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

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

July 2021

Abstract

There are large differences in mean income across US cities, but little is known about the levels 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 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 all debit and credit card transactions, check and ACH payments, and cash withdraws 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 expensive commuting zone enjoy a level of consumption that is about half that of low income residents in the most affordable commuting zone.

In the second part of the analysis, we endogenize income and estimate the standard of living that low-skill and high-skill households can expect in each US commuting zone, once we account for geographical variation both in cost of living and also in 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. 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: High school drop outs moving from the most to the least affordable commuting zone would experience a 23.5% decline in consumption.

We thank Matthias Hoelzlein, Erik Hurst, Matt Turner, and seminar participants at Berkeley, Cornell, Harris, Irvine, Mannheim, NBER Summer Institute, OSUS, Stanford, Toronto for useful comments. Abhisit Jiranaphawiboon provided outstanding research assistance.

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 some 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 large, 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 to 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 most all debit and credit card transactions, check and ACH payments, and cash withdraws 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 both in cost of living and also in 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 important advantages. Our consumption data is comprehensive and includes virtually all purchases conducted by individuals in our sample, and unlike other consumption data, it is not

[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) that uses the CEX to study consumption of the low-income in rich versus poor states.

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

self-reported. It matches well both the measure of aggregate 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 has 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 goods and services purchased by credit card, debit card and ACH, we only observe the value but not the type of purchases when the purchases are paid for by cash or check.

To measure local prices, we build consumer price indexes that vary by commuting zone and by 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 price 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 Nielsen, ACCRA, and the CEX.

The price indices point to large differences in cost of living across commuting zones. The overall cost of living faced by low-income households (post-tax income <\$50,000) in the most expensive city—San Jose, CA—is 44% higher than in the median commuting zone, Cleveland, and 95% higher than the most affordable commuting zone–Presque Isle, ME. 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 experienced by most Americans is between -20% and +20% of the median commuting zone, there are a handful of very expensive cities in the right tail, where cost of living is much higher than the median.[3]

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.

[3]For low-income households, there are 20 commuting zones with cost of living that is more than 20% higher than the median, and 10 where cost of living is more than 30% higher than the median.

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 Diego, CA; San Francisco, CA; and San Jose, CA, with consumption levels between 25% and 28% lower than the median CZ. At the other extreme of the spectrum, examples of commuting zones with high consumption of low-income households are Elizabeth City, NC; Johnstown, PA; and Huntington, WV, with consumption levels in real terms 19-22% higher than the median CZ. Thus, low-income families who live in the most affordable commuting zone in the US enjoy a level of market-based consumption measured in real terms that is about 70% higher that of families with the same income who live in the least affordable commuting zone.

To validate our findings, we replicate the analysis using model-free evidence. We use direct and transparent measures of consumption quantities from Nielsen. The data contains information on the quantities purchased of specific products measured in physical units, unlike our bank data that measures expenditures. For example, we measure kilograms of beer, the number of light bulbs, or pounds of nuts purchased in a year by each Nielsen consumer. There are 823,507 grocery products, divided in 116 product groups. Consistent with our findings on overall consumption, we find that for 79 product group out of 116, Nielsen consumers in expensive commuting zones buy fewer physical units, holding income and demographics fixed. Using our bank account data, we extend this analysis to seven non-grocery product groups by using the count of transactions as a proxy for quantity within homogeneous products. For example, we measure the number of gas purchases, number of trips to the movies or number of Netflix or Cable payments. Consistent with the findings for Nielsen, we find a negative relationship between consumption of non grocery items and local cost of living.

If differences in cost of living across cities are permanent, and if consumers expect to be in their current city for a long time, variation in local price indexes is a pure income effect. We estimate that the elasticity of overall market consumption with respect to local prices is -0.910 (0.009) and -1.020 (0.037) for low- and high-income households, respectively. The permanent income hypothesis suggests that if utility is locally homothetic, the elasticity should be -1. We cannot reject that this is the case for high-income and middle-income households, but we can reject that this is the case for low-income households.

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.[4] Consistent with this possibility, low-income households in expensive commuting zones are found to have a higher incidence of financial distress than low-income households in affordable commuting zones: they have a significantly higher probability of negative saving and a higher probability of paying overdraft fees.

[4] Alternatively, it could reflect the possibility that more low-income households in expensive cities are expecting larger future income gains than high-income households, and therefore do not need to cut consumption by the same amount; or alternatively, more low-income households in expensive cities are expecting to move to affordable cities in the future than high-income households. It could also reflect the existence of government-provided minimum subsistence levels.

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, 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, once we account both for geographical variation in cost of living and also for geographical variation in 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. For each skill group and commuting zone, we compare pre-tax income to post-tax income to consumption.

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. Across all commuting zones, the standard deviation of consumption is half the standard deviation in 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 almost 40%. However, since pre-tax income is so high, even after accounting for taxes and cost of living, college graduates in San Francisco, New York, and Boston retain a high level of standard of living and remain in the top third of the distribution of all cities in terms of market consumption. Among large cities, Houston is the one where consumption is the highest, since it offers good expected income and moderate prices.

Overall, we find that for college graduates, there is no significant 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 -0.039 (0.052). This suggests that college graduates living in an average city with high costs of living enjoy an expected standard of living similar to college graduates with the same observable characteristics living in the average city with low cost of living.

For less skilled households, the picture that emerges is 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 -0.277 (0.030), 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 9.3% decline in the standard of living. Moving from the commuting zone with the lowest cost of living index (Presque Isle) to the commuting zone with the highest index would imply a decline in the standard of living by 16.6%.

The negative relationship between consumption and cost of living is significantly steeper for high school drop outs. The slope is -0.454 (0.037), suggesting that for this group standard of living in expensive commuting zones is much lower than in cheaper commuting zones. For households in this group, moving from Cleveland to San Jose implies a 15.2% decline in the standard of living. Moving from Presque Isle to San Jose implies a 27.2% decline in the standard of living. Among large cities, the ones that offer the highest standard of living to high school drop outs is Buffalo, NY.

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. We find that the gap 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 to explain the difference in our findings between high- and lower-skill households it would need to be the case that lower-skill households have significantly stronger preferences for amenities found in expensive cities. Since expensive cities also tend to be larger and have a higher share of college graduates, 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. 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

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 consumer has multiple bank or credit card accounts within the same bank, then these accounts are linked and we observe them as linked. On the other hand, if a consumer has credit cards or accounts at many banks, we do not observe the individual transaction on these cards. For these multi-banked consumers, 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. Our data provider's proprietary algorithm already attempted to restrict our data to accounts that active, and thus likely to be their main bank accounts.[5] In addition to the active account restriction, we require households to have at least \$10,000 of annual income and \$1,000 of annual expenditure. If these restrictions leave us with households' main bank accounts, we expect to be missing only 7.8% of the average household's

[5] While we do not have the exact code run by our data provider to filter out inactive accounts, attempts to screen out accounts that are rarely active in terms of deposits and withdrawals/debits.

income and expenditures.[6]

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.

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, 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 level 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. More importantly, we cannot observe some government transfers. Our data includes 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 likely to receive in-kind government transfers. Both these considerations 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, for the population that we study the omission of Food Stamps, TANF and housing assistance does not appear to be an important source of bias. In Section 5, we analyze the sensitivity of our estimates to including in our expenditure measure the imputed value of Food Stamps, TANF and housing assistance and find that our main empirical results remain robust.[7]

As a first step in understanding 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 2014 American Community Survey. To make the ACS data comparable to our data, we drop incomes below \$10,000 and run it through TaxSim to calculate post-tax income for each household. 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).[8] 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%

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

[7] 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.

[8] In our empirical analysis, we seek to measure consumption *conditional on income*. The fact that low-income households are slightly under-represented in our data does not in itself invalidate our estimates.

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 our measure of income in each commuting zone to 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.[9] The figure shows a tight relationship between median income measured in the ACS and in our data, with a slope of 0.914 (0.100). Since the ACS is a 5% random sample of the population, measurement error would lead to a slope of less than one, even if both data sources were representative.

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).[10]

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.[11] [12]

Panel A shows that our bank data appear to closely match NIPA. The bank data estimates average household spending at $74,631. NIPA estimates $77,533. The CEX estimates $53,495. The

[9] CZ's 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. As before, we use Taxsim to turn ACS pre-tax income into post-tax income and drop households with income below $10,000.

[10] 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.

[11] 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 share of health care costs that are out-of-pocket, as measured by the Center for Medicare Medical Services (2021). There is an additional difference between the two data in the way NIPA measures housing expenditures. NIPA measures housing expenditure for homeowners by the imputed rental value of their home. Our bank data measures "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 section. It will turn out that both of these definitions lead to similar estimates of housing spending.

[12] There are some small differences in the CEX housing and health spending measures from 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 does not include spending on paying down mortgage principle 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.

bank data mean is only 4% lower than NIPA's mean, while the CEX mean is 31% lower. The CEX a survey-based dataset which is known significantly under report spending (Sabelhaus and Groen, 2000; Aguiar and Bils, 2015).[13] We will return to healthcare in the next 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 Panel A of Figure 3 we plot expenditure shares by category. For this comparison, we restrict all three datasets to types of expenditure that are measured consistently in each. Since healthcare and housing in NIPA are not defined in the same way as in the other datasets, as previously discussed, in Panel A, we drop these categories from all three datasets and calculate household average expenditure shares among the remaining categories in all three datasets.[14] We also aggregate some categories with definitions that don't quite line up across datasets into "other goods" and "other services". This leaves us with 11 consumption categories in each dataset.

Panel A shows that our bank data line up extremely 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 0.51, and the correlation between the CEX and NIPA is 0.63. Bee et al. (2012) show that there is substantial variation in the under reporting rate of consumption across types of spending in the CEX creating poorly measured expenditure shares. We will return to healthcare and housing in the next section, where we discuss panel B.

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 (2016). On average, our data appears to match closely 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 expected to include 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 &

[13]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 these bank-of-the-envelope adjustment suggests were in the right ballpark.

[14]For our bank data, we drop transactions that do not have an expenditure category assigned to them, such as ATM withdrawals and checks. We will returns to these unclassified transactions in the next section.

Fitch, Chipotle, Costco, Dunkin' Donuts, Home Depot, Kroger, Macy's, McDonald's, Nordstrom, Starbucks, Walmart, Whole Foods. Overall, our expenditure data appears to track sales trends well.

As an 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.

One limitation is that our data can identify the exact category only for goods purchased by credit card, debit card and ACH. When the purchases are paid for by cash or check, we observe the value but not the type of purchases. Thus, expenditure on housing cannot be identified with accuracy 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, we do not observe exactly which transactions are those related to housing. Our data puts them into a category called "Unclassified". We come back to this point in Subsection 3.1, where we explain how we deal with this issue.

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.

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 0.922 (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).

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, their exact location 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.[15] 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.

[15] Our data includes the majority of the ten largest banks in the US. These banks' locations are skewed to larger, urban areas.

We classify households into three income groups: 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. 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. More details on the construction of the sample are in Appendix A.

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).[16]

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:[17]

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

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

[16]The BLS provides CPI-U for 19 metro areas, but the index can be used only to compare prices changes over time, not price levels.

[17]While our index is inspired by 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. This prevents us from using the exact Laspeyres index definition used by the BLS. 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.

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.[18] 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 B.

Measuring Prices of Consumption Items. To measure prices $p_{i,j}$, we combine data from Nielsen, ACCRA, and the ACS. First, 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. For these categories, we observe prices at the twelve-digit barcode level. For example, we observe the price paid in each commuting zone for "Cinnamon Toast Crunch - Cereal 2.00 oz" (UPC 001600014154) or "Budweiser Lager Beer - 30pk/12 fl oz Cans" (UPC 018200110306). To measure the local price for each product category, 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 CZ dummies as our estimate of local price for each product category. In this approach, good quality is held constant since we are comparing the price that consumers in different commuting zones pay for of the same twelve-digit barcode level product. Since Nielsen prices are reported before taxes, for goods categories that in a given state are subject to sale tax, we add the relevant sales tax.[19]

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, 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

[18] 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.

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

Balancing". ACCRA reports the price of "Tire Balancing" in each city for a specific type of tires.[20] 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 it is 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 paid by household 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 CZ 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 price does 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 6.5%, 11.0% and 18.0%, respectively.

Measuring Expenditure Shares. Our data classifies non-housing expenditure in 21 high-level consumption categories. These are listed in Appendix Table A2. We make three adjustments.

First, three categories among the 21 in our data are very broad: Grocery, General Merchandise, and Personal Care. To improve precision, we use data from the 2014 Nielsen Consumer Panel Survey to obtain a more refined product definition nested within each of these three categories. For example, Nielsen identifies 17 subcategories within the Personal Care category: Cosmetics; Deodorant, Vitamins, etc. The subcategories are shown in Appendix Table A3.

Second, expenditure on housing cannot be identified with accuracy in our data because many consumers pay their rent with checks and mortgages with bill-pay transfers to banks. As discussed in the Data section above, our data puts them into a category called "Unclassified". We must use an alternative data source to estimate how much each income group spends on housing. To impute housing expenditure, we follow the same methodology used by the BLS, which uses data from the CEX and ACS. For renters, this simply means measuring how much rent is paid on average by each income group. However, for households that own their home, their expenditures on housing need not equal the cost of purchasing one year of housing services in their given CZ. For example, homeowners that do not have a mortgage pay housing costs (property taxes) that are likely much less than cost of purchasing the year of housing services they consume. In contrast, owners paying down a mortgage are likely spending more than the cost of a single year of housing services. Our goal is the measure the expenditure spent on the single year of housing services only. Following

[20] 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–i.e. the number of observations per item—is more limited than the BLS sample size.

the BLS's methodology, we impute the expenditure on housing by homeowners as the rental value of their home if they were to rent it out. To measure average household expenditure on housing, the BLS uses contract rents for renters and rent equivalent for owners. The BLS estimates rental equivalents using data from the CEX and ACS and we follow their approach. (Poole et al. (2005); Bureau of Labor Statistics (2007)).[21] Homeowners in the CEX are asked the rental value of their home if they were to rent it out. This the imputed rental value of their home. We measure imputed rents 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.[22]

Third, we augment out of pocket health 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. We add this extra health care spending to expenditure. Details are in Appendix B.

After these three adjustments, we return to the comparison of our data to the NIPA and CEX. In Panel B of Figure 2 we compare our adjusted bank data average aggregate expenditure against NIPA and the CEX. The NIPA data is in its "raw" format, since we have adjusted our bank data to make the health and housing definitions consistent with those used by NIPA.[23] The adjusted

[21] Bee et al. (2012) show that the CEX Interview survey accurately tracks housing expenditure, when validated against NIPA.

[22] The expenditures we observe in our transactions data include the *actual* expenditures on housing. For homeowners, this amount may be more or less than the imputed rental value on their home. 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 (Aguiar and Bils, 2015).

To measure homeowner households' actual expenditures on housing, we again use the CEX. The CEX reports the amount homeowners spend on mortgage payments, property taxes, and down payments on houses. For renters, their actual housing payments are simply equal to their housing rental payments. We sum the actual payments from renters and owners within each of our three income groups and divide by the sample size to create average actual housing expenditure by income group. We then subtract out the actual spending on housing from our unclassified spending and add back our estimated cost of a year of housing services. Specifically, we subtract our estimate of actual housing expenditure by income group from the part of expenditure in our bank data that is "unclassified". This removes the true outlays on housing by the average household in each income group. Next, we add back the estimated housing expenditure for a year of housing services, as measured by actual rents for renters and imputed rents for owners. We also add the net difference between imputed and actual housing spending to income to reflect the income generation of the housing asset.

After subtracting out the actual housing expenditures from the unclassified transactions in our bank data, we are left with a residual of unclassified spending that is not housing. We assume this unclassified non-housing spending is split across the non-housing categories proportionally the non-housing expenditure shares observed among the transactions that were classified to expenditure categories. This assumes that probability that a non-housing transaction is able to be classified into a consumption category does not vary among the non-housing expenditure categories. Figure 2 already shows that the expenditure shares among the classified transactions lines up very closely to NIPA, validating this assumption. If certain non-housing transaction categories were much more likely to be paid for by cash, check, a merchant whose transaction category was unknown, then those categories' expenditure shares among the transactions that could be classified should be much lower than the NIPA aggregate expenditure shares.

[23] The difference between the NIPA average spending in Figures 2 and 2 is that Figure 2 subtracts out estimate health spending paid by the government and employers.

average household expenditure is \$90,323 in the bank data, which is about 2.5% less than the raw NIPA estimate of \$92,779. 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, substantially below NIPA.[24]

Figure 3b reports average spending by category in the CEX, NIPA, and bank data. The CEX is consistently lower spending than NIPA and our bank data. The correlations of spending across categories is the highest between NIPA and our data at 0.98.[25]

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 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. 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 (Lewbel and Pendakur, 2009) (a generalization of the AIDS demand system (Deaton and Muellbauer, 1980)), 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 B.

The Törnqvist index is a geometric average, where the price of each good is weighted by the

[24]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: 53,495-6149+10895+2382+6284=66,907.

[25]To get a sense of how spatial variation in prices contribute to spatial variation of our overall-income index, column 5 in Appendix Table A2 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 (0.29). Gasoline/Fuel (0.06) and Electronics (0.03) have much lower geographical variation.

average of its local share and the share in the reference city – thus it is "chained".

The CES index assumes consumers within an income group's underlying preferences to have constant elasticity of substitution across all product categories. That elasticity is implicitly inferred by 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 21 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 corrections for variety. To measure local variety we use the number of unique UPC codes sold in each CZ as observed in the Nielsen RMS (store sales) data. For product categories not covered by Nielsen, we use the number of unique merchants that we observed transacted at within each CZ in our bank data. This index requires that we calibrate the substitution elasticity between varieties within a product group. We explore two choices of the elasticity parameter that have been found in the literature: σ: 7 (Montgomery and Rossi, 1999) and 11.5 (Broda and Weinstein, 2010).

In addition, we estimate an approximate Exact Affine Stone Index (EASI) implicit Marshallian demand system as developed by Lewbel and Pendakur (2009). Lewbel and Pendakur (2009) shows that the EASI demand system is a generalization of the popular AIDS demand system (Deaton and Muellbauer, 1980). Like the AIDS demand system, EASI expenditure shares are linear in parameters. However, unlike AIDS which imposes linear Engel curves, EASI demand Engel curves can have any shape over real expenditures. We follow Lewbel and Pendakur (2009)'s methods to estimate an approximate EASI demand model, derive the price index implied by the demand model estimates and use this as our income group-specific price indices.

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 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).

The second PPP index we estimate is the 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.

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, 5 commuting zones around the median, and the 15 least expensive commuting zones based on our baseline indexes. Throughout the paper, we label commuting zones using the name of its largest city, instead of the official commuting zone name. The most expensive commuting zones are San Jose, CA; San Francisco, CA; and San Diego, CA where the low-income price index is 1.440, 1.434, and 1.370, respectively. This implies that prices faced by low-income residents of these cities are 37% to 44% 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 Natchez, MS; Gallup NM; and Presque Isle, ME with price indexes equal to 0.767, 0.746, and 0.738, respectively.

The geographical price differences revealed by our price index differences are economically large. The overall cost of living in San Jose is estimated to be 95% and 52% higher than the cost in Presque Isle for low- and high-income households, respectively.

Figure 7 displays the spatial dispersion of the our price indexes across all 443 commuting zones. The standard deviation equals 0.110 and 0.072 for the low- and high-income group, respectively, indicating that the cost of living for low-income households exhibits higher spatial variation than for high-income households. This 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. The 75-25 and 90-10 differences are 0.124 and 0.238 for low-income and 0.082 and 0.165 for high-income households, respectively.

The figure also 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 Americans face 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 20 commuting zones with cost of living that is more than 20% above the median and 10 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 share of consumption and smaller geographical variation in prices—like 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 0.267 (0.013) and 0.428 (0.011) for high- and low-income households, respectively (Appendix Table A4). Recall that the share of housing in the indexes of high- and low-income households is 0.144 and 0.294, 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 nonhousing 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. For example, the cost of a haircut or a slice of pizza are higher in San Jose than in Cleveland, holding quality constant, because retail space and labor are more expensive in San Jose. That said, local price indices can be well-approximated by using data only on local housing costs, weighted appropriately. The R-squared of the regressions of our price indices on housing rent are all above 0.88.

