WHERE IS STANDARD OF LIVING THE HIGHEST? LOCAL PRICES AND THE
GEOGRAPHY OF CONSUMPTION

Rebecca Diamond
Enrico Moretti

Where is Standard of Living the Highest? Local Prices and the Geography of Consumption
Rebecca Diamond and Enrico Moretti
NBER Working Paper No. 29533
December 2021, Revised January 2023
JEL No. F1,J00,R00

ABSTRACT

Income differences across US cities are well documented, but little is known about the level of standard of living in each city—defined as the amount of market-based consumption that residents are able to afford. We provide estimates of the standard of living by commuting zone for households in a given income or education group, and we study how they relate to local prices. By combining datasets tracking household spendings including bank and credit card transaction data of 5% of US households in 2014 and NiselenIQ data, we measure mean consumption expenditures by commuting zone and income group. To measure local prices, we build income-specific consumer price indices by commuting zone. We uncover vast geographical differences in material standard of living for a given income level. Low income residents in the most affordable commuting zone enjoy a level of consumption that is 95% higher than that of low income residents in the most expensive commuting zone. We then endogenize income and estimate the standard of living that low-skill and high-skill households can expect in each US commuting zone, accounting for geographical variation in both costs of living and expected income. We find that for college graduates, there is essentially no relationship between consumption and cost of living, suggesting that college graduates living in cities with high costs of living—including the most expensive coastal cities—enjoy a standard of living on average similar to college graduates with the same observable characteristics living in cities with low cost of living—including the least expensive Rust Belt cities. By contrast, we find a significant negative relationship between consumption and cost of living for high school graduates and high school drop-outs, indicating that expensive cities offer lower standard of living than more affordable cities. The differences are quantitatively large: High school drop-outs moving from the most to the least affordable commuting zone would experience a 18.5% decline in consumption.

Rebecca Diamond
Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305
and NBER
diamondr@stanford.edu

Enrico Moretti
University of California, Berkeley
Department of Economics
549 Evans Hall
Berkeley, CA 94720-3880
and CEPR
and also NBER
moretti@econ.berkeley.edu

1 Introduction

Over the last three decades there have been increased differences in income among US communities. Economically vibrant cities, like New York, San Francisco, Boston, and Seattle, have experienced fast increases in mean household income. At the same time, less dynamic local labor markets have experienced more limited increases in income and, in some cases, even declines. What is less clear is how the actual standard of living of residents varies across communities. The standard of living of residents of a city—which we define as the amount of market-based consumption households are able to purchase—depends both on the income level that residents can expect there and the local cost of living. While we know that large, expensive cities tend to have jobs that offer higher nominal earnings, and small, affordable cities tend to have jobs that offer lower nominal earnings, we know little about where market-based consumption is the highest. Are residents of dynamic metro areas better or worse off in terms of consumption compared to residents of smaller, economically struggling communities? This lack of information is surprising, because the amount of market-based consumption is a key component of utility and economic well-being. Models employed to study spatial inequality in utility crucially depend on market consumption, but the existing empirical measures are at best indirect proxies (Moretti, 2013; Diamond, 2016; Giannone, 2017; Monte et al., 2018; Fajgelbaum and Gaubert, 2020; Piyapromdee, 2021). Prior work has assumed that housing is the only locally priced good and housing expenditure and housing prices are sufficient to measure local consumption or it has focused on groceries that can be tracked by scanner data (Handbury, 2019; Handbury and Weinstein, 2015).

Despite the fundamental role of consumption for economic well-being, there is limited direct 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 granular expenditures and prices and that are large enough to allow for a detailed geographical analysis.[2]

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

---

[1] A notable exception to this is Bertrand and Morse (2016) who use the CEX to study consumption of the low-income by state. There is a much larger literature on consumption inequality at the national level—for example see Aguiar and Bils (2015); Attanasio and Pistaferri (2016); Meyer and Sullivan (2017)—and the role of prices —see for example Broda et al. (2009); Jaravel (2019); Hornbeck and Moretti (forthcoming).

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

consumption in affordable commuting zones once we account for geographical variation in both cost of living and expected income. Finally, we study the role played by geographic sorting of households in nationwide consumption inequality.

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

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

To measure local prices, we build consumer price indexes that vary by commuting zone and income group. Our baseline price index is a Laspeyres index which mimics the index used by the BLS to estimate the official national CPI. It is a weighted average of the local prices of items consumed by the average household with income-specific weights reflecting the importance of each item in the bundle for consumers of a given income group. We also examine six alternative price indices based on alternative assumptions, including ones that correct for differences in variety and supply across cities (Handbury and Weinstein, 2015; Handbury, 2019). We augment our data with data on local prices of specific goods that we collected from 11 separate data sources.

Our price indexes point to differences in cost of living across commuting zones that are much larger for low-income than high-income households. The overall cost of living faced by low-income households (post-tax income <$50,000) in the most expensive city—San Jose, CA—is 65% higher than in the median commuting zone, Cleveland, and 95% higher than the most affordable commuting zone—Jonesboro, AR. By contrast, we uncover significantly smaller geographical differences for high-income households (post-tax income >$200,000). The wider spatial variation in cost of living faced by low-income households means that expensive cities like San Francisco, New York or Boston are relatively more expensive for low-income households than high-income households. This finding reflects the larger share of expenditure devoted by low-income households to consumption categories with large spatial variation in prices—housing in particular—and it underscores the importance of using local price indexes that vary not just by location but also by income.

2

We also find that the spatial distribution of cost of living is not symmetric, but highly skewed to the right for all income groups. While the cost of living in most cities is between -20% and +20% of the median city, there are a handful of very expensive cities in the right tail, where cost of living is much higher than the median.

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

We find that geographical differences in material standard of living for a given nominal income level are economically large. The three commuting zones with the lowest consumption of low-income households are San Jose, CA; San Francisco, CA; and Honolulu, HI, with consumption levels between 28% and 34% lower than the median commuting zone. At the other extreme of the spectrum, examples of commuting zones with high consumption of low-income households are Elizabeth City, NC; Traverse City, MI; and Champaign, IL, with consumption levels in real terms 16–29% higher than the median commuting zone. The range of consumption levels observed across U.S. communities is quite wide: Low-income families who live in the most affordable commuting zone enjoy a level of market-based consumption measured in real terms that is 95% higher that of families with the same income who live in the least affordable commuting zone.

We estimate that the elasticity of overall market consumption with respect to the local price index is -0.897 (0.009) and -1.069 (0.039) for low- and high-income households, respectively. The fact that low-income consumers cut consumption less than high-income consumers in response to higher local prices could reflect the fact that the former are closer to a minimum subsistence level, and small consumption cuts cost more in terms of utility.[3] Consistent with this possibility, we find that low-income households in expensive commuting zones have a higher incidence of financial distress than low-income households in affordable commuting zones: they are significantly more likely to have negative saving and pay more overdraft fees.

To validate our findings, we replicate the analysis using model free evidence based on direct and transparent measures of consumption quantities from NielsenIQ. The data contain information on the quantities purchased of specific grocery products measured in physical units, unlike our bank data that measure expenditures. For example, we measure the number of cans of beer, the number of light bulbs, or the number of pounds of nuts purchased in a year by each NielsenIQ consumer. There are 823,507 grocery products, divided in 116 product groups. Consistent with our findings on overall consumption, we find that for 78 out of 116 product groups, NielsenIQ consumers in expensive commuting zones buy fewer physical units, holding income and demographics fixed. The elasticity of consumption with respect to the local price index is lower than the elasticity estimated

---

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

for overall consumption. This likely reflects the fact that grocery items are necessities. When faced with higher cost of living, households seem to cut consumption of necessities less than consumption of all other goods. In addition, groceries exhibit less geographic price variation, making them a relative bargain in expensive cities.

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

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

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

For the high-skill group, San Francisco, New York, and Boston are among commuting zones with the highest pre-tax mean incomes. Accounting for taxation and cost of living reduces the purchasing power of households in these cities by over 40%. However, since pre-tax income is so high, even after accounting for cost of living and taxes, college graduates in San Francisco, New York, and Boston retain a high level of standard of living and remain in the top 20% of the distribution of market consumption across all US cities. Among large cities, Pittsburgh has the highest mean consumption, since it offers good expected income and moderate prices.

Overall, we find that for college graduates, there is no significant bi-variate relationship between expected consumption and cost of living. A regression of expected consumption on the local price index across all commuting zones yields a coefficient of 0.017 (0.058). This suggests that college graduates located in cities with high cost of living enjoy an expected standard of living similar to college graduates with the same observable characteristics located in cities with low cost of living. The reason is that for skilled households, expensive cities offer pre-tax incomes high enough to exactly offset the higher cost of living and personal taxes.

For less skilled households, the picture that emerges is markedly different. San Francisco and New York do offer high pre-tax incomes to high school graduates but not high enough to offset cost of living and taxes (Autor, 2019, 2020). 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 -0.187 (0.029), indicating 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 8.3% decline in their standard of living. Moving from the commuting zone with the lowest cost of living index (Jonesboro) to the commuting zone with the highest index would imply a decline in the standard of living by 10.8%.

The negative relationship between consumption and cost of living is significantly steeper for high school drop outs. The slope is -0.364 (0.035), suggesting that for this group standard of living in expensive commuting zones is quantitatively much lower than in cheaper commuting zones. For households in this group, moving from Cleveland to San Jose implies a 16.1% decline in the standard of living. Moving from Jonesboro to San Jose implies a 20.9% decline in the standard of living.

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

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

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

We conclude the paper by studying the effect of geographic sorting of households into high- and

low-cost commuting zones on income and consumption inequality in the US as a whole (Moretti, 2013). We find that in the absence of sorting, the nationwide mean difference between college graduates and high school graduates would be 19 percent smaller for pre-tax income and 14 percent smaller for consumption than the observed difference.

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

2   Data and External Validity

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

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

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

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

---

[4]Our data provider uses a proprietary algorithm to identify accounts that are active and we drop all inactive accounts.

restrictions leave us with households' main bank accounts, we expect to be missing only 7.8% of the average household's income and expenditures.[5]

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

2.1 Measuring and Validating Income

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

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

In order to compare mean income in each commuting zone obtained from our data to the one obtained from the ACS, in Appendix Figure A1 we plot mean log household income by commuting zone from our data on the y-axis against mean log household income from the 2012-2016 pooled tabulated ACS data. The figure shows a tight relationship between median income measured in the ACS and in our data, with a slope of 0.914 (0.100).[6]

Overall, we conclude that our measure of income appears to be generally consistent with other nationally representative data sources.

In interpreting our estimates below, one should keep in mind that our measure of income has two limitations. First, we cannot observe income that is paid in cash and spent in cash, unless the cash is deposited into the bank account before being spent. Second, we cannot observe some government transfers. Our data include income from Social Security, Disability Insurance, and EITC—since

---

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

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

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. These limitations are likely to be important for households with low levels of true income. This consideration further motivate our restriction to households with income above $10,000. We cannot observe very low income households in our data, so we do not attempt to study their consumption.

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

2.2 Measuring and Validating Consumption Expenditure

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

It is important to benchmark our measure of consumption expenditure against other known measures of consumption expenditure. The most accurate data come from the National Income and Product Accounts (NIPA). 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. The definition of expenditure in NIPA is not exactly the same to the one in our data because NIPA includes among health spending the sum of out-of-pocket health expenses and spending paid by insurers and the government, while our bank data only include out-of-pocket spending. To make the two data more comparable, in Panel A we subtract out non-out-of-pocket health spending from the NIPA expenditure.[8]

Panel A shows that mean expenditure in our bank data closely match the one in NIPA. Our data estimate average household spending at $74,631. NIPA estimates are $77,533, which is only 4% higher. The CEX estimates are 31% lower: $53,495. This is probably not too surprising, since the CEX is a survey-based dataset known to significantly under report spending (Sabelhaus and Groen, 2000; Aguiar and Bils, 2015; Sabelhaus et al., 2015).[9] We will return to healthcare in the

---

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

[8]In this Figure, we estimate non-out-of-pocket health spending in NIPA by taking the NIPA-reported spending on healthcare and net health insurance premiums and multiply it by 0.887, the 2014 share of health care costs that are not out-of-pocket, as measured by the Centers for Medicare & Medical Services (CMS, 2021). There are also some small differences between the CEX health spending measures and the measures in our bank data. The CEX health spending includes out-of-pocket spending, as well as payroll deductions towards health insurance premiums, but does not include any contributions towards health costs directly from employers or the government. This should lead to more health spending in the scope of the CEX than our bank data, but less than in the raw NIPA data.

[9]Our bank data have a slightly different sample than the NIPA data. Our data are restricted to households with bank accounts that earn at least $10k of post-tax income, while NIPA includes all households. According to the SCF, our sample has 12% higher income than the average US households, which would lead our data to have a higher household spending level than NIPA. On the other hand, we miss spending out of unlinked bank accounts from other

next sub-section, where we discuss Panel B.

As a second way to probe the quality of our expenditure data, in Figure 3 we compare the fraction of consumption expenditure by mean of payment in our data with estimates from the Federal Reserve (Greene and Schuh, 2016). On average, our data appear to closely match the corresponding fractions in the general population from the Fed report. The value of all credit card, debit card, and ACH transactions accounts for about 70% of all expenditure in our data, with cash and checks accounting for a smaller fraction. In the figure, we also break down the shares by income group. The Fed does not report these estimates by income.

Third, we compare our measures of expenditure for specific publicly traded merchants to corporate sales reported in SEC filings. SEC filings are precise measures of sales for merchants due to the penalties for misreporting, although corporate sales reported in SEC filings do not need to match exactly our measures of expenditure because they include oversea sales while our measures are only based in the US. Figure 4 shows this comparison for Abercrombie & Fitch, Chipotle, Costco, Dunkin' Donuts, Home Depot, Kroger, Macy's, McDonald's, Nordstrom, Starbucks, Walmart, and Whole Foods. Overall, our expenditure data appear to track sales trends well.

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

Overall, we conclude that our measure of consumption expenditures appears to be generally consistent with other nationally representative data sources.

2.3 Measuring and Validating Location

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

Figure 5 plots the log size of our sample in each commuting zone against the log number of households from 2012-2016 American Community Survey (ACS). There appears to be a tight link between our sample size and the corresponding number of households in the ACS, with an R-squared of 0.81. The slope is 1.340 (0.028), indicating that we under-sample rural areas and over-sample larger cities. This likely reflects the geographical presence of the banks in our sample, which includes the majority of the ten largest banks in the US. These banks' locations are skewed

---

banks of multi-banked households. According to the SCF, un-linked bank accounts likely lead us to miss 7.8% of household spending. Combining these offsetting effects suggests we should overestimate spending by about 4%. We end up underestimating spending by 4%, but this bank-of-the-envelope adjustment suggests that we are in the right ballpark.

to larger, urban areas. We use weights to adjust for sample representativeness, where weights are the ratio of the number of households in a given commuting zone in ACS data to the corresponding number in our data. In practice, this weighting does not impact our results.

2.4 Health and Housing Expenditures Adjustments

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

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

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

For housing, we face the additional limitation that the share of total expenditures that is spent on housing cannot be accurately quantified in our data because many consumers pay their rent with checks and mortgages with bill-pay transfers to banks. While the value of these transactions is included in our measure of total expenditure, our data puts them into a category called "Unclassified". To quantify the share of expenditures that is spent on housing and properly measure housing services for homeowners, we adopt the same methodology and same data that the BLS employs in measuring the CPI (Poole et al. (2005); Bureau of Labor Statistics (2007)). Namely, for renters, we estimate housing expenditures by commuting zone and income group using mean contract rent from the ACS. For homeowners, the BLS uses a measure of "rent equivalent" from the CEX, which is defined as the rental value of their home if they were to rent it out.[10] Details

---

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

are in Appendix B.

After these two adjustments, we return to the comparison of our data to the NIPA and CEX. In Panel B of Figure 2 we compare our adjusted average total expenditure against NIPA and the CEX. (The NIPA data is now in its "raw" format, since we have adjusted our bank data to make the expenditure definition consistent with that used by NIPA.) The adjusted average household expenditure is \$85,446 in the bank data, which is very close to the raw NIPA estimate of \$92,779.[11]

2.5 Categories of Consumption Expenditures

The main focus of our our paper is how overall consumption varies across space. We do not aim to study how the shares of specific consumption categories vary across space. However, as an additional way to validate our expenditure data, Figure 6 compares the composition of consumption expenditure in our data with that in NIPA and the CEX. Our data classify each transaction in 20 high-level consumption categories based on the identity of the merchant.[12] We restrict this comparison to types of expenditure that are measured consistently in all three datasets. The health and housing expenditures that we report in the graph for our data are inclusive of the two adjustments described in the previous subsection.

The Figure shows that our data line up closely with the NIPA. The correlation of spending across categories between NIPA and our data is quite high: 0.98. The correlation with the CEX is lower. This is not surprising since Bee et al. (2012) show that there is substantial variation in the underreporting rate of consumption across types of spending in the CEX creating poorly measured expenditure shares.

Taken together, Figures 2 and 6 indicate that our data match well NIPA data, both in terms of overall mean household consumption and in terms of consumption by category.

2.6 Summary Statistics.

We classify households into three income groups based on unadjusted income: low \$10,000-\$50,000; middle \$50,000-\$200,000; and high >\$200,000. In our empirical analysis, we only include commuting

---

against NIPA. We take rent equivalent for each income group from the CEX, pooling the 2012–2016 data. We estimate average rental payments for renters by income group in the 2012–2016 pooled ACS. We then average these together, weighted by homeownership rates, to get total housing expenditure. To avoid double-counting, we subtract out the actual spending on housing from our unclassified spending and add back our estimated cost of a year of housing services.

[11]Quantitatively, the main reason for why our adjusted average in Panel B is higher than the unadjusted average in Panel A is that the former includes health expenditures that are not out-of-pocket. The housing adjustment for homeowners is quantitatively less consequential on average. In the Figure, we adjust the CEX in a similar way to make it more comparable to NIPA by using imputed rents for homeowners and including healthcare spending paid by employers and the government. The adjusted CEX's mean is at \$66,907, still below NIPA. The baseline CEX average household spending is \$53,495. CEX reported homeowner costs are \$6,149 and estimate imputed rent is 10,896. NIPA estimates healthcare paid by employers at \$2,382 per household and \$6,284 paid by the government. Our adjusted CEX household expenditure is thus: 53,495-6,149+10,895+2,382+6,284=66,907.

[12]These categories only exist for goods purchased by credit card, debit card and ACH. When the purchase is paid for by cash or check, we observe the value but not the type of purchase. These transactions are in a category called "Unclassified". To make this Figure, we apportion unclassified expenditure across the categories proportionally to the household's classified expenditure on each category.

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. Panel B is for adjusted expenditures and income —they are both higher due to the addition of health expenditures that are not out-of-pocket.

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).

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) produces 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). In practice, we will find significant differences in the geographical variation of the price indexes for high- and low-income households.

3.1   Baseline Price Index

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

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

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

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

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

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

(A) Measuring Prices of Consumption Items. To measure prices $p_{i,j}$, we combine data from 11 different datasets.

We use price data from the 2014 NielsenIQ Retail Scanner data for seven consumption categories: *Child/Dependent Expenses, Electronics, General Merchandise, Groceries, Hobbies/Entertainment, Office Supplies, and Personal Care.* The Nielsen price data are detailed, as they provide prices of goods at the twelve-digit barcode level (UPC)—the barcode used by grocery stores at checkout. There are 823,507 distinct UPCs. Each UPC is assigned by Nielsen to a product group, an intermediate classification between UPC and high-level consumption category. Some examples of product group are Milk; Books, Magazines; and Vitamins. To measure the average price of UPCs belonging to each product group in each CZ, we run a UPC-level regression where we regress the price of a UPC code in a given CZ on a UPC fixed effect and a dummy for each community zone. The coefficients on the CZ dummies represent the mean price of a given product group in a given CZ, holding constant the mix UPCs. We use these regressions to predict the CZ-specific mean price for each product group, evaluated at the national bundle of UPC items observed in the NielsenIQ dataset. Quality is held constant since we are comparing the price that consumers in different commuting zones pay for the same bundle. Since NielsenIQ prices are reported before taxes, for

---

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

goods categories that in a given state are subject to sale tax, we add the relevant sales tax.[14]

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

For seven consumption categories—*Automotive Expenses; Telecommunications; Healthcare/Medical; Utilities; Gasoline/Fuel; Clothing/Shoes/Jewellery;* and *Restaurants/Dining*—we collected and homogenized data from various different sources, which are listed below. We provide more details on each of the sources in Appendix C. For goods that in a given state are subject to sale tax, we added the relevant sale tax.

