

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

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

January 2026

Abstract

Income differences across US cities are well documented, but little is known about the level of market-based consumption that residents are able to afford. We provide estimates of market consumption by commuting zone for households in a given income or education group, and we study how they relate to local prices. We use data that track all household expenditures for 5% of US households in 2014. To measure prices at the commuting zone $\times$ income level, we build local consumption price indices that aggregate commuting-zone specific prices from over 140 distinct products. We find that geographical differences in local prices are especially large for low-income households. The spatial standard deviation of the price indexes for the low-income group is almost double that one for the high-income group. We then relate the real consumption that low-skill and high-skill households enjoy in each commuting zone to the price index. We find that for college graduates, there is no relationship between real consumption and local prices, suggesting that college graduates living in cities with high prices—including the most expensive coastal cities—enjoy a real consumption on average similar to college graduates with the same observable characteristics living in cities with low prices—including the least expensive Rust Belt cities. By contrast, we find a significant negative relationship between real consumption and local prices for high school graduates and high school drop-outs, indicating that expensive cities offer lower real consumption 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 20.9% decline in real consumption.

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

1 Introduction

Geographic location plays an important role in an individual’s economic well-being. Moving from the 10 lowest paying cities to the 10 highest paying cities leads to a 27% increase in earnings (Card et al., 2021). The average worker is willing to pay 24% of earnings to consume the amenities of the 75th percentile city, relative to the amenities of the 25th percentile city (Diamond, 2016). Expensive cities such as San Francisco and New York have housing prices nearly double those of rust belt cities. Beyond housing costs, we know little about how local prices vary and how they aggregate with housing to lead to geographic differences in consumption.¹ In particular, despite the fundamental role of consumption in economic well-being, there is limited direct empirical evidence on the differences in consumption across cities and how they relate to local prices. While we know that large, high housing price cities tend to have jobs that offer higher nominal earnings, and small, low housing price cities tend to have jobs that offer lower nominal earnings, we know little about where market-based consumption is the highest. Are residents of expensive cities consuming more compared to residents of affordable cities? 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 build local consumption price indices that vary by commuting zone (CZ) and income group. We begin by providing a descriptive analysis of spatial differences in local prices, their sources, and how they vary by income group. We then combine our price indices with observed household expenditure levels to measure spatial differences in market-based consumption in “real terms” for each income group. We relate our measure of real consumption to the relevant local price index to answer two main questions: (a) Is real consumption higher or lower in CZs where local prices are high compared to CZs where local prices are low? (b) Is the relationship between real consumption and local prices the same for high-skill and low-skill households?

Our measure of market consumption—or standard of living—is not a utility metric since it does not include the non-market amenity value of cities. Measuring the total utility value of cities, inclusive of amenities, is an important topic, as reflected by a growing and active literature (Diamond, 2016; Giannone, 2017; Monte et al., 2018; Fajgelbaum and Gaubert, 2020; Piyapromdee, 2021); however, it is not our focus. Instead, we focus on a sub-component of utility: the geographic dispersion in market-based consumption across the income distribution, a topic that has received less attention in the literature. Every paper studying the utility value of cities implicitly takes a stand on market consumption levels in each city to measure utility, with no local price data beyond housing and no generalized consensus on the best way to do this. Some papers assume housing prices do not impact other local good prices, but do model trade costs in the good market (Monte et al., 2018; Rossi-

¹An exception is Bertrand and Morse (2016), who use the CEX to study consumption among the low-income by state. There is a 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); Moretti (2013); Jaravel (2019, 2024); Hornbeck and Moretti (forthcoming).

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

Hansberg et al., 2019).³ Another approach assumes housing is the only local good, while everything else is priced nationally (Ganong and Shoag, 2017; Piyapromdee, 2021). Others allow a share of non-housing consumption to be priced the same as housing, serving as a proxy for many local goods prices, which are higher in areas with high housing-costs (Moretti, 2013; Diamond, 2016; Giannone, 2017; Fajgelbaum and Gaubert, 2020). These differences in calibration and modeling assumptions are quantitatively important and mostly untested. Albouy (2008) shows that the implied amenity values of cities depend crucially on how much weight is placed on local house prices in determining consumption differences across space.