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. This is shown in Appendix Figure A4. In this figure, we use the version of our index that is not income specific, since the BEA index does not vary by income group.

Alternative Indexes and Product Variety. Our findings are 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 the 120 pairwise combinations formed by 16 different indexes, the associated correlation coefficients have a mean of 0.933 and a standard deviation of 0.064. We will return to these indices to assess the robustness of our analysis below.

Appendix Table A7 quantifies the spatial dispersion of all our alternative price indexes and it compares it to our baseline index. For most 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) generally look quite similar to indices that hold prices fixed across income groups (Panel A).

As discussed above, our baseline index does not allow for differences across commuting zones in the variety of products that are locally available. This can potentially matter because using data on on grocery products from Nielsen, Handbury (2019) and Handbury and Weinstein (2015) has 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 replicate this finding with our variety-corrected nested CES index. However, Appendix Table A6 shows that this negative correlation does not hold for the overall price index because many other types of products, such as housing, gas, utilities, clothing, automotive, and healthcare do not exhibit this inverse relationship, even when adjusting for variety. See the Appendix for the full details of this analysis.

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. All of the price indexes for all commuting zones are available for download from our Online Appendix.

4 Geographical Differences in Consumption by Income Group

Consumption expenditures are unevenly distributed over space and the 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. We seek to compare the standard of living experienced by the residents of expensive and affordable commuting zones, *holding nominal income fixed* (Subsection 4.1). We also 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). In the next section, we study geographical differences in consumption if we allow household income to vary across commuting zones as a function of schooling and demographics.

We reiterate that in both this and the next section 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 $\delta_{j,k}$ 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.[26] We run this regression separately by income group and we condition on household income to control for possible income differences within income groups. For example, low-income households in expensive cities may have a higher nominal income than low-income households in affordable cities. We report estimates where consumption is evaluated at post-tax adjusted 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.[27] 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 \$53,288; \$53,064; and \$52,198. Other examples of commuting zones that offer high consumption of low-income households are Mobile, AL and Traverse City, MI. Three cities with the lowest consumption are San Diego, CA; San Francisco; and San Jose, CA. The corresponding values are \$32,607; \$31,457; and \$31,457. Other examples of commuting zones with low consumption are Honolulu, HI (\$33,066); Washington, DC (\$34,475); Seattle, WA (\$34,815); Los Angeles, CA (\$34,003); and New York, NY (\$33,649).

The geographical differences in standard of living are economically large. For low-income households, a comparison between the top 3 commuting zones in the top group and the bottom 3 commuting zones in the bottom group indicates that households in the top group enjoy a level of consumption that is about 70% higher than households with the same income 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

[26] Adjusted post-tax income means that we have included the the imputed net income from housing for homeowners to make it comparable to the expenditure measure we are using.

[27] We report empirical-Bayes shrunken estimates to improve precision. 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.

with the highest consumption for this group are Huntington, WV; Toledo, OH; and Johnstown, PA. They have level of consumption measured in real terms equal to \$298,341; \$297,887; and \$297,162, respectively. Three cities with the lowest consumption for high-income households are San Diego, CA (\$203,070); San Jose, CA (\$201,532); and Honolulu, HI (\$199,125). Other cities in this category include New York, NY (\$213,284); Boston, MA (\$217,363); Seattle, WA (\$210,708); and San Francisco, CA (\$210,708). 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. By contrast, cities with the highest consumption for high-income families are less overlapped with those for low-income families.

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 -0.910 (0.009), -0.980 (0.019), and -1.020 (0.037) for low-, middle-, and high-income households, respectively, suggesting that the consumption of high- and middle-income consumers is more sensitive to local prices than the consumption of low-income consumers.

Note that 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).

There are a number of possible explanations for the higher elasticity of low-income households. First, low-income households are closer to a minimum subsistence level, and small consumption cuts may cost more in terms of utility. This would also imply less 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.[28] These explanations are of course not mutually exclusive.

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 income in more expensive commuting zones. We now use Nielsen 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 Nielsen 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 Nielsen 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. 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).[29] 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 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.[30]

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. For example, for the product group "Beer", we report the average elasticity across all types of beer. Table 3 shows some selected examples. The first row reports results for the consumption of carbonated beverages measured in kilograms per year. The entries indicates that the elasticity of consumption of carbonated beverages with respect to the cost of living index is -0.833 (0.086) and -1.254 (0.168) 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. The top panel of 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

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

[29] 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 C for details

[30] In the Nielsen 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.

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 72% are negative for low-, middle-, and high-income households.[31] Second, the effects appear to be more negative for high income households than low income households, consistent with what we observed for overall consumption. The median values for the low-, middle-, and high-income groups are -0.068, -0.109, and -0.314. Third, the elaticities for all three groups are much smaller than the elasticities estimated for overall consumption above. This likely reflect the fact that most of the consumption items in the Nielsen 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.[32]

Nielsen 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 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 Nielsen, 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 level 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.

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

[32]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 0.073, 0.101, and 0.189.

4.3 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 exceed 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.[33]

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 0.077 (0.011).[34]

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 (where bankruptcy fees are identified as transactions that contain the words "bankrupt"). 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 other 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.

[33] 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.

[34] 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.

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?

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 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. We use 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 of selection on observables.

We do so by estimating expected household income in each commuting zone for each household type, where household types are defined by 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 then map our previous estimates of consumption by income level into estimates of consumption by skill level.[35] For each commuting zone and skill level, we predict mean consumption holding constant the other observable characteristics of the household.[36] We provide more details in Appendix D.

Our 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. A violation

[35] Since ACS income data is pre-tax, we first use Taxsim to transform pre-tax income to household post-tax income, so that it is consistent with the income definition in our bank account data. We then bin our data into 20 income ventiles by commuting zone, and match households in the ACS to income ventile using income bounds and commuting zone. We then take a random draw of expenditures-to-income ratios from our data given income ventile and commuting zone and multiply the drawn ratio with household post-tax income in the ACS to obtain consumption expenditures. This procedure ensures that for a given commuting zone all households in the ACS data are assigned one of our constructed income ventiles for that commuting zone. We use our price index to deflate consumption expenditures for a given income ventile and commuting zone and obtain consumption, like we did in the previous section.

[36] Specifically, we run a household-level regression of consumption (or consumption expenditures, post-tax income, and pre-tax income) on commuting zone indicators and 664 indicators for household types. For each skill group, we then predict mean consumption evaluated at nationwide weighted-average fractions across types. We then aggregate these 664 groups into three nationally representative skill groups based on education.

of this assumption would occur, for example, if households 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 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 three columns report the difference between consumption and pre-tax income in absolute change, percentage change, and percentile change.

Table 5 is for households where the head has a college degree or more. The first three rows show that the Washington, DC; San Jose, CA; and San Francisco, CA are the three CZs where high-skill households have the highest mean pre-tax nominal income: \$127,517; \$125,209; and \$121,903, 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 Washington DC, San Jose, San Francisco, and New York, the corresponding values are \$77,136, \$73,894, \$74,369, and \$74,864. 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 (column 9). In terms of pre-tax income, these four cities are at the 99th or 100th percentile, while in terms of consumption Washington, San Jose, San Francisco, and New York drop to 92nd, 78th, 81st, and 84th, respectively (column 6).

Thus, despite some of the highest costs of living in the US, Washington, 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 may explain why these regions have attracted so many college graduates over the past three decades.)

Boston consumption level is at the 82nd percentile. The Los Angeles and San Diego commuting zones experience larger drops in relative standings. In terms of pre-tax income, these two cities are all at the 96th and 95th percentile, while in terms of consumption percentiles they drop by -61 and -88 to the 35th and 7th percentiles, respectively. Other examples of cities with large negative percentile changes are Portland (-80), Seattle (-78), Minneapolis (-75), Sacramento (-71), Denver (-70), and Providence (-69). By contrast, Cincinnati (+4), Cleveland (+8), Pittsburgh (+15), and Buffalo (+30) 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 (column 8) because local income is relatively high, cost of living is moderate, and there are no state taxes. Philadelphia and Pittsburgh follow closely.[37]

A feature of the table worth noting is that spatial variation in consumption is much smaller 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 Atlanta, GA and Indianapolis, IN. Similarly, consumption in New York is found to be similar to Detroit, MI and Nashville, TN. Overall, across all 443 commuting zones, the standard deviation in pre-tax income is 9,908, while the standard deviation in consumption is only 5,114. 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 to this group. For example, while in terms of pre-tax income, Washington, DC; San Jose, CA; and San Francisco, CA remain at the top, in terms of consumption they drop to the bottom half of the distribution. New York, Boston and Chicago are at the 17th, 27th, and 18th 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 Pittsburgh.[38] Across all 443 commuting zones, the standard deviation in pre-tax income is 5,343, while the standard deviation in consumption is only 3,411, 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, Seattle and Los Angeles is in the bottom 10

[37] The three commuting zones that offer the highest standard of living among all commuting zones in the sample are Edgartown, MA (87,176); Gallup, NM (86,594); and Huntington, WV (85,254).

[38] The three commuting zones that offer the highest standard of living for this group among all commuting zones in the sample are Gallup, NM (68,252); Summersville, WV (63,643); and Sidney, MT (63,422).

percent of the distribution. In these cities 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 half 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 and Pittsburgh, thanks to low prices, as it was the case for high school graduates.[39] Across all 443 commuting zones, the standard deviation in pre-tax income is 3,719, while the standard deviation in consumption is only 3,432.

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 pre-tax income, post-tax income, expenditures, and consumption as a function of the local cost of living index. The figure includes all 443 commuting zones. We note that income, consumption, and local prices are all simultaneously determined. Thus, 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 slope is above 1 for the high-skill group, and below 1 for the middle- and low-skill groups. In particular, the slope is 1.047 (0.061), 0.871 (0.052), and 0.685 (0.067) 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 more than 10.5%—and a less than proportional increase in expected pre-tax income for the low skill group—specifically: an increase of only 6.9%.

We also observe a positive relationship for post-tax income and for consumption expenditures, indicating that all three groups tend to spend more in cities with higher cost of living. 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 slightly positive and not statistically significant at conventional levels: -0.039 (0.052). 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

[39] The three commuting zones that offer the highest standard of living for this group among all commuting zones in the sample are Summersville, WV (61,471); Gallup, NM (60,765); and Presque Isle, ME (59,261).

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 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 -0.277 (0.030), implying that a middle-skill household moving from the median city (Cleveland) to the most expensive city (San Jose) would experience a decline in standard of living by 9.3%. Moving from the city with the lowest cost of living index (Presque Isle) to the city with the highest index (San Jose) would imply a decline in the standard of living by 16.6%.

The negative relationship between consumption and cost of living is significantly steeper for low-skill households, suggesting that for this group standard of living in expensive commuting zones is much lower than in cheaper commuting zones. The slope is -0.454 (0.037), almost double (in absolute value) the one for middle-skill households, implying vast geographical differences in consumption. Moving from Cleveland to San Jose implies a 15.2% decline in the standard of living. Moving from Presque Isle to San Jose implies a 27.2% 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.

In this respect, one limitation of our income data is that it 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 estimated relationships for less skill households. To assess the magnitude of the problem, we analyze the sensitivity of our estimates to including in our income measure 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 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 E for more details. Appendix Table A8 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 -0.039 to -0.046 for high skill, from -0.277 to -0.279 for middle skill, and from -0.454 to -0.447 for low skill. One reason for the robustness of our estimates is that the dollar amounts involved are not large for households with income above $10,000. More fundamentally, 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.

Finally, we turn to within commuting zone inequality. There have been growing concerns in expensive cities about 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. 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 tends to increase with local prices. For example, consumption of college graduates in San Francisco, San Jose, and New York is 1.80, 1.76, and 1.72 times higher than consumption of high school drop-outs. The corresponding ratios in the three cheapest commuting zones, Natchez, MS; Gallup, NM; and Presque Isle, ME are 1.49, 1.43, and 1.34.

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 middle panel shows the difference in mean consumption between high- and low-skill households. The bottom panel shows the difference in mean consumption between middle- and low-skill households. The slopes are 0.237 (0.025), 0.414 (0.029), and 0.177 (0.018), 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.

5.3 Correlation with City Size and College Share and Robustness

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 a modest upward relationship between consumption and city size 0.005 (0.003), indicating that large cities offer slightly higher consumption than small cities. 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 (**?** and Moretti (2013)). More novel is the relationship between consumption and college share. For high-skill households, there is a modest positive relationship between consumption and college share. For middle- and low-skill households, there is a significant negative significant relationship between consumption and city size indicating that residents in cities with many college graduates enjoy standard of living that is lower than residents in cities with fewer college graduates.[40]

Since cities' prices, population, and college shares are all positively correlated, we investigate a multi-variate 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 does not reflect causal estimates, but rather equilibrium relationships.

Estimates in Table 8 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 he 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 particular, 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 city size range from 0.024 to 0.052 for the high-skill group, differential effects from -0.026 to -0.024 for the middle-skill group and differential effects from -0.041 to -0.035 for the low-skill group. Overall, we conclude that our results appear generally robust to the choice of price index.

[40] 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.

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. 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—labelled "Baseline"—indicate that the difference in consumption is smaller than the corresponding difference in pre-tax income.[41] In particular, the pre-tax income difference between college graduates and high school graduates is 0.518, while the corresponding difference in consumption is only 0.384, or 26 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 0.861 and 0.653. Namely, the consumption gap is 24 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 0.48 and 0.81, respectively. However, the consumption gaps would remain basically the same at 0.39 and 0.65, respectively. Interestingly, geographic sorting of households across CZs seems to have no impact on nationwide consumption inequality between skill groups.[42]

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 0.43 and 0.75 and the consumption gaps to 0.34 and 0.59. [43]

[41] 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.

[42] 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.

[43] 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. $\text{share}_t = \frac{\sum_{h\in\bigcup_{j\in J} H_{j,t}} \text{hhwt}_h}{\sum_{h\in\bigcup_{j\in J,t\in T} H_{j,t}} \text{hhwt}_h}$ is the nationwide share of type t and $\text{share}_{j,t} = \frac{\sum_{h\in H_{j,t}} \text{hhwt}_h}{\sum_{h\in\bigcup_{t\in T} H_{j,t}} \text{hhwt}_h}$ is the commuting-zone-j share of type t, where $H_{j,t}$ is the set of households of type t living in commuting zone j. population_j is the number of households in commuting zone j.

Equalizing population lowers high-middle skill and high-low skill consumption in equality by 10% and 9%, 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 16% of high-middle skill income gap, 13% of the high-low skill income gap, and 8% of the middle-low skill income gap. In terms of consumption inequality, these forces explain 10% of the high-middle skill consumption gap, 9% of the high-low skill consumption gap and 9% 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 expensive commuting zone enjoy a level of consumption that is half that of low-income residents in the most affordable 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, we estimate the standard of living that low- and high-skill households can expect in each US commuting zone, once we account both for geographical variation in cost of living and also in expected income. This allows us to answer the question of where in the US a household with a specific skill level can expect the highest standard of living and how it varies with local prices.

We find that for high-skill households, there is 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 similar to college graduates with the same observable characteristics 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 23.5% decline in market consumption.

While we have not included amenities in any of our calculations, prior work (Diamond, 2016) has found amenities to be higher in expensive cities. This suggests if we were to add in the additional utility value of amenities we would find that inequality between high and low skill households would be even larger. Unlike the simplistic Rosen-Roback model, where amenities perfectly offset differences in market consumption across space, amenities are higher in expensive cities, despite the high levels of high-skill consumption. The Rosen-Roback framework also poorly fits the middle and low-skill consumption differences across space. We find that the lowest skill households have the largest consumption differences between expensive and cheap cities. Through the lens of Rosen-Roback, this would indicate at the lowest skilled households have the highest willingness to pay for the amenities available in the most expensive cities. Again, this goes against prior work (Diamond, 2016). A richer model with preference heterogeneity within and across these skill groups is likely what is needed to understand these equilibrium relationships. We leave this to future research.

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.

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.

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*.

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

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.

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.

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

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.

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.

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

Moretti, E. (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).*

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.

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

Figure 1: **Income Distribution: Bank Account Data and ACS**

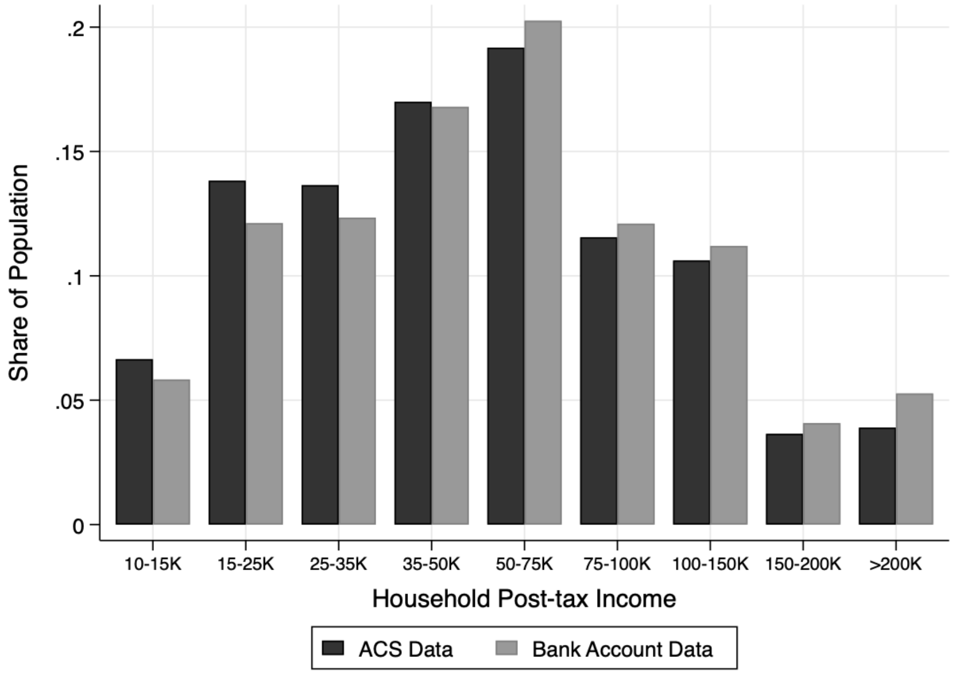

Notes: This figure compares the distribution of households by post-tax income bracket from our bank account data against the distribution from the 2012-2016 ACS micro data. In the ACS data, we use NBER TAXSIM software to calculate income taxes and then subtract it from household pre-tax income, yielding household post-tax income. The corresponding median (mean in parentheses) post-tax income in our data and in ACS data are \$52,955 (\$81,011) and \$48,831 (\$61,618).