– Automotive expenses: We take a weighted average of the cost of buying a car, maintenance and registration. We obtained the price of ten common models of used cars in the most populated zip code in each commuting zone from quickvalues.com—a service provided by Kelley Blue Book. We assume full depreciation after five years (Meyer and Sullivan, 2008), so the amortized cost of a car is 20% of the relevant sale price. To measure car maintenance costs, we use prices from NielsenIQ data. To measure the cost of car registration, we combine 2013 state vehicle registration fees with total motor-vehicle registrations using data from the Federal Highway Administration. To reflect the fact that purchasing a car is the most expensive part of Automotive Expenses, 95% of the the Automotive expenses index is the sum of the amortized cost of a used car plus the registration fee. The remaining 5% of the automotive expenses index is maintenance costs.

– Telecommunications: We obtained data on average cable TV prices by county in 2014 from the Federal Communications Commission (FCC) through a Freedom of Information request.

– Healthcare/Medical: Our price measure comes from the 2014 Healthy Marketplace Index, published by the Health Care Coverage Institute.

– Utilities: We obtained the 2014 price of a kWh of residential electricity by county from the U.S. Energy Information Administration and the price of a gallon of water by county from the American Water Works Association/Raftelis Financial Consultants 2014 Water and Wastewater Rate Survey. Specifically, we add the price of 893 kWh of electricity (the average household consumption of electricity in the US in 2014) to 7,840 gallons of water (the average household consumption of water in the US in 2014) to obtain the average price of utilities for one household in a month.

– Gasoline/Fuel: We scraped the 2014 price of a gallon of Unleaded Regular gas published by

---

[14]Sales tax data are from 2014 from Walczak and Cammenga (2021).

GasBuddy.

– Clothing/Shoes/Jewellery: We purchased data on 2014 prices from ACCRA, which is collected by the Council for Community and Economic Research.

– Restaurants/Dining: we constructed a price index for 14 popular national restaurant chains using data from Pricelisto. For each chain and county, we observe prices at the menu item level, which is a standardized description of a product offered at multiple locations. We regress the price of a menu item on a commuting zone indicator variable and a menu item fixed effect and predict the mean price for each commuting zone by using the coefficient on the commuting zone indicator variable and the nationwide mean of all menu items. Similarly to the Nielsen data above, this procedure holds quality constant by comparing the price that consumers pay for the same mix of menu items in different locations. We then aggregate the prices of the 14 restaurant chains into one number for each CZ by taking a weighted average of the 14 chains, using as weights the amount of money spent nationally on each of the 14 merchants obtained from our card transaction data. The mean price of non-chain restaurants is unobserved. We assume that CZs where chain restaurants are expensive based on our restaurant price index also have expensive non-chain restaurants, and CZs where chain restaurants are affordable have affordable non-chain restaurants. While this assumption may be violated in practice, the share of total household expenditures spent on non-chain restaurants is likely to be too small to induce significant bias.

An alternative approach to measure prices in these seven consumption categories would be to use Council for Community and Economic Research's ACCRA data, since ACCRA provides local price indexes for all seven categories. In a previous version of this paper, we did so and found that the main empirical results are virtually unchanged. Since ACCRA data sources are different from ours (with the exception of Clothing/Shoes/Jewellery), measurement error in ACCRA prices should be independent of measurement error in our prices. The fact that the results based on our sources are similar to the ones based on ACCRA data lends credibility to our estimates.[15]

For the remaining six consumption categories—*Charitable Giving, Education, Financial Fees, Insurance, Printing and Postage, and Travel*—we have no data on geographical variation in prices. We assume that their prices do not vary geographically. This assumption may be violated in practice and the magnitude of any resulting bias is a function of how important these categories are. The sum of expenditure shares of these items for low-, middle-, and high-income households are 7.59%, 11.95%, and 19.57%, respectively.

(B) Measuring Expenditure Shares. The expenditure shares for 21 high-level consumption categories by income group and are listed in Appendix Table A2. As we have shown in Section 2, these shares are reliable because match well the NIPA shares. To get a sense of how spatial

---

[15]The Council for Community and Economic Research markets its price index to private firms considering a relocation. For our purposes, one limitation of the ACCRA price index is that according to the Council for Community and Economic Research web site, it is "appropriate for professional and managerial households in the top income quintile".

variation in prices contribute to spatial variation of our overall-income index, column 5 reports standard deviation of prices across commuting zones. Categories that exhibit large variation in prices include Automotive Expenses (0.482), Housing (0.295), and Telecommunications (0.358). By contrast, Child/Dependent Expenses (0.029) and Gasoline/Fuel (0.053) have much lower geographical variation.

Four categories among the 21 in our data are very broad: Groceries, General Merchandise, Hobbies/Entertainment, and Personal Care. To improve precision, we use data from the 2014 NielsenIQ Consumer Panel Survey to obtain expenditure shares for more refined product definition nested within each of these three categories. For example, NielsenIQ identifies 17 subcategories within the Personal Care category: Cosmetics, Deodorant, Vitamins, etc. The shares for each subcategory are shown in Appendix Table A3. We use these shares to aggregate prices at the product group level to high-level expenditure category prices.

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. Product variety differs across cities, as some goods exist in some commuting zones but not in others, and it been shown to be quantitatively important for measuring local prices (Handbury and Weinstein, 2015; Handbury, 2019).

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

The Törnqvist index is a geometric average of expenditure shares between the nationwide average and the specific CZ. This has been shown to be a second-order approximation to the

16

true price index. The CES index assumes that underlying preferences within an income group have constant elasticity of substitution across all product categories. That elasticity is implicitly inferred from a transformation of the expenditure shares within each CZ. The benefit of this index is that is an "exact" index, not an approximation. The downside is that it is only exact if the true underlying utility function is CES.

The Nested CES allows for more complex substitution patterns between products. There is an elasticity of substitution between the 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 correcting for variety. To measure local variety we use the number of unique UPC codes sold in each CZ as observed in the NielsenIQ RMS (store sales) data. For product categories not covered by NielsenIQ, we use the number of unique merchants that we observed transacted at within each CZ in our bank data. We explore two choices of the elasticity parameter σ that have been found in the literature: 7 (Montgomery and Rossi, 1999) and 11.5 (Broda and Weinstein, 2010).[16]

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

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

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

---

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

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

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

Table 1 shows the 15 most expensive commuting zones, the 5 commuting zones around the median, and the 15 least expensive commuting zones based on our baseline indexes. Throughout the paper, we label each commuting zone using the name of its largest city, instead of the official commuting zone name. For low income families, the most expensive commuting zones are San Jose, CA; San Francisco, CA; and Honolulu, HI where the low-income price index is 1.653, 1.552, and 1.506, respectively. This implies that prices faced by low-income residents of these cities are 65% to 51% higher than prices faced by low-income residents of Cleveland (which has index equal to 1 by construction). Other expensive commuting zones include San Diego, CA; Santa Barbara, CA; and New York, NY. The least expensive commuting zones for low-income residents are Waycross, GA; Batesville, AR; and Jonesboro, AR, with price indexes equal to 0.862, 0.848, and 0.848, respectively.

The geographical price differences revealed by our price index differences are economically large, and this is particularly true for low income families. The overall cost of living in San Jose is estimated to be 95% and 46% higher than the cost in Jonesboro for low- and high-income households, respectively, suggesting that the range of prices that low income families are exposed to is much wider than the range of prices that high income families are exposed.

Figure 7 displays the spatial dispersion of the our price indexes across all 443 commuting zones. Two features of this figure are particularly interesting. First, and most importantly, the cost of living index for low-income households exhibits significantly higher spatial variation than the index for high-income households. The standard deviation equals 0.115 and 0.069 for the low- and high-income group, respectively. The 75-25 and 90-10 percentile differences for low-income households are much larger than the corresponding differences for high-income households. This implies that expensive cities are even more expensive for low income households. This finding reflects the fact that low-income households put higher weights on housing expenditure, which is the item in the consumption basket whose price varies the most across commuting zones.

Second, the figure shows that the distribution is far from symmetric, but highly skewed to the right for all three income groups. While the mass of the distribution is concentrated between 0.8 and 1.2—indicating that most cities have an index that is between -20% and +20% of the median—there are a handful of expensive cities in the right tail, where cost of living is much higher. For low-income families, there are 32 commuting zones with cost of living that is more than 20% above the median and 14 commuting zones with cost of living that is more than 30% above the median. Similar skewness is present for other income groups.

The consumption item that is most responsible for the spatial variation in the cost of living indexes is housing, since its share of consumption is the largest and its price varies over space more than the price of any other goods. By contrast, product categories with lower shares of consumption and smaller geographical variation in prices—Grocery or Electronics, for example—contribute much less to the spatial variation in the indexes. A regression of the log of the index on log rent yields coefficients of 0.241 (0.008) and 0.434 (0.011) for high- and low-income households, respectively (Appendix Table A4), while the share of housing in the indexes of high- and low-income

households is 0.148 and 0.338, respectively. If the only source of geographical variation in prices of consumption items were housing costs, and all other items had the same price nationwide, we would find the coefficients equal to these shares. The fact that the coefficients are higher reflects the fact that the prices of non-housing nontradables tend to be higher in areas with more expensive land. In turn, this reflects the fact that it costs more to produce nontradable goods and services in areas where land is more expensive (Choi and Jo, 2020). For example, the cost of a haircut or a slice of pizza is higher in San Jose than in Cleveland, holding quality constant, because retail space and labor are more expensive in San Jose.[18]

That said, our findings indicate that local price indices can be well-approximated by using data only on local housing costs, weighted appropriately, especially for the low-income group. The R-squared of the regression of our low- and high-income price indexes on housing rent are 0.95 and 0.89, respectively.

Table A5 shows a regression of the price of non housing consumption categories against rent. The price of most categories is positively correlated with rent. One notable exception is the price of telecommunication, which is negatively correlated with rent. This likely reflect the lack of competition in the cable market in rural and more sparsely populated low rent areas compared to more densely populated high rent areas.

Alternative Indexes. These findings are generally similar if we use our alternative price indexes. The reason is that in practice all indexes are highly correlated with one another. This is shown in Appendix Table A6, 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 (A1 to D), the associated correlation coefficients have a mean of 0.884 and a standard deviation of 0.093. The bottom row in the Table shows that our baseline price index has a correlation of 0.93 with the BEA price index, at least for the geographical areas that are covered by the BEA.

Appendix Table A7 quantifies the spatial dispersion of all our alternative price indexes and it compares it to our baseline index. For all indexes except the income group-specific PPP Geary-Khamis and BEA price index, the low-income index exhibits highest variation, followed by the middle-income index and the high-income index, respectively. Indexes that allow different prices across income groups within a location (Panel B) look generally similar to indices that hold prices fixed across income groups (Panel A).

Differences across commuting zones in the variety of products that are locally available are potentially important. Using data on on grocery products from NielsenIQ, Handbury (2019) and Handbury and Weinstein (2015) have shown that correcting for differences across cities in product variety has a large impact on measured prices. Handbury (2019) shows that the correlation of the variety-corrected price index and city income is negative — so that richer cities have lower effective prices — while the price indices that don't include the variety adjustment show a positive correlation

---

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

of city income. We find similar results with our variety-corrected nested CES index applied to Automotive Expenses, Child/Dependent Expenses, Clothing/Shoes/Jewellery, Gasoline/Fuel, General Merchandise, Groceries, Healthcare/Medical, Office Supplies, Personal Care, Restaurants/Dining, and Utilities, as shown in Appendix Table A8. The table also shows that this negative correlation holds for the overall price index.

The Nested CES index is highly correlated with our baseline index, as shown in Appendix Table A6. Below, we find that our empirical findings are not sensitive to using one or the other.

4 Geographical Differences in Consumption by Income Group

Consumption expenditures are unevenly distributed over space and this geographical variation is particularly pronounced for low-income households. The maps in Appendix Figure A4 show mean consumption expenditures in each commuting zone for which we have data. However, the maps are not informative of local consumption. Differences across areas in consumption expenditures reflect not just the quantity of goods consumed by the area residents, but also local prices. It is possible that areas with high expenditures enjoy a lower consumption than areas with low expenditures if local prices are high enough to more than offset the higher expenditures.

We use our price indexes to deflate expenditures and quantify "real" consumption for each commuting zone and income group. This allows us to study how real consumption varies across commuting zones as a function of local cost of living and compare the standard of living experienced by the residents of expensive and affordable commuting zones, *holding nominal income fixed* (Subsection 4.1). We then examine how the quantity consumed of specific grocery items, measured not in dollars but in physical units or weight vary as a function of local prices as "model free" evidence of consumption differences (Subsection 4.2). We also study how measures of financial distress for low-income families vary as a function of cost of living (Subsection 4.3).

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

4.1 Overall Consumption

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

estimate the following model:

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

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

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

The three commuting zones with the highest consumption for low-income households are Elizabeth City, NC; Traverse City, MI; Champaign, IL. They have level of consumption measured in real terms equal to \$47,498; \$43,119; and \$42,832. Other examples of commuting zones that offer high consumption of low-income households are Huntington, WV and Youngstown, OH. Three cities with the lowest consumption are Honolulu, HI; San Francisco, CA; and San Jose, CA. The corresponding values are \$26,457; \$25,781; and \$24,300. Other examples of commuting zones with low consumption are New York, NY (\$28,460); Washington, DC (\$28,361); Seattle, WA (\$30,043); and Los Angeles, CA (\$28,575).

The geographical differences in standard of living are economically large. Low-income households who live in the top commuting zone in the top group enjoy a level of consumption that is 95% higher than households with the same income who live in the bottom commuting zone in the bottom group.

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

The right panel shows the corresponding estimates for high-income households. Three cities with the highest consumption for this group are Toledo, OH; Pittsburgh, PA; and Erie, PA. They have

---

[19]In this table, to limit the role of sample error, we report empirical-Bayes shrunken estimates. In practice, we calculate $\hat{Y}_i^{shrunk} = \omega_i \cdot Mean(\hat{Y}_i) + (1 - \omega_i) \cdot \hat{Y}_i$, where $\omega_i = SE_{\hat{Y}_i}^2 / (Var(\hat{Y}_i) - Mean(SE_{\hat{Y}_i}^2) + SE_{\hat{Y}_i}^2)$. 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.

level of consumption measured in real terms equal to $290,754; $279,005; and $278,635, respectively. Three cities with the lowest consumption for high-income households are San Diego, CA ($189,633); San Jose, CA ($186,819); and Honolulu, HI ($173,899). Other cities in this category include New York, NY ($208,570); Seattle, WA ($208,271); and San Francisco, CA ($195,965). 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 Honolulu, New York, San Jose, and San Francisco are associated with lower consumption by both high- and low-income households.

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

The elasticities of consumption with respect to income-group-specific local prices are economically large, and confirm large differences in the amount of consumption that households can afford in cheap and expensive communities. Specifically, the elasticities are -0.897 (0.009), -0.983 (0.021), and -1.069 (0.039) for low-, middle-, and high-income households, respectively. This means that a 10% higher cost of living index results in a 10.7% lower consumption for high income households, and 9.0% lower consumption for low income households. The effect of cost of living on consumption needs to be interpreted as an income effect, as opposed to a price effect (assuming that most consumers expect to be in their current city for a long time). If utility is locally homothetic, elasticity of consumption with respect to the price index should be -1 — meaning that a 10% higher cost of living index is equivalent to a 10% lower income, implying a 10% lower consumption. We cannot reject that this holds for high-income consumers and for middle-income consumers, while we can reject it for low-income consumers (p-value = 0.0001).

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

---

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

This finding is generally robust to using alternative price indexes. In particular, Appendix Table A9 shows the elasticities of consumption with respect to income-group-specific local prices based on all our alternative price indexes. In every row, the elasticity for high income households remains significantly higher in absolute value than the elasticity for low income households, confirming that irrespective of the specific index used, the consumption of high-income consumers is more sensitive to local prices than the consumption of low-income consumers.

4.2 Consumption of Specific Goods Measured in Physical Units

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

The sample includes 57,627 households in the 2014 NielsenIQ Consumer Panel data with income above $10,000. A product is defined as a twelve-digit UPC. There are 785,609 UPC codes in the data, in 116 product groups. As UPCs may come in different units within a product group, we convert all UPCs within each product group to have the same unit as the most prevalent or "modal unit" within that product group following Allcott et al. (2019).[21] We assign 0 to households that did not purchase that product in 2014. To allow a comparison of the coefficients across products, we divide the household consumption of each UPC by its nationwide income-specific-group mean, and we use this mean-adjusted quantity as the dependent variable. This allows for an elasticity interpretation. We regress the mean-adjusted quantity of a good purchased by a household in a year on the log of the price index controlling for household income; presence of children; type of residence; household size; head's age, gender, race, marital status, education, and employment status.[22]

We run one regression for each product. To summarize the results, we compute the average of the estimated elasticities for each of the 116 product groups in our data. Table 3 shows some examples. The first row reports results for the consumption of carbonated beverages measured in kilograms per year. Entries in this row are the average elasticity across all types of carbonated beverages. They indicate that the elasticity of consumption of carbonated beverages with respect to the cost of living index is -0.889 (0.094) and -1.234 (0.225) 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.

---

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

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

Of course, it is difficult to draw strong conclusions based on selected examples. Figure 9 plots the distribution of all the 116 estimated elasticities—one for each product group—weighted by average household expenditure on these product groups. Three features of the figure are noteworthy. First, the majority of the coefficients are negative, confirming a lower level of consumption in more expensive commuting zones. Of the 116 coefficients, 65%, 64%, and 70% are negative for low-, middle-, and high-income households.[23] Second, the effects appear to be more negative for high income households than low income households. This indicates that low income households in expensive cities cut consumption of grocery goods less than high income households in expensive cities, consistent with what we observed for overall consumption. The median values for the low-, middle-, and high-income groups are -0.088, -0.138, and -0.434. Third, the elasticities for all three groups are much smaller than the elasticities estimated for overall consumption above. This likely reflects the fact that most of the consumption items in the NielsenIQ data are grocery items and many grocery items are necessities. When faced with higher cost of living, households seem to cut consumption of necessities less than consumption of all other goods. In addition, groceries exhibit less geographic price variation, making them a relative bargain in expensive cities.[24]

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

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

Overall, from Figure 9 and Table 4 we conclude that households in more expensive commuting zones tend to have significantly lower levels of consumption of grocery items and some non-grocery items than households with the same income level in less expensive commuting zones. The evidence in this section measures consumption in physical units and does not depend on deflating

---

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

[24]We have run similar regressions relating the price of each goods to the cost of living index, controlling for the same set of household characteristic indicators. We find that goods are more costly in expensive areas than in cheap areas. This is despite the fact that virtually all grocery goods can be considered traded. The median estimated price coefficients on the index for the low-, middle-, and high-income groups are 0.078, 0.113, and 0.180.

expenditures by a price index. This provides some "model free" evidence of dramatic consumption differences across space and generally confirms the evidence on overall consumption in the previous subsection.

4.3 Negative Savings and Measures of Financial Distress

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

Panel A of Figure 10 shows that low-income households in expensive commuting zones have a higher probability of negative savings than low-income households in cheap commuting zones. The slope is 0.079 (0.010).[26]

Of course, having negative savings in a given year does not necessarily imply financial distress since the household may be borrowing from expected future income to smooth consumption across years. A cross-section is poorly suited to draw strong conclusions on the dynamics of consumption. But Panels B and C in the Figure show a positive relationship between the price index and the share of income spent by low income households on overdraft fees (where overdraft fees are identified from entries in bank account statements); and the existence of bankruptcy fees as a proxy for bankruptcy (where bankruptcy fees are identified as transactions that contain the words "bankrupt").[27] Overall, we conclude that for low-income households, the probability of financial distress appears to be higher in more expensive cities, although the magnitude of these estimates is admittedly hard to interpret.

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

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

But of course the income level that a specific household can attain varies significantly across space: for a given level of human capital, some cities offer high labor earnings (and therefore high

---

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

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

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

income), while others offer low labor earnings (and therefore low income). Ultimately, a household's standard of living in a given city is determined by the relation between the income level that it can achieve there and the local cost of living. In equilibrium, local earnings and cost of living are jointly determined and typically, cities that offer higher labor earnings tend to have higher cost of living, while cities that offer lower labor earnings tend to have lower cost of living.

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

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

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

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

or consumption evaluated at nationwide weighted-average fractions across types. We provide more details in Appendix E.

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

5.1 Standard of Living in the Largest 50 Commuting Zones

The maps in Appendix Figure A5 show the geographical distribution of consumption by skill level. Since the maps are not easy to read, Tables 5, 6, and 7 present our main findings for the largest fifty commuting zones, ordered by pre-tax nominal income for each skill group. For each commuting zone, we report estimates of average household adjusted pre-tax income, post-tax income, and consumption. These estimates hold constant the combination of household characteristics that define a type (education, age, gender, race, Hispanic origin, education, marital status, and children). For each variable, we also report its corresponding percentile among all 443 commuting zones in our data (not just the 50 shown in the table). The last two columns report the difference between pre-tax income and consumption, as a percent difference and as a percentile difference.