Just the weight on housing alone is a central parameter for all of these models. Local price variation beyond local housing prices is beyond the scope of all these papers, but it is something that we study directly. The only study that we are aware of that examines non-housing local prices has focused exclusively on groceries that can be tracked by scanner data (Handbury, 2019; Handbury and Weinstein, 2015). We incorporate these grocery prices into our analysis, but we also include local prices and expenditures for all other goods and services.

To measure the value of consumption expenditures, we use a 5% sample of US households’ linked bank and credit card transaction data from 2014. We observe all debit and credit card transactions, check and ACH payments, and cash withdrawals conducted every day in 2014. Relative to existing data sources on consumption, such as the CEX, our combined dataset has important advantages. Our consumption data is comprehensive and includes virtually all purchases conducted by individuals in our sample, and 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 the main consumption categories.

The data has detailed geographical information and, unlike the CEX, a sample size large enough that we have enough observations to cover most commuting zones. Our data, however, have important limitations. The main issues are 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); furthermore, not all accounts can be linked at the family level.

Our baseline price indices are a Laspeyres index—which mimics the index used by the BLS to estimate the official national CPI—and a GEKS-Fisher index, used by the World Bank for cross country purchasing power parity comparisons. We also examine three alternative price indices based on different assumptions, including those that correct for differences in variety and supply of goods across cities (Handbury and Weinstein, 2015; Handbury, 2019). A key aspect of our indexes is that we allow them to be heterogeneous across income groups to approximate taste differences and non-homotheticities across the income distribution. Differences across income groups in the indices turn out to be empirically important.

Our consumption price indices aggregate commuting-zone-specific prices from over 140 distinct product categories and more than a million distinct products. Examples of our data sources include

³These papers do allow other goods and services to have local prices, but they are calibrated to match trade flows of manufacturing goods. They do not use direct data on households’ local expenditures or consumption of non-housing goods, nor do they use data on non-housing local prices.

city-level prices of grocery goods at the twelve-digit barcode level (UPC) from the NielsenIQ Retail Scanner data; average cable TV prices from the Federal Communications Commission (FCC), by county; the prices of the ten most common models of used cars from Kelley Blue Book, by county; the prices of a kWh of residential electricity from the U.S. Energy Information Administration, by county; prices of a gallon of Unleaded Regular gas from GasBuddy, by city; and the price of health services from the Healthy Marketplace Index, by county.

While our micro-data on local prices is spatially granular and covers many goods and services, it does not cover all goods and services in the U.S. economy. Even when they cover a good category, we often have only a fraction of all possible brands. For example, we collected data on the ten most common models of used cars, as mentioned above; however, the number of car models available is much larger. As a result of these limitations, it is important to keep in mind that our price indexes are likely to contain some measurement error when interpreting our estimates. The BEA produces a MSA-specific price index, but it does not allow it to vary by income group. If we aggregate our indexes by combining the income-specific indexes into one, we find that the correlation between the BEA index and our Laspeyres and GEKS-Fisher indexes is 0.93 and 0.92, respectively—a useful diagnostic of the reliability of our indexes.

Empirically, we find that spatial differences in local price indices are significantly larger for low-income households than for high-income households. For example, the local prices faced by low-income households (post tax income <\$50,000) in the most expensive city—San Jose, CA—are 66% higher than in the median commuting zone, Cleveland, OH, whereas the San Jose-Cleveland price gap is only 35% for high-income households (post tax income >\$200,000). The standard deviation of the price indices across all CZs for the low-income group is almost double that of the high-income group. This reflects the larger share of expenditure devoted by low-income households to housing, and it underscores the importance of using local price indexes that vary not just by location but also by income.

We also find that the consumption item most responsible for the spatial variation in our price indexes is housing, because its share of consumption is the largest and its price varies over space more than the prices of any other goods. In addition, the prices of most non-housing goods and services in our data tend to be higher in areas with more expensive housing. Quantitatively, the fraction of spatial variation in the price index that is explained by housing costs is 95% and 90% for low- and high-income households, respectively. An implication of this finding is that local price indices can be well-approximated by using data only on local housing costs, especially for the low-income group. This finding can potentially be useful for researchers who need measures of local prices but do not have access to all the data necessary to build a complete price index.