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

(a) Raw Average Consumption Expenditure

(b) Adjusted Consumption Expenditure

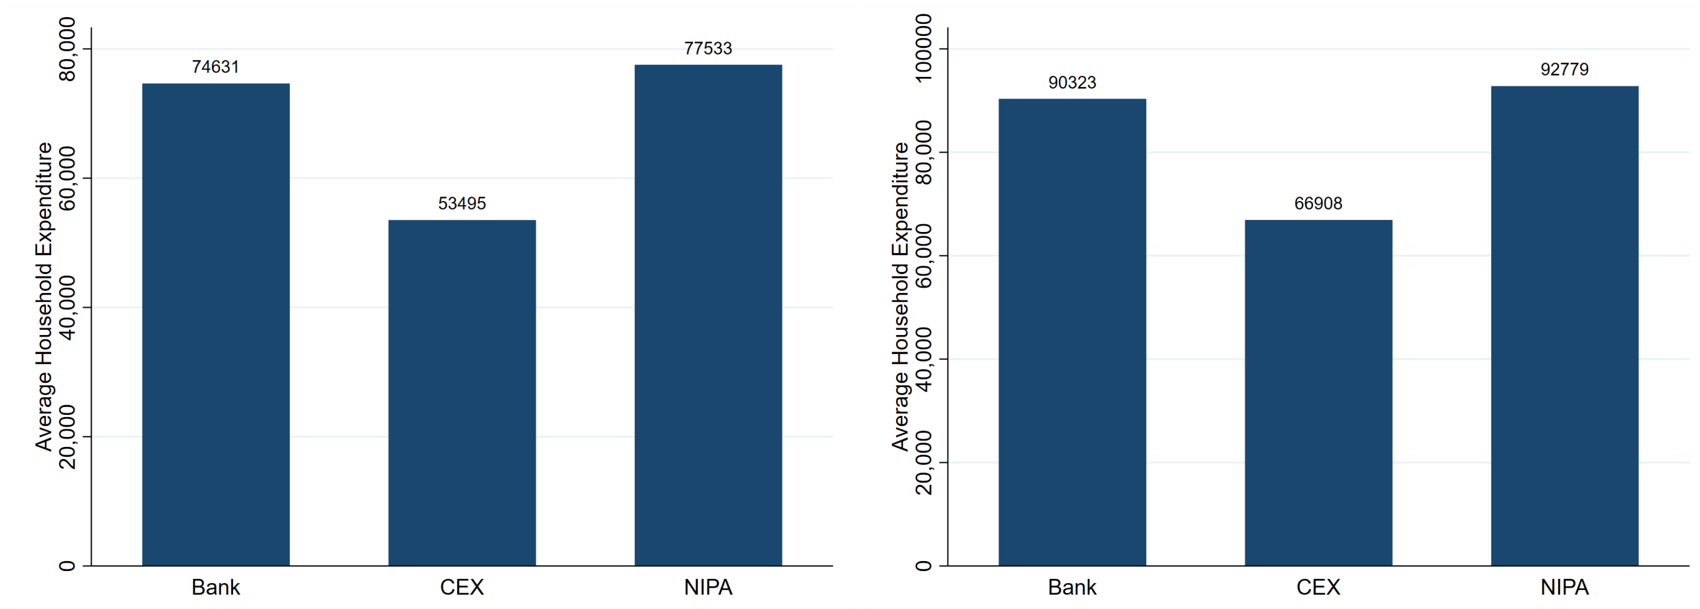

Notes: This figure compares average household annual expenditure in 2014 in the CEX, NIPA, and our data. These three data sources measure expenditure in different ways. Panel (a) compare the "raw" differences in datasets' average household expenditure. The only adjustment done in panel (a) is for NIPA: we subtract out healthcare spending paid by governments and insurance companies 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 share of health care costs that are not out-of-pocket (Center for Medicare Medical Services, 2021). In panel (b), we do more adjustments to the data to make them more even comparable. We adjust the bank data by subtracting out estimated owner housing costs and add back in estimated imputed rental values. We also add in imputed healthcare costs paid by the government and insurers. We adjust the CEX by subtracting out owner housing costs and add back in imputed housing rent (as measured in the CEX data). We then also add in the NIPA reported health insurance premiums paid by employers towards employee health insurance. See text for details.

Figure 3: **Expenditure Shares: Bank Data, NIPA and CEX**

(a) Expenditure Shares (raw)

(b) Expenditure Levels (adjusted)

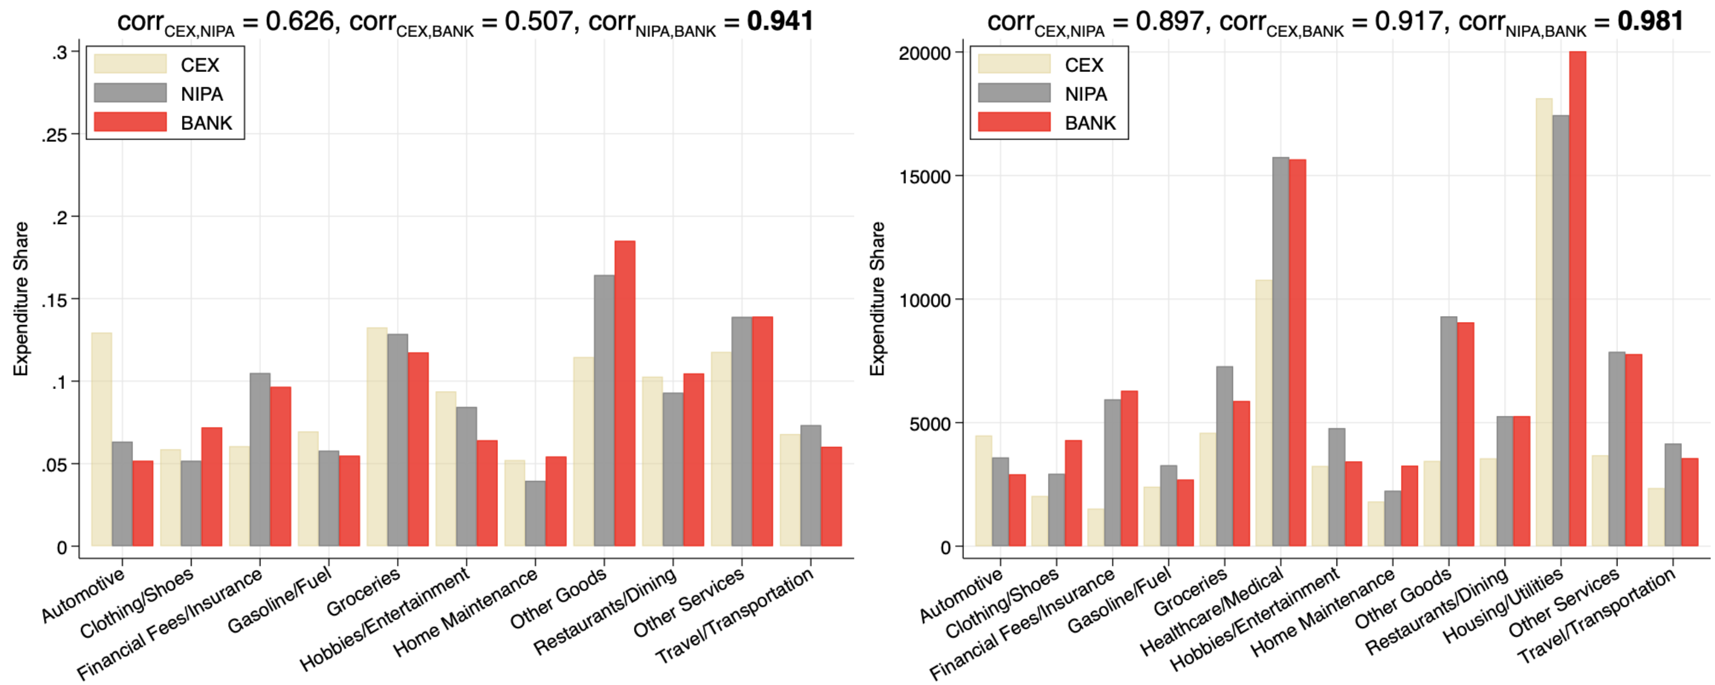

Notes: Panel (a) compares expenditure shares across comparable spending categories across the three datasets, where expenditure from our bank account does not include apportioned unclassified spending. Panel (b) compares average household expenditure levels across spending categories, where we add to the plot our imputed total health charges and housing costs. For CEX, we pool 2012-2016 Interview Survey data to measure annual spending. For NIPA, we use aggregate nationwide personal consumption expenditure in 2014. As spending categories do not align perfectly across the three datasets, we aggregate spending into 13 grouped categories that can be matched. For healthcare/medical spending, NIPA reports total health charges, while CEX and our data only include out-of-pocket health spending. To make feasible comparison, we deflate health spending in CEX and our data by 0.13 —the average ratio of out-of-pocket health spending to total health charges, as reported by the Center for Medicare Medical Services (2021). 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.

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

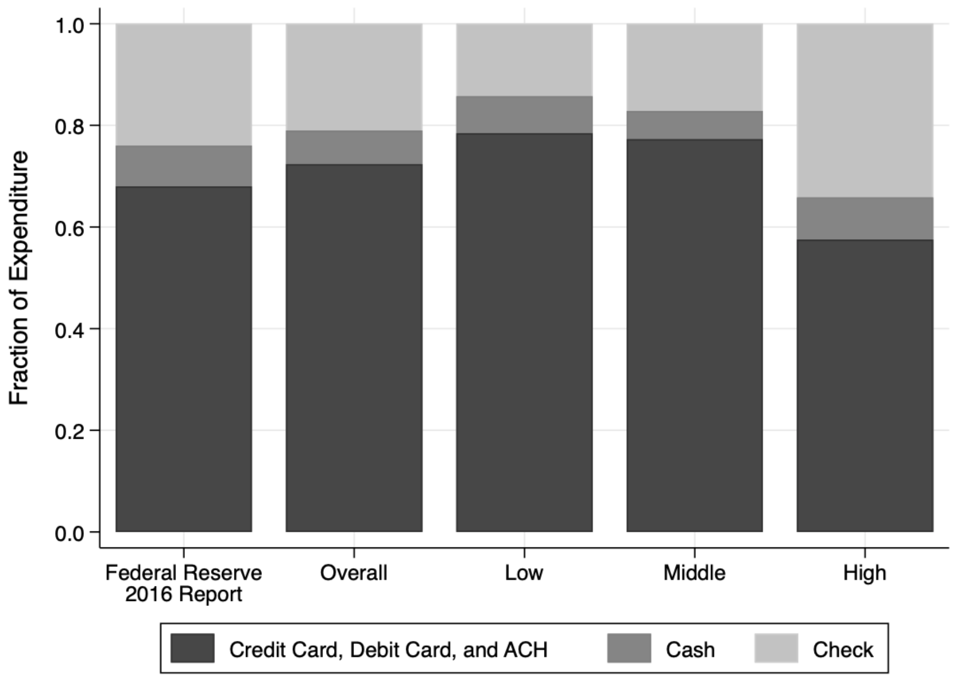

Notes: This figure plots average shares of consumption expenditure by mean of payment from the 2016 Federal Reserve Report and from our data. There are three types of payment including credit card, debit card, and ACH; cash; and check. When calculating averages with our data, we use commuting zone weight to adjust for sampling differences across counties in our data.

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

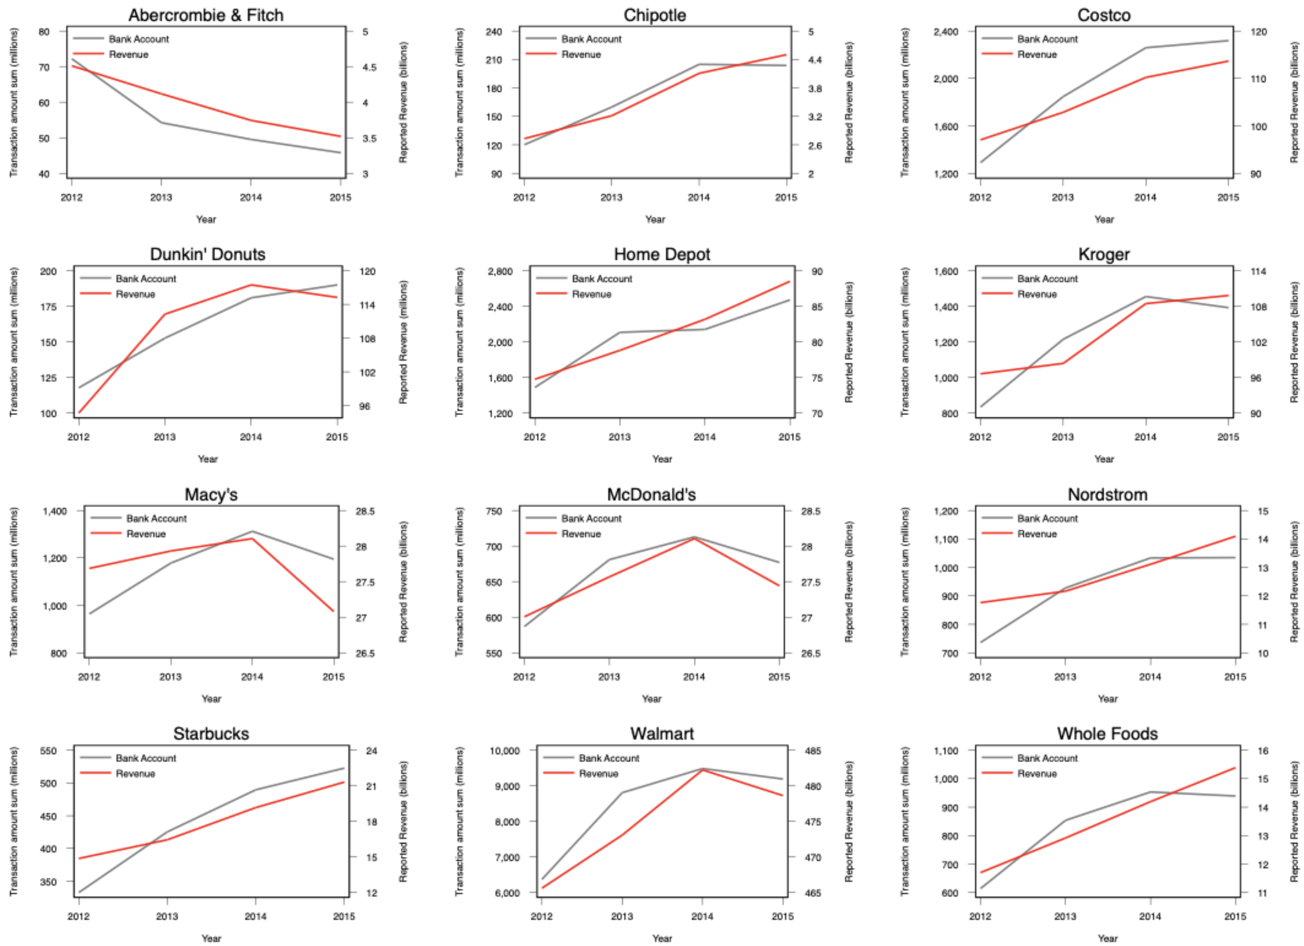

Notes: This figure compares total expenditures from our bank account data (in gray) and reported revenues from SEC 10K filings (in red) from 2012-2015.

Figure 6: **Number of Households: Bank Data vs. ACS**

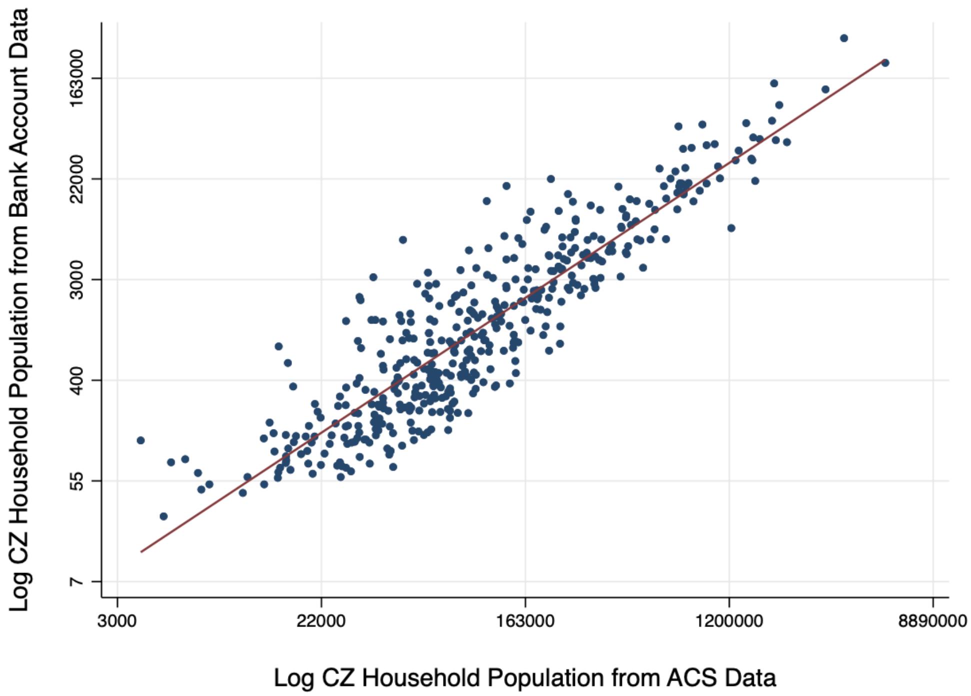

Notes: This figure plots log number of households from our bank account data against log number of households from 2012-2016 ACS data. Each dot is a commuting zone. To make ACS data consistent with our data, we drop households with income less than $10,000 in ACS. Values on both x-axis and y-axis are 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**

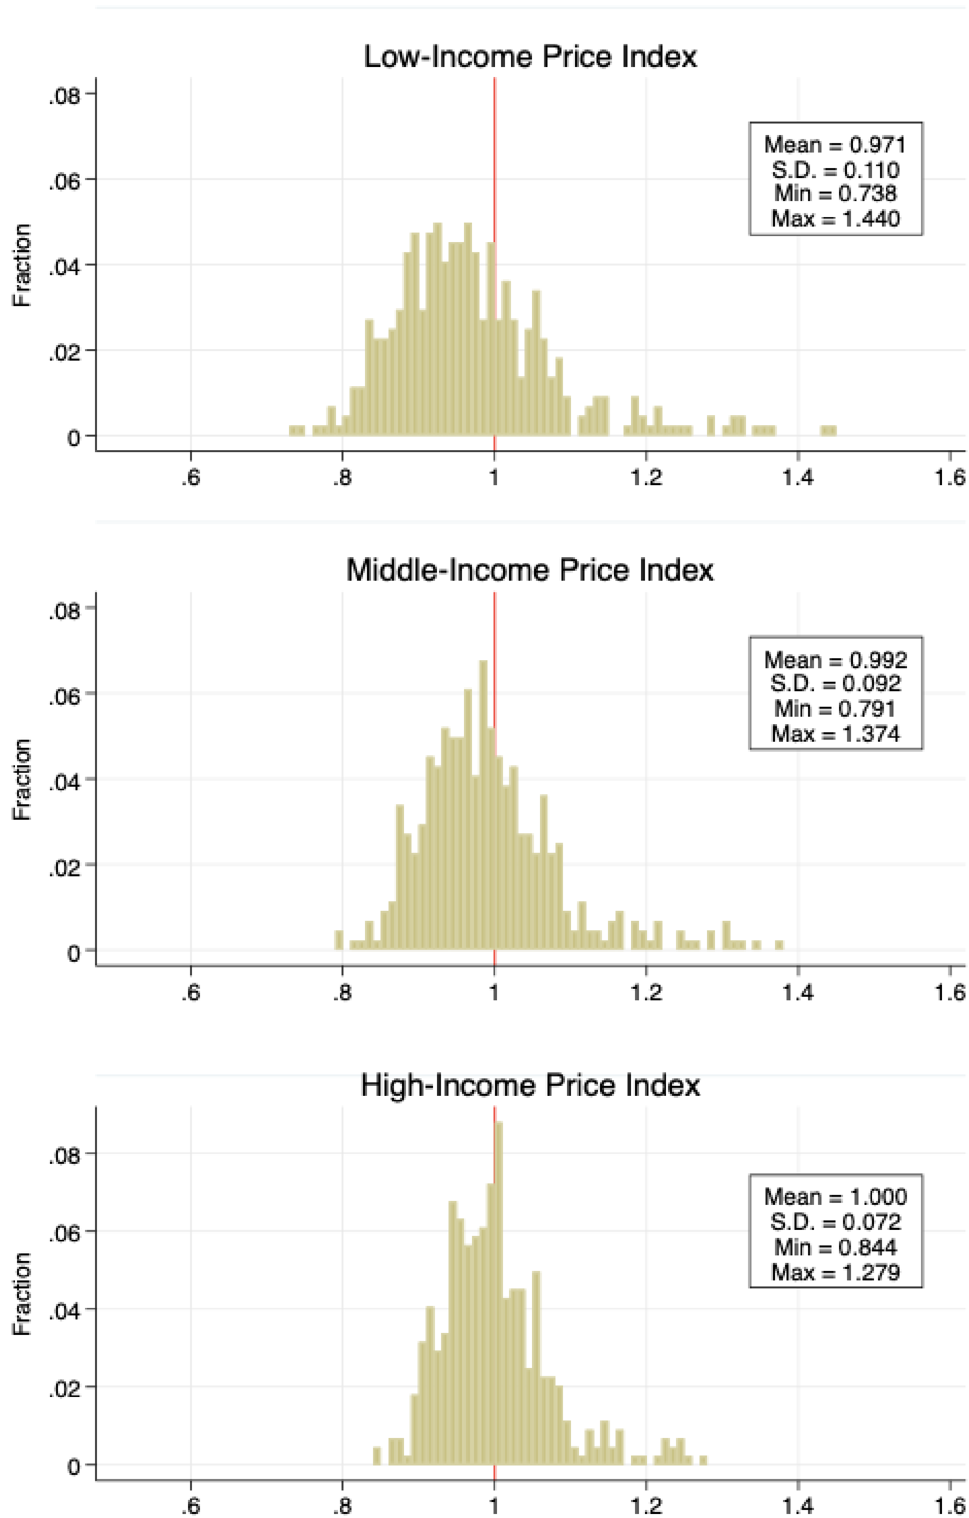

Notes: The level of observation is a commuting zone. $N = 443$.

