.5	–1.149**	18.7
		(0.296)		(0.332)		(0.543)	
Carbonated Beverages	KG	–0.790***	128.5	–0.958***	129.8	–1.277***	122.1
		(0.080)		(0.116)		(0.170)	
Cookies	KG	–0.202**	6.2	–0.328**	6.3	–0.060	5.8
		(0.085)		(0.134)		(0.206)	
Deodorant	KG	–0.190**	0.3	–0.398***	0.4	–0.892***	0.4
		(0.087)		(0.069)		(0.204)	
Eggs	CT	–0.152**	182.8	–0.107	195.8	–0.396**	186.4
		(0.062)		(0.097)		(0.173)	
Housewares, Appliances	CT	–0.852***	2.2	–1.127***	2.4	–1.757***	2.4
		(0.084)		(0.160)		(0.185)	
Kitchen Gadgets	CT	–0.309	39.0	–0.347	55.2	0.465	64.4
		(0.193)		(0.262)		(0.446)	
Laundry Supplies	KG	–0.396***	11.0	–0.439***	12.2	–0.972***	12.0
		(0.093)		(0.126)		(0.177)	
Light Bulbs, Electric Goods	CT	–0.967***	6.7	–1.253***	7.3	–1.914***	7.7
		(0.109)		(0.138)		(0.305)	
Nuts	KG	0.113	3.1	0.207	4.2	–0.478	4.8
		(0.116)		(0.165)		(0.327)	
Pet Food	KG	–0.813***	57.6	–0.848***	55.3	–1.499***	47.7
		(0.128)		(0.178)		(0.383)	
Pizza, Snacks - Frozen	KG	–0.585***	5.8	–0.836***	6.0	–1.141***	5.6
		(0.129)		(0.155)		(0.246)	
Stationery, School Supplies	CT	–0.497***	228.4	–0.686***	278.5	–0.842**	275.6
		(0.169)		(0.228)		(0.405)	
Vegetables - Frozen	KG	–0.181	10.0	–0.258	11.2	–0.600*	10.0
		(0.126)		(0.158)		(0.315)	

Notes: Entries are from regressions of mean-adjusted quantity of consumption on the local price index controlling for household income; household size; age and presence of children; type of residence; household composition; household head's characteristics including age, gender, race, marital status, education, employment status, and education. We average elasticities by product group. The analysis is based on 59,755 households in the 2014 NielsenIQ Consumer Panel data with at least $10,000 annual income. The numbers of households by income group are 26,533; 23,490; and 9,732. Robust standard errors are clustered by commuting zone and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.01.

Table 4: **Elasticity of Consumption wrt Price Index — Bank Account Data**

Product Group	Low Income		Middle Income		High Income	
	$\hat{\beta}$	$\bar{Y}$	$\hat{\beta}$	$\bar{Y}$	$\hat{\beta}$	$\bar{Y}$
Cable and Satellite Services	-0.292	3.3	-0.629***	6.0	0.306	5.5
	(0.181)		(0.212)		(0.282)	
DHL/FedEX/UPS/USPS	-0.224**	2.2	-0.851***	3.8	-1.391***	4.2
	(0.104)		(0.164)		(0.239)	
Gasoline/Fuel	-1.503***	24.1	-1.802***	35.1	-1.113***	23.1
	(0.245)		(0.244)		(0.407)	
Gym/Fitness/Yoga	0.445*	2.2	-0.094	3.4	-0.466	3.6
	(0.265)		(0.293)		(0.449)	
Hulu/Netflix/Google Play	-1.014***	5.1	-1.255***	6.3	-1.301***	4.0
	(0.100)		(0.172)		(0.227)	
Movies	-1.142***	0.7	-1.748***	1.1	-1.866***	0.8
	(0.331)		(0.413)		(0.545)	
Pedicure and Manicure	0.384**	0.7	0.039	1.2	1.715**	1.5
	(0.185)		(0.259)		(0.698)	

Notes: Entries are from regressions of mean-adjusted number of transactions on log price index controlling for log household income, weighting by commuting zone weights. The sample includes 3,000,518 households in our bank account data: 1,368,817 low income, 1,449,978 middle income, and 181,723 high income. Robust standard errors are clustered by commuting zone and are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.01.

Table 5: **Pre-tax Income, Post-tax Income, and Consumption — High Skill**

City	Adj Pre-tax Income		Adj Post-tax Income		Consumption		Difference between Adj Pre-tax Income & Consumption	
	value (1)	pct (2)	value (3)	pct (4)	value (5)	pct (6)	per (7)	pct (8)
1. San Jose, CA	144,255	100	111,641	100	70,692	78	–51%	–22
2. San Francisco, CA	139,677	100	109,189	100	71,580	84	–49%	–16
3. Washington, DC	138,924	100	106,302	100	73,634	91	–47%	–9
4. White Plains, NY	130,986	99	102,581	99	71,645	84	–45%	–15
5. New York, NY	130,858	99	102,670	99	70,416	76	–46%	–23
6. Newark, NJ	128,240	99	101,669	99	68,869	66	–46%	–33
7. Boston, MA	125,393	99	99,579	99	70,511	77	–44%	–22
8. Hartford, CT	124,635	98	98,518	98	71,915	85	–42%	–13
9. Philadelphia, PA	117,464	98	94,085	98	74,055	93	–37%	–5
10. Los Angeles, CA	116,821	97	92,755	97	65,249	37	–44%	–60
11. Baltimore, MD	116,755	97	91,659	97	72,588	88	–38%	–9
12. Houston, TX	116,019	97	94,260	98	77,372	99	–33%	+2
13. San Diego, CA	114,788	97	91,130	96	60,272	9	–47%	–88
14. Chicago, IL	113,035	96	88,987	95	69,418	69	–39%	–27
15. Dallas, TX	110,571	95	90,370	96	73,907	93	–33%	–2
16. Camden, NJ	110,515	95	88,889	95	68,328	61	–38%	–34
17. Seattle, WA	110,324	95	90,559	96	63,277	21	–43%	–74
18. Sacramento, CA	107,972	94	86,244	93	64,073	28	–41%	–66
19. West Palm Beach, FL	107,016	93	88,594	94	71,033	81	–34%	–12
20. Denver, CO	106,992	93	85,690	92	63,667	26	–40%	–67
21. Austin, TX	106,567	93	87,794	94	69,235	68	–35%	–25
22. Providence, RI	106,342	93	86,189	93	63,178	21	–41%	–72
23. Atlanta, GA	104,870	92	82,246	90	69,668	70	–34%	–22
24. Minneapolis, MN	104,150	91	82,655	90	62,667	18	–40%	–73
25. Fort Worth, TX	103,034	90	84,570	91	69,738	71	–32%	–19
26. San Antonio, TX	101,913	90	83,532	91	70,383	75	–31%	–15
27. Detroit, MI	101,724	89	81,104	88	72,611	88	–29%	–1
28. Miami, FL	101,140	89	84,176	91	63,708	26	–37%	–63
29. Cincinnati, OH	100,956	89	81,026	88	75,585	97	–25%	+8
30. Portland, OR	100,790	88	79,496	86	60,439	10	–40%	–78
31. St. Louis, MO	100,604	88	80,161	87	68,047	59	–32%	–29
32. Virginia Beach, VA	100,111	87	79,113	84	59,702	7	–40%	–80
33. Charlotte, NC	99,401	87	78,347	82	67,252	52	–32%	–35
34. Raleigh, NC	99,209	87	78,259	81	66,638	47	–33%	–40
35. Nashville, TN	98,619	86	81,489	89	70,660	77	–28%	–9
36. Columbus, OH	98,271	85	79,464	85	72,731	89	–26%	+4
37. Phoenix, AZ	97,915	85	79,142	84	67,425	54	–31%	–31
38. Cleveland, OH	97,853	85	79,481	85	73,782	92	–25%	+7
39. Pittsburgh, PA	96,952	83	78,747	83	74,591	95	–23%	+12
40. Kansas City, MO	96,426	82	77,272	77	65,306	37	–32%	–45
41. Jacksonville, FL	96,302	82	79,617	86	65,921	41	–32%	–41
42. Las Vegas, NV	96,131	82	79,664	86	66,725	48	–31%	–34
43. Harrisburg, PA	95,268	80	77,443	79	68,338	61	–28%	–19
44. Tampa, FL	94,490	79	78,534	83	64,130	29	–32%	–50
45. Milwaukee, WI	94,398	78	74,748	68	62,659	17	–34%	–61
46. Salt Lake City, UT	94,032	77	75,230	71	62,261	15	–34%	–62
47. Indianapolis, IN	94,020	77	75,804	73	70,064	73	–25%	–4
48. Buffalo, NY	93,700	76	76,218	75	73,641	91	–21%	+15
49. Orlando, FL	91,991	71	76,833	77	62,815	19	–32%	–52
50. Grand Rapids, MI	86,127	47	70,037	44	66,574	46	–23%	–1

Notes: Entries are average adjusted household pre-tax income, adjusted post-tax income, and consumption across the largest 50 commuting zones. For each variable, we report its corresponding unweighted percentile among all 443 CZs in our data.

Table 6: **Pre-tax Income, Post-tax Income, and Consumption — Middle Skill**

City	Adj Pre-tax Income		Adj Post-tax Income		Consumption		Difference between Adj Pre-tax Income & Consumption	
	value (1)	pct (2)	value (3)	pct (4)	value (5)	pct (6)	per (7)	pct (8)
1. San Jose, CA	86,146	100	72,212	100	46,852	25	–46%	–75
2. Washington, DC	84,623	100	68,851	100	48,503	48	–43%	–52
3. San Francisco, CA	83,341	100	69,805	100	46,380	21	–44%	–79
4. New York, NY	76,773	99	64,599	98	45,319	12	–41%	–87
5. White Plains, NY	76,563	98	64,688	98	46,528	23	–39%	–75
6. Newark, NJ	76,409	98	64,938	99	45,115	11	–41%	–87
7. Boston, MA	75,329	98	63,790	97	47,055	29	–38%	–69
8. Hartford, CT	74,281	96	63,107	97	47,154	30	–37%	–66
9. San Diego, CA	74,091	96	62,744	96	42,531	1	–43%	–95
10. Los Angeles, CA	73,512	96	62,276	95	44,771	9	–39%	–87
11. Baltimore, MD	72,726	95	60,526	94	49,606	63	–32%	–32
12. Camden, NJ	71,851	95	61,255	95	48,424	47	–33%	–48
13. Seattle, WA	71,415	95	61,516	95	44,722	8	–37%	–87
14. Philadelphia, PA	70,049	93	59,584	93	48,069	42	–31%	–51
15. Sacramento, CA	69,248	93	59,146	93	45,722	15	–34%	–78
16. Denver, CO	69,234	93	58,584	92	44,740	8	–35%	–85
17. Chicago, IL	68,390	92	57,263	90	46,150	19	–33%	–73
18. Austin, TX	67,899	91	58,739	92	47,678	37	–30%	–54
19. Minneapolis, MN	67,714	91	57,166	90	44,881	10	–34%	–81
20. Houston, TX	67,636	91	58,265	92	50,206	70	–26%	–21
21. Providence, RI	67,442	91	57,972	91	44,318	5	–34%	–86
22. West Palm Beach, FL	66,425	90	57,933	91	47,435	33	–29%	–57
23. Dallas, TX	66,003	89	57,080	90	48,373	46	–27%	–43
24. Fort Worth, TX	65,784	88	56,813	89	48,652	51	–26%	–37
25. Virginia Beach, VA	65,746	88	55,093	81	43,941	3	–33%	–85
26. Portland, OR	65,515	88	55,123	82	43,297	2	–34%	–86
27. Salt Lake City, UT	64,956	87	55,247	83	46,877	26	–28%	–61
28. Phoenix, AZ	64,226	84	54,906	81	48,220	44	–25%	–40
29. Las Vegas, NV	64,119	84	55,457	84	48,211	44	–25%	–40
30. Harrisburg, PA	64,033	83	54,733	80	49,356	60	–23%	–23
31. San Antonio, TX	63,720	83	55,116	82	48,007	41	–25%	–42
32. Atlanta, GA	63,225	81	53,039	74	46,575	23	–26%	–58
33. Miami, FL	63,048	80	55,243	83	43,472	2	–31%	–78
34. St. Louis, MO	62,294	77	53,034	74	46,872	25	–25%	–52
35. Kansas City, MO	61,713	76	52,490	72	46,454	22	–25%	–54
36. Buffalo, NY	61,518	75	53,276	76	53,846	93	–12%	+18
37. Detroit, MI	61,497	75	52,376	71	48,334	46	–21%	–29
38. Nashville, TN	61,248	72	53,436	76	48,308	45	–21%	–27
39. Raleigh, NC	61,010	71	51,498	64	45,681	15	–25%	–56
40. Jacksonville, FL	60,722	70	52,863	73	46,271	19	–24%	–51
41. Cincinnati, OH	60,540	69	51,971	68	49,685	64	–18%	–5
42. Milwaukee, WI	60,254	67	51,182	61	44,473	6	–26%	–61
43. Pittsburgh, PA	60,123	66	51,851	67	50,294	70	–16%	+4
44. Cleveland, OH	59,927	66	51,751	66	48,961	54	–18%	–12
45. Tampa, FL	59,605	64	52,255	70	44,799	10	–25%	–54
46. Orlando, FL	59,602	64	52,355	71	44,436	6	–25%	–58
47. Charlotte, NC	59,538	63	50,429	55	45,272	12	–24%	–51
48. Columbus, OH	59,201	61	51,092	60	48,293	45	–18%	–16
49. Indianapolis, IN	58,724	57	50,161	52	47,727	38	–19%	–19
50. Grand Rapids, MI	55,833	34	48,303	33	47,468	34	–15%	+0

Notes: Entries are average adjusted household pre-tax income, adjusted post-tax income, and consumption across the largest 50 commuting zones. For each variable, we report its corresponding unweighted percentile among all 443 CZs in our data.