Table 5 is for households where the head has a college degree or more. The first three rows show that San Jose, CA; San Francisco, CA; and Washington, DC are the three CZs where high-skill households have the highest expected adjusted pre-tax income: \$143,935; \$139,465; and \$138,555, respectively (column 1). White Plains, NY—a suburb of New York— and New York, NY follow closely. Column 3 reports the corresponding after-tax income obtained by subtracting personal federal and state taxes from column 1. Column 5 shows our estimates of the levels of expected consumption. It quantifies the standard of living that a family with this level of schooling can expect in each commuting zone. For San Jose, San Francisco, Washington DC, and New York, the corresponding values are \$65,873, \$68,213, \$71,218, and \$70,269. The entries in column 5 are substantially lower than column 1 because high-skill residents face a particularly high local cost of living, and, to a lesser degree, because they face high state taxes. But in terms of consumption percentile, the decline for these four cities is modest. In terms of adjusted pre-tax income, these four cities are at the 99th or 100th percentile (column 2), while in terms of consumption San

27

Jose, San Francisco, Washington, DC, and New York drop to 64th, 81st, 94th, and 91st percentile respectively (column 6).

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

Boston consumption level is at the 99th percentile, while the one of Chicago is at the 72nd percentile. Los Angeles and San Diego experience larger drops in relative standings. In terms of pre-tax income, these two cities are at the 98th and 97th percentile respectively, while in terms of consumption percentiles they drop by 53 and 87 to the 45th and 10th percentiles, respectively. Other examples of cities with large negative percentile changes are Portland (-69), Seattle (-51), Minneapolis (-58), Sacramento (-63), and Denver (-65). By contrast, Cincinnati (+11), Cleveland (+13), Pittsburgh (+17), and Buffalo (+10) 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 Pittsburgh because local income is relatively high, cost of living is moderate, and there are no state taxes. Cincinnati follows closely.

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 Charlotte, NC. Similarly, consumption in Washington, DC is found to be similar to Dallas, TX and Baltimore, MD. Overall, across all 443 commuting zones, the standard deviation in adjusted pre-tax income is 10,765, while the standard deviation in consumption is only 4,635—or less than half. This is to be expected if households are at least in part mobile and have a tendency to move toward areas that offer high standard of living.

Table 6 is for households where the head has a high school degree or some college. The picture that emerges is different, in the sense that the most expensive cities appear to offer significant lower consumption than affordable cities. For example, while in terms of adjusted pre-tax income, San Jose, CA and San Francisco, CA remain at the top, in terms of consumption they drop to the bottom quantile of the distribution. New York and Chicago are at the 29th and 16th percentile, respectively. Los Angeles, San Diego, and Denver are in the bottom 15 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 Pittsburgh—due to its low cost of living—followed by Buffalo and Boston.[28] Across all 443 commuting zones, the standard deviation in adjusted pre-tax income is 6,140, while the standard deviation in consumption is only 2,715, 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, Los Angeles, and Denver is in the bottom 10 percent of the distribution. In these cities adjusted pre-tax nominal salaries are higher than in most other commuting zones, but cost of living is so high that low skill residents' standard of living is among the lowest in the nation. Washington, DC and Seattle fare slightly better, although their consumption remains in the bottom quartile of the distribution.

Among the largest 50 commuting zones, the ones that offer the highest standard of living to high school drop outs are Buffalo and Pittsburgh, thanks to low prices.[29] Across all 443 commuting zones, the standard deviation in pre-tax income is 4,692, while the standard deviation in consumption is only 2,720.

5.2 Correlation of Standard of Living with Local Price Indexes

To understand more systematically how standard of living of each skill group varies as a function of local prices, Figure 11 plots household adjusted pre-tax income, post-tax income and consumption as a function of the local cost of living index. The figure includes all 443 commuting zones. Since income, consumption, and local prices are all simultaneously determined, these relationships do not have a causal interpretation. Rather, they need to be interpreted as describing the cross-sectional equilibrium relationship between income, consumption, and local prices.

For all three skill groups, there is a positive correlation between pre-tax income and local cost of living. This is hardly surprising, as more expensive cities have long been known to offer higher equilibrium earnings. Crucially, the elasticity is above 1 for the high-skill group, and below 1 for the middle- and low-skill groups. In particular, the slope is 1.078 (0.062), 0.914 (0.039), and 0.742 (0.047) for high-, middle-, and low-skill households, respectively. This implies that a 10% increase in cost of living is associated with a more than proportional increase in expected pre-tax income for the high skill group —specifically: an increase of 10.8%—and a less than proportional increase in expected pre-tax income for the low skill group—specifically: an increase of only 7.4%.

We also observe a positive relationship for post-tax income. For all three groups, the slopes for post-tax incomes are smaller than the slopes for pre-tax incomes—since expensive cities tend to be located in states with higher income taxes—with the difference in slope larger for the high-skill group than for the low-skill group—due to tax progressivity. The intercept for post-tax income is

---

[28]The three commuting zones that offer the highest standard of living for this group among all commuting zones in the sample are Ketchikan, AK (57,915); Nantucket, MA (57,026); and Oak Bluffs, MA (56,408).

[29]The three commuting zones that offer the highest standard of living for this group among all commuting zones in the sample are Ketchikan, AK (50,732); Summersville, WV (50,222); and Butte, MT (49,274).

lower than the one for pre-tax income—reflecting mean tax burden in the least expensive commuting zones—and the drop in intercept is the largest for high-skill households and minimal for low-skill households—again reflecting tax progressivity.

The findings for consumption are remarkable. For high-skill households, there is essentially no relationship between consumption and cost of living. The coefficient is 0.017 (0.058)—close to zero and not statistically significant at conventional levels. This suggests that college graduates living in cities with high costs of living enjoy a standard of living that is similar to that enjoyed by college graduates with the same observable characteristics living in cities with low cost of living. This appears to be true for the entire range of values observed for the cost of living index, including at the very top of the cost of living distribution. Compared with affordable cities, expensive cities appear to offer incomes high enough to exactly offset the difference in cost of living and personal taxes.

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

The elasticity of consumption with respect to cost of living is -0.187 (0.029), implying that a middle-skill household moving from the median commuting zone (Cleveland) to the most expensive commuting zone (San Jose) would experience a decline in standard of living by 8.3%. Moving from the commuting zone with the lowest cost of living index (Jonesboro) to the commuting zone with the highest index (San Jose) would imply a decline in the standard of living by 10.8%.

The negative relationship between consumption and cost of living is significantly steeper for high school dropouts, suggesting that for this group the standard of living in expensive commuting zones is much lower than in cheaper commuting zones. The slope is -0.364 (0.035), implying vast geographical differences in consumption. Moving from Cleveland to San Jose implies a 15.7% decline in the standard of living. Moving from Jonesboro to San Jose implies a 20.5% 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 (Autor, 2019, 2020). The concern appears particularly serious for the low-skill households, who are increasingly exposed high costs of living and are found to be significantly worse off in terms of market consumption compared to similar households in more affordable areas.

Our findings also have important implications for within-commuting zone inequality. Since consumption of college graduates was found to be uncorrelated with local prices while consumption of less skilled groups was found to decline with local prices, consumption inequality within a commuting zone should increase with local prices. In our data, consumption of college graduates

in San Francisco, San Jose, and New York is 1.96, 1.94, and 1.94 times higher than consumption of high school drop-outs. The corresponding ratios in the three cheapest commuting zones, Jonesboro, AR; Batesville, AR; and Waycross, GA are 1.57, 1.52, and 1.65. The top panel of Figure 12 shows more systematically how the difference in mean consumption between high- and middle-skill households who live in the same commuting zones varies as a function of local cost of living across all commuting zones in the sample. The bottom panel shows the difference in mean consumption between high- and low-skill households. The slopes are 0.204 (0.034) and 0.380 (0.038), 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.[30]

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

In this respect, one possible concern is that our income data misses Food Stamps, TANF, and housing assistance. If low income residents in expensive commuting zones tend to receive more generous transfers than low income residents in affordable commuting zones, this could induce bias in the elasticity of consumption with respect to local prices estimated for less skill households. To assess the magnitude of the problem, we analyze the sensitivity of our estimates of the elasticity of consumption with respect to local prices to including the imputed value of Food Stamps, TANF and housing assistance. For the imputation, we use data from the CPS on the average receipt of Food Stamps, TANF and housing assistance by income level, martial status, number of children and state. Since the CPS has been shown to under-report government transfers, we inflate the reported amounts based on Table 2 in Meyer and Mittag (2019). See Appendix F for more details. Panel B in Appendix Table A10 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 0.017 to 0.008 for high skill, from -0.187 to -0.195 for middle skill, and remains -0.364 for low skill. The main reason for the robustness of our estimates is that Food Stamps, TANF and housing assistance are federal transfers with limited geographical variation and therefore limited effect on our estimated coefficients, which are identified by geographical variation.

Panel C in the same table probes the robustness of our estimates to the adjustments that we made to consumption expenditures. Recall that our expenditure measure includes an adjustment for health expenditures that are not out of pocket; and one for housing costs paid by homeowners. Entries in Panel C show that these two adjustments make the three elasticities more negative.

---

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

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

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

The top panel in Figure 13 presents the results for size, measured by commuting zone population. For all three groups, there is a positive correlation between pre-tax income and population. This is unsurprising, and has been documented by a large literature on the wage premium offered by large cities over small cities and rural areas. What is more interesting is the relationship between consumption and city size. For high-skill households, there is a significant positive relationship between consumption and city size. For middle- and low-skill households, however, there is a significant negative relationship between consumption and city size. Therefore, the standard of living of lower skill households is on average lower in large cities compared to small cities.

The bottom panel focuses on the share of residents with a college degree or more. The positive correlation between pre-tax income and college share is consistent with previous work (Moretti, 2004, 2013). More novel is the relationship between consumption and college share. For high-skill households, there is a significant positive relationship between consumption and college share. For middle- and low-skill households, there is a significant negative significant relationship between consumption and college share indicating that residents in cities with many college graduates enjoy standard of living that is lower than residents in cities with fewer college graduates.

Since cities' prices, population, and college shares are all positively correlated, in Table 8 we investigate a multivariate regression to see which of these characteristics consumption is most related to including the same set of controls used above. We allow the coefficients on each of these three regressors to vary by skill group. In interpreting this table, it is important to keep in mind that local prices, city size, and college share are all simultaneously determined and the Table's entries do not reflect causal estimates, but rather equilibrium relationships. Entries indicate that conditional on city size and college share, the correlation of consumption and the price index is negative for all three groups, with similar elasticities. Interestingly, conditional on prices and college share, the correlation of consumption and city size is positive for the high-skill group, close to zero for the middle-skill group, and negative for the low-skill group. The correlation of consumption and college share appears close to zero for all three groups, but the effects are a bit noisy.

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

Finally, we turn to robustness. In Section 3.3 we found that the alternative price indexes are highly correlated with the baseline index. We now investigate how sensitive are our estimates to the

use of the alternative price indexes. In Appendix Table A11, we report the estimates corresponding to the Table 8 based on all the alternative price indexes. Quantitatively, the coefficients vary somewhat in magnitude, as one might expect. But in most cases, the qualitative picture that emerges is similar to the one in Table 8. In general, we find a negative correlation with the price indexes that is similar across skill groups and a correlation with city size that is positive for the high skill and negative for the less skilled. The coefficients on the price index for the high-skill group range from -0.423 to 0.678, the differential effects for the middle-skill group from 0.045 to 0.176 and differential effects for the low-skill group from -0.080 to 0.124. The coefficients on city size for the high-skill group range from 0.030 to 0.064, differential effects for the middle-skill group from -0.029 to -0.021 and differential effects for the low-skill group from -0.047 to -0.034.

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 65 experience a roughly similar cost of living—namely an index between 0.98 and 1. By contrast, richer households–those with income percentile above 65—are exposed to an index that is exponentially increasing in income.

In this section, we study the implications of geographic sorting into high- and low-cost commuting zones for nationwide inequality (Moretti, 2013). Figure 15 shows the income and consumption differences between high- and middle-skill households (top panel), high- and low-skill households (middle panel), and middle- and low-skill households (bottom panel). The first set of bars—labeled "Baseline"—indicate that the difference in consumption is smaller than the corresponding difference in pre-tax income.[31] In particular, the pre-tax income difference between college graduates and high school graduates is 0.442, while the corresponding difference in consumption is only 0.310, or 30 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 0.672 and 0.469. Namely, the consumption gap is 30 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 0.401 and 0.627, respectively. However, the consumption gaps would remain basically the same at 0.309 and 0.468, respectively. Interestingly, geographic sorting of households across CZs seems to have no impact on nationwide consumption inequality

---

[31]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.

between skill groups.[32]

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 0.359 and 0.571 and the consumption gaps to 0.267 and 0.413.[33] Equalizing population lowers high-middle skill and high-low skill consumption in equality by 14% and 12%, respectively. This is driven by large cities offering higher consumption to the high skill, but lower consumption to middle and low skill households.

Overall, population and demographic sorting across CZs can explain 19% of high-middle skill income gap, 15% of the high-low skill income gap, and 7% of the middle-low skill income gap. In terms of consumption inequality, these forces explain 14% of the high-middle skill consumption gap, 12% of the high-low skill consumption gap and 8% of the middle-low skill consumption gap.

7   Conclusion

We draw two main conclusions. First, we uncover vast geographical differences in material standard of living for a given level of income. Low income residents in the most affordable commuting zone enjoy a level of consumption that is 95% higher that of low income residents in the most expensive commuting zone. When we replicate the analysis focusing on the consumption of specific goods measured in physical units we also find significantly lower consumption in expensive areas.

Second, when we estimate the standard of living that low- and high-skill households can expect in each US commuting zone once we account both for geographical variation in cost of living and also in expected income, we find marked differences between low- and high-skill households. For high-skill households, we find no relationship between expected consumption and cost of living, suggesting that college graduates living in cities with high costs of living enjoy a standard of living generally similar to college graduates living in cities with low cost of living. For high school graduates and high school drop-outs, we find a significant negative relationship between consumption and cost of living, indicating that expensive cities offer lower standard of living than more affordable cities. The differences are quantitatively large. A high school drop-out household moving from the most affordable commuting zone to the most expensive one would experience a 18.5% decline in market consumption.

Establishing the relationship between internal migratory flows and the geography of standard of living in the US, and the precise reasons for the persistence of large difference in standard of living for less educated households should be two primary objectives of future research in this area. Future work should also explore the appropriate model of spatial equilibrium that is consistent

---

[32]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.

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

with our findings. A simplistic version of the Rosen-Roback framework where amenities perfectly offset differences in market consumption across space poorly fits the consumption differences across space that we have uncovered. Through the lens of Rosen-Roback, our finding that the lowest skill households have the largest consumption differences between expensive and cheap cities would indicate that the lowest skilled households have the highest willingness to pay for the amenities available in the most expensive cities. While we have not included amenities in any of our calculations, this possibility appears to be inconsistent with prior work (Diamond, 2016). A richer model with preference heterogeneity within and across these skill groups is likely needed to understand these equilibrium relationships.

References

AGUIAR, M. AND M. BILS (2015): "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review*, 105, 2725–2756.

ALLCOTT, H., R. DIAMOND, J.-P. DUBÉ, J. HANDBURY, I. RAHKOVSKY, AND M. SCHNELL (2019): "Food deserts and the causes of nutritional inequality," *The Quarterly Journal of Economics*, 134, 1793–1844.

AMERICAN WATER WORKS ASSOCIATION AND RAFTELIS FINANCIAL CONSULTANTS, INC. (2022): "2014 Water and Wastewater Rate Survey," .

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

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

ATTANASIO, O. P. AND L. PISTAFERRI (2016): "Consumption Inequality," *Journal of Economic Perspectives*, 30, 3–28.

AUTO REMARKETING (2014): "10 Most-Purchased Models at CarMax," `https://www.autoremarketing.com/trends/10-most-purchased-models-carmax`.

AUTOR, D. (2019): "Work of the Past, Work of the Future," *American Economic Association: Papers and Proceeding*.

——— (2020): "The Faltering Escalator of Urban Opportunity," *Securing Our Economic Future, Melissa S Kearney and Amy Ganz (eds.); Aspen Institute*.

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

BERNARD, D., C. COWAN, T. SELDEN, L. CAI, A. CATLIN, AND S. HEFFLER (2012): "Reconciling medical expenditure estimates from the MEPS and NHEA, 2007," *Medicare & medicaid research review*, 2.

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

BRODA, C., E. LEIBTAG, AND D. E. WEINSTEIN (2009): "The Role of Prices in Measuring the Poor's Living Standards," *Journal of Economic Perspectives*, 23, 77–97.

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

BUREAU OF LABOR STATISTICS (2007): "Handbook of methods," Tech. rep., US Department of Labor, Washington, DC.

CARD, D., , J. ROTHSTEIN, AND M. YI (2021): "Location, location, location," *UC Berkeley*.

CHARLES, K. K., E. HURST, AND N. ROUSSANOV (2009): "Conspicuous Consumption and Race," *The Quarterly Journal of Economics*, 124, 425–467.

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

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

COUTURE, V., C. GAUBERT, J. HANDBURY, AND E. HURST (2019): "Income Growth and the Distributional Effects of Urban Spatial Sorting," *NBER Working Paper No. 26142*.

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

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

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

DEATON, A. AND J. MUELLBAUER (1980): "An Almost Ideal Demand System," *American Economic Review*, 70, 312–326.

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

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

DIAMOND, R., M. J. DICKSTEIN, T. MCQUADE, AND P. PERSSON (2018): "Insurance without Commitment: Evidence from the ACA Marketplaces," Tech. rep., National Bureau of Economic Research.

DYNAN, K. E., J. SKINNER, AND S. P. ZELDES (2004): "Do the Rich Save More?" *Journal of Political Economy*, 112, 397–444.

FAJGELBAUM, P. D. AND C. GAUBERT (2020): "Optimal spatial policies, geography, and sorting," *The Quarterly Journal of Economics*, 135, 959–1036.

FEDERAL COMMUNICATIONS COMMISSION (2022): "Cable Television," `https://www.fcc.gov/media/engineering/cable-television#:~:text=There%20are%20three%20types%20of,service%20a%20subscriber%20can%20buy`.

FEDERAL DEPOSIT INSURANCE CORPORATION (2015): "FDIC National Survey of Unbanked and Underbanked Households, Executive Summary," .

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

GIANNONE, E. (2017): "Skilled-biased technical change and regional convergence," *University of Chicago. Unpublished manuscript*.

GREENE, C. AND S. SCHUH (2016): "The Diary of Consumer Payment Choice," *Federal Reserve Bank of Boston*.

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

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

HEALTH CARE COST INSTITUTE (2022a): "2016 Health Care Cost and Utilization Report: Analytic Methodology," `https://healthcostinstitute.org/images/pdfs/HCCI_2016_Methodology_v1.0_1.23.18.pdf`.

——— (2022b): "2021 Healthy Marketplace Index: Technical Appendix," `https://healthcostinstitute.org/images/pdfs/hmi_2021_technical_appendix.pdf`.

——— (2022c): "Staff," `https://healthcostinstitute.org/staff/`.

HORNBECK, R. AND E. MORETTI (forthcoming): "Estimating Who Benefits From Productivity Growth: Local and Distant Effects of City Productivity Growth on Wages, Rents, and Inequality," *Review of Economics and Statistics*.

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

KELLEY BLUE BOOK (2022): "Frequently Asked Questions," `https://www.kbb.com/faq/used-cars/#question-0`.

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

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

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

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

MEYER, B. D. AND J. X. SULLIVAN (2008): "Changes in the consumption, income, and well-being of single mother headed families," *American Economic Review*, 98, 2221–41.

MONTE, F., S. J. REDDING, AND E. ROSSI-HANSBERG (2018): "Commuting, migration, and local employment elasticities," *American Economic Review*, 108, 3855–90.

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

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

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

PIYAPROMDEE, S. (2021): "The impact of immigration on wages, internal migration, and welfare," *The Review of Economic Studies*, 88, 406–453.

POOLE, R., F. PTACEK, AND R. VERBRUGGE (2005): "Treatment of Owner-Occupied Housing in the CPI," *Paper presented to the Federal Economic Statistics Advisory Committee (FESAC)*.

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

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

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

SABELHAUS, J., D. JOHNSON, S. ASH, D. SWANSON, T. I. GARNER, J. GREENLEES, AND S. HENDERSON (2015): "Is the Consumer Expenditure Survey Representative by Income?" in *Improving the Measurement of Consumer Expenditures*, ed. by C. D. Carroll, T. F. Crossley, and J. Sabelhaus, University of Chicago Press, chap. 8.