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

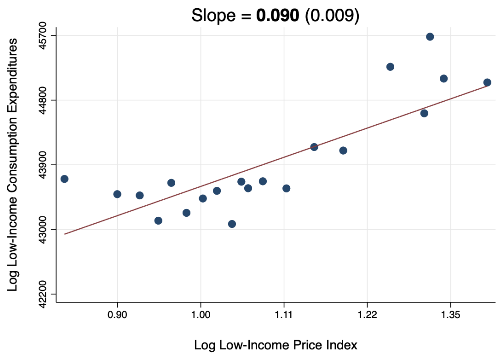

(a) Expenditure, Low Income

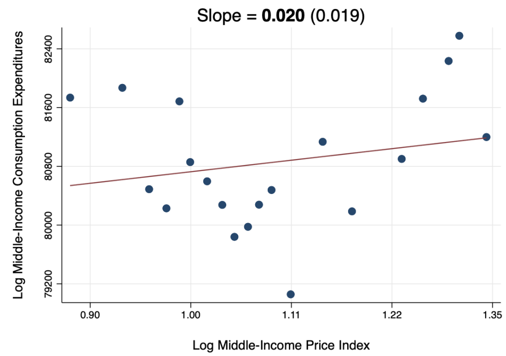

(b) Expenditure, Middle Income

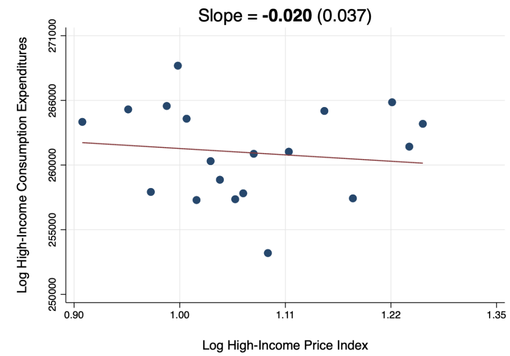

(c) Expenditure, High Income

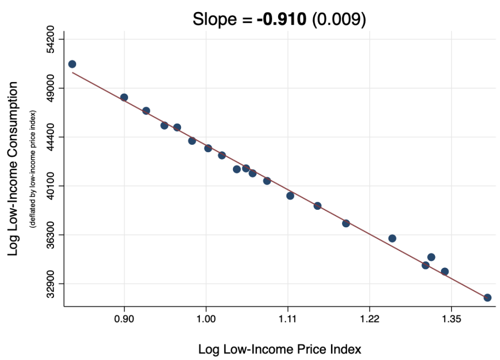

(d) Consumption, Low Income

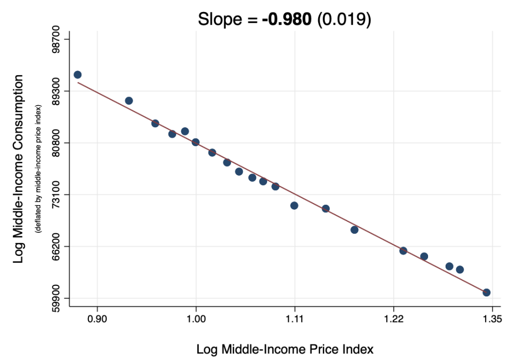

(e) Consumption, Middle Income

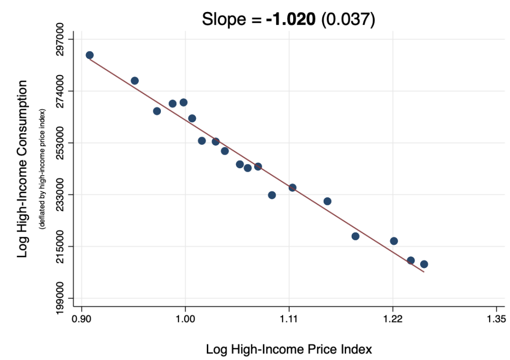

(f) Consumption, High Income

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 Estimates — Nielsen Data**

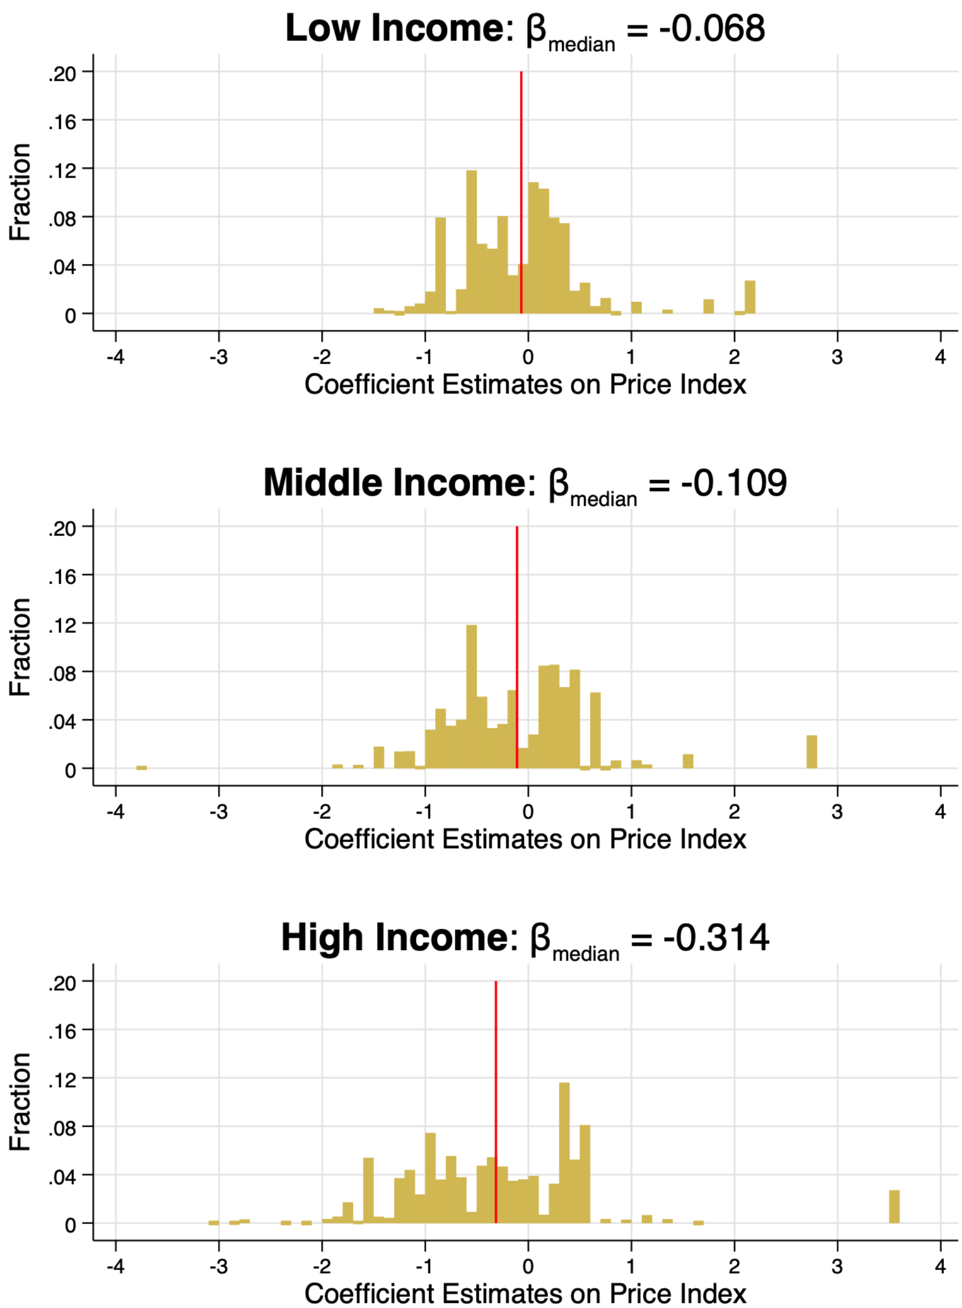

Notes: Each vertical red line denotes the median. We weight by average household expenditure on each product group.

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

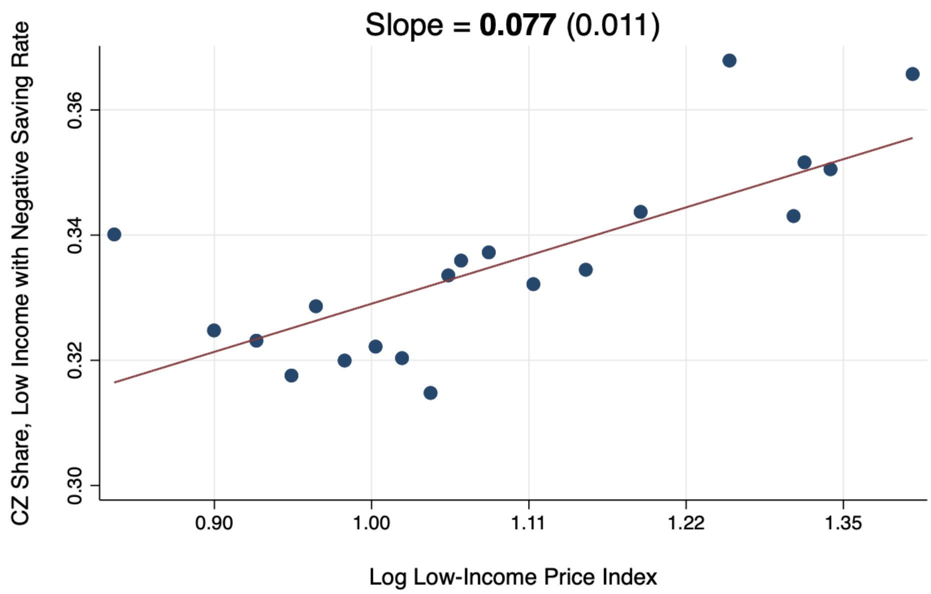

A. Negative Savings

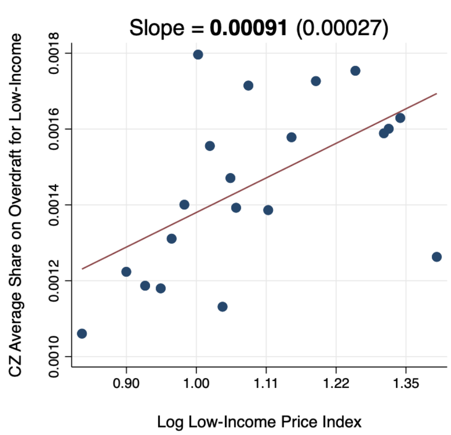

B. Overdraft Fees

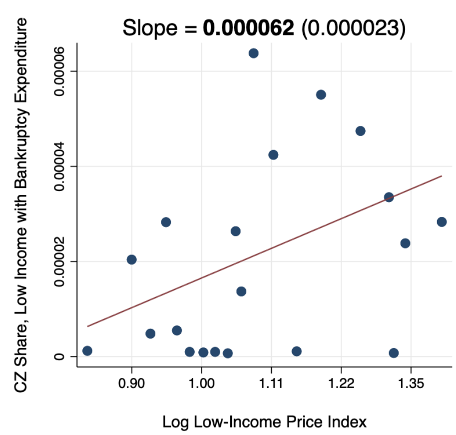

C. Bankruptcy

Figure 11: **Income, Expenditure, and Consumption Against Price Index by Skill Group**

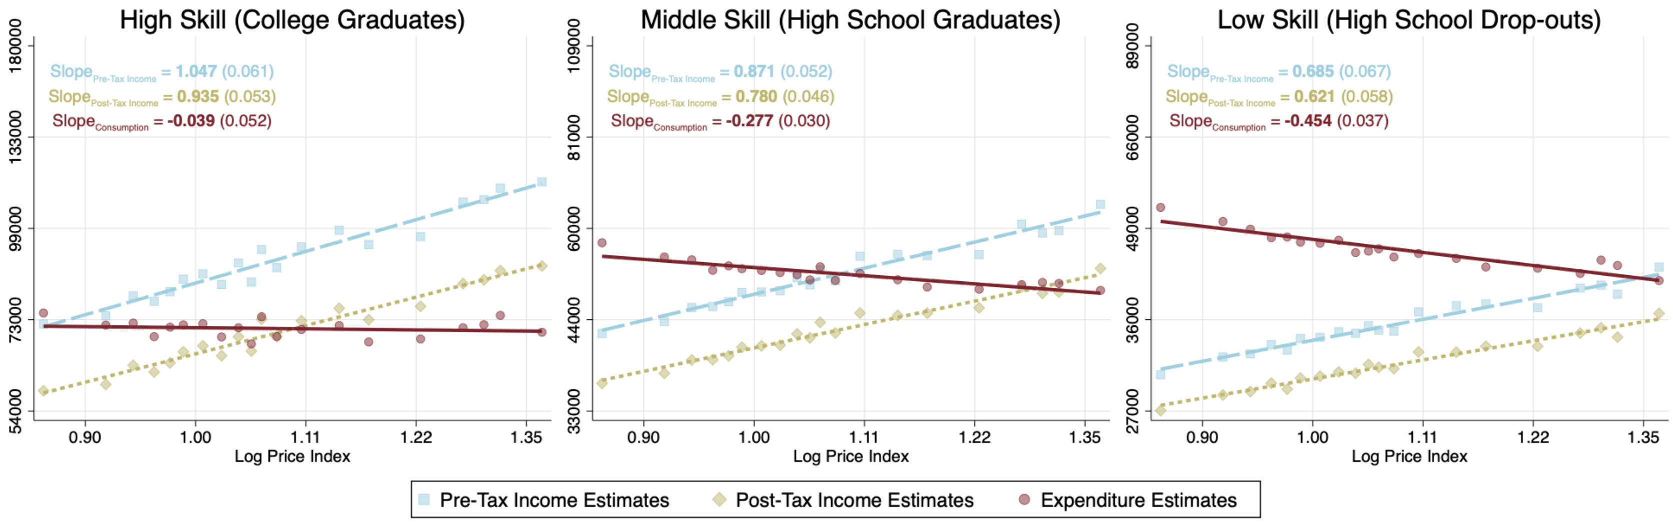

Figure 12: **Inequality in Consumption Within a Commuting Zone**

(a) College Grads vs. High School Grads

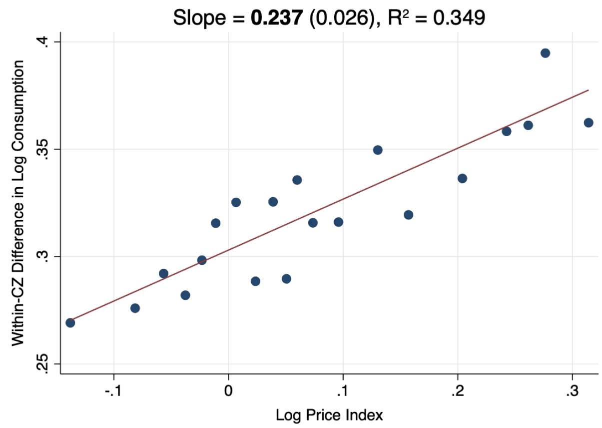

(b) College Grads vs. High School Dropouts

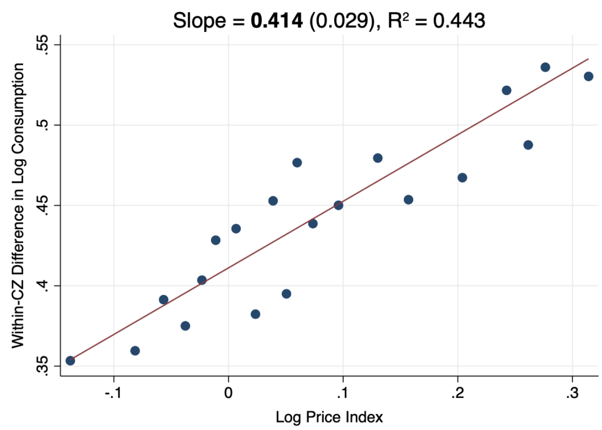

(c) High School Grads vs. High School Dropouts

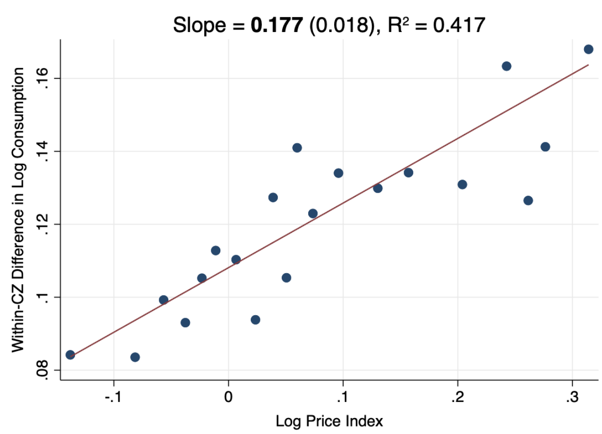

Figure 13: **Income, Expenditure, and Consumption Against City Size or College Share by Skill Group**