Table 7: **Pre-tax Income, Post-tax Income, and Consumption — Low Skill**

City	Adj Pre-tax Income		Adj Post-tax Income		Consumption		Difference between Adj Pre-tax Income & Consumption	
	value (1)	pct (2)	value (3)	pct (4)	value (5)	pct (6)	per (7)	pct (8)
1. San Jose, CA	64,050	100	56,659	100	37,710	6	–41%	–94
2. Washington, DC	63,280	99	54,495	99	39,320	18	–38%	–81
3. San Francisco, CA	61,742	99	54,444	99	36,943	2	–40%	–97
4. Newark, NJ	59,808	98	53,210	98	37,812	6	–37%	–92
5. White Plains, NY	58,297	97	52,144	98	38,143	9	–35%	–88
6. Boston, MA	58,279	97	51,841	98	39,560	21	–32%	–76
7. New York, NY	57,765	96	51,598	97	36,909	2	–36%	–94
8. Camden, NJ	56,672	95	50,727	96	41,378	43	–27%	–52
9. Hartford, CT	56,217	95	50,475	95	38,423	10	–32%	–85
10. Seattle, WA	56,153	95	50,240	95	37,659	5	–33%	–90
11. San Diego, CA	55,018	93	49,286	94	34,453	1	–37%	–92
12. Los Angeles, CA	54,482	93	48,795	93	35,790	1	–34%	–92
13. Denver, CO	54,436	93	48,110	92	37,277	4	–32%	–89
14. Harrisburg, PA	53,826	92	47,511	91	43,495	67	–19%	–25
15. Chicago, IL	53,770	92	47,076	90	39,044	15	–27%	–77
16. Baltimore, MD	53,716	92	47,207	90	39,687	24	–26%	–68
17. Providence, RI	53,183	91	47,869	92	38,061	8	–28%	–83
18. Salt Lake City, UT	53,153	91	47,287	90	41,169	41	–23%	–50
19. Portland, OR	53,140	91	46,658	87	37,417	5	–30%	–86
20. Philadelphia, PA	53,010	91	47,295	91	38,978	14	–26%	–77
21. Minneapolis, MN	52,401	89	46,708	88	37,735	6	–28%	–83
22. Sacramento, CA	51,827	88	46,755	88	36,978	3	–29%	–85
23. West Palm Beach, FL	51,106	84	46,549	87	38,853	12	–24%	–72
24. Las Vegas, NV	51,045	83	45,796	83	41,117	40	–19%	–43
25. Austin, TX	50,871	82	45,931	84	38,495	10	–24%	–72
26. Virginia Beach, VA	50,813	82	44,856	77	37,138	3	–27%	–79
27. Houston, TX	50,001	78	45,103	79	40,508	31	–19%	–47
28. Dallas, TX	49,583	76	44,886	78	39,344	18	–21%	–58
29. Buffalo, NY	49,518	75	45,101	79	47,062	91	–5%	+16
30. Fort Worth, TX	49,237	73	44,423	74	39,381	19	–20%	–54
31. Kansas City, MO	49,037	71	43,590	67	39,571	21	–19%	–50
32. Phoenix, AZ	48,932	70	43,983	71	39,808	25	–19%	–45
33. Miami, FL	48,682	69	44,557	75	35,928	1	–26%	–68
34. Milwaukee, WI	48,593	68	43,319	64	38,629	11	–21%	–57
35. Detroit, MI	48,412	67	43,254	63	40,750	35	–16%	–32
36. St. Louis, MO	48,201	64	43,064	61	39,188	16	–19%	–48
37. Pittsburgh, PA	48,052	63	43,328	64	42,631	59	–11%	–4
38. San Antonio, TX	47,691	60	43,203	62	38,724	12	–19%	–48
39. Nashville, TN	47,591	59	43,394	65	40,615	33	–15%	–26
40. Atlanta, GA	47,546	58	42,384	53	38,060	8	–20%	–50
41. Orlando, FL	47,529	58	43,474	66	37,971	7	–20%	–51
42. Tampa, FL	47,119	55	43,075	61	37,926	7	–20%	–48
43. Jacksonville, FL	46,901	53	42,740	58	39,014	14	–17%	–39
44. Grand Rapids, MI	46,862	52	42,188	50	42,090	51	–10%	–1
45. Cleveland, OH	46,788	51	42,136	50	40,516	32	–13%	–19
46. Cincinnati, OH	46,156	45	41,574	43	40,567	33	–12%	–12
47. Charlotte, NC	46,127	45	41,227	38	38,284	9	–17%	–36
48. Raleigh, NC	45,873	42	41,169	37	38,084	9	–17%	–33
49. Indianapolis, IN	45,779	41	41,039	36	40,248	29	–12%	–12
50. Columbus, OH	45,530	38	41,243	38	39,925	25	–12%	–13

Notes: Entries are average adjusted household pre-tax income, adjusted post-tax income, and consumption across the largest 50 commuting zones. For each variable, we report its corresponding unweighted percentile among all 443 CZs in our data.

Table 8: **Consumption vs. Price Index, City Size, and College Share**

	Log Consumption
Log price index	-0.263***
	(0.075)
Log price index × middle-skill	0.039
	(0.086)
Log price index × low-skill	-0.022
	(0.087)
Log city size	0.027***
	(0.008)
Log population × middle-skill	-0.026***
	(0.009)
Log population × low-skill	-0.040***
	(0.009)
Log college share	0.021
	(0.037)
Log college share × middle-skill	-0.036
	(0.042)
Log college share × low-skill	-0.028
	(0.042)
Middle-skill	-0.063
	(0.144)
Low-skill	-0.036
	(0.145)

Notes: Entries are from a regression of log consumption on log price index, log city size, and log college share all interacted with education group identifiers. The level of analysis is commuting zone × education group. Observations are weighted by commuting zone population. $N = 1329$ *** p<0.01, ** p<0.05, * p<0.10.

Online Appendix

A Household Sample in Bank Account Data

The raw bank account data span 2011 to 2016 and are most populated in 2014, which we choose as our year of focus. To ensure that we have a complete twelve-month coverage for all households in 2014, we keep only households who enter our data during 2011-2013. We have 4,150,659 households in 2014 at the start.

We geocode all physical (i.e., non-online) merchants for which we observe the addresses in our data and take the commuting zone in which each household transact most frequently each year as its annual "modal commuting zone". We drop households for which we do not have sufficient data to identify the modal commuting zone, leaving 3,847,005 households from 703 commuting zones.

For each household, we define annual income as a total dollar amount across all transactions in 2014 paid into bank account as "credit", taking out transfers between accounts and debit income taxes (from both bank and credit or debit card accounts). To identify the transfers, we filter through individual credit transactions in bank account using various keywords

Similarly, we define annual expenditure as a total dollar amount across all transactions in 2014 paid out of bank account as "debit", taking out transfers between accounts and debit income taxes. We also take out transactions that do not reflect consumption realized in the current period such as loans, retirement contributions, and investments

We drop households with missing annual income or annual income less than $10,000, leaving 3,382,105 households. Second, we drop households with missing annual expenditure or annual expenditure less than $1,000, leaving 3,366,135 households.

Among the remaining households, some have high frequencies of small- or medium-sized business transactions (e.g., advertising and marketing, business miscellaneous, employee and officer compensations, paychecks and salaries, and payroll services). Because these households are more likely to be small- or medium-sized businesses rather than family households, we exclude households with spending of these types greater than $500 in 2014. This restriction leaves 3,107,351 households. We further drop households for which we cannot measure spending across different categories precisely. These households are those for which we cannot link their bank accounts with associated card accounts. This restriction leaves 3,013,465 households.

We only keep commuting zones featuring at least three households from each of the three income groups, leaving 3,000,518 households from 443 commuting zones in the the final sample. These commuting zones represent 96.3% of the US population.

B Healthcare and Housing Adjustments

Healthcare. Out-of-pocket (OOP) spending observed in our data does not reflect total health charges. To quantify the total amount of health care expenditures, we turn to Medical Expenditure Survey (MEPS) data, pooling 2012-2016 years. This dataset allows us to measure total expenditures and out-of-pocket spending for both healthcare and pharmacy at the household level. Using the MEPS data, we regress total healthcare (or pharmacy) spending on out-of-pocket healthcare (or pharmacy) spending. Then, we use these relationships to predict total health spending and total pharmacy spending in our data. Finally, we re-calculate our "Healthcare/Medical" spending as a sum of total healthcare charges, total pharmacy charges, and the original non-recreation health

spending.[35]

Specifically, we take the following steps.

First, the MEPS is a household survey that is known to under-report spending relative to NIPA (Bernard et al., 2012). We therefore inflate MEPS spending by a factor of 1.37 to obtain adjusted health spending so that overall healthcare spending matches the reported NIPA healthcare spending. We regress adjusted total healthcare spending on adjusted OOP healthcare spending, controlling for household income and region of the country, and the interactions of these terms. The coefficient on OOP spending is 5.409 (0.539), indicating that each additional dollar of OOP spending (excluding prescriptions) corresponds on average to $5.41 of total expenditure. A similar regression for prescription expenditures yields a coefficient equal to 3.756 (1.219).

Second, we identify all OOP non-drug and drug spending in 2014. We separately regress non-drug health spending and prescription spending on gross annual overall (not just health) expenditure. We control for income group and its interaction with total personal expenditure. We then use this regression to estimate the share of drug and non-drug health expenditure for each individual using their total annual health expenditure and income group.

Finally, for each household in the bank data, we use our estimated coefficients to impute total drug and non-drug health expenditure for a given level of observed OOP health expenditures. The OOP non-drug health expenditure mean is $736 and the total non-drug health expenditure mean is $5,689. The medians are $237 and $4,022. For drug expenditure, the means are $1,111 and $5,547, and the medians are $448 and $2,608.

Housing. We use the same methodology and data employed by the BLS for estimating the housing expenditures used in the CPI (Poole et al. (2005); Bureau of Labor Statistics (2007)).

(a) We estimate average rental payments for renters by income group in the 2012–2016 pooled ACS.

(b) For owners, expenditures on housing need not equal the cost of purchasing one year of housing services. Homeowners that do not have a mortgage pay housing costs that are likely to be lower than cost of purchasing the year of housing services they consume. Owners with a mortgage are likely to spend more than the cost of a single year of housing services. Homeowners who spent more than the imputed rental values of their homes are effectively earning negative income on their housing asset this year. Thus these "excess housing payments" are not actually expenditure on consumption. This excess spending needs to be subtracted out from their spending and income. Homeowners that spend less than the imputed rental values of their homes are earning income from their housing asset. This needs to be added back to their income and expenditure. This adjustment is standard in consumption inequality literature. Following the BLS, for owners we use a measure of "rent equivalent" from the CEX, which is defined as the rental value of their home if they were to rent it out, unfurnished, and without utilities. We take rent equivalent for each income group from the CEX, pooling the 2012–2016 data.

These are the specific steps that we have taken:

First, we measure housing costs using CEX Interview Survey data by pooling 2012–2016 years (centering around 2014). We define two measures of housing costs in the CEX data. The first

[35]Our data also have an additional limitation: our health spending includes health-related transactions that are either recreational or not covered under health insurance such as gym/fitness membership, veterinary services, and vision expenses. To address this issue, we classify "Healthcare/Medical" spending in our data into healthcare, pharmacy, and recreational health-related spending, using relationships among these measures established in Diamond et al. (2018).

measure is "housing costs to be subtracted", which includes contract rent and owner costs (purchase costs, closing costs, mortgage payments, and down payments). The second measure is "housing costs to be added", which includes contract rent and equivalent rent. In the steps below, we take our total expenditure, subtract out the first housing cost measure, and then add back the second housing cost measure to re-define total expenditure for all households.

Second, using the CEX data, we regress our defined housing costs (both measures, separately) on property value, its squared term, post-tax income, and number of rooms separately by region × income group. Then, we use the coefficient estimates to predict our housing measures for owner-occupied units in the ACS data. We can take the rental payments for renters directly from the ACS. Since the ACS has a much larger sample, it can measure the distribution of housing types in each CZ with much more precision. We use the estimated relationship between these housing characteristics and housing expenditures as measured in the CEX, but then apply this relationship to the types of housing and the income observed in the ACS to get a more precise estimates of our housing spending measures at the CZ and income group level.

Third, we need to assign these estimated housing expenditures as measured in the ACS to our bank transaction households. We match households in our bank data to those in the ACS based on income and commuting zone. Specifically, we regress our imputed housing costs in the ACS (both measures, separately) on post-tax income by commuting zone × income group. Then, we use the estimates to predict both types of housing costs for all households in our bank account data. For reference, mean housing costs to be subtracted are $16,924 for overall-income; $12,181 for low-income; $18,812 for middle-income; and $42,249 for high-income households. Mean housing costs to be added are $16,009 for overall; $9,567 for low-income; $17,747 for middle-income; and $57,870 for high-income households.

C Price Indexes

C.1 Measuring Prices of Good and Services

Nielsen Data. We use price data from the 2014 Nielsen Retail Scanner data for six consumption categories: Grocery, General Merchandise, and Personal Care; and three additional categories for which we can find a one-to-one map to a product group in Nielsen: Baby Needs, Electronics, and Office Supplies. The Nielsen data contain all UPCs purchased and recorded by Nielsen-participating households in a given year. We merge in product details (e.g., department, product group, product module, size, and unit) and household characteristic indicators (e.g., household income; household size; age and presence of children; type of residence; household composition; household head's characteristics including age, gender, race, martial status, education, employment status, and education; and the commuting zone they lived in 2014).