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

U.S. ENERGY INFORMATION ADMINISTRATION (2022): "Frequently Asked Questions," `https://www.eia.gov/tools/faqs/faq.php?id=97&t=3`.

WALCZAK, J. AND J. CAMMENGA (2021): "State Business Tax Climate Index," Tax Foundation.

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

Notes: This figure compares the distribution of households post-tax income in our data and the 2012-2016 ACS. In the ACS, we use NBER TAXSIM to calculate income taxes and then subtract it from household pre-tax income, yielding post-tax income. The median (mean) in our data and in the ACS data are $52,956 ($81,011) and $48,837 ($66,452), respectively.

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

(a) Raw Average Consumption Expenditure

(b) Adjusted Consumption Expenditure

Notes: This figure compares average household annual expenditure in 2014 in the CEX, NIPA, and our data. Panel (a) shows average household expenditure from the raw data. The only adjustment done in panel (a) is for NIPA: we subtract healthcare spending paid by insurance companies and the government from the NIPA average total household expenditure. To do this, we take the NIPA total healthcare spending and spending on net health insurance premiums and multiply by it by 0.887, the 2014 share of health care costs that are not out-of-pocket (CMS, 2021). In panel (b), we adjust the data to make health expenditures and housing expenditures more directly comparable. For our bank data, we add healthcare costs paid by the government and insurers; and we adjust housing costs for homeowners. For NIPA, we include all not out-of-pocket expenditure. For the CEX: we adjust housing costs for homeowners. See text for details.

Figure 3: **Consumption Expenditure by Mean of Payment: Bank Data and Fed Data**

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

Figure 4: **Transaction Dollar Amount from Our Data vs. Reported Revenue from SEC Filings**
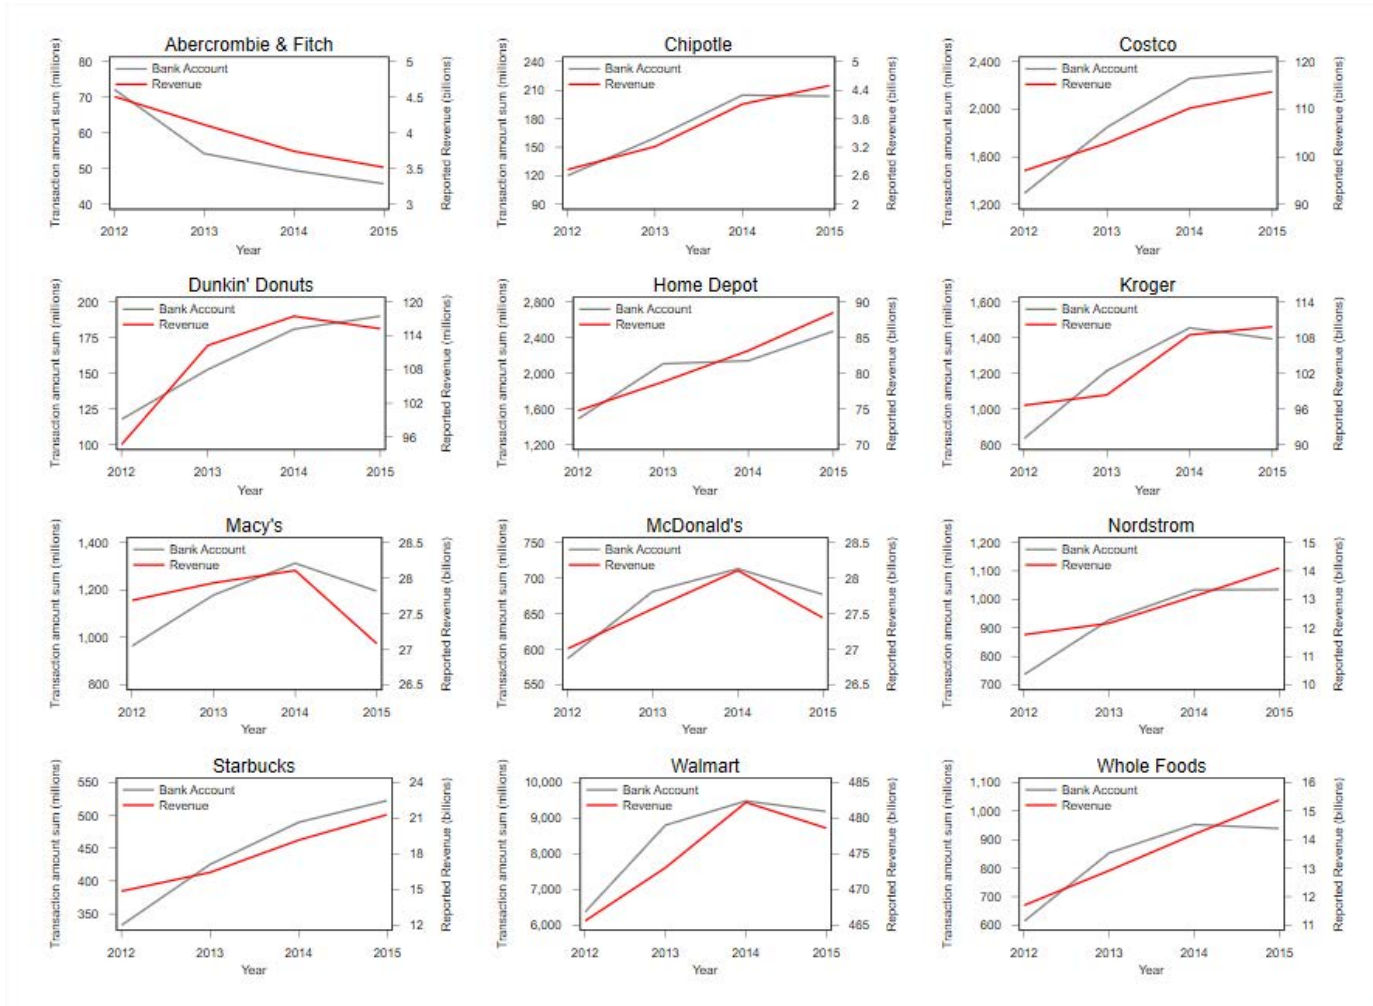
Notes: The gray line is the total value of transactions at each merchant in our bank data, from 2012 to 2015. The red line is the reported revenues from SEC 10K filings from 2012 to 2015.

Figure 5: **Number of Households by Commuting Zone: Bank Data vs. American Community Survey**

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

Figure 6: **Categories of Expenditure: Bank Data, NIPA and CEX**

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

Figure 7: **Spatial Distribution of Price Indexes**
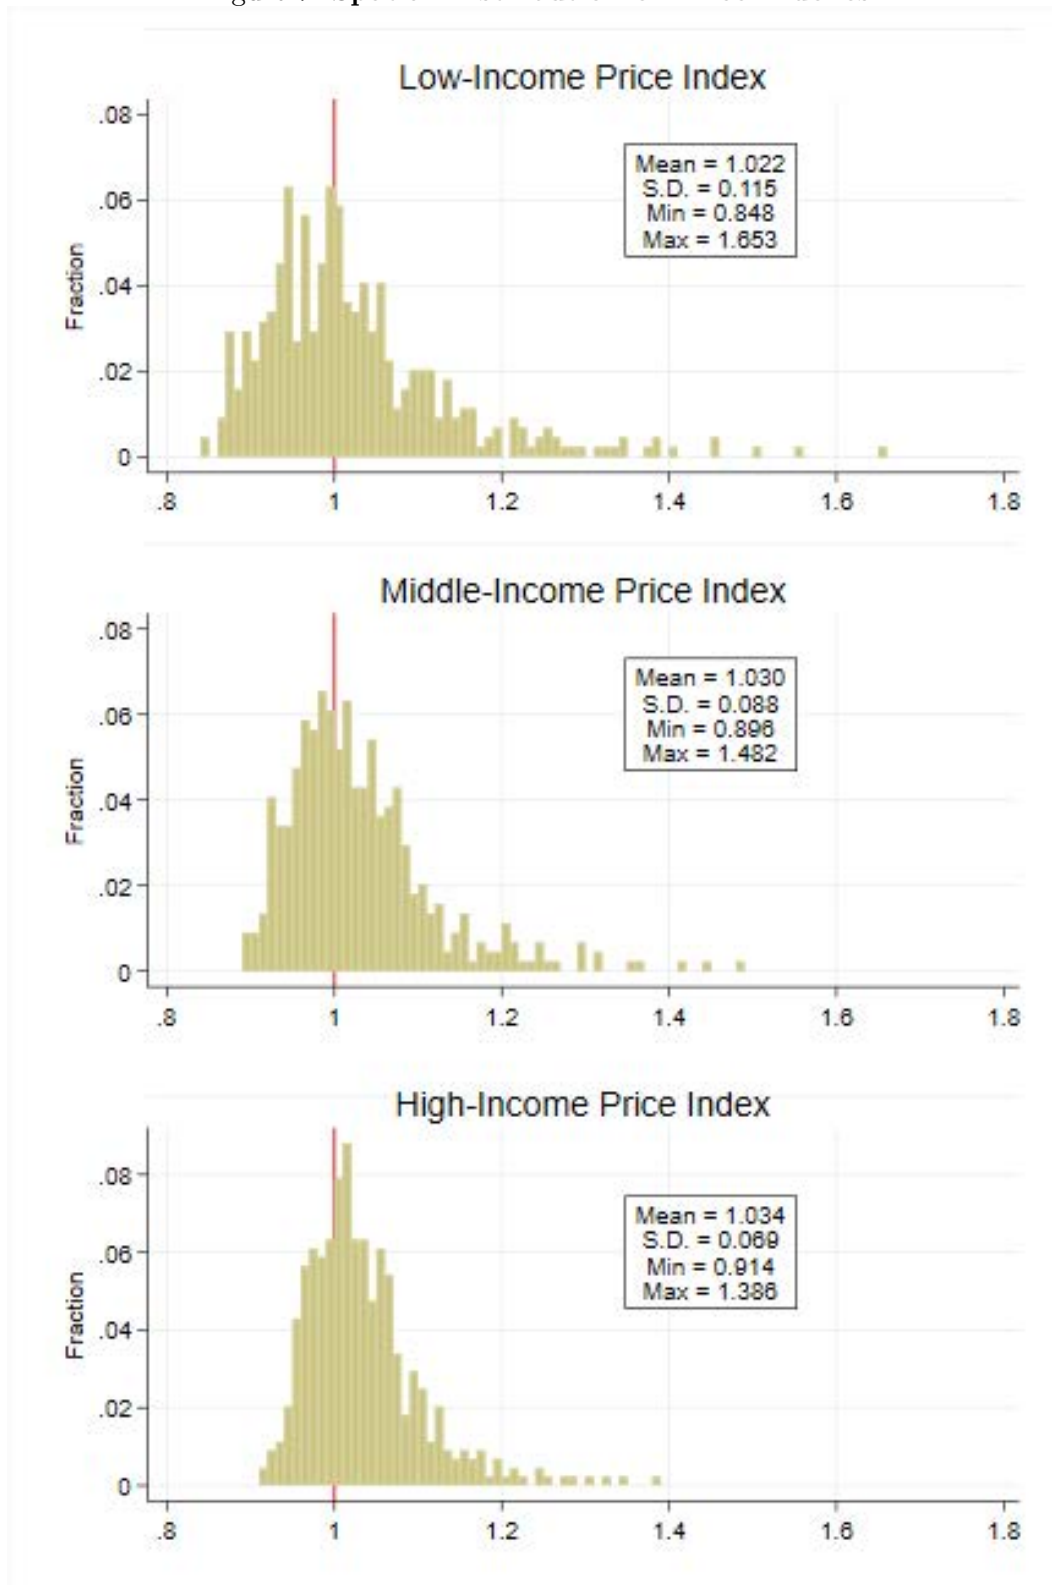
Notes: The level of observation is a commuting zone. $N = 443$.

Figure 8: **Expenditure or Consumption vs. Price Index**
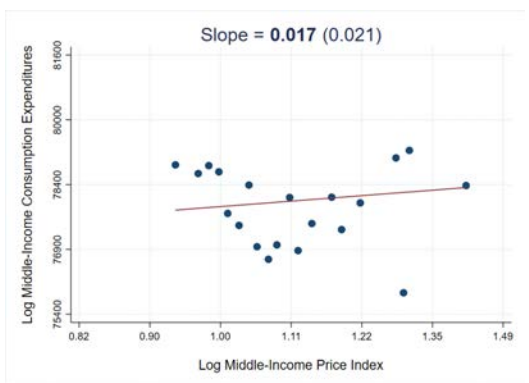
(a) Expenditure, Low Income   (b) Expenditure, Middle Income   (c) Expenditure, High Income

(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. $N = 443$.

47

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

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

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

A. Negative Savings

B. Overdraft Fees

C. Bankruptcy

Notes: Panel A plots the share of low-income households with zero or negative savings as a function of the low income price index. Panels B plots the share of income paid by low-income households on overdraft fees as a function of the low income price index. Panels C plots the share of low-income households who pay bankruptcy fees as a function of the low income price index. Overdraft fees and bankruptcy fees are identified from entries in bank account and credit card statements. Values on x-axis are measured in log scale but we label actual values for easier interpretation. See text for details.

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

Notes: We plot expected adjusted pre-tax income (light blue), adjusted post-tax income (yellow), and consumption (red) on the y-axis against the relevant price index on the x-axis, across 443 commuting zones. Values on both x-axis and y-axis are measured in log scale but we label actual values for easier interpretation.

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

(a) College Grads vs. High School Grads

(b) College Grads vs. High School Dropouts

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

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

(a) **Against City Size**

(b) **Against College Share**

Notes: Panel A: We plot expected adjusted pre-tax income (light blue), adjusted post-tax income (yellow), and consumption (red) against city population, across 443 commuting zones. Panel B: same, but college share is on the x-axis. Values on both x-axis and y-axis are measured in log scale but we label actual values for easier interpretation.

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

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

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

(a) College Grads vs. High School Grads

(b) College Grads vs. High School Dropouts

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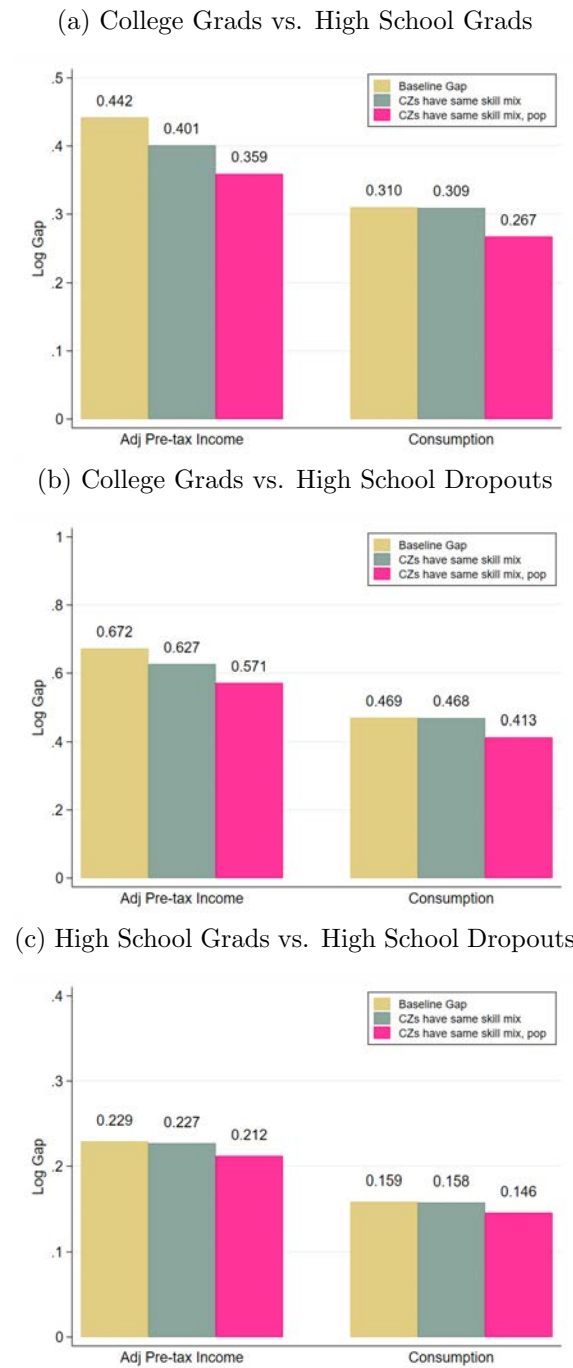
Notes: The Figure shows the adjusted income and consumption differences between high- and middle-skill households (top Panel), high- and low-skill households (Middle Panel), and middle- and low-skill households (bottom panel). The first set of bars refers to the baseline differences. The second next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

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

City	Price Index Low Income	Price Index Middle Income	Price Index High Income
Most Expensive			
San Jose, CA	1.653	1.482	1.349
San Francisco, CA	1.552	1.414	1.307
Honolulu, HI	1.506	1.446	1.386
San Diego, CA	1.459	1.370	1.284
Santa Barbara, CA	1.457	1.357	1.279
New York, NY	1.404	1.311	1.241
Washington, DC	1.384	1.297	1.208
Los Angeles, CA	1.381	1.293	1.215
Anchorage, AK	1.379	1.296	1.252
White Plains, NY	1.344	1.267	1.199
Kapaa, HI	1.341	1.318	1.330
Ketchikan, AK	1.330	1.241	1.220
Newark, NJ	1.326	1.251	1.179
Edison, NJ	1.311	1.228	1.153
Miami, FL	1.290	1.234	1.170
Median			
Twin Falls, ID	1.001	1.004	1.008
Cedar Rapids, IA	1.000	1.032	1.047
Cleveland, OH	1.000	1.000	1.000
Wichita, KS	1.000	1.023	1.044
Gloversville, NY	1.000	0.990	0.998
Least Expensive			
Columbus, MS	0.877	0.923	0.951
Fort Smith, AR	0.877	0.925	0.947
Cape Girardeau, MO	0.876	0.925	0.962
Russellville, AR	0.874	0.913	0.938
Somerset, KY	0.874	0.919	0.948
Harrison, AR	0.873	0.923	0.957
Greenville, MS	0.873	0.924	0.959
Jasper, AL	0.872	0.924	0.956
Pikeville, KY	0.871	0.915	0.940
London, KY	0.868	0.914	0.950
Paducah, KY	0.867	0.897	0.916
Natchez, MS	0.864	0.929	0.962
Waycross, GA	0.862	0.896	0.924
Batesville, AR	0.848	0.898	0.926
Jonesboro, AR	0.848	0.896	0.923

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

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

Low-Income Households			High-Income Households		
City Name (1)	Consumption (2)	Price Index (3)	City Name (4)	Consumption (5)	Price Index (6)
Highest Consumption			**Highest Consumption**		
Elizabeth City, NC	47,498	0.973	Toledo, OH	290,754	0.974
Traverse City, MI	43,119	0.968	Pittsburgh, PA	279,005	1.000
Champaign, IL	42,832	1.015	Erie, PA	278,635	0.995
Huntington, WV	42,127	0.927	Kalamazoo, MI	278,212	1.002
Youngstown, OH	42,054	0.914	Huntington, WV	277,830	0.953
State College, PA	42,015	0.962	Warsaw, IN	277,008	0.980
Kingsport, TN	41,995	0.914	Canton, OH	276,743	0.975
Johnstown, PA	41,840	0.919	Louisville, KY	274,047	0.987
Mobile, AL	41,835	1.007	Cleveland, OH	273,880	1.000
Columbia, MO	41,698	0.928	Lufkin, TX	273,847	0.981
Beaumont, TX	41,566	0.995	Cincinnati, OH	272,882	1.000
Florence, SC	41,321	0.935	Sandusky, OH	272,681	0.999
Morgantown, WV	41,058	0.933	South Bend, IN	271,855	1.001
Lafayette, LA	40,931	0.933	Johnstown, PA	270,797	0.979
Springfield, IL	40,853	0.935	Elizabeth City, NC	269,040	1.012
Median Consumption			**Median Consumption**		
Kalispell, MT	37,035	1.027	Tuscaloosa, AL	239,193	1.020
Lincoln, NE	36,971	1.014	Memphis, TN	239,106	1.003
Lubbock, TX	36,942	1.032	New Orleans, LA	238,852	1.035
Amarillo, TX	36,902	1.037	Gary, IN	238,729	1.028
El Paso, TX	36,894	1.029	Sioux Falls, SD	238,358	1.015
Lowest Consumption			**Lowest Consumption**		
Denver, CO	30,091	1.260	Corpus Christi, TX	210,123	1.091
Edison, NJ	30,059	1.311	Chico, CA	209,872	1.123
Seattle, WA	30,043	1.286	Ogdensburg, NY	209,022	1.044
White Plains, NY	30,007	1.344	Los Angeles, CA	208,901	1.215
Newark, NJ	29,676	1.326	New York, NY	208,570	1.241
Fairbanks, AK	29,659	1.261	Seattle, WA	208,271	1.150
Los Angeles, CA	28,575	1.381	Yuma, AZ	208,119	1.034
New York, NY	28,460	1.404	Santa Barbara, CA	207,731	1.279
Washington, DC	28,361	1.384	Medford, OR	199,553	1.059
Anchorage, AK	27,991	1.379	Anchorage, AK	199,344	1.252
San Diego, CA	27,055	1.459	Olympia, WA	198,274	1.182
Santa Barbara, CA	26,999	1.457	San Francisco, CA	195,965	1.307
Honolulu, HI	26,457	1.506	San Diego, CA	189,633	1.284
San Francisco, CA	25,781	1.552	San Jose, CA	186,819	1.349
San Jose, CA	24,300	1.653	Honolulu, HI	173,899	1.386