We use our price indices to deflate observed expenditures in order to obtain measures of consumption by commuting zone and income group measured in *real* terms. We compare the real consumption levels of households located in expensive cities to those of households with the same income located in affordable cities. We uncover vast differences in real consumption, especially among low-income families. Low-income families who live in the most affordable commuting zone

enjoy a level of consumption measured in real terms that is almost double that of families with the same income who live in the least affordable commuting zone. The corresponding number for high-income households is 66%. The real consumption gap between the most and least expensive cities for low-income households is slightly smaller than the price index gap between these cities because low-income households increase their spending (and thus lower their saving rate) in expensive cities. Higher income households do not have different savings rates across low and high cost CZs.

The analysis up to this point compares the real consumption of residents in expensive and affordable cities while holding their nominal income constant. However, 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. In the main and final part of the analysis, we allow income to vary based on place of residence and measure the real consumption that low- and high-skill households can expect in each US commuting zone, accounting for geographical variation in both local prices and expected income.

We find that for college graduates, there is no significant bi-variate relationship between expected real consumption and local prices. A regression of log expected consumption on log price index across all commuting zones yields a coefficient of 0.014 (0.056). This suggests that college graduates located in cities with high prices enjoy expected real consumption similar to that of college graduates with the same characteristics located in cities with low prices.

The reason is that expensive cities offer incomes high enough to exactly offset the higher prices. College graduates in San Francisco and New York, for example, enjoy a high level of real consumption despite being exposed to local price indexes near the top of the distribution because their expected income is high enough to compensate for the high prices.

For less educated households, the picture that emerges is markedly different. San Francisco and New York offer incomes to high school graduates near the top of the distribution, but not high enough to offset local prices (Autor, 2019, 2020); thus, their real consumption is well below the median. Across all CZs, the elasticity of real consumption of high school graduates with respect to the price index is -0.195 (0.028), indicating that expensive cities offer real consumption that is 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 an 8.7% decline in their real consumption. Moving from the commuting zone with the lowest price index to San Jose would imply a decline in real consumption of 11.3%.

The negative relationship between real consumption and local prices is even steeper for high school drop-outs. The elasticity for this group is -0.373 (0.033). Moving from Cleveland to San Jose implies a 16.6% decline in real consumption. Moving from the cheapest CZ to San Jose implies a 21.7% decline in real consumption.

Since the consumption of college graduates is uncorrelated with local prices, while the consumption of less skilled groups declines with local prices, consumption inequality within a commuting

zone increases significantly with local prices. In particular, we find that the difference in real consumption between high- and low-skill households living in the same commuting zone is much larger in expensive commuting zones than in affordable commuting zones.

Overall, we conclude that the real consumption of less skilled households declines significantly with local prices, while the consumption of college graduates does not. This finding is consistent with the growing concern that high cost cities have become unaffordable for the middle class and low-income households (Autor, 2019, 2020). The concern appears particularly serious for low-skill households, who are increasingly exposed to 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.

An important question for future work is how economically significant differences in real consumption can exist across communities within the US for less skilled households. One possible explanation is that less skilled households are less mobile than highly skilled ones and are less able to arbitrage away geographical differences in consumption. Lower mobility could reflect spatial frictions that impede relocation (such as credit constraints or lack of information); or it could be an optimal choice on the part of households if it reflects idiosyncratic ties to local areas (such as proximity to one’s friends and family). A second possible explanation is the presence of better local non-market amenities in expensive cities, whose utility offsets differences in market consumption for less skilled households. However, for our findings to be explained by amenities, it would have to be the case that high school dropouts have a higher willingness to pay for the amenities of expensive cities like New York and San Francisco compared to college graduates—a possibility that is inconsistent with existing evidence (Diamond, 2016). A third possibility is that expensive cities offer more upward mobility to the children of their less skilled residents. Consistent with this hypothesis, we find that measures of upward mobility from Chetty et al. (2014) are positively correlated with our price index, indicating that less skilled residents might be willing to accept lower consumption at present to offer their children a higher future income.

The remainder of the paper is organized as follows. Section 2 presents the conceptual framework. Sections 3 and 4 describe the data and the price indexes. Section ?? presents descriptive facts about geographical differences in the cost of living and consumption by income group. Section 6 presents our estimates of consumption by skill level. Section 7 concludes.