In 2014 there are 64,717,120 UPC purchases and 823,507 distinct UPCs from 1,100 modules, 116 product groups, and 10 departments. To make it consistent with our household sample in the bank account data, we drop households in Nielsen with 2014 annual income lower than $10,000, leaving 61,903,872 UPC purchases made by 59,756 households in 660 commuting zones. Then, we classify the remaining households into three income groups: low (10K-50K), middle (50K-100K), and high ($\geq$100K; note that the income indicator is top-coded at 100K in 2014). The corresponding numbers of households are 26,534 for low-income, 23,490 for middle-income, and 9,732 for high-income.

We calculate commuting-zone-specific prices at the product group level. In particular, for each product group, we regress log UPC price on commuting zone indicators, UPC fixed effects and

weighting observations by expenditures on the UPCs. We estimate:

$$\log p_{u,j} = \delta_u + \delta_{p(u),j} + \epsilon_{u,j}$$

where $u \in U$ is UPC belonging to product group $p(u) \in P$ purchased in commuting zone $j \in J$. The UPC fixed effects, δ_u, control for quality differences in products consumed in different locations. The estimated coefficient on $\delta_{p,j}$, evaluated at the nationwide shares across all UPCs within a given product group, is used as the conditional mean price of product group p faced by any income group in commuting zone j.

We follow a similar procedure in the case where we allow prices to also vary by income group within the same commuting zone. Specifically, for each product group and income level, we regress log UPC price on commuting zone indicators, absorbing UPC fixed effects and household income group indicators:

$$\log p_{u,j,h} = \delta_u + Y_h + \delta_{p(u),j,k(h)} + \epsilon_{u,j,h}$$

where $k(h) \in \{\text{overall}, \text{low}, \text{middle}, \text{high}\}$ denotes an income group to which household h belongs. The estimated coefficient on $\delta_{p,j,k}$, evaluated at the nationwide shares across UPCs within a given product group and at a fixed nominal income bracket, is used as the conditional mean price of product group p faced by income group k in commuting zone j.

ACCRA Data. We use ACCRA prices for nine consumption categories: Automotive Expenses, Clothing/Shoes/Jewelry, Gasoline/Fuel, Healthcare/Medical, Hobbies/Entertainment, Miscellaneous Services, Restaurants/Dining, Telecommunications, and Utilities.

The geographical unit of analysis is different from commuting zone. Specifically, ACCRA data contain prices of goods and services at the Core-Based Statistical Area (CBSA) level. After correcting one miscoded CBSA from 14460 to 14454, we compute mean prices by CBSA, using city population making up each CBSA as weight. Then, we crosswalk from CBSA to county and to commuting zone and take the mean by commuting zone, using county population as weight. Of 249 CBSAs in the 2014 data: 83% map to exactly one commuting zone; 13% map to two commuting zones; and 4% map to three commuting zones or more. One limitation of ACCRA data is that the raw data only cover 254 commuting zones in 2014. To improve geographical coverage, we impute prices for some missing areas. For a given missing commuting zone, we calculate a population-weighted average price using prices in all "neighbor" counties. A neighbor county is defined as any county with prices available, is contained within a commuting zone that shares borders with the commuting zone being imputed, and is located within a 35 mile radius from centroid to centroid. The assumption is that prices are similar in areas that are at most 35 miles apart. This imputation allows us to cover 326 commuting zones in 2014.

ACS Data. To measure housing costs, we use household-level ACS data. Following the approach used by the BLS to estimate the CPI, we measure housing costs using rental prices.

We use 2012–2016 ACS data (centered at 2014), which include 6,838,804 households. We begin by assigning each household a commuting zone. In the ACS data, we can identify county of residence as long as that county belongs to an MSA; otherwise, the county code is missing. However, information on Public Use Microdata Area (PUMA) is available for all households. To assign each household a commuting zone, we build a crosswalk from state-PUMA to commuting zone by overlaying maps in ArcGIS. Because some PUMAs map to multiple commuting zones, we randomly assign each

household a commuting zone based on a fraction of PUMA population that is made up of that commuting zone such that a commuting zone with a larger population share has a higher assignment probability.

We then estimate mean rents controlling for observable housing characteristics. In particular, we interact the following five housing characteristics to define "housing types" (n): (1) Year the structure was built (before 1950, 1950-1969, 1970-1989, and from 1990 onward); (2) Unit structure (one-family house, multiple family building, and other remaining structures); (3) Number of rooms (at most three rooms, four rooms, five rooms, six to seven rooms, and eight rooms or more); (4) Number of bedrooms (at most one bedroom, two bedrooms, three bedrooms, and four bedrooms or more); and (5) Presence of facilities (having all of the above listed facilities; and lacking at least one facility). There are $N = 192$ types of housing nationwide. We calculate $\bar{s}_{n,j,k}$ or the share of all housing units (owner-occupied and renter-occupied) of type n for income group k in commuting zone j.

For each commuting zone and for each housing type, we calculate mean monthly contract rent among all observed renter-occupied units, using household weights in the ACS data.[36] Then, for a given commuting zone and income group, we calculate mean contract rent across our defined housing types, where we weight each housing type by its relative prevalence within the commuting zone. Specifically, we estimate the commuting-zone-level monthly rents as $\text{rent}_{j,k} = \sum_{n=1}^{N}(\text{rent}_{n,j,k} \times \bar{s}_{n,j,k})$.

C.2 Measuring Expenditure Shares

We closely follow the methodology that the BLS uses to calculate expenditure shares to compute the CPI. An expenditure share on a given item is defined as total consumption expenditure on this item across households divided by total consumption expenditure on all items across households. Specifically, we define income-group-specific nationwide shares and income-and-commuting-zone-specific shares, respectively, as

$$s_{i,k} := \frac{\sum_{h \in \bigcup_{j \in J} H_j} E_{h,i,j(h),k(h)}}{\sum_{i \in I} \sum_{h \in \bigcup_{j \in J} H_j} E_{h,i,j(h),k(h)}} = \frac{\bar{E}_{i,k}}{\sum_{i \in I} \bar{E}_{i,k}}$$

$$s_{i,j,k} := \frac{\sum_{h \in H_j} E_{h,i,j(h),k(h)}}{\sum_{i \in I} \sum_{h \in H_j} E_{h,i,j(h),k(h)}} = \frac{\bar{E}_{i,j,k}}{\sum_{i \in I} \bar{E}_{i,j,k}}$$

where I denotes the set of 22 high-level categories; J denotes the set of commuting zones in our data; and K denotes the set of income groups. H_j is the set of households in commuting zone j. $E_{h,i,j(h),k(h)}$ is the total expenditures on high-level category i of household h belonging to income

[36] Not all housing types are available for all income groups in all commuting zones. For such cases, we use contract rents from 2012-2016 county-level ACS data. For each county, we calculate housing characteristic "fractions". For example, if there are 10,000 rental units in county A such that 9,900 units have complete plumbing facilities and 100 units lack such, the corresponding fractions are 0.99 and 0.01. We do this for all categories within each housing characteristic. Then, we regress log monthly contract rent on commuting zone indicators, controlling for characteristic fractions and using county population as weight. Precisely, we let $p_{housing,c}$ be a median rent in county c. We estimate $\log p_{housing,c} = \delta_{j(c)} + X\beta + \epsilon$, where $j(c)$ is the commuting zone to which county c belongs; and X is a vector of country-level housing characteristic fractions. We predict commuting-zone-level monthly rent, evaluated at the nationwide population-weighted-average characteristic fractions that are the same for all commuting zones, i.e., $\widehat{p_{housing,j}} = exp(\hat{\delta}_j + \bar{X}\hat{\beta})$.

group $k(h)$ and living in commuting zone $j(h)$. For both types of shares, we divide the numerator and the denominator by their corresponding total number of households: $\sum_{j \in J} |H_j|$ for $s_{i,k}$ and $|H_j|$ for $s_{i,j,k}$. Therefore, $\bar{E}_{i,k}$ is household-average expenditure on category i for income group k nationwide and $\bar{E}_{i,j,k}$ is household-average expenditure on category i for income group k in commuting zone j.

Our data classify expenditure in 21 high-level consumption categories. An important limitation is that we can identify categories only for transactions done by credit card or debit card or electronic transfer. Transactions by cash or checks are labelled in our data as "Unclassified" because the identity of the merchant is unknown. However, we note that the distribution of classified expenditures across categories match well the NIPA shares (Section 2).

To calculate household h's expenditure shares, we take its total expenditure (E_h) and then subtract out our imputed housing costs (H_h^-), leaving total non-housing expenditure (N_h): $N_h = E_h - H_h^-$. This non-housing expenditure is the sum of expenditures paid through bank or card accounts—which are assigned to the 21 non-housing categories ($X_{h,i}$ for $i \in I = \{1, \dots 21\}$)—and expenditures paid in cash or checks—which are "Unclassified": $N_h - \sum_{i \in I} X_{h,i}$. Because we cannot identify what types the latter spending consists of, we apportion it back to our focal categories or $\tilde{X}_{h,i} = \frac{X_{h,i}}{\sum_{i \in I} X_{h,i}} \times N_h$: as such, $\sum_{i \in I} \tilde{X}_{h,i} = N_h$. Next, we add back our imputed housing costs (H_h^+ or $X_{h,22}$) to our total non-housing expenditure to re-calculate total expenditure, equivalently, $\tilde{E}_h = H_h^+ + N_h$. For each household, we calculate expenditure shares defined as $s_{h,i} = \frac{\tilde{X}_{h,i}}{\tilde{E}_h}$ for $i = 1, \dots, 22$.

The expenditure shares for each category vary by income group and are listed in Appendix Table A2.

The following three categories—General Merchandise, Groceries, and Personal Care—are very broad. To improve precision, we build expenditure shares of product groups nested within each of these three categories using Nielsen data. In practice, we build $s_{g,k}$ and $s_{g,j,k}$ for $g \in i(g)$ or the set of product groups belonging to a high-level category $i \in \{\text{General Merchandise}, \text{Groceries}, \text{Personal Care}\}$. We calculate expenditure shares by product group by dividing total expenditure for a given product group by total expenditure from all product groups that map to the high-level category considered. We do this separately by income group to obtain income-groups-specific shares. For each of these three high-level categories, we scale down the nested shares so that they sum to the corresponding share relative to the 22 high-level categories.

The shares for each subcategory are shown in Appendix Table A3.

Some of our alternative price indices require measuring expenditure shares at the commuting zone level. We discuss below how we estimate commuting zone-income group-specific expenditure shares.

C.3 Alternative Price Indices

Törnqvist Index The Törnqvist index is a second-order approximation to the true price index for a pair of cities or time periods. We compare the pair of a given city and the nationwide average. The Törnqvist is a geometric means of relative prices, weighted by the average of the CZ-specific

expenditure shares and the nationwide average expenditure shares.

$$P_{j,k}^{\text{Törnqvist}} = \prod_{i \in I} \left(\frac{p_{i,j}}{\overline{p_i}} \right)^{\frac{s_{i,k}+s_{i,j,k}}{2}} \tag{3}$$

CES Price Index The CEX price index is not income group specific, but is an exact cost-of-living index if the true utility function is CES. The elasticity of substitution is implicitly estimated through the transformation of the CZ-specific expenditure shares. The formula is:

$$P_{j,k}^{\text{CES}} = \prod_{i \in I} \left(\frac{p_{i,j}}{\overline{p_i}} \right)^{\omega_{i,j,k}} \tag{4}$$

where (i) $\omega_{i,j,k} = \frac{\mu_{i,j,k}}{\sum_{i \in I} \mu_{i,j,k}}$ and $\mu_{i,j,k} = \frac{s_{i,j,k} - s_{i,k}}{\ln(s_{i,j,k}) - \ln(s_{i,k})}$;

Nested CES: We follow Handbury and Weinstein (2015) in building a nested-CES exact local price index, accounting for variation in local supply of products. We measure the same price index for all three income groups. Just like our main Laspeyres price index, we index the "high-level" product categories by i where I is the set of all high-level product categories. Within each product category i, there are mid-level categories that classify purchases into product groups. These are indexed by im. Only the 3 high-level product categories have products split into mid-level nests (e.g. yogurt versus cheese). This is because the Nielsen data provides this additional level information about the products. These mid-level nests are split based on the product groups assigned to each product by Nielsen. For high-level product categories not covered by Nielsen, these is no mid-level nest. Finally, the lowest level nest measures utility from each individual variety of product. These are indexed by g. For the 3 Nielsen groups, we use UPC codes to identify unique varieties. For the rest of the product categories not covered by Nielsen, we use merchants, as observed in our transaction data to identify a unique variety. Surely most merchants sell a variety of products, but merchant is the most granular data we observe. For most transactions, our data provider as already listed the merchant associated with each transaction. For smaller merchants, this variable is blank in our data. To measure merchants for these additional transactions, we standardize the description string from the transaction by cleaning out text from the bank itself (e.g. remove words like "CHECKCARD PURCHASE"), and other formatting differences across banks to create a text string unique to each merchant. The utility function is:

$$U = (\sum_{i \epsilon I} (C_i)^{\frac{1}{\sigma - 1}})^{\sigma - 1},$$

$$C_i = (\sum_{m \epsilon M_i} (d_{im})^{\frac{1}{\sigma_i - 1}})^{\sigma_i - 1}, \quad d_{im} = (\sum_{g \epsilon G_{im}} (\lambda_{img} c_{img})^{\frac{1}{\sigma_{im} - 1}})^{\sigma_{im} - 1}$$

c_{img} is the quantity of variety g within expenditure category im consumed. M_i is the set of product groups within high-level expenditure category i. For categories not covered by the Nielsen data, there is only a single variety m in the set M_i. G_{im} is the set of varieties within mid-level category im. λ_{igm} measures the quality of variety g within expenditure category im. σ_{im} is the elasticity of substitution between varieties within category im. σ_i is the elasticity of substitution between mid-level product categories m within high-level category i. σ is the elasticity of substitution between high-level categories.