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

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

Product	Unit	Low Income $\hat{\beta}_{index}$	$\bar{Y}$	Middle Income $\hat{\beta}_{index}$	$\bar{Y}$	High Income $\hat{\beta}_{index}$	$\bar{Y}$
Beer	KG	0.541*	21.2	0.377	22.5	−0.623	18.7
		(0.326)		(0.346)		(0.547)	
Carbonated Beverages	KG	−0.889***	128.5	−0.950***	129.8	−1.234***	122.1
		(0.094)		(0.143)		(0.225)	
Cookies	KG	−0.323***	6.2	−0.415***	6.3	−0.298	5.8
		(0.093)		(0.150)		(0.224)	
Deodorant	KG	−0.242**	0.3	−0.422***	0.4	−0.993***	0.4
		(0.102)		(0.076)		(0.246)	
Eggs	CT	−0.164**	182.8	−0.143	195.8	−0.499**	186.4
		(0.068)		(0.106)		(0.210)	
Housewares, Appliances	CT	−0.892***	2.2	−1.056***	2.4	−1.592***	2.4
		(0.095)		(0.170)		(0.235)	
Kitchen Gadgets	CT	−0.514***	39.0	−0.201	55.2	1.054**	64.4
		(0.193)		(0.297)		(0.454)	
Laundry Supplies	KG	−0.435***	11.0	−0.427***	12.2	−0.988***	12.0
		(0.099)		(0.127)		(0.176)	
Light Bulbs, Electric Goods	CT	−0.986***	6.7	−1.303***	7.3	−2.083***	7.7
		(0.122)		(0.153)		(0.309)	
Nuts	KG	0.147	3.1	0.286	4.2	−0.311	4.8
		(0.124)		(0.188)		(0.436)	
Pet Food	KG	−0.704***	57.6	−0.766***	55.3	−1.211**	47.7
		(0.127)		(0.203)		(0.474)	
Pizza, Snacks - Frozen	KG	−0.614***	5.8	−0.805***	6.0	−1.303***	5.6
		(0.150)		(0.196)		(0.283)	
Stationery, School Supplies	CT	−0.538**	228.4	−0.453	278.5	−0.482	275.6
		(0.213)		(0.283)		(0.525)	
Vegetables - Frozen	KG	−0.280**	10.0	−0.446***	11.2	−1.134***	10.0
		(0.133)		(0.164)		(0.279)	

Notes: Entries are from regressions of mean-adjusted quantity of consumption on the local price index controlling for household income; household size; age and presence of children; type of residence; household composition; household head's characteristics including age, gender, race, marital status, education, employment status, and education. We average elasticities by product group. The analysis is based on 57,627 households in the 2014 NielsenIQ Consumer Panel data with at least \$10,000 annual income. The numbers of households by income group are 25,265 low income households, 22,753 middle income households, and 9,609 high income households. Robust standard errors are clustered by commuting zone and reported in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.01$.

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

Product Group	Low Income $\hat{\beta}$	Low Income $\bar{Y}$	Middle Income $\hat{\beta}$	Middle Income $\bar{Y}$	High Income $\hat{\beta}$	High Income $\bar{Y}$
Cable and Satellite Services	-0.169 (0.187)	3.3	-0.444* (0.235)	6.0	0.485** (0.232)	5.5
DHL/FedEX/UPS/USPS	-0.223** (0.108)	2.2	-0.839*** (0.191)	3.8	-1.568*** (0.273)	4.2
Gasoline/Fuel	-1.506*** (0.263)	24.1	-1.796*** (0.274)	35.1	-1.295*** (0.371)	23.1
Gym/Fitness/Yoga	0.451 (0.277)	2.2	-0.022 (0.309)	3.4	-0.601 (0.446)	3.6
Hulu/Netflix/Google Play	-1.028*** (0.114)	5.1	-1.222*** (0.217)	6.3	-1.320*** (0.306)	4.0
Movies	-1.044*** (0.309)	0.7	-1.585*** (0.399)	1.1	-1.850*** (0.543)	0.8
Pedicure and Manicure	0.434** (0.193)	0.7	0.172 (0.281)	1.2	1.368 (0.843)	1.5

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

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

City	Adj Pre-tax Income		Adj Post-tax Income		Consumption		Difference between Adj Pre-tax Income & Consumption	
	value (1)	pct (2)	value (3)	pct (4)	value (5)	pct (6)	per (7)	pct (8)
1. San Jose, CA	143,935	100	111,323	100	65,873	64	−54%	−36
2. San Francisco, CA	139,465	100	108,985	100	68,213	81	−51%	−19
3. Washington, DC	138,555	100	105,959	100	71,218	94	−49%	−6
4. White Plains, NY	130,657	99	102,273	99	73,231	98	−44%	−1
5. New York, NY	130,567	99	102,409	99	70,269	91	−46%	−8
6. Newark, NJ	127,971	99	101,405	99	72,203	97	−44%	−2
7. Boston, MA	125,160	99	99,353	99	74,637	99	−40%	+0
8. Hartford, CT	124,391	98	98,270	98	73,374	98	−41%	+0
9. Philadelphia, PA	117,250	98	93,887	98	73,889	98	−37%	+0
10. Los Angeles, CA	116,548	98	92,502	97	63,636	45	−45%	−53
11. Baltimore, MD	116,400	97	91,300	97	70,880	93	−39%	−4
12. Houston, TX	115,878	97	94,099	98	74,223	99	−36%	+2
13. San Diego, CA	114,430	97	90,781	96	58,110	10	−49%	−87
14. Chicago, IL	112,847	96	88,815	95	66,708	72	−41%	−24
15. Dallas, TX	110,199	95	90,026	96	71,271	95	−35%	+0
16. Camden, NJ	110,058	95	88,452	95	67,492	77	−39%	−18
17. Seattle, WA	110,021	95	90,254	96	63,426	44	−42%	−51
18. Sacramento, CA	107,682	94	85,967	93	61,643	31	−43%	−63
19. West Palm Beach, FL	106,739	94	88,345	95	70,364	92	−34%	−2
20. Denver, CO	106,517	93	85,254	92	61,173	28	−43%	−65
21. Austin, TX	106,408	93	87,620	94	66,236	67	−38%	−26
22. Providence, RI	106,211	93	86,032	93	66,708	72	−37%	−21
23. Atlanta, GA	104,703	92	82,065	90	66,212	67	−37%	−25
24. Minneapolis, MN	103,540	91	82,160	90	61,825	33	−40%	−58
25. Fort Worth, TX	103,097	91	84,586	91	69,415	88	−33%	−3
26. San Antonio, TX	101,637	90	83,251	91	67,866	79	−33%	−11
27. Detroit, MI	101,437	89	80,855	88	70,554	93	−30%	+4
28. Miami, FL	100,863	89	83,902	91	63,168	42	−37%	−47
29. Cincinnati, OH	100,696	88	80,800	88	74,741	99	−26%	+11
30. Portland, OR	100,467	88	79,194	84	59,850	19	−40%	−69
31. St. Louis, MO	100,425	88	79,977	87	69,001	86	−31%	−2
32. Virginia Beach, VA	99,661	87	78,719	83	60,230	21	−40%	−66
33. Charlotte, NC	99,067	87	78,048	82	65,891	64	−33%	−23
34. Raleigh, NC	98,909	87	77,992	81	64,868	55	−34%	−32
35. Nashville, TN	98,219	86	81,093	89	67,715	79	−31%	−7
36. Columbus, OH	98,128	85	79,257	85	70,899	94	−28%	+9
37. Cleveland, OH	97,734	85	79,357	86	73,597	98	−25%	+13
38. Phoenix, AZ	97,683	85	78,919	83	65,266	59	−33%	−26
39. Pittsburgh, PA	96,849	83	78,636	83	75,072	100	−22%	+17
40. Kansas City, MO	96,175	82	77,044	78	62,918	39	−35%	−43
41. Jacksonville, FL	95,879	81	79,192	84	64,678	53	−33%	−28
42. Las Vegas, NV	95,795	81	79,354	85	63,129	41	−34%	−40
43. Harrisburg, PA	95,194	80	77,365	80	67,119	75	−29%	−5
44. Tampa, FL	94,224	79	78,267	82	62,850	39	−33%	−40
45. Milwaukee, WI	93,989	78	74,364	68	60,988	26	−35%	−52
46. Indianapolis, IN	93,709	77	75,500	74	67,355	76	−28%	−1
47. Salt Lake City, UT	93,632	77	74,806	71	60,771	25	−35%	−52
48. Buffalo, NY	93,120	75	75,703	75	68,813	85	−26%	+10
49. Orlando, FL	91,686	71	76,526	77	59,229	16	−35%	−55
50. Grand Rapids, MI	85,593	44	69,629	42	64,879	56	−24%	+12

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

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

City	Adj Pre-tax Income		Adj Post-tax Income		Consumption		Difference between Adj Pre-tax Income & Consumption	
	value (1)	pct (2)	value (3)	pct (4)	value (5)	pct (6)	per (7)	pct (8)
1. San Jose, CA	85,629	100	71,716	100	42,774	7	−50%	−93
2. Washington, DC	84,271	100	68,508	99	46,893	56	−44%	−44
3. San Francisco, CA	82,953	100	69,442	100	44,129	18	−47%	−82
4. New York, NY	76,419	98	64,267	98	44,968	29	−41%	−69
5. White Plains, NY	76,109	98	64,263	98	47,002	58	−38%	−40
6. Newark, NJ	75,980	98	64,535	98	46,719	54	−39%	−44
7. Boston, MA	74,897	97	63,377	98	49,094	85	−34%	−12
8. Hartford, CT	73,803	96	62,655	97	47,528	65	−36%	−31
9. San Diego, CA	73,711	96	62,379	96	41,060	1	−44%	−95
10. Los Angeles, CA	73,176	96	61,955	96	43,553	13	−40%	−83
11. Baltimore, MD	72,310	95	60,126	95	48,342	78	−33%	−17
12. Camden, NJ	71,391	95	60,820	95	47,833	71	−33%	−24
13. Seattle, WA	71,095	95	61,199	95	44,459	23	−37%	−72
14. Philadelphia, PA	69,680	94	59,233	93	47,395	63	−32%	−31
15. Sacramento, CA	68,895	93	58,806	93	43,646	14	−37%	−79
16. Denver, CO	68,816	93	58,198	92	42,876	8	−38%	−85
17. Chicago, IL	68,122	92	57,008	91	43,943	16	−35%	−76
18. Austin, TX	67,500	92	58,390	93	45,505	35	−33%	−57
19. Minneapolis, MN	67,367	91	56,850	90	44,101	18	−35%	−73
20. Houston, TX	67,335	91	57,978	92	47,898	72	−29%	−19
21. Providence, RI	67,257	91	57,787	92	46,833	55	−30%	−36
22. West Palm Beach, FL	66,122	89	57,660	91	46,639	53	−29%	−36
23. Dallas, TX	65,811	89	56,882	90	46,578	51	−29%	−38
24. Fort Worth, TX	65,370	88	56,438	88	47,829	70	−27%	−18
25. Virginia Beach, VA	65,266	87	54,644	81	43,886	16	−33%	−71
26. Portland, OR	65,156	87	54,797	82	42,522	6	−35%	−81
27. Salt Lake City, UT	64,492	86	54,824	83	45,683	39	−29%	−47
28. Phoenix, AZ	63,831	84	54,528	81	46,504	50	−27%	−34
29. Las Vegas, NV	63,776	84	55,126	85	45,451	35	−29%	−49
30. Harrisburg, PA	63,756	84	54,457	80	48,414	79	−24%	−5
31. San Antonio, TX	63,417	83	54,831	83	46,183	46	−27%	−37
32. Atlanta, GA	63,022	82	52,832	74	44,346	22	−30%	−60
33. Miami, FL	62,731	81	54,941	84	42,609	7	−32%	−74
34. St. Louis, MO	62,051	78	52,781	74	47,411	64	−24%	−14
35. Buffalo, NY	61,267	75	53,018	76	49,784	89	−19%	+14
36. Detroit, MI	61,262	75	52,152	71	46,885	56	−23%	−19
37. Kansas City, MO	61,253	74	52,070	71	44,296	20	−28%	−54
38. Nashville, TN	60,820	72	53,066	76	45,947	43	−24%	−29
39. Raleigh, NC	60,618	71	51,125	62	44,512	23	−27%	−48
40. Jacksonville, FL	60,337	69	52,496	72	44,773	26	−26%	−43
41. Cincinnati, OH	60,225	68	51,678	67	48,915	84	−19%	+16
42. Milwaukee, WI	60,009	67	50,951	60	43,039	8	−28%	−59
43. Pittsburgh, PA	59,785	66	51,543	67	50,502	93	−16%	+27
44. Cleveland, OH	59,693	66	51,533	66	48,784	82	−18%	+16
45. Tampa, FL	59,254	64	51,919	70	43,629	14	−26%	−50
46. Charlotte, NC	59,175	64	50,106	54	43,866	16	−26%	−48
47. Orlando, FL	59,157	63	51,923	70	41,910	3	−29%	−60
48. Columbus, OH	58,869	61	50,794	59	46,828	55	−20%	−6
49. Indianapolis, IN	58,429	56	49,888	52	45,646	37	−22%	−19
50. Grand Rapids, MI	55,427	32	47,939	32	46,088	46	−17%	+14

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

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

City	Adj Pre-tax Income		Adj Post-tax Income		Consumption		Difference between Adj Pre-tax Income & Consumption	
	value (1)	pct (2)	value (3)	pct (4)	value (5)	pct (6)	per (7)	pct (8)
1. San Jose, CA	63,533	99	56,160	99	33,915	1	−47%	−98
2. Washington, DC	62,939	99	54,147	99	37,953	21	−40%	−78
3. San Francisco, CA	61,246	99	53,981	98	34,668	2	−43%	−97
4. Newark, NJ	59,284	98	52,713	98	38,718	32	−35%	−66
5. White Plains, NY	58,102	97	52,004	98	38,371	27	−34%	−70
6. Boston, MA	57,837	97	51,418	97	41,252	72	−29%	−25
7. New York, NY	57,312	96	51,175	97	36,307	7	−37%	−89
8. Camden, NJ	56,357	95	50,430	95	40,291	57	−29%	−38
9. Hartford, CT	55,984	95	50,251	95	39,302	40	−30%	−55
10. Seattle, WA	55,750	94	49,863	95	37,193	14	−33%	−80
11. San Diego, CA	54,723	94	48,998	94	32,993	0	−40%	−94
12. Los Angeles, CA	54,138	93	48,470	93	34,777	2	−36%	−91
13. Denver, CO	54,040	93	47,762	93	35,522	4	−34%	−89
14. Harrisburg, PA	53,430	93	47,150	91	42,160	81	−21%	−12
15. Baltimore, MD	53,382	93	46,881	90	39,034	36	−27%	−57
16. Chicago, IL	53,352	92	46,670	89	36,736	10	−31%	−82
17. Portland, OR	52,918	92	46,480	89	36,885	12	−30%	−80
18. Salt Lake City, UT	52,758	91	46,856	90	40,177	55	−24%	−36
19. Providence, RI	52,689	91	47,364	92	39,406	42	−25%	−49
20. Philadelphia, PA	52,662	91	46,960	91	38,101	23	−28%	−68
21. Minneapolis, MN	52,339	90	46,676	90	36,545	9	−30%	−81
22. Sacramento, CA	51,454	88	46,386	88	35,227	3	−32%	−85
23. Las Vegas, NV	50,948	86	45,687	84	38,624	30	−24%	−56
24. West Palm Beach, FL	50,651	84	46,103	86	38,186	25	−25%	−59
25. Virginia Beach, VA	50,452	83	44,510	77	37,512	17	−26%	−66
26. Austin, TX	50,334	82	45,395	83	36,376	8	−28%	−74
27. Houston, TX	49,516	78	44,638	79	38,264	25	−23%	−53
28. Dallas, TX	49,199	77	44,505	77	37,565	17	−24%	−60
29. Buffalo, NY	49,067	76	44,658	79	43,136	89	−12%	+13
30. Fort Worth, TX	48,893	74	44,095	74	38,332	26	−22%	−48
31. Kansas City, MO	48,757	73	43,304	67	37,966	21	−22%	−52
32. Phoenix, AZ	48,574	70	43,625	71	38,398	27	−21%	−43
33. Detroit, MI	48,335	67	43,161	67	39,685	47	−18%	−20
34. Miami, FL	48,259	67	44,159	75	35,331	3	−27%	−64
35. Milwaukee, WI	48,209	67	42,942	64	36,843	12	−24%	−55
36. St. Louis, MO	48,009	65	42,888	63	39,802	49	−17%	−16
37. Pittsburgh, PA	47,582	62	42,860	63	42,550	84	−11%	+22
38. San Antonio, TX	47,240	60	42,767	62	37,337	14	−21%	−46
39. Atlanta, GA	47,226	60	42,068	54	36,325	7	−23%	−53
40. Orlando, FL	47,121	59	43,073	65	35,613	4	−24%	−55
41. Nashville, TN	47,075	58	42,939	64	37,832	20	−20%	−38
42. Grand Rapids, MI	46,681	55	42,030	53	41,073	69	−12%	+14
43. Tampa, FL	46,680	55	42,642	61	36,679	10	−21%	−45
44. Cleveland, OH	46,526	52	41,890	52	40,168	54	−14%	+2
45. Jacksonville, FL	46,440	51	42,368	58	37,927	21	−18%	−30
46. Cincinnati, OH	45,751	44	41,160	42	39,771	49	−13%	+5
47. Charlotte, NC	45,680	44	40,803	38	37,187	13	−19%	−31
48. Raleigh, NC	45,363	40	40,682	36	36,755	11	−19%	−29
49. Columbus, OH	45,225	38	40,975	40	38,547	30	−15%	−8
50. Indianapolis, IN	45,177	37	40,494	33	37,384	15	−17%	−22

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

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

	Log Consumption
Log price index	-0.289***
	(0.077)
Log price index × middle-skill	0.044
	(0.087)
Log price index × low-skill	-0.048
	(0.089)
Log city size	0.032***
	(0.008)
Log population × middle-skill	-0.026***
	(0.009)
Log population × low-skill	-0.039***
	(0.009)
Log college share	0.046
	(0.040)
Log college share × middle-skill	-0.034
	(0.045)
Log college share × low-skill	-0.027
	(0.046)
Middle-skill	-0.061
	(0.150)
Low-skill	-0.054
	(0.150)

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

Online Appendix

A Household Sample in Bank Account Data

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

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

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

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

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

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

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

B Healthcare and Housing Adjustments

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

spending.[34]

Specifically, we take the following steps.

First, the MEPS is a household survey that is known to under-report spending relative to NIPA (Bernard et al., 2012). We therefore inflate MEPS spending by a factor of 1.32 to obtain adjusted health spending so that overall healthcare spending matches the reported NIPA healthcare spending. We regress adjusted total healthcare spending on adjusted OOP healthcare spending, controlling for household income and region of the country, and the interactions of these terms. The coefficient on OOP spending is 5.232 (0.522), indicating that each additional dollar of OOP spending (excluding prescriptions) corresponds on average to $5.23 of total expenditure. A similar regression for prescription expenditures yields a coefficient equal to 3.633 (1.179).

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

Finally, for each household in the bank data, we use our estimated coefficients to impute total drug and non-drug health expenditure for a given level of observed OOP health expenditures. The OOP non-drug health expenditure mean is $889 and the total non-drug health expenditure mean is $6,021. The medians are $259 and $3,992. For drug expenditure, the means are $1,324 and $6,109, and the medians are $485 and $2,690.