(a) **Against City Size**

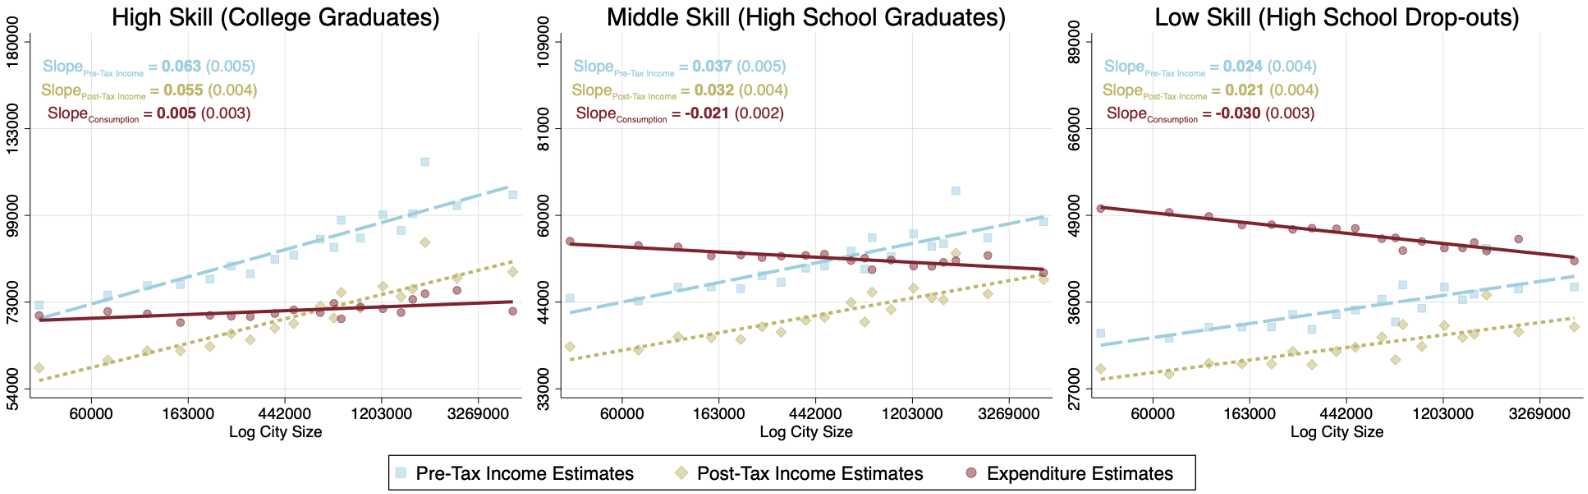

(b) **Against College Share**

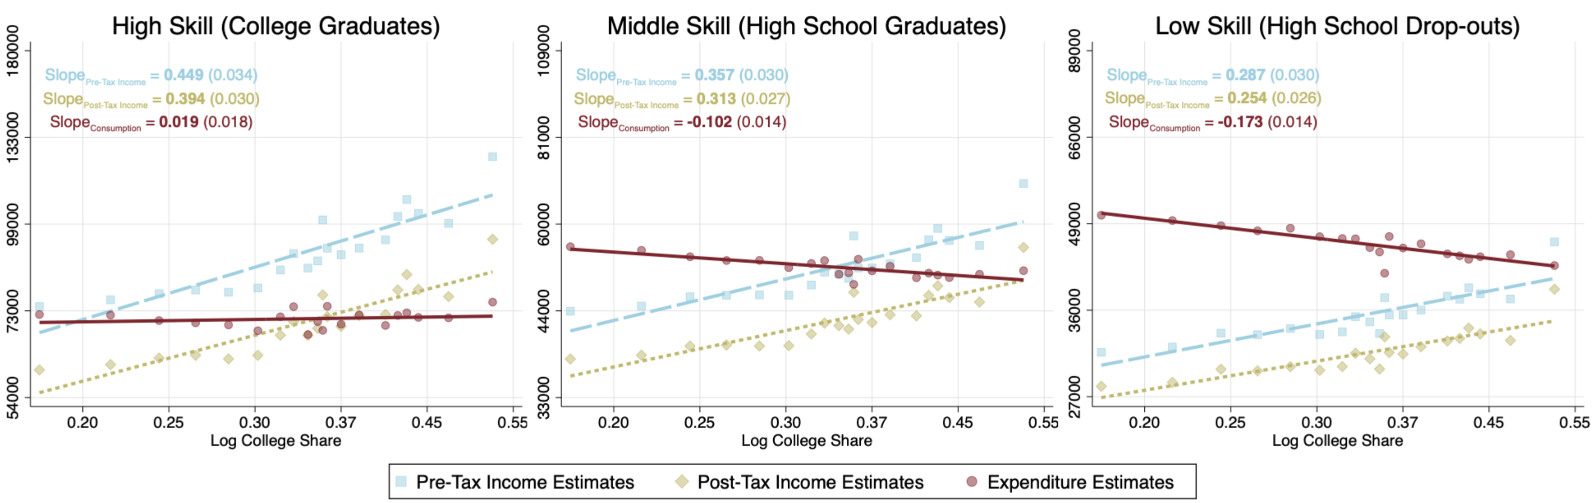

Figure 14: **Mean Price Index by Income Percentile**

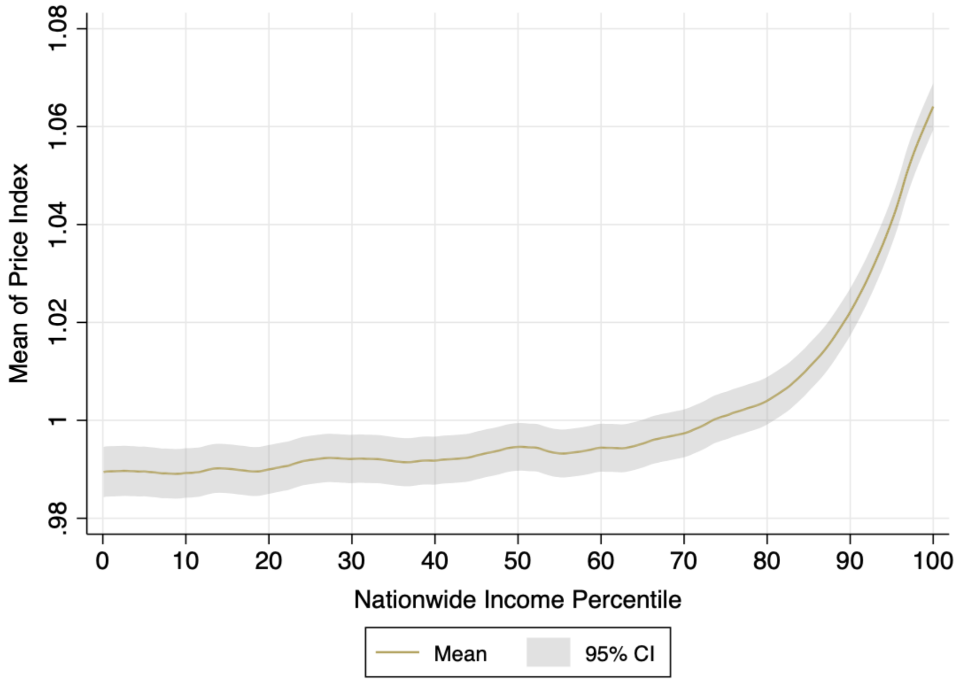

Figure 15: **Pre-tax Income and Consumption Inequality at the National Level**

(a) College Grads vs. High School Grads

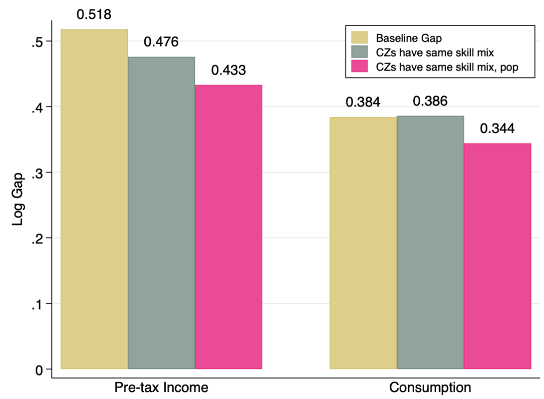

(b) College Grads vs. High School Dropouts

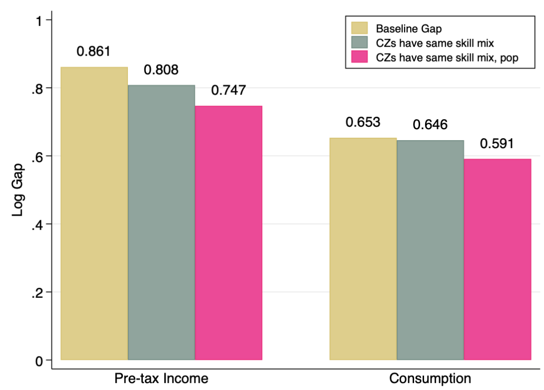

(c) High School Grads vs. High School Dropouts

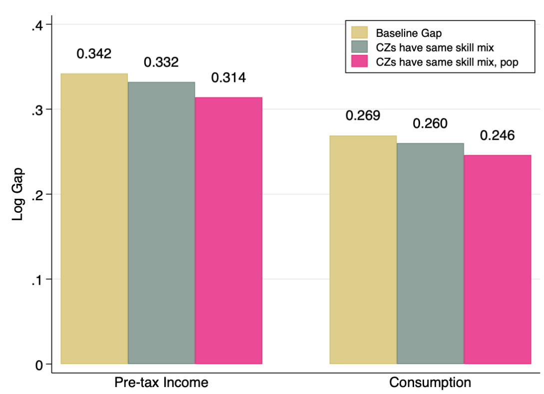

Table 1: **Commuting Zones by Price Index**

City	Price Index Overall Income (1)	Price Index Low Income (2)	Price Index Middle Income (3)	Price Index High Income (4)
Most Expensive				
San Jose, CA	1.399	1.440	1.374	1.279
San Francisco, CA	1.382	1.434	1.347	1.243
San Diego, CA	1.338	1.370	1.320	1.228
Honolulu, HI	1.328	1.359	1.310	1.234
New York, NY	1.323	1.347	1.307	1.244
Newark, NJ	1.310	1.324	1.303	1.253
White Plains, NY	1.305	1.317	1.300	1.250
Santa Barbara, CA	1.296	1.322	1.281	1.213
Edison, NJ	1.290	1.302	1.284	1.236
Los Angeles, CA	1.281	1.313	1.261	1.183
Boston, MA	1.266	1.285	1.253	1.229
Washington, DC	1.263	1.290	1.248	1.167
Hartford, CT	1.242	1.242	1.245	1.221
Seattle, WA	1.235	1.258	1.219	1.170
Miami, FL	1.223	1.235	1.219	1.166
Median				
Montrose, CO	1.000	0.991	1.009	1.014
Manhattan, KS	1.000	0.989	1.013	1.008
Cleveland, OH	1.000	1.000	1.000	1.000
Des Moines, IA	1.000	0.993	1.007	1.000
Omaha, NE	0.999	0.994	1.004	1.007
Least Expensive				
Somerset, KY	0.844	0.820	0.863	0.901
Batesville, AR	0.842	0.818	0.862	0.900
North Platte, NE	0.842	0.818	0.861	0.900
Union City, TN	0.839	0.814	0.858	0.897
Bluefield, WV	0.838	0.814	0.858	0.897
Elkins, WV	0.832	0.807	0.852	0.892
Beckley, WV	0.832	0.806	0.851	0.892
Summersville, WV	0.824	0.798	0.845	0.886
Greenville, MS	0.811	0.784	0.832	0.876
Marquette, MI	0.810	0.783	0.831	0.875
Pikeville, KY	0.810	0.783	0.831	0.875
London, KY	0.800	0.772	0.821	0.868
Natchez, MS	0.795	0.767	0.817	0.864
Gallup, NM	0.775	0.746	0.798	0.849
Presque Isle, ME	0.768	0.738	0.791	0.844

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	53,288	0.833	Huntington, WV	298,341	0.910
Johnstown, PA	53,064	0.829	Toledo, OH	297,887	0.976
Elizabeth City, NC	52,198	0.991	Johnstown, PA	297,162	0.907
Mobile, AL	51,749	0.918	Kalamazoo, MI	296,426	0.959
Traverse City, MI	50,723	0.929	Erie, PA	295,842	0.957
McAllen, TX	50,701	0.860	South Bend, IN	293,801	0.945
Florence, SC	50,689	0.866	Warsaw, IN	292,060	0.949
Youngstown, OH	50,584	0.867	Canton, OH	291,013	0.949
Charleston, WV	50,481	0.859	Cincinnati, OH	284,045	0.985
State College, PA	50,336	0.914	Pittsburgh, PA	283,000	1.011
Beaumont, TX	49,761	0.938	Fort Wayne, IN	282,879	0.952
Sunbury, PA	49,686	0.875	Bloomington, IN	281,911	0.944
Dothan, AL	49,674	0.854	McAllen, TX	281,820	0.900
Kingsport, TN	49,343	0.886	Sandusky, OH	281,599	0.994
Cadillac, MI	49,200	0.879	Youngstown, OH	280,910	0.933
Median Consumption			**Median Consumption**		
San Angelo, TX	43,779	0.947	Hagerstown, MD	250,747	1.031
Lebanon, NH	43,751	1.014	Columbia, SC	250,732	1.005
Nashville, TN	43,751	0.986	Utica, NY	250,441	1.030
New Orleans, LA	43,689	1.048	Boise City, ID	250,344	1.027
Prescott, AZ	43,683	1.023	Killeen, TX	250,292	0.985
Lowest Consumption			**Lowest Consumption**		
Hartford, CT	36,169	1.242	Boston, MA	217,363	1.229
Anchorage, AK	35,878	1.219	Olympia, WA	217,239	1.094
Boston, MA	35,615	1.285	Providence, RI	217,165	1.199
Seattle, WA	34,815	1.258	White Plains, NY	216,791	1.250
White Plains, NY	34,700	1.317	Edison, NJ	215,904	1.236
Washington, DC	34,475	1.290	Virginia Beach, VA	215,862	1.125
Edison, NJ	34,366	1.302	Yuma, AZ	213,959	1.039
Los Angeles, CA	34,003	1.313	New York, NY	213,284	1.244
Newark, NJ	33,788	1.324	Newark, NJ	212,308	1.253
New York, NY	33,649	1.347	San Francisco, CA	210,846	1.243
Santa Barbara, CA	33,628	1.322	Seattle, WA	210,708	1.170
Honolulu, HI	33,066	1.359	Medford, OR	205,997	1.059
San Diego, CA	32,607	1.370	San Diego, CA	203,070	1.228
San Francisco, CA	31,531	1.434	San Jose, CA	201,532	1.279
San Jose, CA	31,457	1.440	Honolulu, HI	199,125	1.234

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

Product	Unit	Low Income		Middle Income		High Income	
		$\hat{\beta}_{index}$	$\bar{Y}$	$\hat{\beta}_{index}$	$\bar{Y}$	$\hat{\beta}_{index}$	$\bar{Y}$
Carbonated Beverages	KG	-0.833***	128.5	-0.972***	129.8	-1.254***	122.1
		(0.086)		(0.118)		(0.168)	
Beer	KG	0.424	21.2	0.089	22.5	-1.205**	18.7
		(0.315)		(0.336)		(0.539)	
Cookies	KG	-0.219**	6.2	-0.330**	6.3	-0.067	5.8
		(0.090)		(0.133)		(0.203)	
Deodorant	KG	-0.206**	0.3	-0.416***	0.4	-0.899***	0.4
		(0.093)		(0.071)		(0.203)	
Eggs	CT	-0.163**	182.8	0.109	195.8	-0.379**	186.4
		(0.066)		(0.098)		(0.169)	
Housewares, Appliances	CT	-0.913***	2.2	-1.164***	2.4	-1.789***	2.4
		(0.090)		(0.159)		(0.183)	
Kitchen Gadgets	CT	-0.328	39.0	-0.374	55.2	0.441	64.4
		(0.210)		(0.266)		(0.448))	
Laundry Supplies	KG	-0.447***	11.0	-0.467***	12.2	-0.995***	12.0
		(0.099)		(0.129)		(0.174)	
Light Bulbs, Electric Goods	CT	-1.043***	6.7	-1.283***	7.3	-1.866***	7.7
		(0.116)		(0.137)		(0.299)	
Nuts	KG	0.135	3.1	0.223	4.2	-0.459	4.8
		(0.124)		(0.166)		(0.324)	
Pet Food	KG	-0.877***	57.6	-0.861***	55.3	-1.529***	47.7
		(0.138)		(0.179)		(0.378)	
Pizza, Snacks - Frozen	KG	-0.591***	5.8	-0.826***	6.0	-1.116***	5.6
		(0.140)		(0.156)		(0.244)	
Stationery, School Supplies	CT	-0.502***	228.4	-0.701***	278.5	-0.858**	275.6
		(0.176)		(0.223)		(0.402)	
Vegetables - Frozen	KG	-0.181	10.0	-0.240	11.2	-0.563*	10.0
		(0.135)		(0.160)		(0.310)	

Notes: Entries 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. The analysis is based on 59,755 households in the 2014 Nielsen 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**

Transaction Type	Low Income		Middle Income		High Income	
	$\hat{\beta}$	$\bar{Y}$	$\hat{\beta}$	$\bar{Y}$	$\hat{\beta}$	$\bar{Y}$
Hulu/Netflix/Google Play	-1.078***	5.1	-1.283***	6.3	-1.311***	4.0
	(0.109)		(0.174)		(0.228)	
Movies	-1.247***	0.7	-1.793***	1.1	-1.895***	0.8
	(0.362)		(0.421)		(0.549)	
Cable and Satellite Services	-0.310	3.3	-0.634***	6.0	0.318	5.5
	(0.196)		(0.217)		(0.285)	
DHL/FedEX/UPS/USPS	-0.241**	2.2	-0.869***	3.8	-1.390***	4.2
	(0.112)		(0.166)		(0.241)	
Gasoline/Fuel	-1.597***	24.1	-1.834***	35.1	-1.125***	23.1
	(0.263)		(0.250)		(0.411)	
Gym/Fitness/Yoga	0.462	2.2	-0.109	3.4	-0.486	3.6
	(0.283)		(0.298)		(0.452)	
Pedicure and Manicure	0.385*	0.7	0.028	1.2	1.710**	1.5
	(0.202)		(0.265)		(0.707)	

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**

	Pre-tax Income		Post-tax Income		Consumption		Difference between Pre-tax Income & Consumption		
	value (1)	pct (2)	value (3)	pct (4)	value (5)	pct (6)	value (7)	per (8)	pct (9)
1. Washington, DC	127,517	100	95,339	100	77,136	92	-50,381	-40%	-8
2. San Jose, CA	125,209	100	93,554	100	73,894	78	-51,315	-41%	-22
3. San Francisco, CA	121,903	100	92,230	100	74,369	81	-47,534	-39%	-19
4. White Plains, NY	115,260	99	87,653	99	75,238	85	-40,023	-35%	-14
5. New York, NY	112,974	99	85,802	99	74,864	84	-38,110	-34%	-15
6. Newark, NJ	112,105	99	86,267	99	72,831	71	-39,274	-35%	-28
7. Hartford, CT	109,037	99	83,555	99	75,370	85	-33,667	-31%	-14
8. Boston, MA	107,898	98	82,688	98	74,463	82	-33,434	-31%	-16
9. Baltimore, MD	105,688	98	80,930	98	76,180	88	-29,507	-28%	-10
10. Houston, TX	103,844	98	82,521	98	81,636	98	-22,208	-21%	+0
11. Philadelphia, PA	103,025	97	80,126	97	78,137	94	-24,888	-24%	-3
12. Los Angeles, CA	101,616	97	78,282	96	68,314	35	-33,302	-33%	-62
13. Chicago, IL	100,957	96	77,409	95	72,620	68	-28,337	-28%	-28
14. San Diego, CA	100,381	96	77,355	95	62,913	7	-37,467	-37%	-89
15. Seattle, WA	98,665	96	79,269	97	65,748	18	-32,917	-33%	-78
16. Dallas, TX	98,546	96	78,784	97	77,706	93	-20,840	-21%	-3
17. Camden, NJ	97,448	95	76,252	95	71,844	62	-25,604	-26%	-33
18. Sacramento, CA	94,964	94	73,776	93	66,513	23	-28,451	-30%	-71
19. Denver, CO	94,761	94	73,878	94	66,597	24	-28,164	-30%	-70
20. Austin, TX	94,704	93	76,254	95	72,941	72	-21,763	-23%	-21
21. Atlanta, GA	94,505	93	72,321	91	73,020	72	-21,485	-23%	-21
22. Minneapolis, MN	93,540	93	72,478	91	65,772	18	-27,768	-30%	-75
23. West Palm Beach, FL	92,583	92	74,709	94	74,915	84	-17,669	-19%	-8
24. Providence, RI	92,564	92	72,849	92	66,534	23	-26,031	-28%	-69
25. Fort Worth, TX	91,967	91	73,849	93	73,525	75	-18,442	-20%	-16
26. San Antonio, TX	91,475	91	73,390	93	73,428	74	-18,047	-20%	-17
27. Virginia Beach, VA	90,804	91	70,113	88	62,345	6	-28,459	-31%	-85
28. Detroit, MI	90,667	91	70,484	89	75,868	87	-14,799	-16%	-4
29. Cincinnati, OH	90,153	90	70,591	90	78,691	94	-11,462	-13%	+4
30. St. Louis, MO	89,730	89	69,669	87	71,277	58	-18,452	-21%	-31
31. Portland, OR	89,388	89	68,475	84	63,302	9	-26,086	-29%	-80
32. Raleigh, NC	88,977	88	68,415	84	69,852	47	-19,125	-21%	-41
33. Charlotte, NC	88,938	88	68,340	84	70,201	51	-18,738	-21%	-37
34. Nashville, TN	88,205	87	71,383	90	74,840	83	-13,365	-15%	-4
35. Columbus, OH	87,579	86	69,108	86	75,956	87	-11,623	-13%	+1
36. Milwaukee, WI	87,293	86	67,924	82	65,442	16	-21,851	-25%	-70
37. Phoenix, AZ	87,244	85	68,878	85	70,407	53	-16,837	-19%	-32
38. Kansas City, MO	86,786	84	67,906	82	68,358	36	-18,428	-21%	-48
39. Miami, FL	86,458	84	70,107	88	68,175	34	-18,283	-21%	-50
40. Cleveland, OH	86,170	83	68,191	83	76,797	91	-9,373	-11%	+8
41. Las Vegas, NV	85,082	81	69,050	86	69,965	49	-15,117	-18%	-32
42. Jacksonville, FL	85,008	80	68,733	85	70,111	50	-14,896	-18%	-30
43. Pittsburgh, PA	84,669	80	66,818	78	78,721	95	-5,949	-7%	+15
44. Indianapolis, IN	84,159	79	66,219	77	73,441	74	-10,718	-13%	-5
45. Harrisburg, PA	83,551	77	66,048	77	72,228	66	-11,324	-14%	-11
46. Tampa, FL	83,229	77	67,645	80	68,334	36	-14,895	-18%	-41
47. Salt Lake City, UT	82,460	73	64,174	66	65,661	17	-16,799	-20%	-56
48. Orlando, FL	80,285	66	65,534	75	66,440	23	-13,845	-17%	-43
49. Buffalo, NY	80,021	65	63,062	62	78,928	95	-1,093	-1%	+30
50. Grand Rapids, MI	75,883	47	60,290	43	69,276	43	-6,607	-9%	-4

Notes: Entries are average household pre-tax income, 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. The last three columns report the difference between consumption and pre-tax income in absolute, percentage, and percentile changes.