As shown by Handbury and Weinstein (2015), the price index EPI_j for CZ j that accounts for

variation in access to local variety can be written as:

$$EPI_j = \prod_i [CEPI_{ij} VA_{ij}]^{w_{ij}},$$

where:

$$CEPI_{ij} = \prod_{g \epsilon G_{ji}} \left(\frac{P_{gj}}{P_g} \right)^{w_{gj}},$$

$$VA_{ij} = \prod_{i \epsilon I, m \epsilon M_i} s_{imj}^{\frac{w_{imj}}{1-\sigma_{im}}},$$

$$P_g = \frac{\sum_j E_{gj}}{\sum_j \frac{E_{gj}}{P_{gj}}}, \qquad s_{imj} = \frac{\sum_{g \epsilon G_{jim}} \sum_{j \epsilon J} E_{gj}}{\sum_{g \epsilon G_{im}} \sum_{j \epsilon J} E_{gj}}.$$

$CEPI_{ij}$ measures the contribution of the local prices P_{gj} relative to national average prices P_g for each variety g to the price index for CZ j, among G_{ji}, the set of varieties within product category i available for sale in CZ j. VA_{ij} represents the variety adjustment to differences in varieties available in each CZ j. s_{imj} measures the share of nationwide sales that are available among the variety for sale in CZ j within product category im. E_{gj} is the total expenditure on variety g in CZ j. G_{jim} is the set of varieties for sale in CZ j in product category im. w_{ij}, w_{gj}, and w_{imj} are the Sato-Vartia weights and are defined as follows:

$$w_{ij} = \frac{\frac{sh_{ij} - sh_i}{\ln sh_{ij} - \ln sh_i}}{\sum_{i' \epsilon I} \left(\frac{sh_{i'j} - sh_{i'}}{\ln sh_{i'j} - \ln sh_{i'}} \right)} \qquad w_{gj} = \frac{\frac{sh_{gj} - sh_g}{\ln sh_{gj} - \ln sh_g}}{\sum_{m \epsilon M_i} \sum_{g' \epsilon G_{im}} \left(\frac{sh_{g'j} - sh_{g'}}{\ln sh_{g'j} - \ln sh_{g'}} \right)},$$

$$w_{imj} = \frac{\frac{sh_{mj} - sh_m}{\ln sh_{mj} - \ln sh_m}}{\sum_{m' \epsilon M_i,} \left(\frac{sh_{m'j} - sh_{m'}}{\ln sh_{m'j} - \ln sh_{m'}} \right)}, w_{imj} = 1 \text{ for non-nielsen categories.}$$

$$sh_{ij} = \frac{\sum_{m \epsilon M_i, g \epsilon \{G_i, G_{im}\}} E_{gj}}{\sum_{i \epsilon I} \sum_{m' \epsilon M_i, g' \epsilon \{G_i, G_{im}\}} E_{g'j}}, \qquad sh_i = \frac{\sum_{m \epsilon M_i, g \epsilon \{G_i, G_{im}\}} E_g}{\sum_{i \epsilon I} \sum_{m' \epsilon M_i, g' \epsilon \{G_i, G_{im}\}} E_{g'}},$$

$$sh_{gj} = \frac{E_{gj}}{\sum_{g' \epsilon G_i} E_{g'j}}, \qquad sh_g = \frac{E_g}{\sum_{g' \epsilon G_i} E_{g'}},$$

$$sh_{mj} = \frac{\sum_{g \epsilon G_{im}} E_{gj}}{\sum_{m' \epsilon M_i, g' \epsilon \ G_{im}} E_{g'j}}, \qquad sh_m = \frac{\sum_{g \epsilon G_{im}} E_g}{\sum_{m' \epsilon M_i, g' \epsilon \ G_{im}} E_{g'}}.$$

E_g is national total expenditure on variety g.

For housing, we assume there is only one variety and it's available everywhere. For products with price data not from Nielsen, we assume all varieties within a high-level product category i have the same local price, as measured by the average price we use in our Laspeyres index for each product category.

Geary-Khamis PPP Index The Geary-Khamis index is a Paasche index that compares the local prices in a given CZ to nationwide average prices. The weights on the relative prices differences between the CZ and the nationwide average are equal the focal CZ's expenditure shares. This is the method used by the BEA to estimate local price indices. A desirable property of The Geary-

Khamis index is that preserves aggregation. Thus, the Geary-Khamis index is a weighted average of Geary-Khamis indices for each sub-component of consumption (e.g. housing or restaurants). It is measured as:

$$P_{j,k}^{\text{Geary-Khamis}} = \frac{\sum_{i \in I}(p_{i,j} \cdot \mathrm{q}_{i,j,k})}{\sum_{i \in I}(\pi_{i,k} \cdot \mathrm{q}_{i,j,k})} \tag{5}$$

where $\pi_{i,k} = \sum_{j \in J} \frac{p_{i,j} \cdot \mathrm{q}_{i,j,k}}{P_{j,k}^{\text{Geary-Khamis}} \cdot \sum_{j\prime \in J} \mathrm{q}_{i,j\prime,k}}$.

GEKS-Fischer PPP Index GEKS-Fisher index. A Fisher index is the Geometric mean of a Laspeyres and Paasche price index for a given pair of cities. The Fisher index is a second-order approximation for the true price index. However, the standard Fisher index is only defined for pairs of cities, and it is not transitive. This means the Fisher index between cities A and B, multiplied by the Fisher index between cities B and C does not equal the Fisher index between cities A and C. The GEKS-Fisher index uses these pairwise Fisher indices to estimate price indices that impose this transitivity. This is implicitly done by an OLS regression of pairwise log Fisher indices on the difference of CZ specific fixed effects for these CZ pairs. The CZ fixed effects are the GEKS-Fisher indices and thus impose transitivity. Instead of running this regression explicitly, the analytic formula below solves for the estimated regression fixed effects directly.

$$P_{j,k}^{\text{GEKS-Fischer}} = \left(\prod_{j\prime \in J} P_{j_0,j\prime,k}^{\text{Fischer}} P_{j\prime,j,k}^{\text{Fischer}} \right)^{\frac{1}{|J|}} \tag{6}$$

and (iii) $P_{j_1,j_2,k}^{\text{Fischer}} = \sqrt{P_{j_1,j_2,k}^{\text{Laspeyres}} / P_{j_2,j_1,k}^{\text{Laspeyres}}}$ and $P_{j_1,j_2,k}^{\text{Laspeyres}} = \sum_{i \in I} \left(\mathrm{s}_{i,j_1,k} \cdot \frac{\mathrm{p}_{i,j_2}}{\mathrm{p}_{i,j_1}} \right)$.

EASI Demand System. Based on Lewbel and Pendakur (2009), the implicit Marshallian budget shares are given by $\mathbf{s} = \sum_{r=0}^{5} \mathbf{b_r} y^r + \mathbf{B}\mathbf{p}y + \epsilon$, where, in the context of our work, $\mathbf{s}$ is a vector of expenditure shares across our focal spending categories; $\mathbf{p}$ is a vector of prices (normalized to 1 for our baseline commuting zone); $y = (x - \mathbf{p}'\mathbf{s})/(1 - \mathbf{p}'\mathbf{B}\mathbf{p}/2)$ is an implicit utility function; and x is our measure of consumption expenditures for each household. This EASI demand model allows expenditure shares to have flexible price effects through $\mathbf{B}$ and non-linear Engel curve shapes through $\mathbf{b_r}$ for r=0,1,2,3,4,5. We follow the authors in estimating this model using 3SLS to account for endogeneity that results from having $\mathbf{s}$ appear in y, specifically, by instrumenting y^r and $\mathbf{p}y$ by $\tilde{y}^r$ and $\mathbf{p}\tilde{y}$, where $y = x - \mathbf{p}'\mathbf{s}$ and $\tilde{y} = x - \mathbf{p}'\bar{\mathbf{s}}$.

To calculate an income-group-specific price index implied by the exact EASI demand model, we first predict $\tilde{\mathbf{s}}_\mathbf{h}$ for each household if it were to face average prices. These predicted shares as a function of real consumption expenditures are shown in Appendix Figure A4. Next, we mean-collapse to calculate $\tilde{\mathbf{s}}_\mathbf{j}$ at the commuting zone level separately by income group. Finally, for each commuting zone and each income group, we construct a Stone index by mean-collapsing prices across all categories using these shares, precisely, $\mathbf{p}_\mathbf{j}'\tilde{\mathbf{s}}_\mathbf{j} = \sum_i (p_{ij} \cdot \tilde{s}_{ij})$.

D Consumption in Physical Units in Nielsen Data

Here, we describe the Nielsen data used in Section 4.2. Since UPCs for a given product group can come in different units, we identify the most prevalent unit or "modal unit" within each product group. We seek to convert non-modal units to the modal unit for each product group: this procedure allows us to aggregate a quantity of UPCs consumed by each household for each product group,

since all UPCs within the same product group are measured in the same unit.

For each product group, we first convert ounce, pound, milliliter, liter, and quart to kilogram, assuming density of water (1,000 kg/m^3). When direct conversion is not possible (e.g., from count or square foot to kilogram), we assume the log of quantity has the same underlying distribution across different units within the product group being considered. We compute z scores for each unit-specific distribution and then equate z scores based the non-modal-unit distributions with z scores based on the modal-unit distribution. Finally, we convert all non-modal units to the modal unit within each product group. Specifically, for a given $q_{nonmodal}$, we solve for q_{modal} satisfying $\frac{q_{modal}-\mu_{modal}}{\sigma_{modal}} = \frac{q_{nonmodal}-\mu_{nonmodal}}{\sigma_{nonmodal}}$, where μ_{modal} and $\mu_{nonmodal}$ denote a given product group's mean quantity measured in modal unit and nonmodal unit, respectively, and σ_{modal} and $\sigma_{nonmodal}$ denote the corresponding standard deviations. We also truncate extreme values at the minimum and maximum quantities within the modal-unit distribution.

We combine the 116 product-group-level files that we have dealt with modal unit adjustment above. We sum-collapse modal-unit-adjusted UPC quantities by household $\times$ product group. We assign 0 to if a household did not buy any UPC for a given product group.

E Estimating Consumption by Education Group

Here, we describe in detail the data and the methodology used in Sections 5 to estimate consumption by commuting zone and education group.

We augment our data with the pooled 2012-2016 ACS data, which include 6,838,804 households. We assign each household a commuting zone. Since household income in our bank account data is post-tax and household income in the ACS data is pre-tax, we calculate household post-tax income in the ACS data using the NBER TAXSIM software. Specifically, for each household, we input into the software its pre-tax income and information on state, number of dependents, martial status, age of household head and spouse, and wages of household head and spouse (if exists). We always use joint filing for households with the spouse present and use single filing otherwise. We subtract state taxes, federal taxes, and social securities (these are outputs from the software) from household pre-tax income to obtain household post-tax income. To make households in the ACS data consistent with those in our bank account data, we drop households with missing post-tax income, households with post-tax income less than \$10,000, and households not belonging to the 443 commuting zones identified in our data. These restrictions together leave 5,302,154 households in the ACS data.

With these data in hand, we take the following steps:

Step 1: We define household types. We interact the following household characteristics to define types:

1. Age — based on mean age of household head and spouse (if exists):
   - Less than 30 years old — 446,250 (8.42%)
   - From 30 to less than 45 years old — 1,249,376 (23.56%)
   - From 45 to less than 65 years old — 1,647,023 (31.06%)
   - At least 65 years old — 1,959,505 (36.96%)
2. Gender — based on a composition of household head and spouse (if exists):
   - Household head is male OR both are male — 959,606 (18.10%)

- Household head is female OR both are female 1,486,558 (28.04%)
- One person is male and the other person is female 2,855,990 (53.86%)

3. Race — based on a composition of household head and spouse (if exists):
   - Household head is white OR both are white 4,190,909 (79.04%)
   - At least one person is nonwhite 1,111,245 (20.96%)

4. Hispanic Origin — based on a composition of household head and spouse (if exists):
   - At least one person has Hispanic origin 381,620 (7.20%)
   - None has Hispanic origin within the household 4,920,534 (92.80%)

5. Education — based on a composition of household head and spouse (if exists):
   - Both are ≥college OR household head is ≥college 1,455,299 (27.45%)
   - One is ≥college AND the other is <college 683,094 (12.88%)
   - Both are ≥highsch<college OR head is ≥highsch<college 2,527,382 (47.67%)
   - One is ≥highsch<college AND the other is <highsch 247,666 (4.67%)
   - Both are <highsch OR household head is <highsch 388,713 (7.33%)

6. Marital Status — based on a composition of household head and spouse (if exists):
   - Married 2,878,074 (54.28%)
   - Non-married 2,424,080 (45.72%)

7. Number of Children — based on whether the household head has at least one child:
   - At least one child within the household 2,084,155 (39.31%)
   - No children within the household 3,217,999 (60.69%)

Step 2: We assign each household in the ACS data an estimated expenditure value from our bank account data. In particular:

- For each commuting zone, we calculate income ventiles: that is, we identify $v = 1, 2, ..., 20$ for each commuting zone $j \in J$.
- We calculate expenditure-to-income ratios (R) for all households within each $j \times v$ bucket. At this stage, we have created a map from income ventile range within each commuting zone to a pool of observed expenditures-to-income ratios in our bank account data.
- For each household h in the ACS data, we identify a commuting zone $\times$ income ventile in our bank account data to which h belongs. We take a random draw of expenditure-to-income ratios, allowing repetition. Let us assume that the sampled value for a specific household is $\tilde{R_h}$. To calculate expenditure for this household, we multiply its post-tax income and the pooled ratio, i.e., $expenditure_h = income_h \times \tilde{R_h}$. This procedure allows us to go from household post-tax income in the ACS to its corresponding commuting zone $\times$ income ventile in our data, take a random draw of observed expenditure-to-income ratios, and then compute expenditure.
- Finally, to calculate consumption, we deflate this expenditure value by the corresponding income-group-specific price index of the commuting zone to which this household belongs.