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

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

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

These are the specific steps that we have taken:

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

---

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

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

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

Third, we need to assign these estimated housing expenditures as measured in the ACS to our bank transaction households. We match households in our bank data to those in the ACS based on income and commuting zone. Specifically, we regress our imputed housing costs in the ACS (both measures, separately) on post-tax income by commuting zone × income group. Then, we use the estimates to predict both types of housing costs for all households in our bank account data. When housing costs to be subtracted exceed unclassified spending in the bank account data, we adjust the housing cost to be subtracted to equal the unclassified spending. For reference, mean housing costs to be subtracted are $16,539 across all income groups; $9,360 for low-income; $18,791 for middle-income; and $52,705 for high-income households. Mean housing costs to be added are $18,650 across all income groups; $13,290 for low-income; $20,457 for middle-income; and $44,632 for high-income households.

C   Price Indexes

We construct price indexes for the 443 commuting zones covered by our linked bank and credit card transaction data. The company that provided us with the bank and card transactions data has categorized expenditures into 20 high-level categories, and we use these as a guide to the categories that we select for our price index. We measure the prices of goods and services belonging to these 20 high level categories plus housing. The categories are listed in Appendix Table A2. To obtain a price index for each CZ and income group, we combine the prices of these 21 categories using their relative expenditure shares. We measure expenditure shares by income group and, in some specifications, by income group–commuting zone.

Here we describe in detail how we measure prices and expenditure shares.

C.1   Measuring Prices of Good and Services

Child/Dependent Expenses; Electronics; General Merchandise; Groceries; Hobbies /Entertainment; Office Supplies; and Personal Care. For these seven consumption categories, we use price data from the 2014 Nielsen Retail Scanner data The Nielsen data contain all UPCs purchased and recorded by Nielsen-participating households in a given year. We merge in product details (e.g., department, product group, product module, size, and unit) and household characteristic indicators (e.g., household income; household size; age and presence of children; type

of residence; household composition; household head's characteristics including age, gender, race, martial status, education, employment status, and education; and the commuting zone they lived in 2014).

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

We calculate commuting-zone-specific prices at the product group level. In particular, for each product group, we regress UPC price on commuting zone indicators with UPC fixed effects, weighting observations by household spending on the UPC. We estimate:

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

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 $\delta_{p,j}$ added to δ_u 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 UPC price on commuting zone indicators, absorbing UPC fixed effects and household income group indicators:

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

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 $\delta_{p,j,k}$ added to δ_u 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.

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

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

We then estimate mean rents controlling for observable housing characteristics. In particular, we interact the following five housing characteristics to define "housing types" (n): (1) Year the

4

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

For each commuting zone and for each housing type, we calculate mean monthly contract rent among all observed renter-occupied units, using household weights in the ACS data.[35] 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 $\text{rent}_{j,k} = \sum_{n=1}^{N}(\text{rent}_{n,j,k} \times \bar{s}_{n,j,k})$.

Automotive Expenses. To measure automotive expenses, we combine three separate data sources: car registration prices from the Federal Highway Administration; used car prices from Kelley Blue Book; and maintenance costs from Nielsen IQ.

We estimate the cost of used cars using quickvalues.com, a service provided by Kelley Blue Book that provides historical data on the price of a used car in a particular zip code. For each commuting zone, we look up the Fair Purchase Price in the most populated zip code. Kelley Blue book defines the Fair Purchase Price as:

> This is the price that Kelley Blue Book has determined people like you are typically paying a dealer for a used car with typical mileage in good condition or better. This price is based on actual used-car transactions and adjusted regularly as market conditions happen to change (Kelley Blue Book, 2022).

We select the ten used car models to be broadly representative of the used car market in 2014. We include the most popular pickup truck (Ford F150), SUV (Ford Escape), and Sedan/Coupe (Nissan Altima) at used car retailer CarMax in 2014 (Auto Remarketing, 2014). We also ensure that our selection of cars covers models bought by customers of different ages. The Nissan Altima was the most popular 2014 car for Generation X and Millenials; baby boomers favored the Toyota Camry, which we also include in our price index. See Appendix Table A12 for the full specifications of the ten cars that we used to construct the price index.

We get 2013 state-level car registration prices by combining two datasets from the Federal Highway Administration. Total receipts for vehicle registration fees comes from the state motor-vehicle and motor-carrier tax receipts (MV-2). Total motor vehicle registrations comes from state motor-vehicle

---

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

registrations (MV-1). We divide the total registration fees by the number of registrations in each state to estimate a state-level registration cost for every state and Washington, D.C.. We use the cost for 2013 instead of 2014 because MV-2 is not available for 2014. We then adjust for inflation using the BLS inflation calculator to obtain the 2013 registration cost measured in 2014 US dollars.

We use Nielsen data to get the costs of car maintenance in five UPCs. Refer to the Nielsen section for details on how this data was extracted.

We aggregate the price of purchasing a car, the price of registration and maintenance into one number we take a weighted average of the three prices. To reflect the fact that purchasing a car is the most expensive part of automotive expenses, we set 95% of the index is the amortized cost of a used car plus one registration fee and 5% of the automotive expenses index to be auto maintenance costs. We assume full depreciation after five years (Meyer and Sullivan, 2008), so the amortized cost of a used car is 20% of the cost recorded.

Gasoline/Fuel. We download graphs that contain historical gas prices from GasBuddy. GasBuddy is a crowdsourcing platform in which users can report gas prices in exchange for rewards. We use `https://apps.automeris.io/wpd/` to scrape the graphs and get point estimates for the price of gas on a given day. Beginning with 9/1/2014, we record the price of gas every 9–10 days until 12/22/2014 to obtain 39 data points for the price of gas. We measure the price of gas on the same day for all cities and states. We take the average price of gas for a city or state to be the mean of these 39 prices. We crosswalk the city to a commuting zone using the city name, and we check with GasBuddy when multiple cities have the same name. We obtain an average gas price for each commuting zone by taking the mean of the price of gas in each city in the commuting zone with the 2000 city population as an analytic weight. This procedure gives us a gas price for 143 commuting zones which are covered by the bank account transactions data. For the remaining 300 commuting zones, we use the state-level gas price.

Healthcare/Medical. We get the price for healthcare from the Healthy Marketplace Index (HMI), produced by the Health Care Cost Institute (HCCI). The Health Care Cost Institute is "an independent, non-profit organization with leading health care claims datasets that enable research, policy and journalism" (Health Care Cost Institute, 2022c). Its price data comes from de-identified health care claims data from around 40 million Americans (Health Care Cost Institute, 2022a). The HMI gives the price of several health care services by location. It also contains an overall health price index, which is the weighted average of various inpatient, outpatient, and professional services. In particular, the HCCI is a weighted average of inpatient claims (100 DRG service codes), outpatient claims (500 CPT codes), and professional claims (500 CPT codes). The specific service codes in the 2014 index are the most frequently occurring codes in 2017 claims data. The overall price index captures 86.0% of claims and 63.4% of total spending (allowed amount) in 2014. The weights assigned to these items are determined by the national claims in each service category. Critically, this means that the weights assigned to the various services do not change by CBSA; the overall price index is the cost of the same basket of goods in different CBSAs.

The overall price index is reported as percent deviations from the national median. Consequently, we invert the index reported in the HMI to obtain a dollar cost of the basket in each CBSA and state (Health Care Cost Institute, 2022b). The raw data cover 121 CBSAs and 42 states as well as Washington, D.C.. We crosswalk CBSA prices to commuting zones and use the population in each CBSA-commuting zone intersection as a weight to combine 121 CBSA prices to get 140 commuting zone prices.

Telecommunications. We collect the price of cable TV in 2014 using a Freedom of Information

request of the Federal Communications Commission's (FCC) Cable Price Survey (Form 333). We use the "basic service price" as our measure of cable TV price. Basic service is defined by the FCC as follows:

> Basic service is the lowest level of cable service a subscriber can buy. It includes, at a minimum, all over-the-air television broadcast signals carried pursuant to the must-carry requirements of the Communications Act, and any public, educational, or government access channels required by the system's franchise agreement. It may include additional signals chosen by the operator. Basic service is generally regulated by the local franchising authority (the local or state entity empowered by Federal, State, or local law to grant a franchise to a cable company to operate in a given area) (Federal Communications Commission, 2022).

We begin with cable TV prices for 778 providers across the United States. For each provider, we have the name of the county (and state) in which they are located. First, we take a simple mean of the 778 cable TV providers at the county level to get 556 county prices. Next, we merge the counties with a county-cz crosswalk, and drop 47 counties which we cannot match to a commuting zone. Finally, we get commuting zone prices by taking an average of the county cable prices using the population count in 2000 to weight the counties within a commuting zone.

Clothing/Shoes/Jewellery. We use ACCRA prices for Clothing/Shoes/Jewelry prices. We purchase clothing prices from ACCRA Cost of Living Index (COLI), which is collected by the Council for Community and Economic Research. From these data, we select three items and one service – we take the price of these as representative of the local price of clothing, shoes, and jewellery: Boys' jeans: Blue Denim jeans, regular, relaxed or loose fit, sizes 8–20; Men's dress shirt: Cotton/polyester, pinpoint weave, long sleeves; Women's slacks: At least 95% cotton, twill khakis, sizes 4–14; Dry cleaning: Cost of cleaning man's two-piece suit. For clothing, ACCRA publishes its own expenditure weights by year, which we use to aggregate prices for our four products.

The data are at the core-based statistical area (CBSA) level. We have 281 annual data points from 251 different CBSAs. After correcting one miscoded CBSA from 14460 to 14454, we obtain a simple mean of the price of each of the four products for each CBSA, and crosswalk these data to commuting zones. We obtain average commuting zone price using the population in each CBSA-CZ intersection as a weight to get the price of each item of clothing and dry cleaning by commuting zone. We are able to get a commuting zone price in 254 commuting zones (this number is greater than the number of CBSAs because some CBSAs are present in multiple commuting zones). Not all of these commuting zones are covered in the paper; however, we do use all commuting zones are part of the imputation procedure described below.

Restaurants/Dining. We take restaurant prices from Pricelisto, a crowdsourcing website that records local prices for a number of consumption amenities, from restaurants and gyms to salons and flu shots. We use a dataset of menu items from 20 popular restaurant chains in 2019 that covers 6,861 zip codes. We were unable to access price data for a date earlier than 2019. We use the BLS inflation calculator for food away from home to get prices in 2014 dollars. We crosswalk a zip code to a commuting zone. When a zip code covers multiple commuting zones, we assign it a weight equal to the share of residential addresses of the zip code in that particular commuting zone.

We drop observations for two restaurant chains that we cannot find in our bank and card data (which we use for relative restaurant shares). For the remaining 18 restaurant chains, we regress

price on a CZ indicator and include a menu item fixed effect. We estimate:

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

where $u \in U$ is a standardized product description of restaurant chain $r(u) \in R$ purchased in commuting zone $j \in J$. The menu item fixed effects, δ_u, control for differences in products consumed in different locations. The estimated coefficient on $\delta_{r(u),j}$ added to δ_u evaluated at the nationwide shares across all products within a given restaurant is used as the conditional mean price of restaurant $r(u)$ faced by any income group in commuting zone j. We weight each observation using the crosswalk generated weight.

We save all restaurants with unimputed prices that cover at least 100 commuting zones. In restricting our sample to restaurants with at least 100 commuting zones, we drop four restaurant chains. Thus, the final sample has commuting zone prices for 14 restaurant chains.

To aggregate the price of the 14 restaurant chains, we weight each restaurant chain based on its share of expenditures. To do so, we return to our bank account data. In our 5% sample of bank and credit card transactions, we can observe merchant name. We therefore use the amount of money spent on each of the 14 merchants by income group and use these as weights to aggregate the chain specific prices into one aggregate.

Utilities. For utility prices, we combine the average monthly costs for a household's water and electricity consumption.

We take water prices from the American Water Works Association/Raftelis Financial Consultants 2014 Water and Wastewater Rate Survey. The survey was collected in the second and third quarters of 2014 to reflect prices on January 1, 2014, and contains responses from 318 water service providers. The survey identifies residential usage at 7,480 gallons per month as a "key usage rate"; we choose to measure the price of water as the price of this amount of water (American Water Works Association and Raftelis Financial Consultants, Inc., 2022).

The dataset is at the water utility level, but it contains information on the state, city, and county of the utility. First, we create a state-level price by taking an average of the cost of 7,480 monthly gallons of water using the total number of residential accounts held by that utility as the analytic weight. There are three states not covered: we impute state prices for these using the imputation procedure described below. We then return to the list of utilities and manually extract the names of all cities and counties covered by each utility. In total, we are able to connect the 318 water service providers to 587 cities and separately to 293 counties. We crosswalk these cities and counties to commuting zones and take a simple mean of the price of all utilities that are present in a commuting zone as the price. Importantly, each utility can only be counted once in the commuting zone price, regardless of the number of cities that we record it covering in the commuting zone; our utility coverage is incomplete, so we do not attempt to weight the various utilities within a commuting zone.

We take residential electricity prices from the U.S. Energy Information Administration (EIA). In particular, we use a dataset which reports forms EIA-861 (schedules 4A and 4D) and EIA-861S. The EIA reports that the average household electricity consumption was 893 kWh per month in 2020 (U.S. Energy Information Administration, 2022). We therefore take the cost of electricity to be represented by the price of 893 kWh.

We begin with 2,124 electricity providers and drop one because it has a duplicated name (Tri-

County Electric Coop, Inc), so it cannot be matched to a zip code. We download lists of all zip codes covered by each utility from the EIA and crosswalk the utility name to a zip code. We take a simple mean of the price of all utilities present in a zip code to get the electricity price in 39,805 zip codes. Finally, we crosswalk the zip codes to commuting zones and use the number of residential addresses that each zip code has in a commuting zone as analytic weights for the mean electricity price. We have electricity prices for 704 commuting zones.

To aggregate the price of water and electricity into one number, we add the price of 893 kWh of electricity to 7,840 gallons of water to estimate the average price of utilities for one household in a month.

Imputation Procedure. We have described the data we use for prices in a number of high-level categories, as well as how we connected prices at different levels of geography (eg. CBSA and city) to commuting zones. However, for the majority of the datasets described above we are unable to obtain raw prices for all 443 commuting zones we examine in our paper. For each commuting zone covered by our bank account data, we assign a price according to the following procedure:

1. We first check whether we have any unimputed prices at the commuting zone level.
2. We then check whether we have prices given at the state level. For car registration, gas, and healthcare, we have state-specific prices which we did not impute. We use these state prices where cz-specific prices are missing.
3. For commuting zones still missing a price, we calculate an imputed commuting zone price by taking a simple mean of all neighboring commuting zones for which we have an unimputed commuting zone price.
4. When the unimputed state prices do not cover all states, we use these to impute state-level prices for the missing states by taking a simple mean of the price in all neighboring states. For Hawaii, we use the California price; for Alaska, we use the Washington price. We use the imputed state price where state- and cz-specific prices are missing.
5. Next, we calculate our own state-level price by taking a simple average of the unimputed prices in all commuting zones within a state, and use this for commuting zones still missing a price.
6. Finally, we use our state-level prices to calculate an imputed state price for commuting zones in states that are still missing a price in that category. For Hawaii, we use the California price; for Alaska, we use the Washington price. (For the Outback Steakhouse price in Alaska, we use our imputed Washington price.)

Using this procedure, we are able to assign all 443 commuting zones a state-level price. See Appendix Table A13 for a summary of the number of commuting zones in each price imputation category. No imputation is needed for housing costs, as the American Community Survey covers all the commuting zones that we study in the paper.

Correlations Between the Prices of Consumption Categories. Appendix Table A15 reports all pairwise correlations among all components of the price index.

C.2 Measuring Expenditure Shares

We closely follow the methodology that the BLS uses to calculate expenditure shares to compute the CPI. An expenditure share on a given item is defined as total consumption expenditure on this item across households divided by total consumption expenditure on all items across households.

Specifically, we define income-group-specific nationwide shares and income-and-commuting-zone-specific shares, respectively, as

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

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

where I denotes the set of 21 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 $s_{i,j,k}$. Therefore, $\bar{E}_{i,k}$ is household-average expenditure on category i for income group k nationwide and $\bar{E}_{i,j,k}$ is household-average expenditure on category i for income group k in commuting zone j.

Our data classify expenditure in 20 high-level consumption categories. An important limitation is that we can identify categories only for transactions done by credit card or debit card or electronic transfer. Transactions by cash or checks are labelled in our data as "Unclassified" because the identity of the merchant is unknown. However, we note that the distribution of classified expenditures across categories match well the NIPA shares (Section 2).

To calculate household h's expenditure shares, we take its total expenditure (E_h) and then subtract out our imputed housing costs (H_h^-), leaving total non-housing expenditure (N_h): $N_h = E_h - H_h^-$. This non-housing expenditure is the sum of expenditures paid through bank or card accounts—which are assigned to the 20 non-housing categories ($X_{h,i}$ for $i \in I = \{1, \dots 21\}$, $i \neq 13$)—and expenditures paid in cash or checks—which are "Unclassified": $N_h - \sum_{i \in I} X_{h,i}$. Because we cannot identify what types the latter spending consists of, we apportion it back to our focal categories or $\tilde{X}_{h,i} = \frac{X_{h,i}}{\sum_{i \in I} X_{h,i}} \times N_h$: as such, $\sum_{i \in I} \tilde{X}_{h,i} = N_h$. Next, we add back our imputed housing costs (H_h^+ or $X_{h,13}$) 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 $i = 1, \dots, 21$.

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

Four of these categories– General Merchandise, Groceries, Hobbies/Entertainment, and Personal Care –are very broad and can be broken down into finer categories to improve precision. We use the Nielsen Consumer Panel Price Data to obtain product group-specific expenditure shares within the broader category. 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 $i \in \{\text{General Merchandise}, \text{Groceries}, \text{Hobbies/Entertainment}, \text{Personal Care}\}$. 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 21 high-level categories.

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

Some of our alternative price indices require measuring expenditure shares at the commuting zone level. For this, we repeat the procedure above but at the commuting zone level, with a few changes. Importantly, the Nielsen consumer data do not cover all products in all commuting zones. Therefore, for commuting zones with no expenditure on any products within a high-level category, we impute total spending for each product group using the same method as we used for missing prices. For high-level categories with expenditure on some products, we don't impute any spending. For example, we would impute spending in all categories for a commuting zone with no expenditure on "Books, Magazines", Pet Care, or "Toys, Sporting Goods" (the three components of Hobbies/Entertainment). However, we would not impute any spending for a commuting zone with expenditure on Pet Care but no spending on either "Books, Magazines" or "Toys, Sporting Goods". We do this for all income groups and separately for low-, middle-, and high-income households. We use unimputed and imputed spending to calculate expenditure shares by commuting zone and income group within the high-level categories. We impute shares in 370 (5.22%) out of 7,088 category-income group-commuting zone groups. Out of these 370 groups, 242 are high-income, 53 are middle-income, 47 are low-income, and 28 are overall.

C.3 Alternative Price Indices

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

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

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

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

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

Nested CES: We follow Handbury and Weinstein (2015) in building a nested-CES exact local price index, accounting for variation in local supply of products. We measure the same price index for all three income groups. Just like our main Laspeyres price index, we index the "high-level" product categories by i where I is the set of all high-level product categories. Within each product category i, there are mid-level categories that classify purchases into product groups. These are indexed by im. Only the 3 high-level product categories have products split into mid-level nests (e.g. yogurt versus cheese). This is because the Nielsen data provides this additional level information about the products. These mid-level nests are split based on the product groups assigned to each product by Nielsen. For high-level product categories not covered by Nielsen, these is no mid-level nest. Finally, the lowest level nest measures utility from each individual variety of product. These

are indexed by g. For the 3 Nielsen groups, we use UPC codes to identify unique varieties. For the rest of the product categories not covered by Nielsen, we use merchants, as observed in our transaction data to identify a unique variety. Surely most merchants sell a variety of products, but merchant is the most granular data we observe. For most transactions, our data provider as already listed the merchant associated with each transaction. For smaller merchants, this variable is blank in our data. To measure merchants for these additional transactions, we standardize the description string from the transaction by cleaning out text from the bank itself (e.g. remove words like "CHECKCARD PURCHASE"), and other formatting differences across banks to create a text string unique to each merchant. The utility function is:

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

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

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

As shown by Handbury and Weinstein (2015), the price index EPI_j for CZ j that accounts for variation in access to local variety can be written as:

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

where:

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

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

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

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

weights and are defined as follows:

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

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

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

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

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

E_g is national total expenditure on variety g.

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

Geary-Khamis PPP Index The Geary-Khamis index is a Paasche index that compares the local prices in a given CZ to nationwide average prices. The weights on the relative prices differences between the CZ and the nationwide average are equal the focal CZ's expenditure shares. This is the method used by the BEA to estimate local price indices. A desirable property of The Geary-Khamis index is that preserves aggregation. Thus, the Geary-Khamis index is a weighted average of Geary-Khamis indices for each sub-component of consumption (e.g. housing or restaurants). It is measured as:

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

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

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.