2 Conceptual Framework: Market Consumption and Local Amenities

Consumption of market goods is a fundamental contributor to the utility of a household. Quantifying geographical differences in market consumption is arguably an important step in ultimately understanding geographical differences in utility. Non-market local amenities, such as clean air, pleasant weather, or crime, also play a role in household utility. Local amenities are difficult to measure in a comprehensive way. From an empirical point of view, this presents a challenge, since in equilibrium, geographical differences in the consumption of market goods likely correlate with the desirability of local amenities (Roback, 1982). In this section, we formalize the concept and definition of geographic variation in market consumption and clarify what assumptions are needed

to identify it empirically in a setting with local amenities.

Consider household i , of type k , in year t , with current wealth W_{it-1} and living in commuting zone j_{t-1} . Each year, the household chooses a CZ to live in j_t , a vector of consumption quantities across all goods and services sold, $\mathbf{C}_{it} = (C_{1it}, \dots, C_{nit}, \dots, C_{Nit})$ based on the following value function maximization:

$$V_{it}(W_{it-1}, j_{t-1}) = \max_{\mathbf{C}_{it}, j_t} U_{kt}(\mathbf{C}_{it}, a_{it}, A_{jt}, j_{t-1}, j_t) + \beta E_t(V_{it+1}(W_{it}, j_t)) \quad (1)$$

$$\text{subject to } W_{it} = W_{it-1} + I_{ijt} - \mathbf{P}_{jt}\mathbf{C}_{it}.$$

The per-period flow utility depends on the consumption vector $\mathbf{C}_{it}$, a vector of idiosyncratic tastes for each CZ a_{it} , a vector of non-market amenities A_{jt} , and the origin and destination CZs (j_{t-1}, j_t) to reflect possible moving costs. The household enters year t with wealth W_{it-1} , earns income in CZ j of I_{ijt} and purchases its preferred consumption vector $\mathbf{C}_{it}$ at prices $\mathbf{P}_{jt} = (P_{1jt}, \dots, P_{njt}, \dots, P_{Njt})$, leaving it with wealth W_{it} .

The model in equation (1) is very general and places no restrictions on the functional form of the per-period utility function, income levels, amenities, idiosyncratic preferences, or how expectations are formed about future prices. In this full generality, the concept of a consumption “level” or a consumption index that is comparable across CZs is ill-defined. The reason is that the amenity vector of a CZ may interact with the marginal utility of consumption for a subset of consumption goods, so that one cannot compare consumption levels across space without also incorporating this vector of amenities. For example, consider the case where CZs with cold weather lead households to prefer consuming hot chocolate, while CZs with hot weather lead households to prefer consuming ice-cream. In this case, the preference ordering of consumption vectors varies across CZs and one cannot make cross-sectional comparisons of consumption. Only cross-sectional comparisons of utility would be well-defined.

We assume that the per-period utility function U_{kt} can be written such that there exists a function $F_k(C_{it})$ that maps the vector of consumption C_{it} to a scalar value such that: $U_{kt}(\mathbf{C}_{it}, a_{it}, A_{jt}, j_{t-1}, j_t) = U_{kt}(F_k(\mathbf{C}_{it}), a_{it}, A_{jt}, j_{t-1}, j_t)$. Here $F_k(\mathbf{C}_{it})$ is the consumption index, implying that that consumption sub-component of utility can be measured and aggregated into an index, regardless of the CZ of residence. This restriction rules out the possibility that sub-components of the consumption vector interact with the amenities or idiosyncratic tastes in the utility function, as described above. While this does place some restrictions on the utility function, essentially all prior work on spatial differences in utility makes this restriction or stronger functional form restrictions. Monte et al. (2018) and Redding and Weinstein (2020) assume that $F_k(\mathbf{C}_{it})$ is a CES aggregate of the components of the consumption vector $\mathbf{C}_{it}$; Giannone (2017) assumes that $F_k(\mathbf{C}_{it})$ is Stone-Geary; Diamond (2016) and Piyapromdee (2021) assume that $F_k(C_{it})$ is Cobb-Douglas; and Fajgelbaum and Gaubert (2020) make this same restriction.