Step 3: We calculate pre-tax income, post-tax income, consumption expenditure, and consumption estimates by skill level and commuting zone following the below steps:

- We define three skill levels based on the education level of a household head: (i) "high-skill" households in which the household head obtained a four-year college degree or higher; (ii) "middle-skill" households in which the household head finished high school but did not obtain a four-year college degree; and (iii) "low-skill" households in which the household head did not finish high school. The corresponding numbers of households by skill level are 1,882,956; 2,916,322; and 502,876.

- For each skill level $s \in S = \{\text{high}, \text{middle}, \text{low}\}$, we calculate commuting-zone-level value, evaluated at the nationwide skill-group-specific shares across household types that are the same for all commuting zones. In practice, we estimate

$$\log Y_{h,j(h),s(h)} = \delta_{Y_{j,s}} + 1_{h,t(h),s(h)} \times \beta_s + \epsilon_{h,j(h),t(h),s(h)}$$

  where $Y \in \{\text{pre-tax income}, \text{post-tax income}, \text{expenditure}, \text{consumption}\}$. For each household h, $j(h)$ denotes commuting zone; $s(h)$ denotes skill group; and $t(h)$ denotes household type. Finally, we calculate $exp(\widehat{\delta_{Y_{j,s}}} + \bar{1}_{t,s} \times \widehat{\beta_s})$, where $\bar{1}_{t,s}$ is a vector of nationwide-average shares across all household types for skill level s.

F Government Transfers

Our income data do not include housing subsidies, food stamp and TANF. Here we describe how we impute the value of these three types of government assistance, which we add to our measure of consumption expenditures in a robustness check.

First, for housing subsidies, we use the 2013 American Housing Survey (AHS) data. We restrict the household sample to subsidized renters having non-missing rents and positive income. We then construct a housing subsidy to income ratio and regress it on the interaction of region, household size, whether the household head has a spouse, and whether there is at least one child in the household.

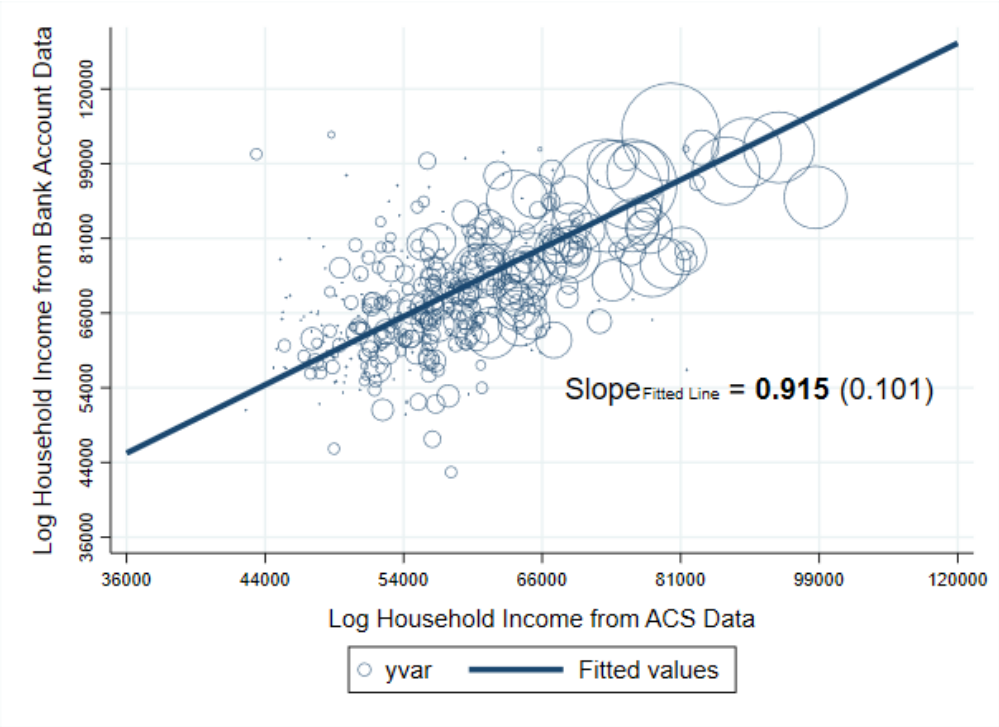
Second, we use the 2014 Survey of Income and Program Participation (SIPP) data, which contain information on dollar amounts of food stamp, TANF, and TGA each household member receives in a given month. We begin by combining all four quarterly survey datasets in 2014: observations are at the household-person-month level. Next, we collapse data at the household level by summing up values across all household members across all months. Then, we use the coefficient estimates from the AHS regression described above to predict housing subsidy for all households in SIPP. We define three government assistance categories: "housing subsidy", "food stamp", and "other public assistance", which consists of TANF and TGA.

Third, because government assistance in both AHS and SIPP data is likely to be under-reported by participating households, we perform the adjustment proposed by Meyer and Mittag (2019). They calculate numbers to scale up these three government assistance measures by income to federal poverty level, and we use their numbers.

We add up these three measures for each households, and divide the sum by household income. We then regress the government assistance to income ratio on the interaction of household size, presence of spouse, presence of children, income group indicator, and state. We then use the resulting coefficient estimates to predict government assistance for all households in 2012-2016

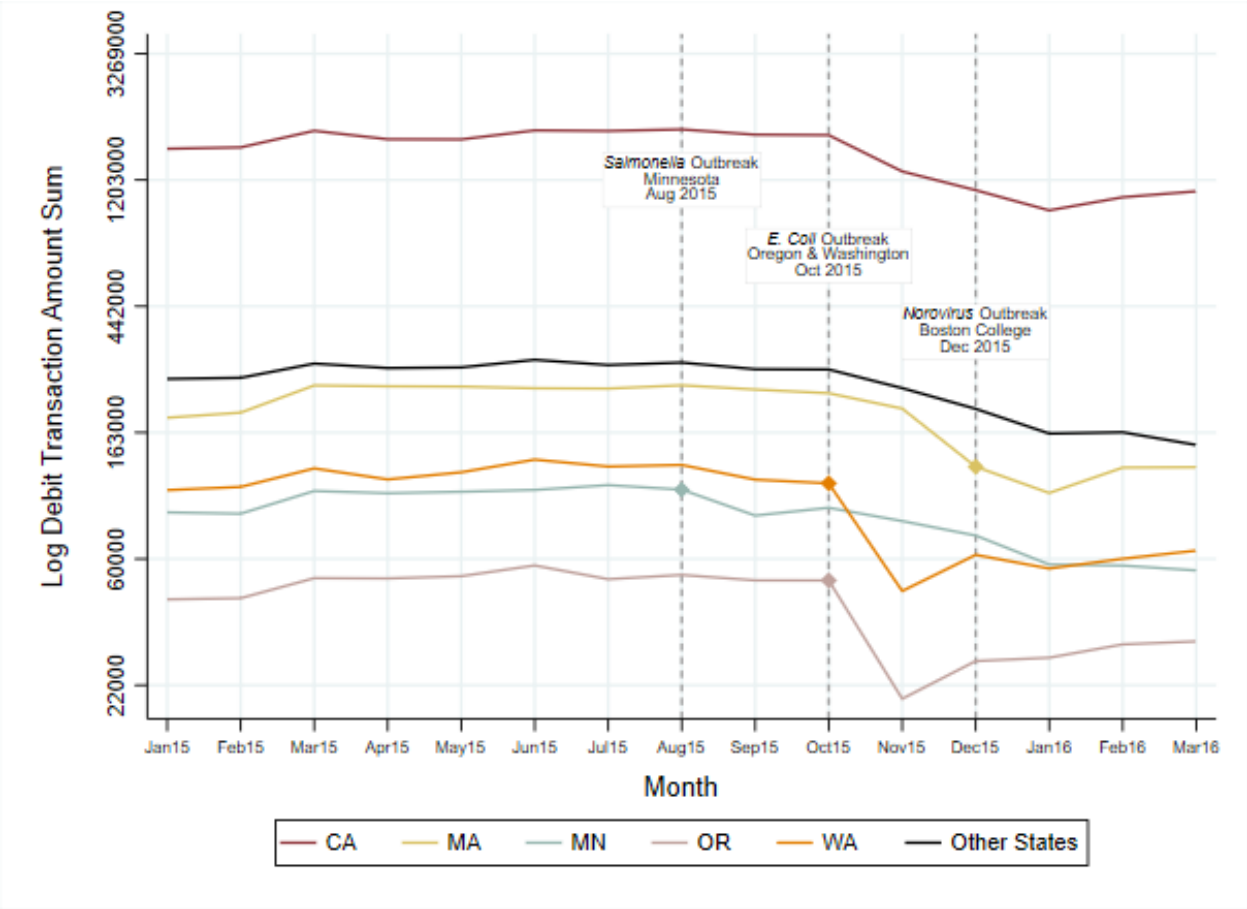
ACS data and then add this imputed measure to our measures of consumption expenditures and income.

Appendix Figure A1: **Mean Income by Commuting Zone: Our Data vs. ACS**



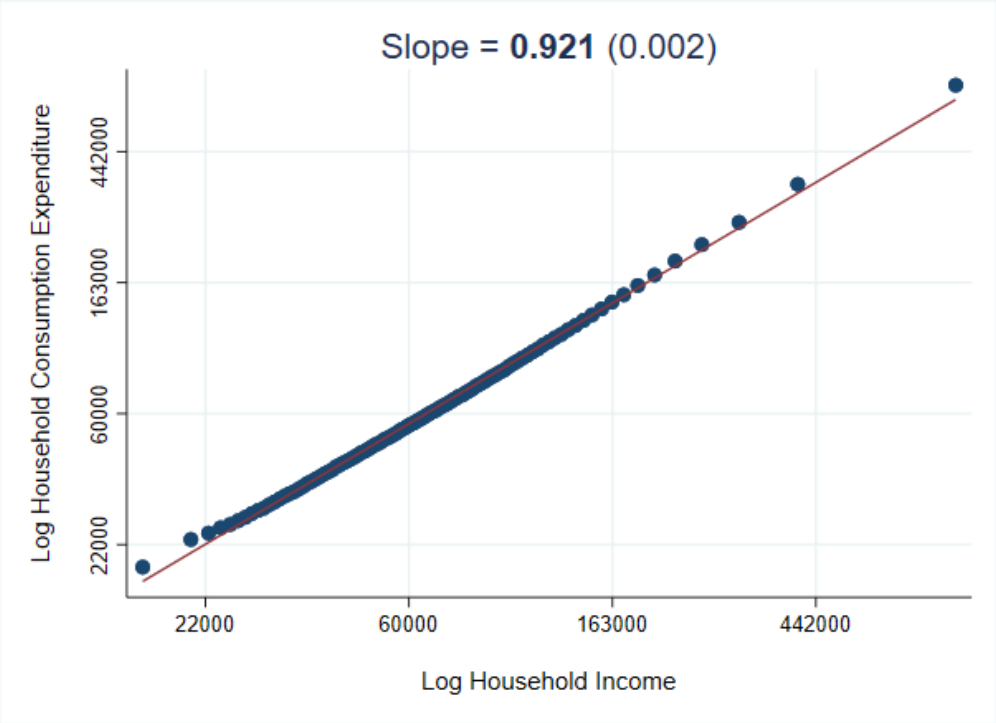
Notes: Observations are at the commuting zone level. ACS data are from 2012-2016. Household income in our data is post-tax. To obtain post-tax income in the ACS data, we subtract from household pre-tax income the income taxes calculated using the NBER TAXSIM software. We weight observations by their corresponding commuting zone population. Values on both axes are in a log scale, but we label actual values for easier interpretation.