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

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

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); $y = (x - \mathbf{p'}\mathbf{s})/(1 - \mathbf{p'Bp}/2)$ 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 $y = x - \mathbf{p'}\mathbf{s}$ and $\tilde{y} = x - \mathbf{p'}\mathbf{\bar{s}}$.

To calculate an income-group-specific price index implied by the exact EASI demand model, we first predict $\mathbf{\tilde{s}_h}$ for each household if it were to face average prices. These predicted shares as a function of real consumption expenditures are shown in Appendix Figure A6. 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, $\mathbf{p'_j}\mathbf{\tilde{s}_j} = \sum_i (p_{ij} \cdot \tilde{s}_{ij})$.

D Consumption in Physical Units in Nielsen Data

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

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

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

E Estimating Consumption by Education Group

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

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

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

Step 1: We define household types. We <u>interact</u> 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 $s \in S = \{\text{high, middle, low}\}$, 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 $Y \in \{\text{pre-tax income, post-tax income, expenditure, consumption}\}$. 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.

16

F   Government Transfers

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

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

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

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

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

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

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

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

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

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

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

Appendix Figure A4: **Map of Consumption Expenditure, By Income Group**
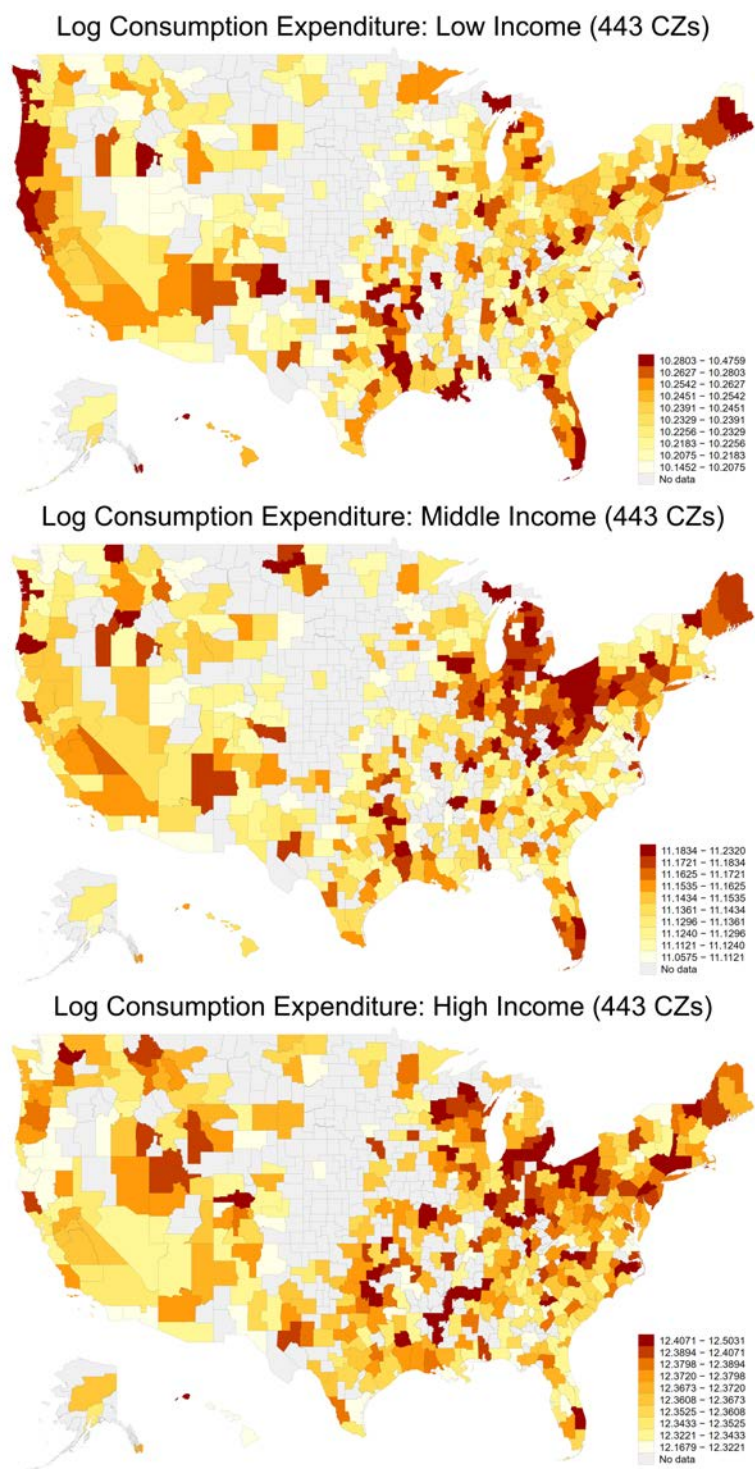
Log Consumption Expenditure: Low Income (443 CZs)

Log Consumption Expenditure: Middle Income (443 CZs)

Log Consumption Expenditure: High Income (443 CZs)

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

Appendix Figure A5: **Map of Consumption, by Skill Level**

A. High Skill
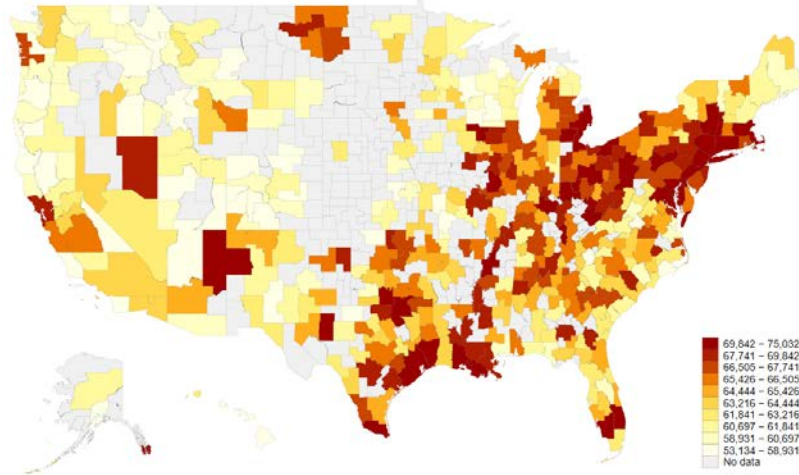
B. Middle Skill
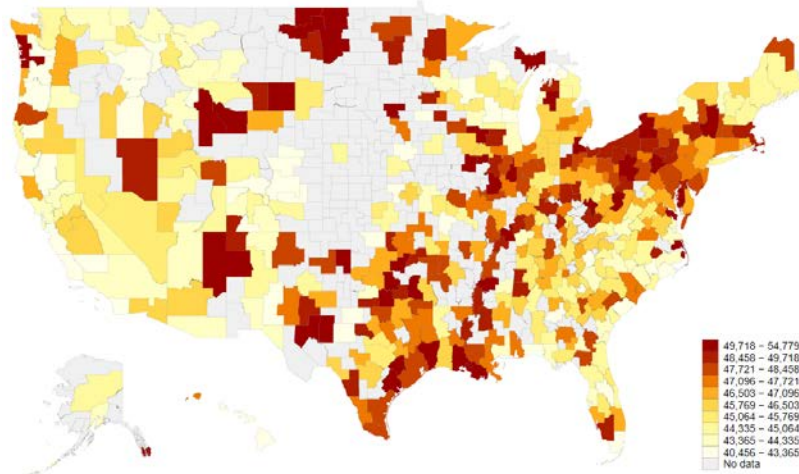
C. Low Skill

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

Appendix Figure A6: **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	47,750.00	27,652.95	71,200.55	251,517.69
Panel B. Adjusted Measures of Income and Expenditure				
Post-tax Income				
Mean	92,171.30	39,536.41	100,881.41	483,806.94
Median	62,146.05	38,997.29	89,993.14	305,953.72
Expenditure				
Mean	85,791.79	39,799.98	91,895.27	441,625.25
Median	56,645.06	36,872.89	79,645.93	267,712.03
Number of Commuting Zones	443	443	443	443
Number of Households	3,000,518	1,368,817	1,449,978	181,723

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

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

	Expenditure Shares				Price
	Overall	Low Income	Middle Income	High Income	Standard Deviation
	(1)	(2)	(3)	(4)	(5)
Automotive Expenses	2.89%	2.17%	3.42%	4.47%	0.482
Charitable Giving	0.37%	0.22%	0.47%	0.76%	0
Child/Dependent Expenses	0.43%	0.26%	0.58%	0.68%	0.029
Clothing/Shoes/Jewellery	4.10%	3.24%	4.66%	6.72%	0.251
Electronics	1.07%	0.91%	1.20%	1.32%	0.135
Education	0.88%	0.70%	0.95%	1.95%	0
Financial Fees	1.17%	0.72%	1.41%	2.94%	0
Gasoline/Fuel	3.31%	3.07%	3.58%	2.88%	0.053
General Merchandise	9.88%	8.63%	11.29%	8.25%	0.237
Groceries	6.95%	5.86%	8.00%	7.13%	0.156
Healthcare/Medical	14.46%	18.73%	10.78%	9.67%	0.137
Hobbies/Entertainment	3.55%	2.92%	4.01%	5.00%	0.169
Housing	27.77%	33.80%	23.25%	14.75%	0.295
Insurance	4.36%	3.33%	5.01%	7.75%	0
Office Supplies	0.24%	0.20%	0.27%	0.39%	0.055
Personal Care	1.20%	1.00%	1.34%	1.62%	0.197
Printing and Postage	0.30%	0.28%	0.31%	0.26%	0
Restaurants/Dining	6.29%	5.83%	6.73%	6.36%	0.232
Telecommunications	4.92%	4.06%	5.66%	5.81%	0.358
Travel	3.22%	2.34%	3.80%	5.91%	0
Utilities	2.66%	1.73%	3.27%	5.39%	0.239

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

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

	Low Income	Medium Income	High Income
Groceries			
Baby Food	0.21%	0.24%	0.25%
Baked Goods - Frozen	0.41%	0.37%	0.31%
Baking Mixes	0.43%	0.40%	0.35%
Baking Supplies	0.55%	0.56%	0.53%
Beer	1.40%	1.35%	1.23%
Bread, Baked Goods	4.02%	3.80%	3.57%
Breakfast Food	0.84%	0.98%	1.02%
Breakfast Food, Frozen	0.59%	0.60%	0.58%
Butter, Margarine	0.92%	0.85%	0.76%
Candy	2.82%	2.60%	2.49%
Carbonated Beverages	3.39%	3.01%	2.76%
Cereal	1.87%	1.91%	1.83%
Charcoal, Logs	0.11%	0.12%	0.12%
Cheese	3.24%	3.48%	3.55%
Coffee	1.89%	2.02%	2.21%
Condiments, Gravies, Sauces	1.38%	1.43%	1.41%
Cookies	1.22%	1.16%	1.10%
Cottage Cheese, Sour Cream	0.59%	0.60%	0.60%
Crackers	0.85%	0.88%	0.89%
Desserts, Fruits, Toppings	0.33%	0.35%	0.34%
Desserts, Gelatins, Syrup	0.47%	0.46%	0.42%
Detergents	1.33%	1.41%	1.45%
Disposable Diapers	0.26%	0.35%	0.36%
Dough Products	0.36%	0.36%	0.32%
Dressings, Salads, Prepared Foods	5.98%	6.28%	6.33%
Eggs	0.85%	0.81%	0.78%
Flour	0.15%	0.14%	0.14%
Fresh Meat	0.78%	0.72%	0.66%
Fresh Produce	5.44%	6.47%	7.51%
Fresheners, Deodorizers	0.46%	0.42%	0.38%
Fruit - Canned	0.42%	0.35%	0.29%
Fruit, Dried	0.33%	0.39%	0.44%
Gum	0.18%	0.21%	0.22%
Household Cleaners	0.61%	0.65%	0.70%
Household Supplies	0.70%	0.74%	0.75%
Ice	0.02%	0.02%	0.02%
Ice Cream	1.62%	1.49%	1.40%
Jams, Jellies, Spreads	0.68%	0.65%	0.62%
Juice, Drinks - Canned-Bottled	2.01%	2.10%	2.15%
Juice, Drinks - Frozen	0.09%	0.09%	0.08%
Laundry Supplies	0.69%	0.74%	0.75%
Liquor	1.03%	1.23%	1.50%
Milk	2.75%	2.62%	2.52%
Nuts	1.03%	1.20%	1.38%
Packaged Meats - Deli	3.75%	3.75%	3.51%
Packaged Milk, Modifiers	0.76%	0.67%	0.61%
Paper Products	3.38%	3.42%	3.38%
Pasta	0.37%	0.40%	0.41%
Pet Food	4.72%	4.23%	3.84%
Pickles, Olives, Relish	0.38%	0.38%	0.37%
Pizza, Snacks - Frozen	1.16%	1.07%	1.00%
Prepared Food - Dry Mixes	0.98%	0.93%	0.83%
Prepared Food - Ready-to-Serve	1.07%	0.94%	0.88%
Prepared Foods - Frozen	3.14%	2.88%	2.71%
Puddings, Dessert - Dairy	0.07%	0.07%	0.07%
Salad Dressings, Mayo, Toppings	0.76%	0.74%	0.66%
Seafood, Canned	0.44%	0.40%	0.39%
Shortening, Oil	0.56%	0.54%	0.55%
Snacks	3.56%	3.73%	3.77%
Snacks, Spreads, Dips - Dairy	0.30%	0.36%	0.46%
Soap, Bath Additives	0.65%	0.73%	0.77%
Soft Drinks - Non-Carbonated	1.16%	1.15%	1.22%
Soup	1.12%	1.09%	1.04%
Spices, Seasoning, Extracts	0.52%	0.53%	0.55%
Sugar, Sweeteners	0.53%	0.45%	0.39%
Table Syrups, Molasses	0.14%	0.15%	0.14%
Tea	0.72%	0.75%	0.79%

Tobacco	3.22%	1.78%	1.12%
Unprepared Meat, Poultry, Seafood	5.79%	6.16%	6.45%
Vegetables - Canned	0.97%	0.95%	0.88%
Vegetables - Frozen	1.04%	1.02%	0.92%
Vegetables, Grains - Dried	0.37%	0.39%	0.40%
Wine	1.04%	1.47%	2.34%
Wrapping Materials, Bags	0.80%	0.84%	0.84%
Yeast	0.00%	0.00%	0.00%
Yogurt	1.18%	1.44%	1.63%
General Merchandise			
Batteries, Flashlights	12.64%	11.73%	10.82%
Canning, Freezing Supplies	1.25%	1.15%	0.90%
Cookware	4.26%	4.07%	3.73%
Floral, Gardening	8.22%	9.11%	9.54%
Glassware, Tableware	4.59%	4.60%	4.69%
Hardware, Tools	6.48%	6.65%	6.31%
Housewares, Appliances	28.52%	28.69%	28.94%
Insecticides, Pesticides, Rodenticides	5.76%	4.77%	4.12%
Kitchen Gadgets	9.66%	10.41%	10.25%
Light Bulbs, Electric Goods	12.72%	12.77%	14.72%
Party Needs	0.27%	0.21%	0.19%
Photographic Supplies	3.30%	3.58%	3.66%
Seasonal	0.62%	0.55%	0.48%
Sewing Notions	0.34%	0.36%	0.35%
Shoe Care	0.26%	0.26%	0.24%
Soft Goods	1.11%	1.08%	1.07%
Hobbies/Entertainment			
Books, Magazines	7.94%	6.94%	6.22%
Pet Care	91.32%	92.46%	93.36%
Toys, Sporting Goods	0.74%	0.61%	0.42%
Personal Care			
Cosmetics	3.49%	4.12%	4.60%
Cough and Cold Remedies	6.56%	6.89%	6.89%
Deodorant	2.00%	2.39%	2.52%
Diet Aids	0.72%	0.88%	0.91%
Ethnic Haba	0.12%	0.09%	0.07%
Feminine Hygiene Products	0.39%	0.37%	0.36%
First Aid	2.47%	2.38%	2.37%
Fragrances - Women	1.17%	1.23%	1.41%
Grooming Aids	1.20%	1.30%	1.31%
Hair Care	6.32%	7.16%	7.65%
Medications, Remedies, Health Aids	41.22%	36.09%	31.72%
Men's Toiletries	0.44%	0.52%	0.53%
Oral Hygiene	6.27%	6.69%	7.24%
Sanitary Protection	1.88%	1.99%	1.98%
Shaving Needs	2.20%	2.84%	3.41%
Skin Care Preparations	5.01%	6.34%	7.75%
Vitamins	18.52%	18.71%	19.27%

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

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

	Low Income (1)	Middle Income (2)	High Income (3)
Log Monthly Rent	0.434***	0.330***	0.241***
	(0.011)	(0.009)	(0.008)
R^2	0.950	0.939	0.893

Notes: All columns use a log-log specification. We use commuting zone population as regression weight. Robust standard errors are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.10. N = 443.

Appendix Table A5: **Regression of Components of Main Price Index Against Rent**

	Automotive Expenses (1)	Child Expenses (2)	Clothing/Shoes/Jewellery (3)	Electronics (4)	Gasoline/Fuel (5)	General Merchandise (6)	Groceries (7)
Log Rent	-0.017*	0.036***	0.136***	0.018***	0.120***	0.042***	0.023*
	(0.010)	(0.010)	(0.035)	(0.006)	(0.019)	(0.009)	(0.012)
R^2	0.023	0.198	0.084	0.050	0.343	0.319	0.063

	Healthcare/Medical (8)	Hobbies/Entertainment (9)	Office Supplies (10)	Personal Care (11)	Restaurants/Dining (12)	Telecommunications (13)	Utilities (14)
Log Rent	0.182***	0.017*	0.085***	0.068***	0.094***	-0.169***	0.366***
	(0.033)	(0.010)	(0.015)	(0.016)	(0.013)	(0.064)	(0.070)
R^2	0.230	0.022	0.428	0.347	0.456	0.040	0.355

Notes: N = 443.

Appendix Table A6: **Price Index Correlations**

	A1. Combined index, Laspeyres	A2. Combined index, Laspeyres (CEX weights)	A3. Combined index, CES	A4. Combined index, Nested CES σ=11.5	A5. Combined index, Nested CES σ=7	A6. Combined index, PPP, Geary-Khamis	A7. Combined index, PPP, GEKS-Fischer	B1. Combined index, Laspeyres	B2. Combined index, Laspeyres (CEX weights)	B3. Combined index, CES	B4. Combined index, Nested CES σ=11.5	B5. Combined index, Nested CES σ=7	B6. Combined index, PPP, Geary-Khamis	B7. Combined index, PPP, GEKS-Fischer	C. EASI demand based index	D. BEA index
A1. Combined index, Laspeyres	1															
A2. Combined index, Laspeyres (CEX weights)	.81	1														
A3. Combined index, CES	.99	.81	1													
A4. Combined index, Nested CES σ=11.5	.92	.79	.93	1												
A5. Combined index, Nested CES σ=7	.81	.72	.81	.96	1											
A6. Combined index, PPP, Geary-Khamis	.99	.81	.99	.93	.82	1										
A7. Combined index, PPP, GEKS-Fischer	.99	.81	.99	.92	.81	.98	1									
B1. Combined index, Laspeyres	.99	.78	.98	.90	.78	.97	.98	1								
B2. Combined index, Laspeyres (CEX weights)	.80	.91	.79	.75	.67	.77	.80	.78	1							
B3. Combined index, CES	.99	.78	.99	.91	.79	.98	.98	.99	.78	1						
B4. Combined index, Nested CES σ=11.5	.92	.77	.92	.99	.95	.93	.92	.92	.74	.92	1					
B5. Combined index, Nested CES σ=7	.81	.70	.81	.95	.99	.82	.81	.80	.65	.80	.96	1				
B6. Combined index, PPP, Geary-Khamis	.98	.78	.99	.92	.80	.99	.98	.98	.76	.99	.93	.81	1			
B7. Combined index, PPP, GEKS-Fischer	.98	.79	.97	.90	.78	.96	.98	.98	.80	.98	.91	.79	.97	1		
C. EASI demand based index	.99	.80	.99	.92	.81	.99	.98	.99	.78	.99	.93	.82	.99	.98	1	
D. BEA index	.93	.73	.93	.84	.72	.91	.92	.93	.75	.92	.84	.72	.91	.91	.92	1

Notes: N = 302 for BEA x Nested CES indexes. N = 305 for all other correlations with BEA index. N = 412 for all other correlations with Nested CES indexes. N = 443 otherwise. Correlation matrix of all alternative price indices. See text for detailed definition of each index.