Assuming $F_k(\mathbf{C}_{it})$ is homothetic, one can write the consumption expenditure functions as:

$$\begin{aligned}
& \min_{\mathbf{C}_{it}} \mathbf{P}_{jt} \mathbf{C}_{it} \text{ such that: } F_k(\mathbf{C}_{it}) = C_{ikjt} \\
& = C_{ikjt} * \min_{\mathbf{C}_{it}} \mathbf{P}_{jt} \frac{\mathbf{C}_{it}}{C_{ikjt}} \\
& = e_k(\mathbf{P}_{jt}, 1) * C_{ikjt}
\end{aligned} \tag{2}$$

C_{ikjt} is the value of the consumption utility achieved, $C_{ikjt} = F_k(\mathbf{C}_{ijt}^*)$ when the household chooses its optimal consumption vector, $\mathbf{C}_{ijt}^*$, to maximize equation (1). The price index of CZ j for households of type k is thus defined as: $P_{jk} = e_k(\mathbf{P}_{jk}, 1)$, where $e_k(\mathbf{P}_{jt}, 1)$ is the consumption expenditure function that minimizes expenditure subject to achieving a consumption utility value of 1.⁴

Now consider a household of type k in CZ l : they maximize their value function to achieve the level of consumption utility C_{iklt} . Combining the first and third lines in equation (2) gives:

$$\begin{aligned}
P_{lk} * C_{iklt} &= \mathbf{P}_{lt} \mathbf{C}_{ilt}^* \\
\frac{P_{jk}}{P_{lk}} P_{lk} * C_{iklt} &= \mathbf{P}_{lt} \mathbf{C}_{ilt}^* \\
P_{jk} * C_{iklt} &= \mathbf{P}_{lt} \mathbf{C}_{ilt}^* \frac{P_{jk}}{P_{lk}}
\end{aligned} \tag{3}$$

Equation (3) tells us that one can take the total expenditure of household i living in CZ l , $\mathbf{P}_{lt} \mathbf{C}_{ilt}^*$, multiplied by the ratio of the price indices of CZ j and CZ l to measure the cost household i would need to spend in CZ j to achieve the same level of consumption utility as what they received in CZ l . Normalizing the price index in CZ j to 1, this gives us the value of the consumption utility function: $F_k(\mathbf{C}_{ilt}^*) = C_{ikjt}$. We refer to this as the standard of living of household i in CZ l .

In this model, amenities can influence consumption decisions in two key ways. First, we allow for arbitrary substitutability and complementarity of the aggregate consumption index with each sub-component of the amenity vector A_{jt} and each household's vector of idiosyncratic tastes for each CZ a_{it} . Prior work on spatial differences in utility either assume utility derived from consumption and amenities is additively separable (Diamond, 2016; Moretti, 2013) or multiplicative (Rossi-Hansberg et al., 2019; Monte et al., 2018). Second, the desirability of local amenities clearly plays a key role in choosing a CZ of residence, and this may change the demand for consumption of goods and services in that CZ. A clear implication of CZ demand driven by high quality amenities is that housing prices are likely to be higher high amenities CZs. Our model explicitly allows for this force to impact consumption demand and local prices in equilibrium. Despite the relationship between local prices and local amenities levels, our framework shows that a consumption index within the broader per-period utility function is well-defined.

⁴If preferences are non-homothetic, then the price index depends on the utility value at which the expenditure function is minimized and the concept of a single price index for type k households is ill-defined. We accommodate non-homotheticities across the income distribution by allowing income groups to have different tastes. We do not incorporate non-homothetic effects driven by differences in local prices across space. We can allow for this, but it complicates how one interprets spatial differences in consumption.

The key insight that we harness in the following sections is that we can measure the total expenditure ($\mathbf{P}_{it}\mathbf{C}_{ilt}^*$) of each household living in their chosen CZs, which they selected by solving equation (1). If total expenditure can be observed in data, we can deflate these total expenditures by the relative price indices between the CZ of residence of household i and a focal normalizing CZ, j . The main implication for our purposes is that within households of type k , comparing these deflated consumption expenditures across CZs measures the difference in per-period consumption utility these households are receiving across space. Spatial differences in consumption utility can be measured using data that capture local consumption expenditures and local price indices, without having to model the broader utility or value function that depends on the amenities and determines location choice.