Appendix Figure A2: **Episodes of Changes in Sales of Chipotle After Outbreaks**



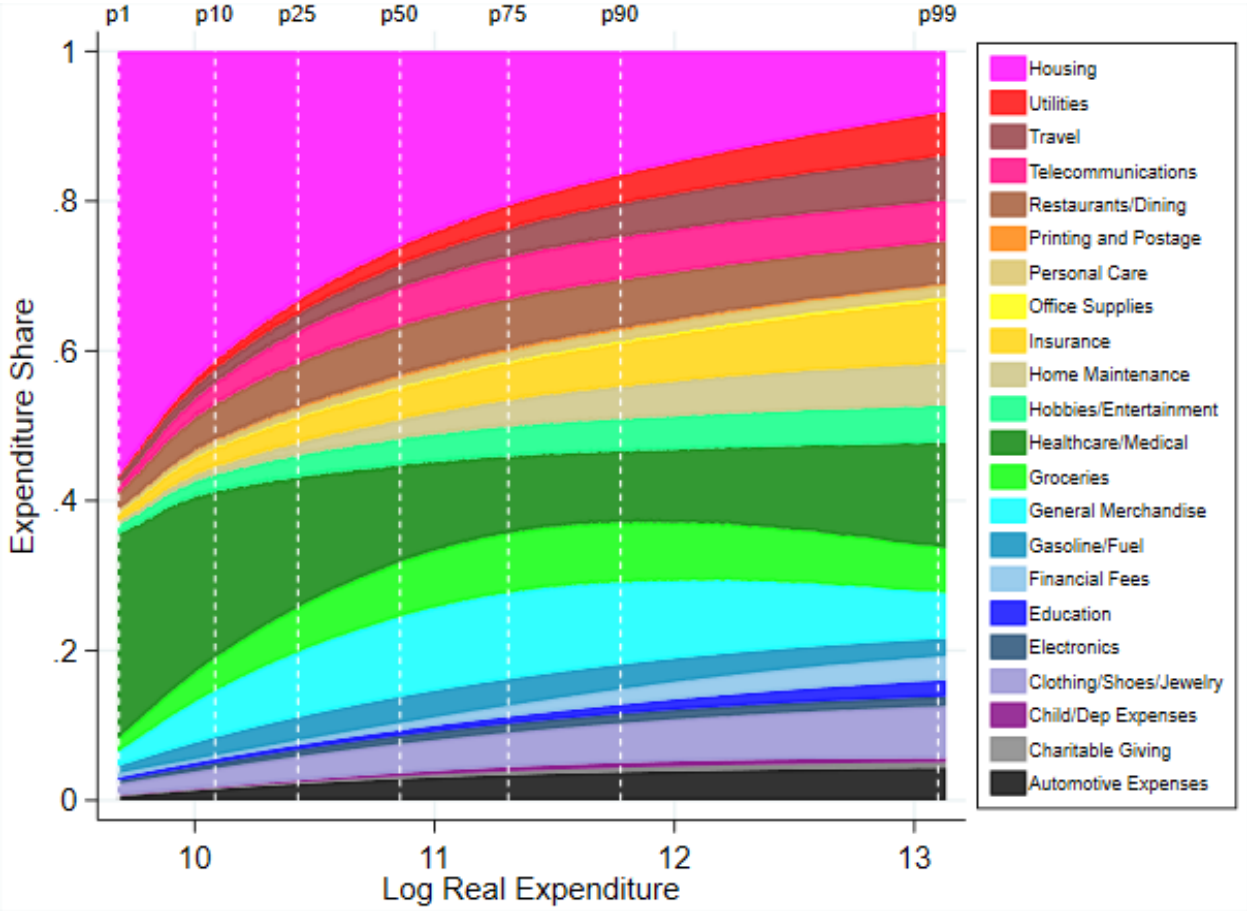
Notes: This figure shows changes in expenditures in Chipotle stores observed in our data after the *Salmonella* outbreak in Minnesota in August 2015; the *E. Coli* outbreak in Oregon and Washington in October 2015; and the Norovirus outbreak in Boston in December 2015. The dash lines indicate the months during which these outbreaks occurred.

Appendix Figure A3: **Consumption Expenditure vs. Income**



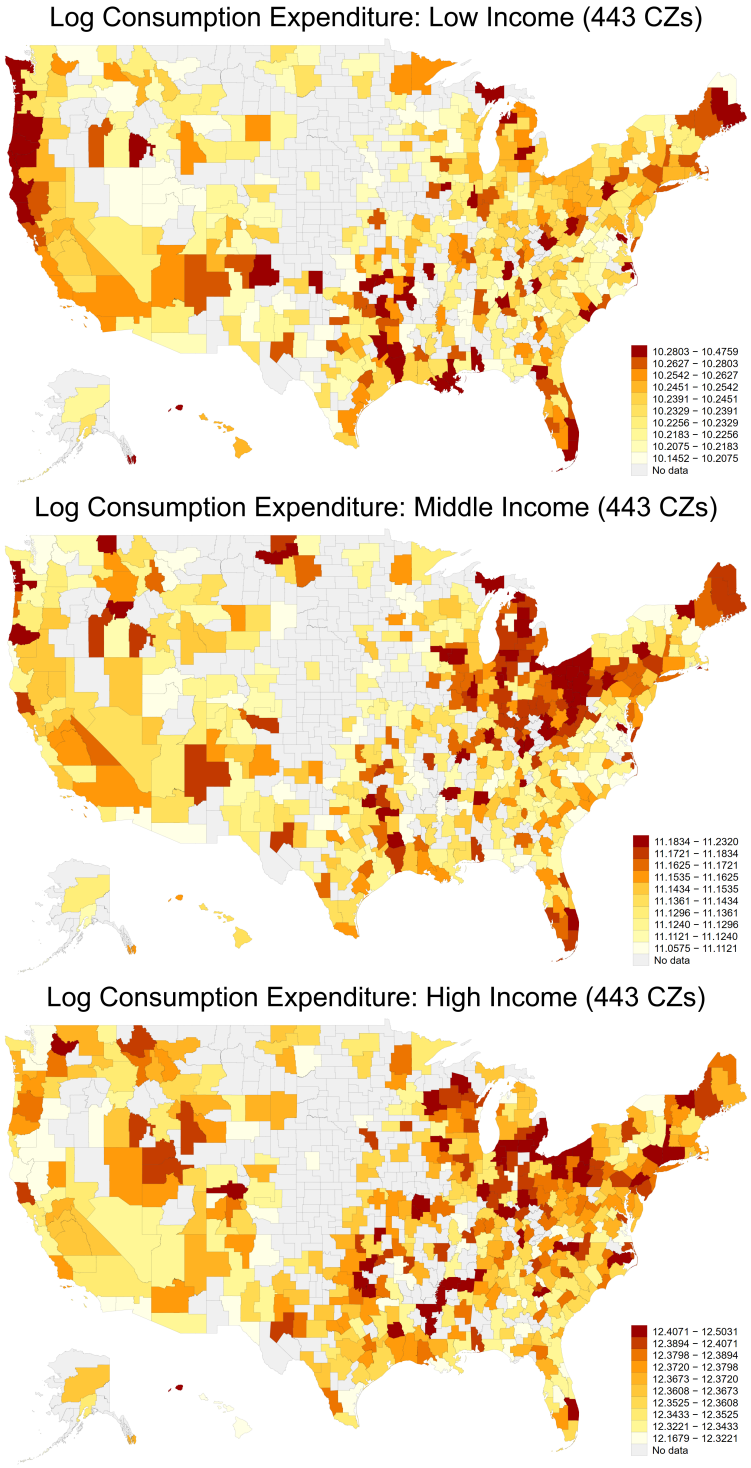
Notes: The sample includes all households in our sample and uses commuting zone weights. Values on both x-axis and y-axis are in measured in log scale, but we label actual values for easier interpretation. N = 3,000,518 households.

Appendix Figure A4: **EASI Demand Expenditure Shares as vs. Real Expenditure**



Notes: The expenditure shares across our focal spending categories displayed in the figure have been recovered from an exact EASI demand estimation, following a framework in Lewbel and Pendakur (2009). In this model, we allow each household's expenditure shares to be a function of polynomials of real expenditure (up to degree 5), local prices, and the interaction of real expenditure and local prices. We then predict shares for all households if they were to face average prices nationwide. The dash lines indicate conventional percentiles of real expenditure across all households in our data. We provide a discussion of EASI demand analysis in the Appendix.

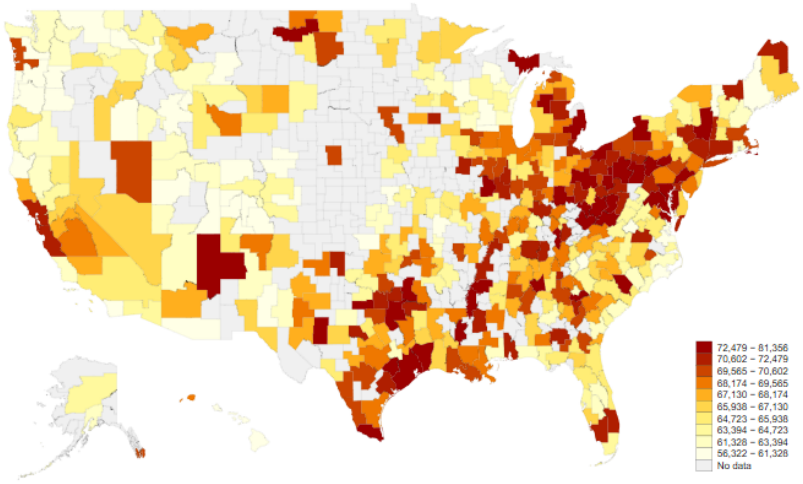
Appendix Figure A5: **Map of Consumption Expenditure, By Income Group**



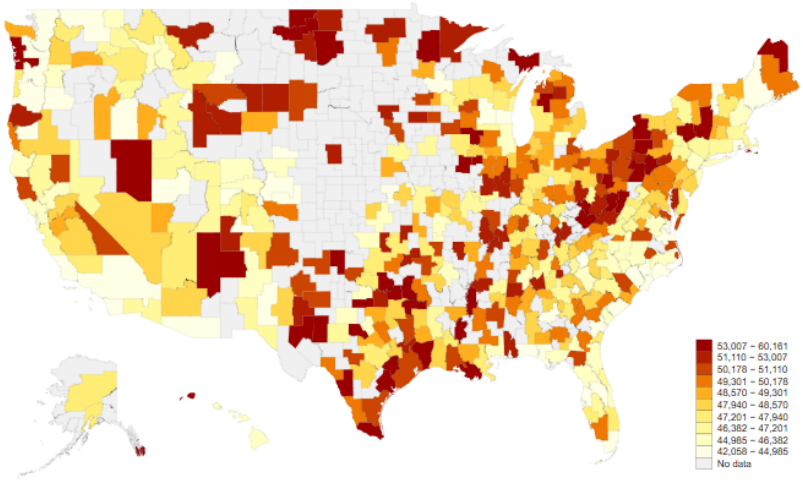
Notes: In this Figure, to limit the role of sample error, we report empirical-Bayes shrunken estimates.

Appendix Figure A6: **Map of Consumption, by Skill Level**

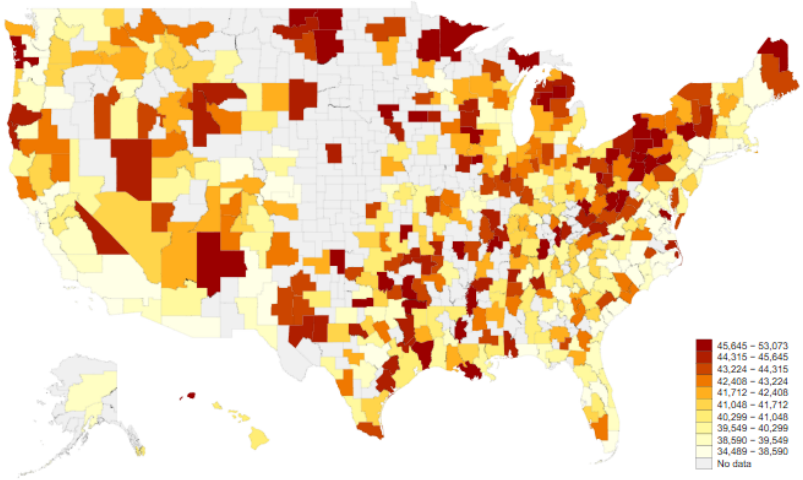
A. High Skill



B. Middle Skill



C. Low Skill



Notes: In this Figure, to limit the role of sample error, we report empirical-Bayes shrunken estimates.

Appendix Table A1: **Summary Statistics**

	All	Low Income	Middle Income	High Income
	(1)	(2)	(3)	(4)
Panel A. Raw Measures of Income and Expenditure				
Post-tax Income				
Mean	81,010.77	29,638.88	91,121.28	448,699.56
Median	52,955.79	29,495.21	81,021.33	288,099.91
Expenditure				
Mean	74,631.26	29,902.46	82,135.14	406,517.88
Median	47,750.03	27,652.95	71,200.55	251,517.69
Panel B. Adjusted Measures of Income and Expenditure				
Post-tax Income				
Mean	92,370.05	39,820.80	101,050.11	483,512.03
Median	62,390.10	39,280.70	90,181.38	305,713.72
Expenditure				
Mean	85,990.54	40,084.38	92,063.97	441,330.31
Median	56,898.91	37,182.11	79,799.09	267,485.66
Number of Commuting Zones	443	443	443	443
Number of Households	3,000,518	1,368,817	1,449,978	181,723

Notes: Panel A summarizes raw post-tax income and expenditure in our bank account data. Panel B summarizes adjusted post-tax income and adjusted expenditure, where we add imputed non out-of-pocket health spending and housing cost adjustments. See text for details.