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

	Overall Income				Low Income				Middle Income				High Income			
	mean	sd	min	max	mean	sd	min	max	mean	sd	min	max	mean	sd	min	max
A. Combined indexes:																
same prices for all households																
- Laspeyres	1.026	0.099	0.875	1.555	1.022	0.115	0.848	1.653	1.030	0.088	0.896	1.482	1.034	0.069	0.914	1.386
- Laspeyres (CEX weights)	0.991	0.033	0.936	1.194	0.990	0.034	0.933	1.197	0.993	0.033	0.937	1.196	0.995	0.031	0.942	1.194
- CES	1.017	0.092	0.867	1.491	1.015	0.109	0.845	1.622	1.020	0.081	0.877	1.427	1.025	0.062	0.905	1.304
- Nested CES (Feenstra, $\sigma = 7$)	1.119	0.088	0.931	1.376	1.103	0.093	0.936	1.523	1.123	0.089	0.910	1.400	1.103	0.079	0.956	1.356
- Nested CES (Feenstra, $\sigma = 11.5$)	1.070	0.067	0.914	1.332	1.064	0.086	0.909	1.520	1.072	0.067	0.893	1.342	1.063	0.054	0.944	1.246
- PPP Geary-Khamis	1.019	0.087	0.854	1.450	1.019	0.105	0.840	1.601	1.022	0.078	0.856	1.395	1.016	0.061	0.831	1.271
- PPP GEKS-Fischer	1.016	0.079	0.882	1.437	1.015	0.099	0.817	1.573	1.009	0.072	0.857	1.378	1.011	0.067	0.750	1.253
B. Combined indexes:																
income-group-specific prices																
- Laspeyres	1.027	0.100	0.872	1.534	1.023	0.103	0.851	1.490	1.017	0.078	0.830	1.391	0.983	0.053	0.839	1.226
- Laspeyres (CEX weights)	1.001	0.029	0.947	1.171	1.003	0.033	0.943	1.243	0.999	0.028	0.910	1.144	0.990	0.027	0.911	1.124
- CES	1.017	0.093	0.853	1.477	1.015	0.098	0.840	1.432	1.006	0.075	0.774	1.364	0.965	0.064	0.582	1.171
- Nested CES (Feenstra, $\sigma = 7$)	1.116	0.087	0.943	1.377	1.114	0.086	0.923	1.436	1.105	0.085	0.867	1.375	1.067	0.071	0.938	1.326
- Nested CES (Feenstra, $\sigma = 11.5$)	1.068	0.066	0.926	1.320	1.075	0.078	0.897	1.430	1.057	0.064	0.850	1.303	1.036	0.051	0.902	1.203
- PPP Geary-Khamis	1.018	0.088	0.829	1.438	1.017	0.096	0.668	1.433	1.008	0.081	0.411	1.355	0.987	0.083	0.348	1.217
- PPP GEKS-Fischer	1.025	0.084	0.882	1.449	1.014	0.094	0.773	1.474	1.013	0.072	0.841	1.329	0.968	0.071	0.562	1.167
EASI demand index	1.021	0.083	0.865	1.424	1.024	0.101	0.830	1.537	1.017	0.072	0.821	1.402	0.979	0.059	0.765	1.227
BEA price index	1.040	0.073	0.874	1.359	1.040	0.073	0.874	1.359	1.040	0.073	0.874	1.359	1.040	0.073	0.874	1.359

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

Appendix Table A8: **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	0.046	-0.068***	-0.119***	-0.021	-0.071**
	(0.026)	(0.007)	(0.013)	(0.024)	(0.024)
B. Child/Dependent Expenses					
Log Median CZ Income	0.058**	-0.048***	-0.084***	0.012	-0.023
	(0.019)	(0.004)	(0.008)	(0.018)	(0.018)
C. Clothing/Shoes/Jewellery					
Log Median CZ Income	0.161*	-0.040***	-0.069***	0.121	0.092
	(0.080)	(0.004)	(0.007)	(0.079)	(0.078)
D. Electronics					
Log Median CZ Income	-0.001	-0.329***	-0.576***	-0.328***	-0.573***
	(0.012)	(0.030)	(0.053)	(0.034)	(0.056)
E. Gasoline/Fuel					
Log Median CZ Income	0.157***	-0.029***	-0.050***	0.128***	0.106***
	(0.032)	(0.003)	(0.005)	(0.031)	(0.031)
F. General Merchandise					
Log Median CZ Income	0.099**	-0.341***	-0.596***	-0.243***	-0.498***
	(0.033)	(0.039)	(0.069)	(0.050)	(0.075)
G. Groceries					
Log Median CZ Income	0.056***	-0.101***	-0.177***	-0.046*	-0.122***
	(0.015)	(0.012)	(0.020)	(0.021)	(0.028)
H. Healthcare/Medical					
Log Median CZ Income	0.231**	-0.039***	-0.069***	0.191*	0.162*
	(0.079)	(0.004)	(0.007)	(0.078)	(0.078)
I. Hobbies/Entertainment					
Log Median CZ Income	-0.002	-0.297***	-0.519***	-0.301***	-0.525***
	(0.023)	(0.035)	(0.062)	(0.041)	(0.065)
J. Office Supplies					
Log Median CZ Income	0.039**	-0.026***	-0.046***	0.013	-0.007
	(0.015)	(0.003)	(0.005)	(0.016)	(0.017)
K. Personal Care					
Log Median CZ Income	0.076***	-0.322***	-0.564***	-0.247***	-0.490***
	(0.016)	(0.040)	(0.070)	(0.038)	(0.067)
L. Restaurants/Dining					
Log Median CZ Income	0.128***	-0.075***	-0.131***	0.053*	-0.003
	(0.027)	(0.006)	(0.011)	(0.024)	(0.023)
M. Telecommunications					
Log Median CZ Income	-0.354**	-0.020***	-0.035***	-0.374**	-0.388**
	(0.126)	(0.002)	(0.003)	(0.126)	(0.126)
N. Utilities					
Log Median CZ Income	0.495***	-0.252***	-0.442***	0.243*	0.054
	(0.118)	(0.027)	(0.048)	(0.110)	(0.107)
Overall					
Log Median CZ Income	0.368***	-0.092***	-0.160***	0.278***	0.212***
	(0.034)	(0.009)	(0.016)	(0.031)	(0.031)

Notes: This table decomposes the impact of the price effect versus the supply of variety in the nested CES price index across each sub-component of the price index. Standard errors are clustered at the commuting zone level. See text for details.

Appendix Table A9: **Consumption vs. Price Index – Robustness**

Index	Low-Income	Middle-Income	High-Income
A1. Laspeyres	−0.897***	−0.983***	−1.069***
	(0.009)	(0.021)	(0.039)
A2. Laspeyres (CEX)	−0.648***	−0.890***	−1.083***
	(0.028)	(0.059)	(0.086)
A3. CES	−0.893***	−0.980***	−1.084***
	(0.009)	(0.023)	(0.043)
A4. Nested CES (11.5)	−0.861***	−0.960***	−1.096***
	(0.011)	(0.030)	(0.067)
A5. Nested CES (7)	−0.855***	−0.963***	−1.113***
	(0.012)	(0.029)	(0.060)
A6. PPP, Geary-Khamis	−0.890***	−0.983***	−1.108***
	(0.009)	(0.025)	(0.046)
A7. PPP, GEKS-Fischer	−0.880***	−0.980***	−1.108***
	(0.010)	(0.026)	(0.045)
B1. Laspeyres	−0.886***	−0.989***	−1.058***
	(0.012)	(0.025)	(0.056)
B2. Laspeyres (CEX)	−0.623***	−0.876***	−0.986***
	(0.037)	(0.068)	(0.105)
B3. CES	−0.882***	−0.985***	−0.987***
	(0.012)	(0.025)	(0.055)
B4. Nested CES (11.5)	−0.847***	−0.974***	−1.080***
	(0.014)	(0.034)	(0.078)
B5. Nested CES (7)	−0.844***	−0.978***	−1.109***
	(0.015)	(0.032)	(0.069)
B6. PPP, Geary-Khamis	−0.895***	−0.986***	−1.042***
	(0.017)	(0.022)	(0.051)
B7. PPP, GEKS-Fischer	−0.874***	−0.988***	−1.059***
	(0.013)	(0.028)	(0.048)
C. EASI Demand System	−0.888***	−0.981***	−1.038***
	(0.011)	(0.025)	(0.055)
D. BEA Price Parities	−0.840***	−0.983***	−1.028***
	(0.015)	(0.024)	(0.038)

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

Appendix Table A10: **Consumption Against Price Index – Robustness**

	College Graduates (1)	High School Graduates (2)	High School Dropouts (3)
A. Baseline			
Log price index	0.017	−0.187***	−0.364***
	(0.058)	(0.029)	(0.035)
B. Income Includes Imputed Food Stamps, TANF, and Housing Assistance			
Log price index	0.008	−0.195***	−0.364***
	(0.059)	(0.031)	(0.036)
C. Consumption Does Not Include Any Imputation			
Log price index	−0.052	−0.248***	−0.438***
	(0.062)	(0.037)	(0.049)

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

Appendix Table A11: **Consumption vs Price Index, City Size and College Share – Robustness**

	A1. Laspeyres	A2. Laspeyres (CEX WT)	A3. CES	A4. Nested CES (σ=11.5)	A5. Nested CES (σ=7)	A6. PPP, Geary-Khamis	A7. PPP, GEKS-Fischer	B1. Laspeyres
Log index	-0.289***	0.318	-0.258***	-0.332***	-0.423***	-0.299***	-0.178**	-0.127
	(0.077)	(0.208)	(0.082)	(0.100)	(0.089)	(0.089)	(0.088)	(0.087)
Log index × middle-skill	0.044	0.157	0.046	0.062	0.062	0.062	0.051	0.042
	(0.087)	(0.269)	(0.092)	(0.114)	(0.101)	(0.101)	(0.100)	(0.102)
Log index × low-skill	-0.048	0.083	-0.058	-0.046	-0.024	-0.048	-0.063	-0.074
	(0.089)	(0.266)	(0.095)	(0.114)	(0.101)	(0.104)	(0.104)	(0.104)
Log city size	0.032***	0.046***	0.031***	0.047***	0.056***	0.034***	0.032***	0.032***
	(0.008)	(0.007)	(0.008)	(0.009)	(0.009)	(0.008)	(0.007)	(0.008)
Log population × middle-skill	-0.026***	-0.022***	-0.026***	-0.026***	-0.026**	-0.027***	-0.026***	-0.025***
	(0.009)	(0.008)	(0.009)	(0.010)	(0.010)	(0.009)	(0.008)	(0.009)
Log population × low-skill	-0.039***	-0.035***	-0.039***	-0.041***	-0.043***	-0.039***	-0.038***	-0.037***
	(0.009)	(0.008)	(0.009)	(0.009)	(0.010)	(0.009)	(0.008)	(0.009)
Log college share	0.046	0.113***	0.052	0.091**	0.106***	0.074*	0.055	0.041
	(0.040)	(0.035)	(0.040)	(0.041)	(0.040)	(0.042)	(0.038)	(0.041)
Log college share × middle-skill	-0.034	-0.014	-0.035	-0.037	-0.037	-0.037	-0.032	-0.035
	(0.045)	(0.043)	(0.046)	(0.047)	(0.047)	(0.050)	(0.045)	(0.048)
Log college share × low-skill	-0.027	-0.011	-0.028	-0.038	-0.042	-0.032	-0.026	-0.018
	(0.046)	(0.043)	(0.048)	(0.047)	(0.047)	(0.050)	(0.046)	(0.050)
Middle-skill	-0.061	-0.087	-0.059	-0.073	-0.075	-0.058	-0.064	-0.075
	(0.150)	(0.138)	(0.150)	(0.170)	(0.178)	(0.152)	(0.141)	(0.152)
Low-skill	-0.054	-0.083	-0.055	-0.038	-0.019	-0.059	-0.063	-0.070
	(0.150)	(0.134)	(0.150)	(0.170)	(0.178)	(0.152)	(0.141)	(0.153)

	B2. Laspeyres (CEX WT)	B3. CES	B4. Nested CES (σ=11.5)	B5. Nested CES (σ=7)	B6. PPP, Geary-Khamis	B7. PPP, GEKS-Fischer	C. EASI Demand System	D. BEA Price Parities
Log index	0.678***	-0.116	-0.262**	-0.380***	-0.177	-0.082	-0.238**	-0.165
	(0.220)	(0.092)	(0.110)	(0.098)	(0.111)	(0.095)	(0.097)	(0.108)
Log index × middle-skill	0.070	0.050	0.122	0.124	0.121	0.059	0.045	0.176
	(0.309)	(0.108)	(0.129)	(0.115)	(0.152)	(0.111)	(0.111)	(0.125)
Log index × low-skill	0.018	-0.073	0.018	0.049	0.027	-0.068	-0.080	0.124
	(0.301)	(0.110)	(0.128)	(0.113)	(0.168)	(0.114)	(0.114)	(0.127)
Log city size	0.045***	0.030***	0.055***	0.064***	0.030***	0.032***	0.034***	0.034***
	(0.007)	(0.008)	(0.009)	(0.010)	(0.008)	(0.007)	(0.008)	(0.008)
Log population × middle-skill	-0.021***	-0.026***	-0.029***	-0.028***	-0.029***	-0.025***	-0.026***	-0.024***
	(0.008)	(0.009)	(0.010)	(0.011)	(0.011)	(0.008)	(0.009)	(0.009)
Log population × low-skill	-0.034***	-0.037***	-0.046***	-0.047***	-0.041***	-0.037***	-0.039***	-0.036***
	(0.008)	(0.009)	(0.010)	(0.010)	(0.011)	(0.008)	(0.009)	(0.009)
Log college share	0.095***	0.047	0.100**	0.120***	0.080*	0.048	0.067	0.042
	(0.032)	(0.042)	(0.043)	(0.042)	(0.043)	(0.039)	(0.042)	(0.040)
Log college share × middle-skill	-0.016	-0.034	-0.041	-0.040	-0.034	-0.034	-0.036	-0.039
	(0.040)	(0.048)	(0.050)	(0.049)	(0.051)	(0.045)	(0.048)	(0.046)
Log college share × low-skill	-0.016	-0.019	-0.036	-0.039	-0.021	-0.020	-0.027	-0.034
	(0.039)	(0.050)	(0.050)	(0.050)	(0.053)	(0.047)	(0.050)	(0.047)
Middle-skill	-0.103	-0.072	-0.034	-0.047	-0.037	-0.074	-0.064	-0.098
	(0.135)	(0.151)	(0.179)	(0.186)	(0.174)	(0.139)	(0.155)	(0.147)
Low-skill	-0.101	-0.072	0.017	0.022	-0.023	-0.070	-0.063	-0.101
	(0.129)	(0.152)	(0.179)	(0.187)	(0.182)	(0.140)	(0.156)	(0.148)

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

Appendix Table A12: **Specifications of Used Cars used in Automotive Expenses Index**

#	Value as of	Vehicle Type	Type	Year	Make	Model	Trim
1	7/1/2014	Automobile	Used	2012	Nissan	Altima	2.5 Sedan 4D
2	7/1/2014	Automobile	Used	2011	Toyota	Camry	Sedan 4D
3	7/1/2014	Automobile	Used	2010	Honda	Accord	LX Sedan 4D
4	7/1/2014	Automobile	Used	2009	Chevrolet	Malibu	LS Sedan 4D
5	7/1/2014	Automobile	Used	2008	Honda	Civic	LX Sedan 4D
6	7/1/2014	Automobile	Used	2012	Ford	F150 Regular Cab	XL Pickup 2D 6 1/2 ft
7	7/1/2014	Automobile	Used	2011	Chevrolet	Silverado 1500 Regular Cab	LT Pickup 2D 6 1/2 ft
8	7/1/2014	Automobile	Used	2010	Ford	Escape	XLT Sport Utility 4D
9	7/1/2014	Automobile	Used	2009	Ford	Edge	SE Sport Utility 4D
10	7/1/2014	Automobile	Used	2008	Hyundai	Sonata	SE Sedan 4D

#	Engine	Trans	Drivetrain	Ext. Color	Mileage	Condition	Other options
1	4-Cyl, 2.5 Liter	Automatic, CVT	FWD	Black	32,500	Good	*Leave as defaults*
2	4-Cyl, 2.5 Liter	Automatic, 6-Spd w/Overdrive	FWD	Black	43,500	Good	*Leave as defaults*
3	4-Cyl, PZEV i-VTEC, 2.4L	Automatic, 5-Spd w/Overdrive	FWD	Black	56,500	Good	*Leave as defaults*
4	4-Cyl, 2.4 Liter	Auto, 4-Spd Overdrive	FWD	Black	71,000	Good	*Leave as defaults*
5	4-Cyl, VTEC, 1.8 Liter	Automatic, 5-Spd w/Overdrive	FWD	Black	85,000	Good	*Leave as defaults*
6	V6, Flex Fuel, 3.7 Liter	Automatic, 6-Spd	2WD	Black	32,500	Good	*Leave as defaults*
7	V8, Flex Fuel, 4.8 Liter	Automatic, 4-Spd w/Overdrive	2WD	Black	43,500	Good	*Leave as defaults*
8	4-Cyl, 2.5 Liter	Automatic, 6-Spd w/Overdrive	2WD	Black	56,500	Good	*Leave as defaults*
9	V6, 3.5 Liter	Automatic, 6-Spd w/Overdrive	2WD	Black	71,000	Good	*Leave as defaults*
10	4-Cyl, 2.4 Liter	Automatic, 5-Spd w/Overdrive & Shiftronic	FWD	Black	85,000	Good	*Leave as defaults*

Notes: These are the full specifications of the cars for which we obtained prices using quickvalues.com.

Appendix Table A13: **Count of Price Origin of Components of Price Index**

	Unimputed CZ price	Umimputed state price	Imputed CZ price using neighboring CZs	Imputed state price using unimputed states	Imputed state price using unimpted CZs	Imputed state price using imputed states
Automotive Expenses						
Auto Maintenance	438	0	0	0	0	5
Car Registration	0	443	0	0	0	0
Used Car	443	0	0	0	0	0
Child/Dependent Expenses	436	0	1	0	1	5
Clothing/Shoes/Jewellery	236	0	188	0	19	0
Electronics	437	0	0	0	1	5
Gasoline/Fuel	143	300	0	0	0	0
General Merchandise	434	0	3	0	1	5
Groceries	321	0	111	0	6	5
Healthcare/Medical	132	253	32	2	11	13
Hobbies/Entertainment	437	0	0	0	1	5
Office Supplies	438	0	0	0	0	5
Personal Care	434	0	3	0	1	5
Restaurants/Dining	30	0	199	0	132	82
Telecommunications	238	0	190	0	15	0
Utilities						
Electricity	442	0	1	0	0	0
Water	169	0	211	0	57	6

Notes: This table shows the price origin of components of the price index. For categories with multiple products combined into one index in the table – Clothing, Groceries, General Merchandise, Hobbies/Entertainment, Personal Care, and Restaurants/Dining – we report the price origin of the least precisely measured product in that commuting zone. Concretely, this means that all grocery products are unimputed in 321 commuting zones; there are 111 commuting zones in which at least one grocery product price is imputed using neighboring commuting zone prices.

Appendix Table A15: **Correlations Among Components of Main Price Index**

	Housing	Automotive Expenses	Child/Dependent Expenses	Clothing/Shoes/Jewellery	Electronics	Gasoline/Fuel	General Merchandise	Groceries	Healthcare/Medical	Hobbies/Entertainment	Office Supplies	Personal Care	Restaurants/Dining	Telecommunications	Utilities
Housing	1														
Automotive Expenses	-.166	1													
Child/Dependent Expenses	.492	.191	1												
Clothing/Shoes/Jewellery	.280	-.106	.134	1											
Electronics	.216	.140	.506	.145	1										
Gasoline/Fuel	.637	-.217	.325	.146	.188	1									
General Merchandise	.602	.164	.929	.179	.565	.441	1								
Groceries	.287	.136	.525	.136	.300	.143	.526	1							
Healthcare/Medical	.500	-.111	.252	.166	.181	.579	.356	.0602	1						
Hobbies/Entertainment	.136	.366	.239	.106	.299	-.126	.234	.208	-.111	1					
Office Supplies	.677	.0507	.773	.267	.495	.537	.885	.567	.394	.139	1				
Personal Care	.610	.123	.736	.219	.340	.392	.756	.605	.327	.173	.830	1			
Restaurants/Dining	.714	-.0637	.645	.275	.373	.656	.748	.415	.450	.136	.749	.613	1		
Telecommunications	-.208	.00396	-.00144	-.248	.00453	-.196	-.0796	-.142	-.0975	.0362	-.212	-.126	-.250	1	
Utilities	.610	-.107	.464	.334	.199	.594	.540	.369	.470	-.115	.671	.626	.677	-.277	1

Notes: N = 443. Correlation matrix of all non-constant components of main price index.