3 Data

3.1 Household Data on Income and Consumption Expenditures

The source of our household data is a firm that provides financial software to banks. The data is 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. The selection of our sample is based on the banks with which the firm that provided the data works. 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. Details on the construction of the sample are in the 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. Ganong and Noel (2019) deal with this problem using the same approach. In addition, we require accounts to have at least \$10,000 of annual income and \$1,000 of annual expenditures. If these

restrictions leave us with the main bank accounts of households, we expect to miss 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.

3.1.1 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. 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 generally follows the ACS distribution. 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 \$48,837, 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.

For our analysis to be valid, it is important that our income data is representative not only at the national level but also at the local level. Appendix Figure A1 shows a tight relationship between CZ median income measured in the ACS and in our data, with a slope of 0.914 (0.100).⁶

Our income and spending measures have 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. In principle, this is a problem for our estimates, especially if cash earnings are more prominent in expensive cities and the elasticity of cash earnings to local housing prices is large. In practice, only 7% of workers get paid in cash in the US (Foster et al, 2023), and they are mostly concentrated among unbanked individuals (Appendix of FDIC, 2023) and low-income workers. Since our sample excludes all unbanked and households with income below \$10,000, we suspect that any bias from cash payments, if it exists, is likely to be modest in our context.

Second, we cannot observe some government transfers. Our data include income from Social Security, Disability Insurance, and EITC—since these transfers are deposited into the household’s

⁵We miss 0% of data for 70% of the sample and 26% of data for 30% of the sample.

⁶Measurement error stemming from imperfect geographical matching and standard sampling error would lead to a slope of less than one, even if both data sources were representative. Commuting zones are not available in the ACS, so we cross-walk using PUMAs, which sometimes span CZ borders, which leads to measurement error.

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. In practice, the omission of Food Stamps, TANF, and housing assistance does not appear to be an important source of bias. In Section 6, 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.

3.1.2 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 \$79,223. 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 next sub-section, where we discuss Panel B.

As a second way to probe the quality of our expenditure data, in Figure A2 we compare the fraction of consumption expenditure by mean of payment in our data with estimates from the

⁷We 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.

⁸In 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.

⁹Our bank data have a slightly different sample than the NIPA data. Our data are restricted to households with bank accounts that earn at least \$10k of post-tax income, while NIPA includes all households. According to the SCF, our sample has 12% higher income than the average US households, which would lead our data to have a higher household spending level than NIPA. On the other hand, we miss spending out of unlinked bank accounts from other banks of multi-banked households. According to the SCF, un-linked bank accounts likely lead us to miss 7.8% of household spending. Combining these offsetting effects suggests we should overestimate spending by about 4%. We end up underestimating spending by 6%. Thus, this back-of-the-envelope calculation suggests that we are in the right ballpark.

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.

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

3.1.3 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.

Appendix Figure A4 plots the log size of our sample in each commuting zone against the log number of households from 2012-2016 American Community Survey (ACS). There appears to be a tight link between our sample size and the corresponding number of households in the ACS, with an R-squared of 0.81. The slope is 1.340 (0.028), indicating that we under-sample rural areas and over-sample larger cities. This likely reflects the geographical presence of the banks in our sample, which includes the majority of the ten largest banks in the US. These banks’ locations are skewed to larger, urban areas. We use weights to adjust for sample representativeness, where weights are the ratio of the number of households in a given commuting zone in ACS data to the corresponding number in our data. In practice, weighting does not impact our results.

3.1.4 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. A 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.

The solution typically adopted by the literature is to measure the cost of the housing and health services consumed (see, for example, Aguiar and Bils (2015)). In practice, this means (a) adjusting the household’s healthcare and housing expenditures to represent the cost of services consumed (b) adjusting the household income to include these housing rents (since one can think of the home owner as paying rent to himself) and to include the employers and employee contributions toward health insurance costs.