Appendix Table A2: **High-Level Category Expenditure Shares**

	Expenditure Shares				Price
	Overall Income (1)	Low Income (2)	Middle Income (3)	High Income (4)	Standard Deviation (5)
Automotive Expenses	2.76%	2.09%	3.25%	4.18%	0.149
Charitable Giving	0.35%	0.21%	0.45%	0.72%	0
Child/Dependent Expenses	0.41%	0.25%	0.55%	0.64%	0.021
Clothing/Shoes/Jewellery	3.91%	3.13%	4.43%	6.29%	0.260
Electronics	1.02%	0.87%	1.14%	1.22%	0.029
Education	0.85%	0.68%	0.90%	1.85%	0
Financial Fees	1.12%	0.70%	1.36%	2.82%	0
Gasoline/Fuel	3.16%	2.97%	3.40%	2.67%	0.062
General Merchandise	9.45%	8.34%	10.73%	7.69%	0.219
Groceries	6.63%	5.66%	7.59%	6.65%	0.157
Healthcare/Medical	14.58%	19.03%	10.77%	9.43%	0.320
Hobbies/Entertainment	3.39%	2.82%	3.80%	4.66%	0.340
Home Maintenance/Improvement	2.65%	1.63%	3.38%	5.08%	0.195
Insurance	4.17%	3.22%	4.76%	7.29%	0
Office Supplies	0.23%	0.19%	0.25%	0.36%	0.032
Personal Care	1.14%	0.97%	1.27%	1.51%	0.195
Printing and Postage	0.28%	0.27%	0.30%	0.25%	0
Restaurants/Dining	6.02%	5.65%	6.39%	5.92%	0.253
Telecommunications	4.70%	3.92%	5.39%	5.45%	0.149
Travel	3.08%	2.27%	3.61%	5.54%	0
Utilities	2.53%	1.67%	3.10%	5.03%	0.133
Housing	27.57%	33.48%	23.16%	14.73%	0.296

Notes: Columns 1-4 report expenditure shares. Column 5 reports the standard deviation in price across commuting zones.

Appendix Table A3: **Expenditure Shares within Nielsen Product Groups**

	Low Income	Medium Income	High Income
Groceries			
Baby Food	0.21%	0.23%	0.25%
Baked Goods - Frozen	0.40%	0.36%	0.31%
Baking Mixes	0.42%	0.39%	0.35%
Baking Supplies	0.54%	0.55%	0.53%
Beer	1.37%	1.33%	1.21%
Bread, Baked Goods	3.96%	3.74%	3.51%
Breakfast Food	0.83%	0.96%	1.01%
Breakfast Food, Frozen	0.58%	0.59%	0.57%
Butter, Margarine	0.91%	0.84%	0.75%
Candy	2.77%	2.56%	2.45%
Carbonated Beverages	3.33%	2.96%	2.72%
Cereal	1.84%	1.88%	1.80%
Charcoal, Logs	0.11%	0.12%	0.11%
Cheese	3.19%	3.43%	3.50%
Coffee	1.86%	1.99%	2.18%
Condiments, Gravies, Sauces	1.36%	1.41%	1.39%
Cookies	1.20%	1.14%	1.08%
Cottage Cheese, Sour Cream	0.58%	0.59%	0.59%
Crackers	0.84%	0.87%	0.87%
Desserts, Fruits, Toppings	0.32%	0.34%	0.34%
Desserts, Gelatins, Syrup	0.46%	0.46%	0.42%
Detergents	1.30%	1.39%	1.43%
Disposable Diapers	0.26%	0.35%	0.35%
Dough Products	0.35%	0.35%	0.31%
Dressings, Salads, Prepared Foods	5.89%	6.18%	6.23%
Eggs	0.84%	0.79%	0.77%
Flour	0.15%	0.14%	0.14%
Fresh Meat	0.77%	0.71%	0.65%
Fresh Produce	5.35%	6.37%	7.40%
Fresheners, Deodorizers	0.45%	0.41%	0.38%
Fruit - Canned	0.41%	0.34%	0.28%
Fruit, Dried	0.33%	0.38%	0.43%
Gum	0.18%	0.20%	0.21%
Household Cleaners	0.60%	0.64%	0.69%
Household Supplies	0.69%	0.73%	0.74%
Ice	0.02%	0.02%	0.02%
Ice Cream	1.60%	1.46%	1.38%
Jams, Jellies, Spreads	0.66%	0.64%	0.61%
Juice, Drinks - Canned-Bottled	1.98%	2.07%	2.12%
Juice, Drinks - Frozen	0.09%	0.09%	0.08%
Laundry Supplies	0.68%	0.73%	0.74%
Liquor	1.02%	1.21%	1.47%
Milk	2.70%	2.58%	2.48%
Nuts	1.02%	1.19%	1.36%
Packaged Meats - Deli	3.69%	3.69%	3.46%
Packaged Milk, Modifiers	0.75%	0.66%	0.60%
Paper Products	3.33%	3.37%	3.33%
Pasta	0.37%	0.39%	0.40%
Pet Care	1.61%	1.52%	1.52%
Pet Food	4.65%	4.17%	3.78%
Pickles, Olives, Relish	0.37%	0.38%	0.37%
Pizza, Snacks - Frozen	1.14%	1.05%	0.99%
Prepared Food - Dry Mixes	0.97%	0.92%	0.82%
Prepared Food - Ready-to-Serve	1.05%	0.92%	0.87%
Prepared Foods - Frozen	3.09%	2.83%	2.67%
Puddings, Dessert - Dairy	0.07%	0.07%	0.07%
Salad Dressings, Mayo, Toppings	0.75%	0.73%	0.65%
Seafood, Canned	0.43%	0.40%	0.38%
Shortening, Oil	0.55%	0.54%	0.54%
Snacks	3.50%	3.67%	3.71%
Snacks, Spreads, Dips - Dairy	0.30%	0.35%	0.45%
Soap, Bath Additives	0.64%	0.72%	0.76%
Soft Drinks - Non-Carbonated	1.14%	1.13%	1.20%
Soup	1.10%	1.07%	1.02%
Spices, Seasoning, Extracts	0.52%	0.53%	0.54%
Sugar, Sweeteners	0.52%	0.44%	0.38%
Table Syrups, Molasses	0.14%	0.15%	0.14%

Tea	0.71%	0.74%	0.78%
Tobacco	3.17%	1.75%	1.11%
Unprepared Meat, Poultry, Seafood	5.70%	6.07%	6.35%
Vegetables - Canned	0.95%	0.94%	0.86%
Vegetables - Frozen	1.02%	1.01%	0.90%
Vegetables, Grains - Dried	0.36%	0.38%	0.39%
Wine	1.02%	1.45%	2.31%
Wrapping Materials, Bags	0.79%	0.83%	0.83%
Yeast	0.00%	0.00%	0.00%
Yogurt	1.16%	1.42%	1.60%
General Merchandise			
Automotive	6.00%	4.99%	4.09%
Batteries, Flashlights	11.51%	10.88%	10.18%
Books, Magazines	2.66%	2.10%	1.76%
Canning, Freezing Supplies	1.14%	1.07%	0.85%
Cookware	3.88%	3.77%	3.50%
Floral, Gardening	7.49%	8.45%	8.97%
Glassware, Tableware	4.18%	4.27%	4.41%
Hardware, Tools	5.90%	6.17%	5.94%
Housewares, Appliances	25.98%	26.60%	27.21%
Insecticides, Pesticides, Rodenticides	5.25%	4.42%	3.88%
Kitchen Gadgets	8.80%	9.66%	9.64%
Light Bulbs, Electric Goods	11.59%	11.85%	13.84%
Party Needs	0.24%	0.19%	0.18%
Photographic Supplies	3.01%	3.32%	3.44%
Seasonal	0.57%	0.51%	0.45%
Sewing Notions	0.31%	0.34%	0.33%
Shoe Care	0.23%	0.24%	0.22%
Soft Goods	1.01%	1.00%	1.00%
Toys, Sporting Goods	0.25%	0.18%	0.12%
Personal Care			
Cosmetics	3.49%	4.12%	4.60%
Cough and Cold Remedies	6.56%	6.89%	6.89%
Deodorant	2.00%	2.39%	2.52%
Diet Aids	0.72%	0.88%	0.91%
Ethnic Haba	0.12%	0.09%	0.07%
Feminine Hygiene Products	0.39%	0.37%	0.36%
First Aid	2.47%	2.38%	2.37%
Fragrances - Women	1.17%	1.23%	1.41%
Grooming Aids	1.20%	1.30%	1.31%
Hair Care	6.32%	7.16%	7.65%
Medications, Remedies, Health Aids	41.22%	36.09%	31.72%
Men's Toiletries	0.44%	0.52%	0.53%
Oral Hygiene	6.27%	6.69%	7.24%
Sanitary Protection	1.88%	1.99%	1.98%
Shaving Needs	2.20%	2.84%	3.41%
Skin Care Preparations	5.01%	6.34%	7.75%
Vitamins	18.52%	18.71%	19.27%

Notes: Columns 1-3 report expenditure shares by income group. Created using 61,903,872 purchases in Nielsen 2014 data.

Appendix Table A4: **Price Index vs. Rent**

	Low Income (1)	Middle Income (2)	High Income (3)
Log Monthly Rent	0.462***	0.360***	0.269***
	(0.011)	(0.010)	(0.013)
R^2	0.965	0.951	0.888

Notes: All columns use a log-log specification. We use commuting zone population as regression weight. Robust standard errors are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

Appendix Table A5: **Price Index Correlations**

	A1. Combined index, Laspeyres	A2. Combined index, Laspeyres (CEX weights)	A3. Combined index, CES	A4. Combined index, Nested CES σ=11.5	A5. Combined index, Nested CES σ=7	A6. Combined index, PPP, Geary-Khamis	A7. Combined index, PPP, GEKS-Fischer	B1. Combined index, Laspeyres	B2. Combined index, Laspeyres (CEX weights)	B3. Combined index, CES	B4. Combined index, Nested CES σ=11.5	B5. Combined index, Nested CES σ=7	B6. Combined index, PPP, Geary-Khamis	B7. Combined index, PPP, GEKS-Fischer	C. EASI demand based index	D. BEA index
A1. Combined index, Laspeyres	1															
A2. Combined index, Laspeyres (CEX weights)	.99	1														
A3. Combined index, CES	.97	.97	1													
A4. Combined index, Nested CES σ=11.5	.97	.97	.95	1												
A5. Combined index, Nested CES σ=7	.97	.97	.95	1.0	1											
A6. Combined index, PPP, Geary-Khamis	.99	.99	.97	.98	.98	1										
A7. Combined index, PPP, GEKS-Fischer	.99	.99	.97	.98	.98	.99	1									
B1. Combined index, Laspeyres	.99	.98	.96	.96	.96	.98	.98	1								
B2. Combined index, Laspeyres (CEX weights)	.98	.99	.96	.96	.96	.98	.98	.99	1							
B3. Combined index, CES	.96	.95	.99	.94	.94	.96	.96	.97	.96	1						
B4. Combined index, Nested CES σ=11.5	.97	.97	.95	.99	.99	.97	.97	.97	.97	.95	1					
B5. Combined index, Nested CES σ=7	.97	.97	.95	.99	.99	.97	.97	.97	.97	.95	1.0	1				
B6. Combined index, PPP, Geary-Khamis	.99	.98	.97	.97	.97	.99	.99	.99	.98	.97	.98	.98	1			
B7. Combined index, PPP, GEKS-Fischer	.99	.98	.97	.97	.97	.99	.99	.99	.99	.97	.98	.98	.99	1		
C. EASI demand based index	.99	.98	.96	.97	.96	.99	.99	.99	.98	.96	.97	.97	.99	.99	1	
D. BEA index	.93	.92	.90	.88	.88	.91	.92	.93	.92	.89	.88	.88	.91	.92	.92	1

Notes: N = 443. Correlation matrix of all alternative price indices. See text for detailed definition of each index.

Appendix Table A6: **Spatial Dispersion – Alternative Price Indexes**

	Overall Income				Low Income				Middle Income				High Income			
	mean	sd	min	max	mean	sd	min	max	mean	sd	min	max	mean	sd	min	max
A. Combined indexes: same prices for all households																
- Laspeyres	0.985	0.105	0.783	1.429	0.974	0.119	0.749	1.491	0.993	0.094	0.810	1.385	1.001	0.073	0.856	1.283
- Laspeyres (CEX weights)	0.987	0.110	0.777	1.454	0.978	0.117	0.757	1.482	0.995	0.113	0.783	1.470	1.012	0.106	0.801	1.441
- CES	0.960	0.094	0.794	1.358	0.946	0.112	0.751	1.445	0.965	0.091	0.778	1.310	0.977	0.068	0.824	1.239
- Nested CES (Feenstra, $\sigma = 7$)	0.971	0.071	0.837	1.253	0.974	0.092	0.815	1.405	0.968	0.072	0.820	1.267	0.969	0.053	0.803	1.174
- Nested CES (Feenstra, $\sigma = 11.5$)	0.972	0.071	0.839	1.259	0.975	0.093	0.816	1.408	0.970	0.073	0.822	1.273	0.970	0.053	0.825	1.176
- PPP Geary-Khamis	0.971	0.091	0.797	1.343	0.964	0.106	0.761	1.415	0.979	0.081	0.823	1.302	0.980	0.063	0.763	1.223
- PPP GEKS-Fischer	0.977	0.093	0.797	1.363	0.969	0.107	0.763	1.425	0.985	0.083	0.823	1.319	0.988	0.063	0.814	1.225
B. Combined indexes: income-group-specific prices																
- Laspeyres	0.985	0.106	0.761	1.401	0.978	0.106	0.737	1.357	0.979	0.087	0.795	1.287	0.958	0.061	0.422	1.136
- Laspeyres (CEX weights)	0.987	0.112	0.754	1.425	0.983	0.105	0.746	1.370	0.975	0.104	0.762	1.348	0.935	0.084	0.273	1.161
- CES	0.960	0.097	0.775	1.341	0.955	0.101	0.739	1.335	0.956	0.085	0.771	1.253	0.946	0.062	0.544	1.129
- Nested CES (Feenstra, $\sigma = 7$)	0.969	0.072	0.817	1.236	0.986	0.085	0.795	1.311	0.954	0.071	0.775	1.211	0.955	0.057	0.542	1.133
- Nested CES (Feenstra, $\sigma = 11.5$)	0.971	0.073	0.819	1.241	0.987	0.085	0.797	1.314	0.956	0.072	0.778	1.217	0.957	0.057	0.546	1.136
- PPP Geary-Khamis	0.971	0.092	0.777	1.323	0.967	0.099	0.736	1.326	0.964	0.084	0.504	1.244	0.953	0.076	0.350	1.154
- PPP GEKS-Fischer	0.978	0.094	0.776	1.341	0.972	0.098	0.747	1.317	0.971	0.083	0.656	1.255	0.950	0.070	0.365	1.132
EASI demand index	0.977	0.083	0.839	1.328	0.981	0.100	0.810	1.431	0.977	0.078	0.838	1.312	0.942	0.060	0.553	1.110
BEA price index	1.024	0.076	0.874	1.359	1.024	0.076	0.874	1.359	1.024	0.076	0.874	1.359	1.024	0.076	0.874	1.359

Notes: N = 443. Summary statistics for all alternative price indices across CZs. See text for detailed definition of each index.