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

	Pre-tax Income		Post-tax Income		Consumption		Difference between Pre-tax Income & Consumption		
	value (1)	pct (2)	value (3)	pct (4)	value (5)	pct (6)	value (7)	per (8)	pct (9)
1. Washington, DC	72,181	100	57,103	100	52,509	42	-19,672	-27%	-58
2. San Jose, CA	67,755	100	54,751	100	50,175	17	-17,580	-26%	-83
3. San Francisco, CA	66,410	99	53,771	100	49,870	14	-16,540	-25%	-85
4. Baltimore, MD	60,830	98	49,248	98	53,703	58	-7,126	-12%	-40
5. Newark, NJ	60,165	98	49,575	98	49,591	12	-10,574	-18%	-86
6. New York, NY	59,126	97	48,092	95	50,185	17	-8,941	-15%	-80
7. White Plains, NY	59,021	97	48,199	96	50,758	22	-8,263	-14%	-75
8. San Diego, CA	58,962	97	48,400	96	45,900	0	-13,062	-22%	-97
9. Seattle, WA	58,616	96	49,260	98	47,801	3	-10,814	-18%	-93
10. Hartford, CT	58,528	96	48,178	95	51,356	27	-7,172	-12%	-69
11. Los Angeles, CA	58,234	95	47,831	95	48,394	6	-9,840	-17%	-89
12. Boston, MA	57,873	95	47,144	94	51,317	27	-6,556	-11%	-68
13. Camden, NJ	57,265	94	47,366	94	52,986	49	-4,279	-7%	-45
14. Denver, CO	56,273	93	46,232	93	48,131	5	-8,143	-14%	-88
15. Minneapolis, MN	55,913	93	45,995	92	48,345	6	-7,569	-14%	-87
16. Philadelphia, PA	55,539	93	45,710	91	53,007	49	-2,532	-5%	-44
17. Austin, TX	55,492	92	46,878	94	51,777	32	-3,715	-7%	-60
18. Houston, TX	55,464	92	46,664	93	54,911	70	-553	-1%	-22
19. Chicago, IL	55,389	92	44,929	89	50,288	18	-5,101	-9%	-74
20. Sacramento, CA	55,024	91	45,649	91	49,164	10	-5,861	-11%	-81
21. Virginia Beach, VA	54,796	91	44,715	89	47,335	1	-7,461	-14%	-90
22. Fort Worth, TX	54,210	90	45,764	91	52,844	46	-1,366	-3%	-44
23. Dallas, TX	54,179	90	45,781	91	52,817	45	-1,361	-3%	-45
24. Portland, OR	52,985	88	43,180	84	46,761	1	-6,224	-12%	-87
25. Salt Lake City, UT	52,730	88	43,544	85	50,837	23	-1,893	-4%	-65
26. Providence, RI	52,697	87	43,912	86	48,637	7	-4,060	-8%	-80
27. Las Vegas, NV	52,694	87	44,519	88	51,992	35	-703	-1%	-52
28. San Antonio, TX	52,428	87	44,361	88	52,394	41	-34	-0%	-46
29. West Palm Beach, FL	52,388	86	44,507	88	51,928	34	-460	-1%	-52
30. Atlanta, GA	52,381	86	42,769	81	50,603	21	-1,778	-3%	-65
31. Phoenix, AZ	52,354	86	43,574	86	51,915	34	-439	-1%	-52
32. Harrisburg, PA	51,319	83	42,532	81	54,331	65	3,012	6%	-18
33. St. Louis, MO	50,837	81	42,069	78	51,181	25	344	1%	-56
34. Kansas City, MO	50,728	80	42,067	78	50,356	19	-372	-1%	-61
35. Milwaukee, WI	50,125	78	41,630	75	48,097	4	-2,028	-4%	-74
36. Nashville, TN	50,026	77	42,741	81	53,381	55	3,355	7%	-22
37. Detroit, MI	49,944	77	41,442	73	52,794	44	2,851	6%	-33
38. Raleigh, NC	49,721	75	40,784	69	50,000	15	279	1%	-60
39. Cincinnati, OH	49,367	73	41,359	72	53,989	61	4,622	9%	-12
40. Jacksonville, FL	49,115	72	41,741	75	50,798	23	1,683	3%	-49
41. Miami, FL	48,861	70	41,698	75	48,085	4	-775	-2%	-66
42. Cleveland, OH	48,358	67	40,726	68	53,331	54	4,973	10%	-13
43. Charlotte, NC	48,254	67	39,698	57	49,466	11	1,213	3%	-56
44. Columbus, OH	48,107	66	40,504	65	52,970	48	4,863	10%	-18
45. Indianapolis, IN	47,995	64	39,913	60	51,939	35	3,944	8%	-29
46. Tampa, FL	47,399	58	40,563	66	49,561	12	2,162	5%	-46
47. Buffalo, NY	47,093	56	39,508	54	59,694	94	12,601	27%	+38
48. Orlando, FL	47,038	56	40,312	64	49,313	11	2,275	5%	-45
49. Pittsburgh, PA	46,825	54	39,098	50	55,547	76	8,721	19%	+22
50. Grand Rapids, MI	44,789	36	37,767	37	51,966	35	7,177	16%	-1

Notes: Entries are average household pre-tax income, 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. The last three columns report the difference between consumption and pre-tax income in absolute, percentage, and percentile changes.

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

	Pre-tax Income		Post-tax Income		Consumption		Difference between Pre-tax Income & Consumption		
	value (1)	pct (2)	value (3)	pct (4)	value (5)	pct (6)	value (7)	per (8)	pct (9)
1. Washington, DC	48,785	100	40,873	100	43,740	7	-5,044	-10%	-93
2. San Jose, CA	45,160	99	38,725	99	42,049	3	-3,111	-7%	-96
3. San Francisco, CA	43,903	99	37,570	99	41,305	1	-2,598	-6%	-98
4. Seattle, WA	41,905	97	36,650	98	41,366	1	-539	-1%	-96
5. Newark, NJ	41,751	97	36,145	97	43,377	5	1,625	4%	-92
6. Baltimore, MD	40,115	96	34,395	94	45,559	22	5,444	14%	-74
7. Denver, CO	39,944	95	34,411	95	41,985	2	2,041	5%	-93
8. Camden, NJ	39,888	95	34,791	96	46,937	35	7,050	18%	-60
9. Salt Lake City, UT	39,428	94	34,191	93	46,476	29	7,048	18%	-65
10. White Plains, NY	39,287	94	34,262	94	44,658	14	5,371	14%	-80
11. Chicago, IL	39,215	93	33,284	89	44,305	9	5,090	13%	-84
12. Boston, MA	39,146	93	33,609	92	45,124	17	5,979	15%	-76
13. Harrisburg, PA	39,050	93	33,419	90	50,306	71	11,257	29%	-22
14. Portland, OR	39,022	93	33,290	89	42,016	2	2,994	8%	-91
15. Hartford, CT	38,997	92	34,181	93	44,524	13	5,527	14%	-79
16. Las Vegas, NV	38,539	91	33,820	92	45,686	24	7,147	19%	-67
17. New York, NY	38,438	90	33,423	90	43,417	6	4,980	13%	-84
18. San Diego, CA	38,391	90	33,498	91	39,352	0	961	3%	-90
19. Minneapolis, MN	38,379	89	33,568	91	43,158	5	4,779	12%	-84
20. Los Angeles, CA	38,167	88	33,291	89	40,882	1	2,715	7%	-87
21. Virginia Beach, VA	37,905	88	32,668	86	42,431	3	4,526	12%	-85
22. Philadelphia, PA	36,834	84	31,884	81	45,711	24	8,876	24%	-60
23. Austin, TX	36,667	83	32,429	85	43,858	7	7,191	20%	-76
24. Kansas City, MO	36,494	81	31,706	79	45,354	20	8,860	24%	-61
25. Fort Worth, TX	36,482	81	32,210	84	45,086	17	8,604	24%	-64
26. Providence, RI	36,393	81	31,802	81	43,613	6	7,220	20%	-75
27. Houston, TX	36,300	80	32,034	82	46,446	29	10,146	28%	-51
28. Milwaukee, WI	36,233	80	31,724	79	43,659	7	7,426	20%	-73
29. Sacramento, CA	36,152	80	31,855	81	42,409	3	6,257	17%	-77
30. Dallas, TX	36,107	79	31,978	81	45,180	18	9,072	25%	-61
31. West Palm Beach, FL	35,637	74	31,712	79	43,985	8	8,347	23%	-66
32. Detroit, MI	35,365	72	30,974	70	46,705	33	11,340	32%	-39
33. Phoenix, AZ	35,171	70	30,877	69	45,712	24	10,540	30%	-46
34. St. Louis, MO	35,171	70	30,713	67	44,841	14	9,670	27%	-56
35. Atlanta, GA	34,648	65	30,183	60	44,097	8	9,449	27%	-57
36. Nashville, TN	34,626	65	30,992	71	48,012	48	13,386	39%	-17
37. San Antonio, TX	34,429	63	30,555	65	44,918	15	10,488	30%	-48
38. Cleveland, OH	34,070	58	30,057	59	46,367	28	12,297	36%	-30
39. Grand Rapids, MI	33,882	57	29,947	58	48,821	57	14,939	44%	+0
40. Jacksonville, FL	33,734	55	30,123	59	45,258	19	11,524	34%	-36
41. Orlando, FL	33,713	55	30,194	61	43,955	8	10,242	30%	-47
42. Cincinnati, OH	33,483	53	29,577	53	46,517	30	13,034	39%	-23
43. Tampa, FL	33,289	51	29,829	56	44,257	9	10,968	33%	-42
44. Indianapolis, IN	33,025	48	28,885	43	45,731	25	12,706	38%	-23
45. Buffalo, NY	33,007	48	29,396	51	54,594	93	21,587	65%	+45
46. Miami, FL	32,952	46	29,437	51	42,425	3	9,473	29%	-43
47. Columbus, OH	32,901	45	29,218	49	46,205	28	13,304	40%	-17
48. Charlotte, NC	32,800	44	28,472	38	44,337	10	11,536	35%	-34
49. Pittsburgh, PA	32,796	44	28,678	41	49,821	65	17,026	52%	+21
50. Raleigh, NC	32,437	40	28,372	37	44,382	10	11,945	37%	-30

Notes: Entries are average household pre-tax income, 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. The last three columns report the difference between consumption and pre-tax income in absolute, percentage, and percentile changes.

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

	Log Consumption
Log price index	-0.289***
	(0.080)
Log price index × middle-skill	0.023
	(0.093)
Log price index × low-skill	-0.063
	(0.094)
Log city size	0.028***
	(0.008)
Log city size × middle-skill	-0.026***
	(0.009)
Log city size × low-skill	-0.037***
	(0.009)
Log college share	0.021
	(0.038)
Log college share × middle-skill	-0.040
	(0.044)
Log college share × low-skill	-0.041
	(0.044)
Middle-skill	-0.028
	(0.151)
Low-skill	0.006
	(0.151)

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 both enter our data during 2011-2013 and exit during 2015-2016 and then keep only the 2014 data for these households. 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[44]. We identify income taxes by filtering through debit transactions from both bank and card accounts using tax-related words[45].

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[46].

Because we observe infrequent transaction activities especially for households with extremely low annual income or annual expenditure, we deal with these inactive accounts into two ways. First, 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.

[44] "wire", "online banking transfer", "agent assisted transfer", "transfer to imma", "transfer to sav", "transfer to ch", "transfer to money market", "transfer from imma", "transfer from sav", "transfer from ch", "transfer from money market", "sweep", "bank transfer", "online transfer to", "online transfer from", "internal transfer", "save as you go transfer", "usaa funds transfer", "transfer ppd", "transfer out", "schedule transfer", "transfer to account", "xfer to acct", "phone transfer", "citibank global transfer", "fund ", "usaa chk-intrnt", "hsbc online-transfer", "bankrock", "citibank xfer", and "keep the change"

[45] "tax pmt", "franchise tax", "tax pymt", "irs", "tax ref", "tax reb", "tax deduc", "city tax", "county tax", "tax collect", and "tps tax".

[46] "student ln", "masviaach", "retirement contributions", "loan", "prosper", "lending club", "advance crd", "home equity", "payday adv", "ez money", "cash america", "cash land", "advance america", "loc", "line of credit", "revolver", "revolving", "brokerage credit", "brokerage debit", "brokerage misc", and "us treas".

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 Construction of Price Indexes

Here we describe the details of our price index construction.

B.1 Measuring Prices

We combine price data from three sources.

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. 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/Jewelries, 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 uses to estimate the CPI, we measure housing costs using rental prices. Rental prices are a better measure of the user cost of housing than home values because home values reflect not just the user cost of a unit, but also any expectations of appreciation or depreciation.

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) that is 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.[47] 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 $\text{rent}_{j,k} = \sum_{n=1}^{N}(\text{rent}_{n,j,k} \times \bar{s}_{n,j,k})$.

B.2 Measuring Expenditure Shares

We closely follow methodologies the BLS uses to calculate expenditure shares as a part of their local CPIs. 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-groups-specific 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 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 $s_{i,j,k}$ to obtain expressions on the right-hand side. 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. Many of our alternative price indices require measuring expenditure shares at the commuting zone level. Our main Laspeyres index only requires nationwide expenditure shares. We discuss below how we impute commuting zone-income group-specific expenditure share that can then be aggregated to nationwide income group-specific expenditure shares.

Non-Housing Expenditure. To measure expenditure of the 22 non-housing categories listed in Appendix Table A2 we use our bank account data. For each household, we sum all the expenditures within each category. This leaves the remaining unclassified transactions that include atm cash, checks, rents, mortgage payments, and other unclassified transactions. These unclassified transactions include housing expenditure, as well as spending on the 22 non-housing categories. We assign

[47] 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., $\widehat{p_{housing,j}} = exp(\widehat{\delta}_j + \bar{X}\widehat{\beta})$.

these unclassified transactions as follow.

Housing Expenditure Following the CPI's methods, we will use CEX and ACS data to estimate housing expenditures. First, we measure housing costs using CEX Interview Survey data by pooling 2012-2016 years (centering around 2014). Consistent with BLS's housing cost definition used in its CPI-U, our definition of housing costs is contract rent for renters plus equivalent rent for owners (or how much owners think their properties would rent for monthly, unfurnished, and without utilities). We define two measures of housing costs in the CEX data. The first 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 $\times$ 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 will 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 $\times$ income group. Then, we use the coefficient 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.

Adjusting Healthcare Spending We adjust for "Healthcare/Medical" spending in our bank account data. There are two issues with this category. One issue is that 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. The other issue is that our non-recreational health spending covered under health insurance is out-of-pocket spending and thus does not reflect total health charges, thus health consumption is underestimated. To address the first 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). To address the second issue, 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 (pharmacy) spending on out-of-pocket healthcare (pharmacy) spending by region and income group. 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.

Final Expenditure Shares Finally, we calculate household-level expenditure shares. For each household h, we take its total expenditure (E_h) and then subtract out our imputed housing costs to be subtracted (H_h^-), leaving total non-housing expenditure (N_h), precisely, $N_h = E_h - H_h^-$. Within this total non-housing expenditure, we identify spending from our focal 21 non-housing categories paid through bank or card accounts ($X_{h,i}$ for $i \in I = \{1, \dots 21\}$) and remaining spending either paid in cash or checks or outside the 21 categories, e.g., "Unclassified", precisely, $N_h - \sum_{i \in I} X_{h,i}$. 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, our non-housing spending now includes only spending from these categories such that $\sum_{i \in I} \tilde{X}_{h,i} = N_h$. Next, we add back our imputed housing costs to be added (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 $i = 1, \dots, 22$.

Detailed Expenditure Shares from Nielsen 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 23 high-level categories.

B.3 Alternative Price Indices

Here we provide the details of all of our 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:

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

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

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:

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

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

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

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

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

E_g is national total expenditure on variety g.

For housing, we assume there is only one variety and its 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 $\pi_{i,k} = \sum_{j\in J}\frac{p_{i,j}\cdot \mathrm{q}_{i,j,k}}{P_{j,k}^{\text{Geary-Khamis}}\cdot\sum_{j'\in J}\mathrm{q}_{i,j',k}}$.

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 $\mathbf{s} = \sum_{r=0}^{5} \mathbf{b_r} y^r + \mathbf{Bp}y + \epsilon$, 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); $y = (x - \mathbf{p's})/(1 - \mathbf{p'Bp}/2)$ 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 $y = x - \mathbf{p's}$ and $\tilde{y} = x - \mathbf{p'\bar{s}}$.

To calculate an income-group-specific price index implied by the exact EASI demand model, we first predict $\mathbf{\tilde{s}_h}$ for each households if it were to face average prices. These predicted shares as a function of real consumption expenditures are shown in Appendix Figure A7. Next, we mean-collapse to calculate $\mathbf{\tilde{s}_j}$ 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, $\mathbf{p_j'\tilde{s}_j} = \sum_i (p_{ij} \cdot \tilde{s}_{ij})$.

C 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.

D 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, as described above (Section B). 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.

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.

E Government Transfers

In this section, we describe how we impute total government assistance, which we add to our measure of consumption expenditures as 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. We save the resulting coefficient estimates for imputing housing subsidies below.

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 above AHS regression 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 adjustments following an approach in Meyer and Mittag (2019). They calculate numbers to scale up these three government assistance measures by income to federal poverty level, and we precisely use their numbers. We add up these three measures for each households to define total government assistance.

Finally, we define a total government assistance to income ratio and regress it 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 total government assistance for all households in 2012-2016 ACS data and then add this imputed measure to our consumption expenditures.