Appendix Table A7: **Variety Effect Decomposition**

Dependent variable:	Log Nested CES Price Index				
Effect decomposition:	price only	variety only		price and variety	
Elasticity parameter (σ):	–	11.5	7	11.5	7
	(1)	(2)	(3)	(4)	(5)
A. Automotive Expenses					
Log Median CZ Income	0.379***	-0.008***	-0.014***	0.371***	0.365***
	(0.088)	(0.001)	(0.001)	(0.088)	(0.088)
B. Child/Dependent Expenses					
Log Median CZ Income	0.068***	-0.008***	-0.013***	0.060***	0.054**
	(0.017)	(0.001)	(0.001)	(0.017)	(0.016)
C. Clothing/Shoes/Jewelry					
Log Median CZ Income	0.232**	-0.006***	-0.011***	0.226**	0.222**
	(0.076)	(0.001)	(0.001)	(0.076)	(0.076)
D. Electronics					
Log Median CZ Income	0.003	-0.004***	-0.006***	-0.001	-0.003
	(0.011)	(0.000)	(0.001)	(0.011)	(0.011)
E. Gasoline/Fuel					
Log Median CZ Income	0.165***	-0.005***	-0.009***	0.160***	0.156***
	(0.036)	(0.001)	(0.001)	(0.036)	(0.036)
F. General Merchandise					
Log Median CZ Income	0.134***	0.015	0.026	0.149***	0.161***
	(0.030)	(0.009)	(0.015)	(0.033)	(0.036)
G. Groceries					
Log Median CZ Income	0.075***	-0.000	-0.000	0.075***	0.075***
	(0.014)	(0.001)	(0.002)	(0.014)	(0.014)
H. Healthcare/Medical					
Log Median CZ Income	0.139***	-0.010***	-0.018***	0.129***	0.121**
	(0.037)	(0.001)	(0.002)	(0.037)	(0.037)
I. Hobbies/Entertainment					
Log Median CZ Income	0.428***	-0.019***	-0.033***	0.409***	0.394***
	(0.082)	(0.002)	(0.003)	(0.081)	(0.081)
J. Home Maintenance/Improvement					
Log Median CZ Income	0.089	-0.045***	-0.079***	0.044	0.011
	(0.131)	(0.004)	(0.008)	(0.130)	(0.130)
K. Office Supplies					
Log Median CZ Income	0.045**	-0.005***	-0.008***	0.041**	0.037*
	(0.015)	(0.001)	(0.001)	(0.015)	(0.015)
L. Personal Care					
Log Median CZ Income	0.072***	-0.013***	-0.023***	0.059***	0.050***
	(0.015)	(0.004)	(0.006)	(0.014)	(0.014)
M. Restaurants/Dining					
Log Median CZ Income	0.067	-0.029***	-0.051***	0.038	0.016
	(0.051)	(0.003)	(0.005)	(0.050)	(0.049)
N. Telecommunications					
Log Median CZ Income	0.112	-0.004***	-0.006***	0.108	0.106
	(0.065)	(0.000)	(0.001)	(0.065)	(0.065)
O. Utilities					
Log Median CZ Income	0.317***	-0.022***	-0.038***	0.295***	0.278***
	(0.077)	(0.003)	(0.005)	(0.077)	(0.077)
Overall					
Log Median CZ Income	0.404***	-0.007***	-0.013***	0.397***	0.391***
	(0.035)	(0.001)	(0.002)	(0.034)	(0.034)

Notes: This table decomposes the impact of the price effect versus the supply of variety in the nested CES price

index across each sub-component of the price index. Standard errors are clustered at the commuting zone level. See text for details.

Appendix Table A8: **Consumption vs. Price Index – Robustness**

Index	Low-Income	Middle-Income	High-Income
A1. Laspeyres	–0.900***	–0.978***	–1.016***
	(0.009)	(0.019)	(0.037)
A2. Laspeyres (CEX)	–0.896***	–0.981***	–1.020***
	(0.009)	(0.016)	(0.026)
A3. CES	–0.904***	–0.982***	–1.008***
	(0.010)	(0.020)	(0.039)
A4. Nested CES (11.5)	–0.873***	–0.964***	–0.976***
	(0.011)	(0.023)	(0.049)
A5. Nested CES (7)	–0.872***	–0.964***	–0.973***
	(0.011)	(0.024)	(0.049)
A6. PPP, Geary-Khamis	–0.890***	–0.973***	–1.012***
	(0.010)	(0.022)	(0.043)
A7. PPP, GEKS-Fischer	–0.890***	–0.974***	–1.014***
	(0.010)	(0.022)	(0.042)
B1. Laspeyres	–0.887***	–0.985***	–0.945***
	(0.012)	(0.022)	(0.056)
B2. Laspeyres (CEX)	–0.881***	–0.986***	–0.968***
	(0.012)	(0.019)	(0.042)
B3. CES	–0.893***	–0.986***	–0.932***
	(0.012)	(0.023)	(0.056)
B4. Nested CES (11.5)	–0.858***	–0.973***	–0.891***
	(0.013)	(0.026)	(0.062)
B5. Nested CES (7)	–0.857***	–0.973***	–0.885***
	(0.013)	(0.026)	(0.062)
B6. PPP, Geary-Khamis	–0.880***	–0.980***	–0.966***
	(0.012)	(0.022)	(0.049)
B7. PPP, GEKS-Fischer	–0.878***	–0.980***	–0.955***
	(0.012)	(0.023)	(0.051)
C. EASI Demand System	–0.886***	–0.978***	–0.984***
	(0.011)	(0.022)	(0.055)
D. BEA Price Parities	–0.846***	–0.973***	–1.024***
	(0.015)	(0.023)	(0.034)

Notes: Both Consumption and the price index are in logs. This table reports the bi-variate regression coefficient of a regression of log consumption on log price index, across all alternative price index definitions. All price indices labeled with "A" use uniform prices across income groups within CZ. Price indices labeled with "B" use income group specific prices within each CZ. N = 443. See text for additional details.

Appendix Table A9: **Consumption Against Price Index – Robustness**

	College Graduates (1)	High School Graduates (2)	High School Dropouts (3)
A. Baseline			
Log price index	–0.032	–0.237***	–0.391***
	(0.047)	(0.026)	(0.032)
B. Income Includes Imputed Food Stamps, TANF, and Housing Assistance			
Log price index	–0.040	–0.242***	–0.389***
	(0.048)	(0.027)	(0.034)
C. Consumption Does Not Include Any Imputation			
Log price index	–0.103**	–0.302***	–0.462***
	(0.051)	(0.034)	(0.045)

Notes: N = 443. Panel A reports our main results of the relationship between consumption an local price index by skill group. Panel B reports these estimates when imputed government transfer program expenditure is added into expenditure. Panel C reports this estimate using our "raw" expenditure data that does not adjust housing and healthcare to accurately track total expenditure on healthcare and the rental equivalent spending on housing. See text for details.

Appendix Table A10: **Consumption vs Price Index, Population, College Share**

	A1. Laspeyres	A2. Laspeyres (CEX WGT)	A3. CES	A4. Nested CES (σ=11.5)	A5. Nested CES (σ=7)	A6. PPP, Geary-Khamis	A7. PPP, GEKS-Fischer	B1. Laspeyres
Log index	-0.263***	-0.341***	-0.276***	-0.132	-0.128	-0.218**	-0.208**	-0.082
	(0.075)	(0.065)	(0.083)	(0.094)	(0.096)	(0.085)	(0.082)	(0.085)
Log index × middle-skill	0.039	0.068	0.041	0.042	0.042	0.045	0.044	0.040
	(0.086)	(0.074)	(0.097)	(0.110)	(0.112)	(0.099)	(0.095)	(0.101)
Log index × low-skill	-0.022	0.039	-0.030	-0.002	-0.002	-0.022	-0.020	-0.041
	(0.087)	(0.074)	(0.099)	(0.112)	(0.114)	(0.100)	(0.096)	(0.104)
Log city size	0.027***	0.025***	0.023***	0.028***	0.028***	0.028***	0.028***	0.026***
	(0.008)	(0.007)	(0.008)	(0.007)	(0.007)	(0.007)	(0.008)	(0.008)
Log population × middle-skill	-0.026***	-0.024***	-0.026***	-0.025***	-0.025***	-0.026***	-0.026***	-0.025***
	(0.009)	(0.008)	(0.009)	(0.008)	(0.008)	(0.008)	(0.008)	(0.009)
Log population × low-skill	-0.040***	-0.038***	-0.040***	-0.040***	-0.040***	-0.040***	-0.040***	-0.038***
	(0.009)	(0.008)	(0.009)	(0.008)	(0.008)	(0.009)	(0.009)	(0.009)
Log college share	0.021	0.016	0.043	0.039	0.039	0.036	0.031	0.007
	(0.037)	(0.032)	(0.036)	(0.037)	(0.037)	(0.040)	(0.038)	(0.038)
Log college share × middle-skill	-0.036	-0.029	-0.034	-0.033	-0.033	-0.035	-0.035	-0.037
	(0.042)	(0.037)	(0.043)	(0.044)	(0.044)	(0.048)	(0.045)	(0.043)
Log college share × low-skill	-0.028	-0.019	-0.025	-0.030	-0.030	-0.027	-0.028	-0.019
	(0.042)	(0.037)	(0.042)	(0.043)	(0.043)	(0.048)	(0.045)	(0.045)
Middle-skill	-0.063	-0.067	-0.050	-0.075	-0.075	-0.058	-0.061	-0.080
	(0.144)	(0.132)	(0.142)	(0.138)	(0.137)	(0.146)	(0.144)	(0.146)
Low-skill	-0.036	-0.041	-0.031	-0.046	-0.046	-0.031	-0.034	-0.059
	(0.145)	(0.132)	(0.141)	(0.139)	(0.139)	(0.148)	(0.146)	(0.148)

	B2. Laspeyres (CEX WGT)	B3. CES	B4. Nested CES (σ=11.5)	B5. Nested CES (σ=7)	B6. PPP, Geary-Khamis	B7. PPP, GEKS-Fischer	C. BEA Price Parities	D. EASI Demand System
Log index	-0.139*	-0.146	0.029	0.032	-0.106	-0.069	-0.143	-0.116
	(0.073)	(0.089)	(0.098)	(0.100)	(0.093)	(0.091)	(0.100)	(0.103)
Log index × middle-skill	0.048	0.044	0.063	0.064	0.057	0.055	0.043	0.158
	(0.087)	(0.107)	(0.118)	(0.121)	(0.111)	(0.108)	(0.117)	(0.121)
Log index × low-skill	-0.004	-0.044	0.011	0.012	-0.019	-0.022	-0.044	0.118
	(0.088)	(0.111)	(0.123)	(0.125)	(0.114)	(0.111)	(0.121)	(0.124)
Log city size	0.024***	0.023***	0.027***	0.028***	0.026***	0.026***	0.026***	0.032***
	(0.007)	(0.008)	(0.007)	(0.007)	(0.008)	(0.008)	(0.008)	(0.008)
Log population × middle-skill	-0.024***	-0.026***	-0.024***	-0.024***	-0.025***	-0.025***	-0.026***	-0.024***
	(0.008)	(0.009)	(0.008)	(0.008)	(0.009)	(0.009)	(0.009)	(0.008)
Log population × low-skill	-0.036***	-0.038***	-0.039***	-0.039***	-0.038***	-0.037***	-0.039***	-0.038***
	(0.008)	(0.009)	(0.008)	(0.008)	(0.009)	(0.009)	(0.009)	(0.009)
Log college share	0.001	0.035	0.022	0.023	0.028	0.018	0.034	0.029
	(0.033)	(0.037)	(0.036)	(0.036)	(0.041)	(0.039)	(0.040)	(0.039)
Log college share × middle-skill	-0.029	-0.032	-0.030	-0.030	-0.035	-0.035	-0.035	-0.035
	(0.038)	(0.045)	(0.044)	(0.044)	(0.050)	(0.046)	(0.046)	(0.045)
Log college share × low-skill	-0.009	-0.016	-0.020	-0.020	-0.022	-0.020	-0.023	-0.024
	(0.038)	(0.045)	(0.043)	(0.043)	(0.050)	(0.047)	(0.047)	(0.046)
Middle-skill	-0.087	-0.060	-0.084	-0.084	-0.077	-0.080	-0.070	-0.086
	(0.133)	(0.144)	(0.134)	(0.133)	(0.148)	(0.146)	(0.151)	(0.141)
Low-skill	-0.068	-0.050	-0.060	-0.059	-0.060	-0.062	-0.056	-0.067
	(0.132)	(0.144)	(0.135)	(0.135)	(0.150)	(0.148)	(0.154)	(0.144)

Notes: Mutivariate consumption regressions on price index, log CZ population, and CZ college share interacted with skill.