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

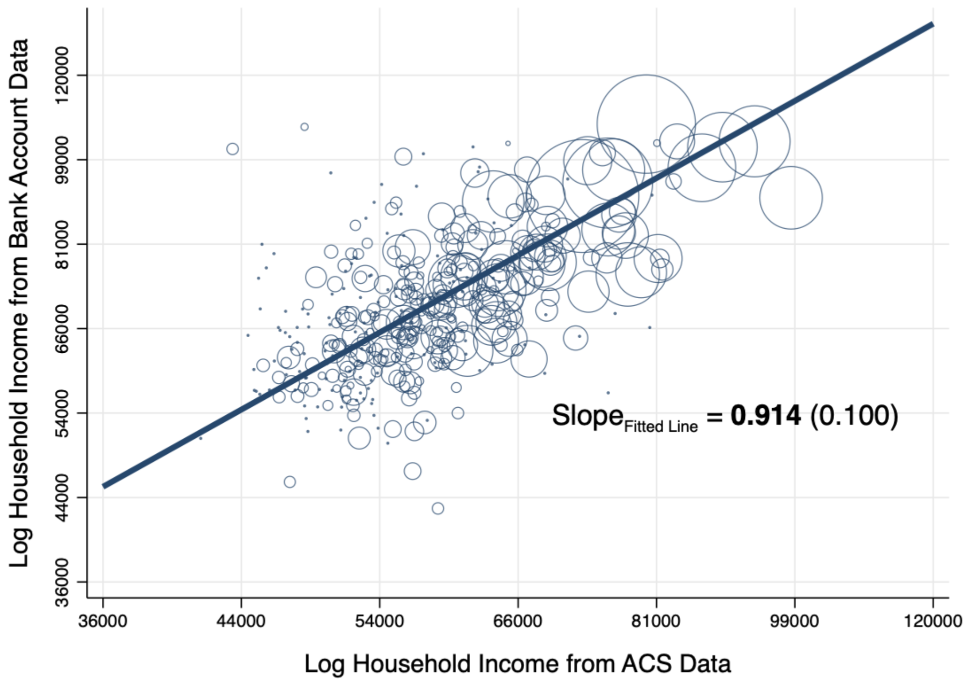

Notes: Observations are at the commuting zone level. ACS data are from 2012-2016. Household income in our data is post-tax. To obtain post-tax income in the ACS data, we subtract from household pre-tax income the income taxes calculated using the NBER TAXSIM software. We weight observations by their 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**

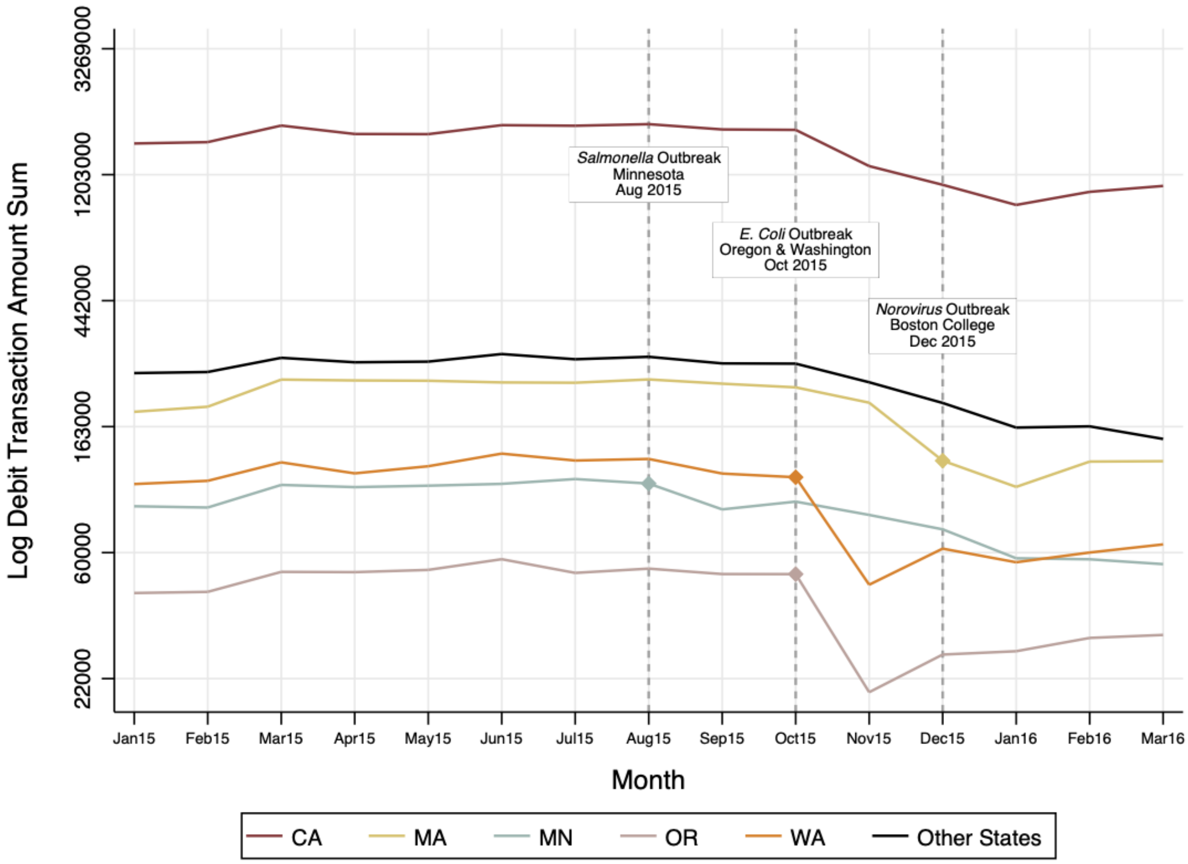

Notes: This figure shows changes in expenditures in Chipotle stores observed in our data after the *Salmonella* outbreak in Minnesota in August 2015; the *E. Coli* outbreak in Oregon and Washington in October 2015; and the Norovirus outbreak in Boston in December 2015. The dash lines indicate the months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

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

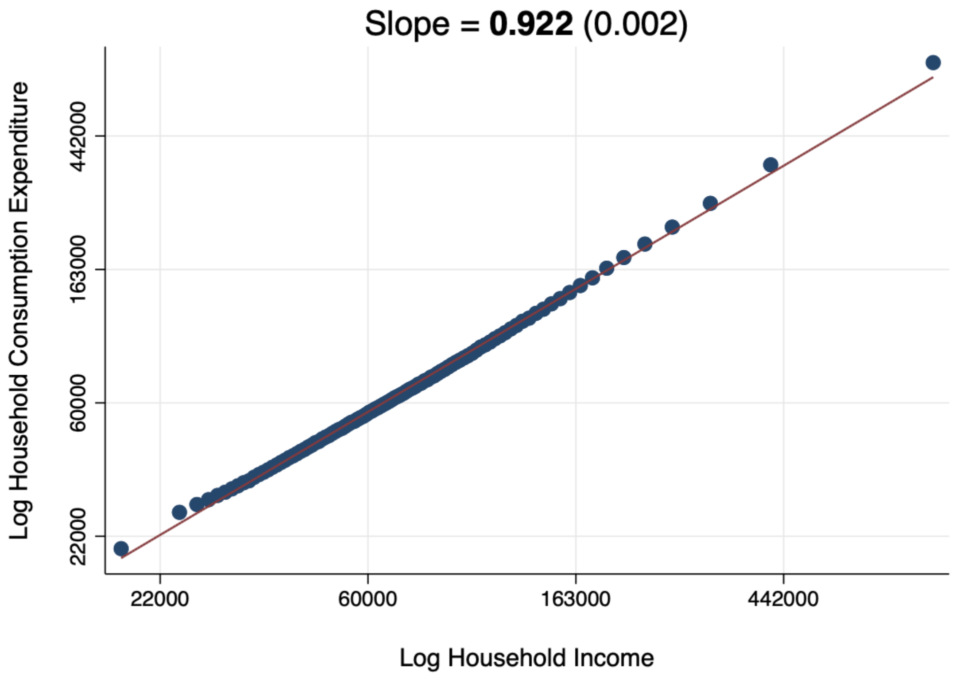

Notes: The sample includes all households in our sample and uses commuting zone weight as weight. Values on both x-axis and y-axis are in measured in log scale, but we label actual values for easier interpretation. N = 3,000,518 households.

Appendix Figure A4: **Comparison of Price Index with BEA Index**

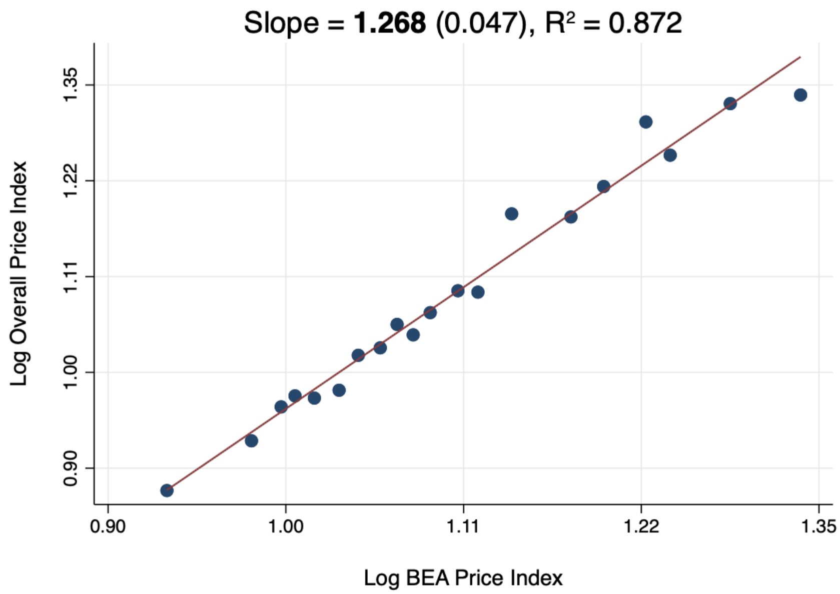

Notes: Observations are at the CZ level. $N = 443$. We use CZ population as regression weight.

Appendix Figure A5: **Consumption Expenditure Across CZs, Low vs. High Income**

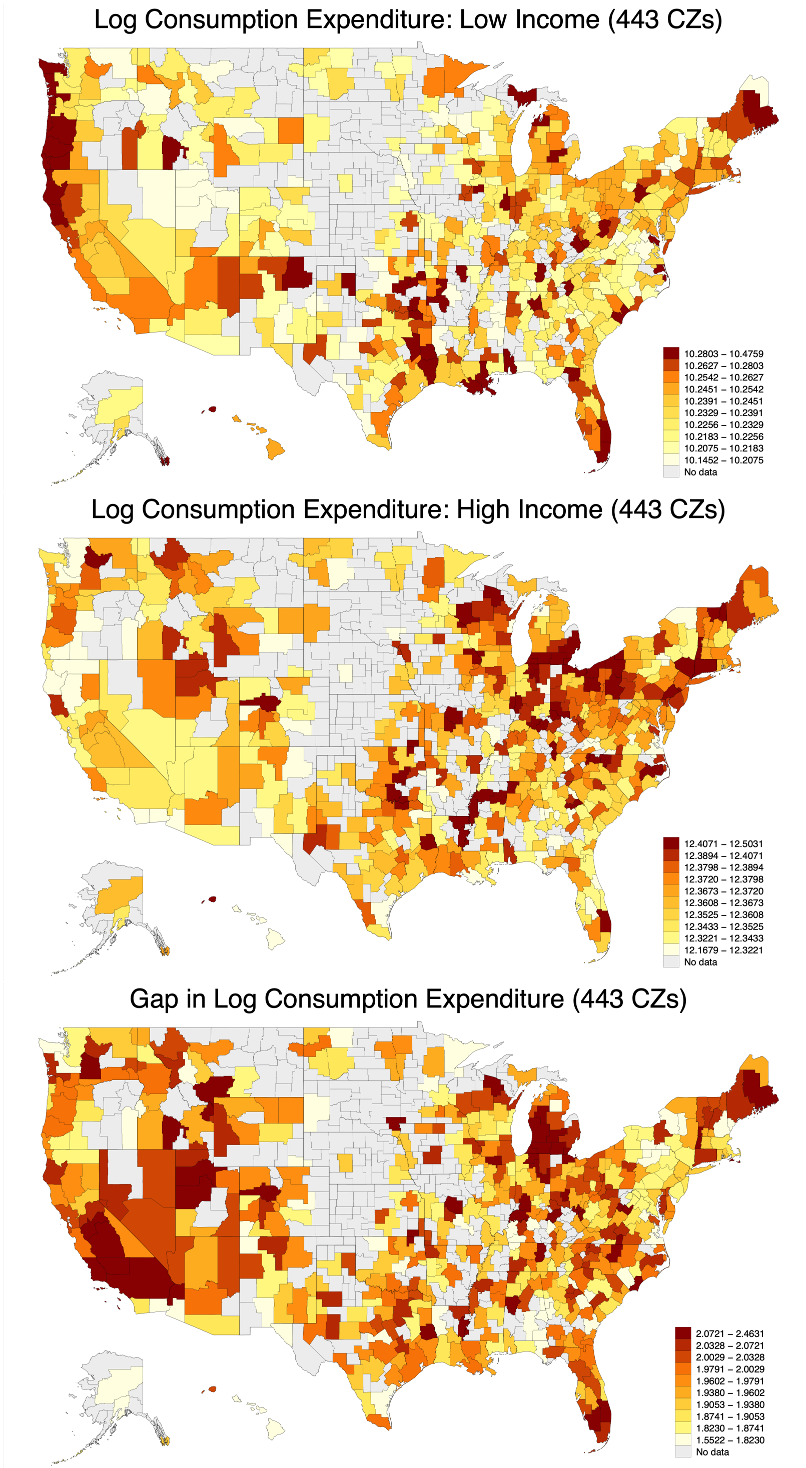

Notes: This figure shows log consumption expenditure for low-income households (left) and high-income households (middle) and the difference between these two groups (right). We shrink estimates so that small commuting zones with less precise estimates are shrunk toward the nationwide mean. We use the following formula: $\hat{Y}_i^{shrinkage} = w_i \cdot Mean(\hat{Y}_i) + (1 - w_i) \cdot \hat{Y}_i$, where $w_i = {SE_{\hat{Y}_i}}^2 / (Var(\hat{Y}_i) - Mean({SE_{\hat{Y}_i}}^2) + {SE_{\hat{Y}_i}}^2)$.

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

A. High Skill

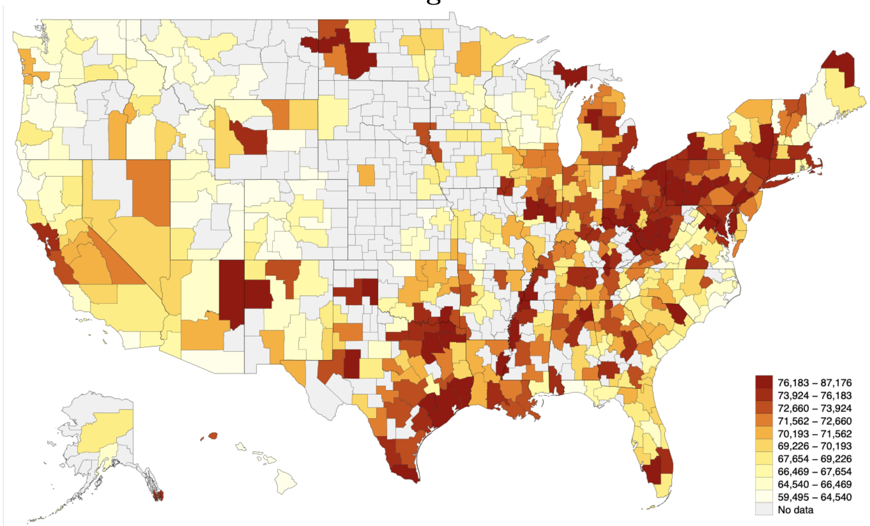

B. Middle Skill

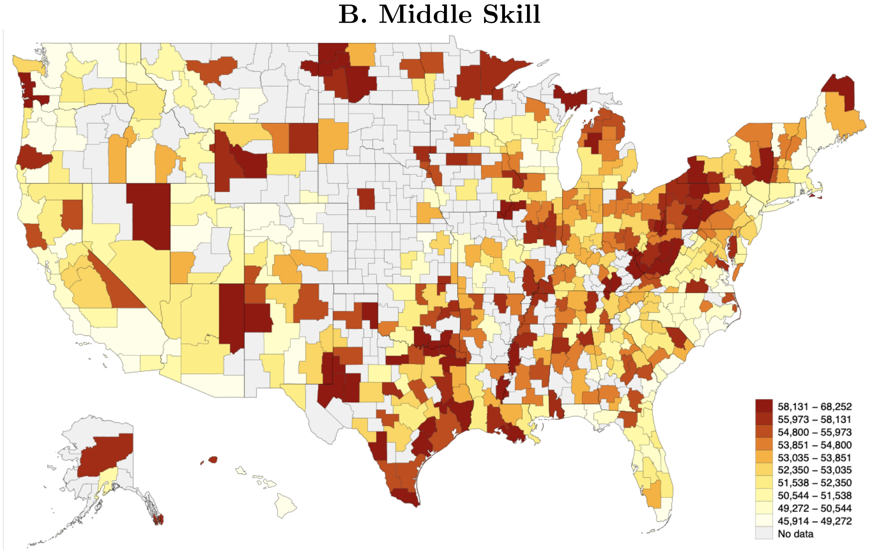

C. Low Skill

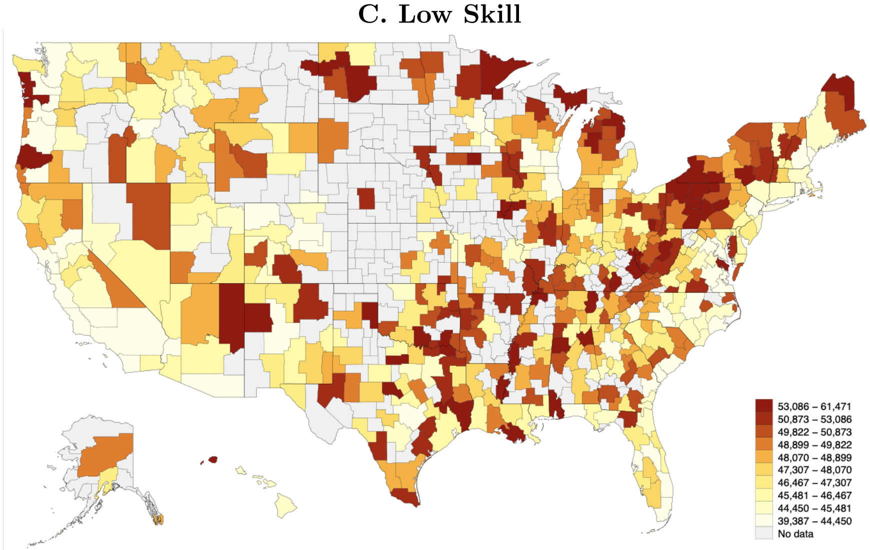

Appendix Figure A7: **EASI Demand Expenditure Shares as vs. Real Expenditure**

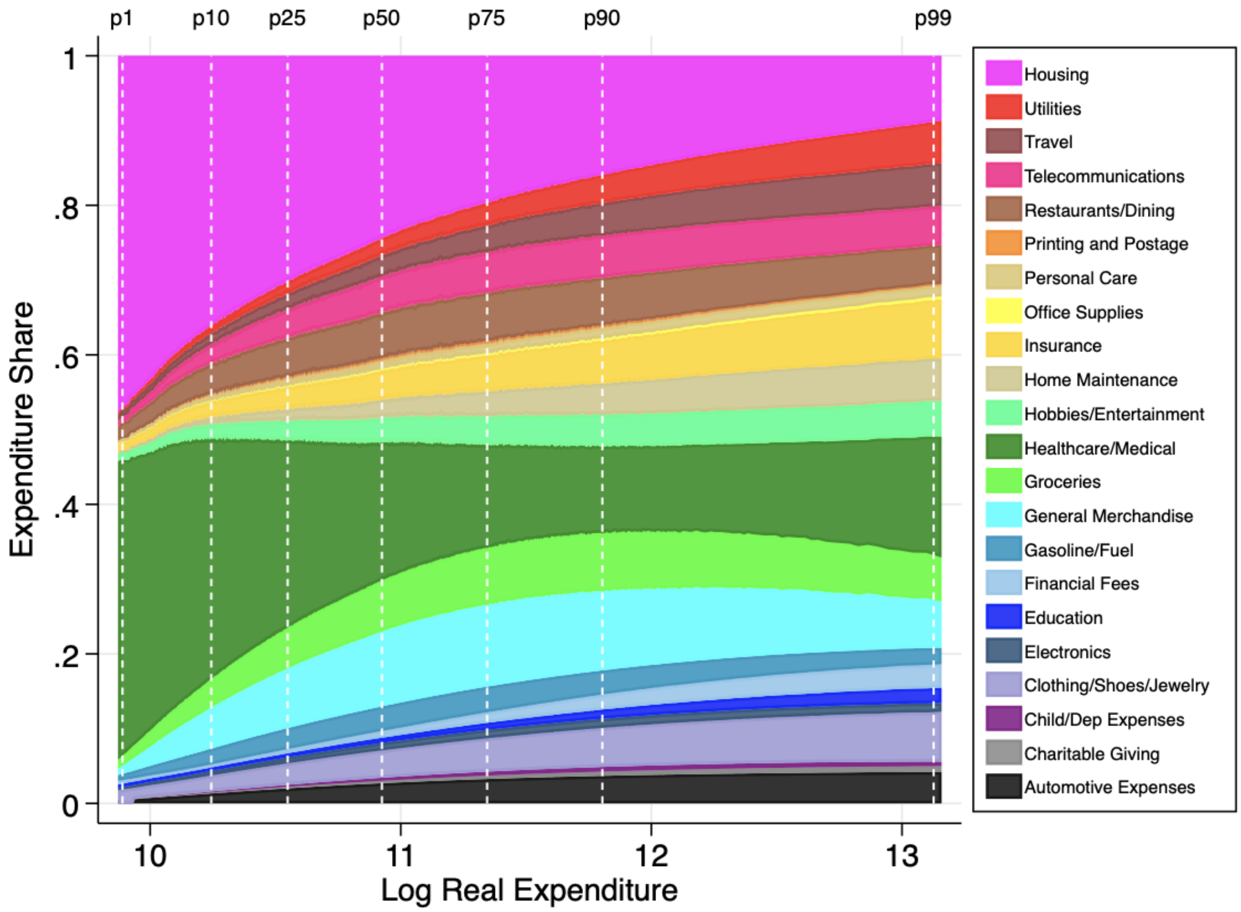

Notes: The expenditure shares across our focal spending categories displayed in the figure have been recovered from an exact EASI demand estimation, following a framework in Lewbel and Pendakur (2009). In this model, we allow each household's expenditure shares to be a function of polynomials of real 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 Table A1: **Summary Statistics**

	Overall Income (1)	Low Income (2)	Middle Income (3)	High Income (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	96,702.82	44,696.75	104,046.84	495,068.62
Median	65,727.21	43,916.22	92,858.72	310,952.81
Expenditure				
Mean	90,323.31	44,960.32	95,060.70	452,886.91
Median	60,585.61	41,764.11	82,429.19	272,304.97
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 across all households in our bank account data. Panel B summarizes adjusted post-tax income and expenditure, where for each household we add to both measures in Panel A our imputed non out-of-pocket health spending and housing cost adjustments. We provide a detailed discussion of these imputations in Appendix B.2

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	2.60%	1.88%	3.15%	4.11%	0.149
Charitable Giving	0.33%	0.19%	0.43%	0.71%	0
Child/Dependent Expenses	0.39%	0.22%	0.54%	0.63%	0.021
Clothing/Shoes/Jewellery	3.68%	2.79%	4.29%	6.18%	0.260
Education	0.80%	0.61%	0.88%	1.82%	0
Electronics	0.96%	0.78%	1.10%	1.20%	0.029
Financial Fees	1.07%	0.62%	1.32%	2.77%	0
Gasoline/Fuel	2.95%	2.65%	3.29%	2.63%	0.062
General Merchandise	8.86%	7.45%	10.39%	7.55%	0.219
Groceries	6.22%	5.04%	7.34%	6.53%	0.157
Healthcare/Medical	20.38%	28.19%	13.71%	11.10%	0.320
Hobbies/Entertainment	3.18%	2.52%	3.68%	4.59%	0.340
Home Maintenance/Improvement	2.51%	1.46%	3.28%	5.00%	0.195
Insurance	3.92%	2.87%	4.61%	7.14%	0
Office Supplies	0.22%	0.17%	0.25%	0.36%	0.032
Personal Care	1.07%	0.86%	1.23%	1.49%	0.195
Printing and Postage	0.27%	0.24%	0.29%	0.24%	0
Restaurants/Dining	5.62%	5.03%	6.19%	5.82%	0.253
Telecommunications	4.42%	3.51%	5.21%	5.34%	0.149
Travel	2.91%	2.03%	3.50%	5.45%	0
Utilities	2.39%	1.49%	3.00%	4.92%	0.133
Housing	25.25%	29.41%	22.33%	14.42%	0.294

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

	Low Income (1)	Middle Income (2)	High Income (3)
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%

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

	Low Income (1)	Middle Income (2)	High Income (3)
Log Monthly Rent	0.428***	0.354***	0.267***
	(0.011)	(0.010)	(0.013)
R^2	0.952	0.947	0.885

Notes: All columns employ a log-log specification and use commuting zone population as regression weight. Robust standard errors are reported in parentheses. N = 443, *** 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=7	A5. Combined index, Nested CES σ=11.5	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	B5. Combined index, Nested CES σ=7	B4. Combined index, Nested CES σ=11.5	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=7	.84	.84	.80	1												
A5. Combined index, Nested CES σ=11.5	.94	.93	.91	.97	1											
A6. Combined index, PPP, Geary-Khamis	.99	.98	.97	.85	.94	1										
A7. Combined index, PPP, GEKS-Fischer	.99	.99	.97	.85	.94	.99	1									
B1. Combined index, Laspeyres	.99	.98	.97	.81	.92	.98	.98	1								
B2. Combined index, Laspeyres (CEX weights)	.98	.99	.96	.81	.91	.97	.98	.99	1							
B3. Combined index, CES	.96	.96	.99	.77	.88	.96	.96	.97	.96	1						
B5. Combined index, Nested CES σ=7	.85	.84	.81	.99	.96	.85	.85	.83	.83	.79	1					
B4. Combined index, Nested CES σ=11.5	.94	.93	.91	.95	.99	.94	.94	.93	.93	.90	.96	1				
B6. Combined index, PPP, Geary-Khamis	.99	.98	.97	.83	.93	.99	.99	.99	.98	.97	.84	.94	1			
B7. Combined index, PPP, GEKS-Fischer	.99	.98	.97	.82	.93	.99	.99	.99	.99	.97	.84	.94	.99	1		
C. EASI demand based index	.98	.98	.96	.82	.92	.98	.98	.98	.98	.96	.83	.93	.99	.99	1	
D. BEA index	.93	.92	.90	.74	.83	.91	.92	.93	.92	.89	.74	.84	.91	.92	.92	1

Appendix Table A6: **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***	-0.066***	-0.116***	0.313***	0.263**
	(0.088)	(0.007)	(0.012)	(0.084)	(0.082)
B. Child/Dependent Expenses					
Log Median CZ Income	0.067***	-0.047***	-0.082***	0.020	-0.015
	(0.017)	(0.004)	(0.007)	(0.016)	(0.016)
C. Clothing/Shoes/Jewelry					
Log Median CZ Income	0.232**	-0.039***	-0.068***	0.193**	0.164*
	(0.076)	(0.004)	(0.007)	(0.075)	(0.074)
D. Electronics					
Log Median CZ Income	0.003	-0.020***	-0.034***	-0.017	-0.032*
	(0.011)	(0.002)	(0.003)	(0.012)	(0.012)
E. Gasoline/Fuel					
Log Median CZ Income	0.165***	-0.028***	-0.050***	0.137***	0.116***
	(0.036)	(0.003)	(0.005)	(0.035)	(0.033)
F. General Merchandise					
Log Median CZ Income	0.134***	-0.334***	-0.584***	-0.199***	-0.449***
	(0.030)	(0.038)	(0.067)	(0.048)	(0.072)
G. Groceries					
Log Median CZ Income	0.075***	-0.126***	-0.221***	-0.051**	-0.146***
	(0.014)	(0.014)	(0.025)	(0.019)	(0.027)
H. Healthcare/Medical					
Log Median CZ Income	0.139***	-0.039***	-0.068***	0.100**	0.071*
	(0.037)	(0.004)	(0.007)	(0.036)	(0.035)
I. Hobbies/Entertainment					
Log Median CZ Income	0.428***	-0.055***	-0.096***	0.373***	0.332***
	(0.082)	(0.005)	(0.009)	(0.080)	(0.079)
J. Home Maintenance/Improvement					
Log Median CZ Income	0.089	-0.069***	-0.120***	0.020	-0.031
	(0.131)	(0.006)	(0.010)	(0.130)	(0.129)
K. Office Supplies					
Log Median CZ Income	0.045**	-0.026***	-0.045***	0.019	0.000
	(0.015)	(0.003)	(0.005)	(0.016)	(0.018)
L. Personal Care					
Log Median CZ Income	0.072***	-0.322***	-0.564***	-0.250***	-0.491***
	(0.015)	(0.040)	(0.069)	(0.037)	(0.065)
M. Restaurants/Dining					
Log Median CZ Income	0.067	-0.074***	-0.130***	-0.008	-0.063
	(0.051)	(0.006)	(0.011)	(0.049)	(0.048)
N. Telecommunications					
Log Median CZ Income	0.112	-0.020***	-0.034***	0.092	0.078
	(0.065)	(0.002)	(0.003)	(0.065)	(0.064)
O. Utilities					
Log Median CZ Income	0.317***	-0.249***	-0.435***	0.068	-0.118
	(0.077)	(0.027)	(0.047)	(0.072)	(0.075)

Appendix Table A7: **Alternative Price Indexes: Spatial Dispersion**

	count	mean	sd	min	p1	p5	p10	p25	p50	p75	p90	p95	p99	max
Nested CES Price Index (σ=7)	443	1.055	0.071	0.886	0.930	0.955	0.971	0.999	1.050	1.100	1.146	1.194	1.237	1.289
Nested CES Price Index (σ=11.5)	443	1.019	0.062	0.876	0.905	0.938	0.950	0.975	1.011	1.055	1.094	1.131	1.215	1.256
Laspeyres Price Index, Overall Income	443	0.982	0.099	0.768	0.810	0.856	0.874	0.917	0.969	1.028	1.087	1.174	1.323	1.399
Laspeyres Price Index, Low Income	443	0.971	0.110	0.738	0.783	0.832	0.851	0.898	0.956	1.022	1.089	1.188	1.347	1.440
Laspeyres Price Index, Middle Income	443	0.992	0.092	0.791	0.831	0.875	0.890	0.931	0.982	1.034	1.092	1.169	1.307	1.374
Laspeyres Price Index, High Income	443	1.000	0.072	0.844	0.868	0.906	0.917	0.951	0.993	1.033	1.082	1.142	1.243	1.279
BEA Price Index	443	1.024	0.076	0.860	0.885	0.920	0.936	0.977	1.014	1.060	1.107	1.178	1.285	1.359

	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	0.982	0.099	0.768	1.399	0.971	0.110	0.738	1.440	0.992	0.092	0.791	1.374	1.000	0.072	0.844	1.279
- Laspeyres (CEX weights)	0.987	0.110	0.754	1.456	0.979	0.116	0.733	1.483	0.996	0.112	0.759	1.471	1.013	0.105	0.794	1.442
- CES	0.959	0.089	0.787	1.337	0.947	0.102	0.753	1.403	0.965	0.088	0.764	1.303	0.978	0.067	0.833	1.236
- Nested CES (Feenstra, $\sigma = 7$)	1.055	0.071	0.886	1.289	1.038	0.076	0.893	1.368	1.063	0.074	0.869	1.301	1.045	0.061	0.923	1.250
- Nested CES (Feenstra, $\sigma = 11.5$)	1.019	0.062	0.876	1.256	1.010	0.076	0.870	1.368	1.023	0.064	0.859	1.280	1.013	0.048	0.902	1.172
- PPP Geary-Khamis	0.969	0.085	0.784	1.317	0.961	0.097	0.752	1.372	0.978	0.078	0.807	1.293	0.979	0.062	0.777	1.220
- PPP GEKS-Fischer	0.974	0.088	0.783	1.336	0.965	0.098	0.753	1.379	0.984	0.080	0.807	1.310	0.987	0.062	0.825	1.222
B. Combined indexes: income-group-specific prices																
- Laspeyres	0.983	0.101	0.760	1.373	0.973	0.099	0.738	1.321	0.978	0.086	0.759	1.280	0.956	0.066	0.277	1.133
- Laspeyres (CEX weights)	0.987	0.112	0.744	1.426	0.982	0.104	0.735	1.371	0.975	0.104	0.718	1.349	0.932	0.089	0.146	1.173
- CES	0.959	0.091	0.776	1.321	0.953	0.093	0.747	1.308	0.956	0.083	0.759	1.248	0.944	0.069	0.247	1.127
- Nested CES (Feenstra, $\sigma = 7$)	1.053	0.070	0.898	1.318	1.049	0.070	0.895	1.303	1.048	0.070	0.828	1.293	1.028	0.070	0.682	1.288
- Nested CES (Feenstra, $\sigma = 11.5$)	1.017	0.062	0.882	1.239	1.020	0.070	0.852	1.284	1.009	0.062	0.818	1.225	0.996	0.059	0.510	1.192
- PPP Geary-Khamis	0.969	0.087	0.775	1.298	0.961	0.091	0.742	1.299	0.963	0.082	0.513	1.236	0.952	0.079	0.217	1.152
- PPP GEKS-Fischer	0.975	0.089	0.775	1.316	0.966	0.090	0.748	1.289	0.970	0.081	0.664	1.247	0.948	0.074	0.227	1.134
C. EASI demand index	0.977	0.081	0.842	1.359	0.979	0.093	0.828	1.387	0.979	0.077	0.848	1.354	0.943	0.082	0.024	1.217
D. BEA price index	1.024	0.076	0.860	1.359	1.024	0.076	0.860	1.359	1.024	0.076	0.860	1.359	1.024	0.076	0.860	1.359

Appendix Table A8: **Consumption Against Price Index, City Size, or College Share by Skill Group**

Panel A. Skill-Specific Consumption vs. Price Index

	College Graduates (1)	High School Graduates (2)	High School Dropouts (3)
A1. Consumption Includes Imputed Health Spending and Housing Cost Adjustment (baseline)			
Log price index	-0.039	-0.277***	-0.454***
	(0.052)	(0.030)	(0.037)
A2. Consumption Does Not Include Any Imputation			
Log price index	-0.082	-0.274***	-0.433***
	(0.052)	(0.036)	(0.044)
A3. Consumption Includes Imputed Health Spending, Housing Cost Adjustment, and Government Assistance			
Log price index	-0.046	-0.279***	-0.447***
	(0.054)	(0.033)	(0.040)

Panel B. Skill-Specific Consumption vs. City Size

	College Graduates (1)	High School Graduates (2)	High School Dropouts (3)
B1. Consumption Includes Imputed Health Spending and Housing Cost Adjustment (baseline)			
Log city size	0.005	-0.021***	-0.030***
	(0.003)	(0.002)	(0.003)
B2. Consumption Does Not Include Any Imputation			
Log city size	0.010***	-0.014***	-0.025***
	(0.004)	(0.003)	(0.004)
B3. Consumption Includes Imputed Health Spending, Housing Cost Adjustment, and Government Assistance			
Log city size	0.003	-0.022***	-0.031***
	(0.003)	(0.003)	(0.003)

Panel C. Skill-Specific Consumption vs. College Share

	College Graduates (1)	High School Graduates (2)	High School Dropouts (3)
C1. Consumption Includes Imputed Health Spending and Housing Cost Adjustment (baseline)			
Log college share	0.019	-0.102***	-0.173***
	(0.018)	(0.014)	(0.014)
C2. Consumption Does Not Include Any Imputation			
Log college share	0.021	-0.079***	-0.139***
	(0.022)	(0.024)	(0.026)
C3. Consumption Includes Imputed Health Spending, Housing Cost Adjustment, and Government Assistance			
Log college share	0.014	-0.106***	-0.172***
	(0.018)	(0.014)	(0.015)

Appendix Table A9: **Consumption vs. Price Index, Population, and College Share**

	A1. Laspeyres (1)	A2. Laspeyres (CEX) (2)	A3. CES (3)	A4. Nested CES (11.5) (4)	A5. Nested CES (7) (5)	A6. PPP, Geary-Khamis (6)	A7. PPP, GEKS-Fischer (7)	B1. Laspeyres (8)
Log index	-0.289***	-0.367***	-0.297***	-0.212**	-0.287***	-0.248***	-0.238***	-0.117
	(0.080)	(0.067)	(0.090)	(0.094)	(0.088)	(0.094)	(0.089)	(0.089)
Log index × middle-skill	0.023	0.037	0.029	0.030	0.032	0.030	0.028	0.013
	(0.093)	(0.076)	(0.106)	(0.109)	(0.103)	(0.110)	(0.104)	(0.107)
Log index × low-skill	-0.063	-0.023	-0.059	-0.033	-0.022	-0.062	-0.063	-0.100
	(0.094)	(0.076)	(0.108)	(0.113)	(0.106)	(0.111)	(0.105)	(0.108)
Log city size	0.028***	0.025***	0.024***	0.040***	0.049***	0.029***	0.029***	0.027***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Log population × middle-skill	-0.026***	-0.025***	-0.026***	-0.026***	-0.026***	-0.026***	-0.026***	-0.025***
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Log population × low-skill	-0.037***	-0.036***	-0.037***	-0.040***	-0.041***	-0.038***	-0.037***	-0.035***
	(0.009)	(0.008)	(0.009)	(0.009)	(0.010)	(0.009)	(0.009)	(0.009)
Log college share	0.021	0.014	0.042	0.058	0.076**	0.037	0.031	0.010
	(0.038)	(0.034)	(0.037)	(0.039)	(0.039)	(0.041)	(0.039)	(0.039)
Log college share × middle-skill	-0.040	-0.033	-0.039	-0.039	-0.040	-0.039	-0.039	-0.038
	(0.044)	(0.038)	(0.044)	(0.045)	(0.045)	(0.049)	(0.046)	(0.046)
Log college share × low-skill	-0.041	-0.033	-0.042	-0.050	-0.053	-0.042	-0.042	-0.029
	(0.044)	(0.038)	(0.043)	(0.044)	(0.044)	(0.048)	(0.046)	(0.046)
Middle-skill	-0.028	-0.029	-0.019	-0.031	-0.028	-0.022	-0.026	-0.042
	(0.151)	(0.140)	(0.147)	(0.158)	(0.165)	(0.153)	(0.151)	(0.154)
Low-skill	0.006	0.006	0.005	0.026	0.042	0.012	0.008	-0.012
	(0.151)	(0.139)	(0.145)	(0.160)	(0.169)	(0.153)	(0.152)	(0.155)

	B2. Laspeyres (CEX) (9)	B3. CES (10)	B4. Nested CES (11.5) (11)	B5. Nested CES (7) (12)	B6. PPP, Geary-Khamis (13)	B7. PPP, GEKS-Fischer (14)	C. BEA Price Parities (15)	D. EASI Demand System (16)
Log index	-0.170**	-0.175*	-0.075	-0.170*	-0.145	-0.110	-0.216**	-0.162
	(0.074)	(0.095)	(0.098)	(0.092)	(0.101)	(0.097)	(0.106)	(0.104)
Log index × middle-skill	0.009	0.023	0.040	0.042	0.031	0.027	0.032	0.109
	(0.088)	(0.115)	(0.118)	(0.110)	(0.121)	(0.116)	(0.125)	(0.123)
Log index × low-skill	-0.079	-0.086	-0.034	-0.020	-0.075	-0.082	-0.076	0.021
	(0.087)	(0.117)	(0.122)	(0.115)	(0.123)	(0.117)	(0.126)	(0.124)
Log city size	0.024***	0.025***	0.042***	0.052***	0.027***	0.027***	0.030***	0.032***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Log population × middle-skill	-0.024***	-0.026***	-0.025***	-0.025***	-0.025***	-0.025***	-0.026***	-0.025***
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Log population × low-skill	-0.034***	-0.035***	-0.039***	-0.040***	-0.035***	-0.035***	-0.037***	-0.036***
	(0.008)	(0.009)	(0.009)	(0.010)	(0.009)	(0.009)	(0.009)	(0.009)
Log college share	0.001	0.036	0.049	0.072*	0.030	0.021	0.036	0.028
	(0.034)	(0.039)	(0.039)	(0.039)	(0.042)	(0.040)	(0.040)	(0.041)
Log college share × middle-skill	-0.030	-0.036	-0.036	-0.036	-0.038	-0.038	-0.041	-0.038
	(0.039)	(0.046)	(0.046)	(0.046)	(0.051)	(0.048)	(0.047)	(0.047)
Log college share × low-skill	-0.019	-0.031	-0.040	-0.043	-0.034	-0.032	-0.040	-0.036
	(0.039)	(0.045)	(0.045)	(0.045)	(0.050)	(0.047)	(0.047)	(0.048)
Middle-skill	-0.047	-0.029	-0.041	-0.039	-0.040	-0.042	-0.031	-0.045
	(0.139)	(0.150)	(0.155)	(0.164)	(0.154)	(0.153)	(0.158)	(0.152)
Low-skill	-0.017	-0.011	0.013	0.028	-0.015	-0.016	-0.003	-0.013
	(0.138)	(0.148)	(0.159)	(0.169)	(0.155)	(0.154)	(0.159)	(0.155)