We follow the literature, and make these two adjustments. Specifically, for health expenditures, we use the Medical Expenditure Panel Survey to 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.¹⁰ For housing, we face the additional limitation that the share of total expenditures that is spent on housing cannot be accurately quantified in our data because many consumers pay their rent with checks and mortgages with bill-pay transfers to banks. While the value of these transactions is included in our measure of total expenditure, our data puts them into a category called “Unclassified”. To quantify the share of expenditures that is spent on housing and properly measure housing services for homeowners, we adopt the same methodology and same data that the BLS employs in measuring the CPI (Poole et al. (2005); Bureau of Labor Statistics (2007)). Namely, for renters, we estimate housing expenditures by commuting zone and income group using mean contract rent from the ACS. For homeowners, the BLS uses a measure of “rent equivalent” from the CEX, which is defined as the rental value of their home if they were to rent it out.¹¹ We also add the expected owner-occupied rents back to income. Details are in Appendix B.

After these two adjustments, we return to the comparison of our data to the NIPA and CEX. In Panel B of Figure 2, we compare our adjusted average total expenditure against NIPA and the CEX. (The NIPA data is now in its “raw” format, since we have adjusted our bank data to make the 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 \$92,965.¹²

3.1.5 Validating Categories of Consumption Expenditures

The main focus of 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 addi-

¹⁰Using the Medical Expenditure Panel Survey we estimate the relationship between total health care expenditure and out of pocket spending. We use the estimate 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.

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

¹²In this Panel, 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: $53,495 - 6,149 + 10,895 + 2,382 + 6,284 = 66,907$.

tional way to validate our expenditure data, Figure A5 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.¹³ 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 A5 indicate that our data match well NIPA data, both in terms of overall mean household consumption and in terms of consumption by category.

3.1.6 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 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.

3.2 Data on Local Prices and Expenditure Shares

To estimate our price indexes, we need data on local prices and expenditure shares.

(A) Prices of Consumption Items. To measure commuting-zone specific prices, we combine data on more than a million distinct products belonging to 140 distinct product categories from 11 different datasets. Here we outline the data sources. We provide more details in Appendix C.

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 NielsenIQ 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 NielsenIQ 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

¹³These 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.

zone. The coefficients on the CZ dummies represent the mean price of a given product group in a given CZ, holding constant the mix of 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 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. We generalize the BLS method to allow for 381 different types of housing to be available in each commuting zone, each with its own price, instead of a single housing good. This allows spatial variation in housing prices to be heterogeneous across different housing types. Appendix Table ?? shows that physical housing consumption is quite different across the income distribution. The 23% of the high income group’s housing expenditure is on single family homes built since 1990 with eight or more rooms and four or more bedrooms. The low income’s housing expenditure on this type is only 2.1%. The low income spend 4.6% of their housing expenditure on apartments built in the 1970s with three or fewer rooms and zero or one bedrooms. The high-income group spends only 0.6% of their housing budget on this type. The housing types differ by observables in the ACS (number of rooms, number of bedrooms, the number of housing units in the building, the year the building was built, and basic amenities). Houses are assets, and their prices reflect both the user cost as well as expectations of future appreciation. We measure monthly rent for each of these 381 housing types by the mean contract rent reported in the ACS for each type of housing in each CZ. Housing rents reflect the flow cost of housing services consumed in a year. See Appendix C for more details.

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

- 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%

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

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. This is a Laspeyres-style index based on 4.8 billion private-sector health insurance claims.

- 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 GasBuddy. This is crowd-sourced gas price data at the station level.

- Clothing/Shoes/Jewelry: 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 Pricelist. 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 NielsenIQ 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 spatial differences in chain restaurant pricing is representative of spatial differences in non-chain restaurant pricing.

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 6.36%, 9.59%, and 15.34%, respectively.

(B) Expenditure Shares. We follow the methodology that the BLS uses to calculate expenditure shares to compute the CPI. The expenditure shares for 21 high-level consumption categories by income group are from our bank data and are listed in Appendix Table A2. Recall that these shares match well the NIPA shares (Section 3). Five categories among the 21 in our data are very

broad: Housing, 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 Groceries, General Merchandise, Hobbies/Entertainment, and Personal Care. 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. For housing, we allow for 381 different housing types based on observed characteristics in the ACS. Expenditure shares for these 381 differ across our three income groups to allow for heterogeneous consumption of housing qualities.

4 Local Cost of Living Indexes

We build 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. 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. While the Bureau of Economic Analysis (BEA) produces annual estimates of local prices that cover most 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 (Handbury, 2019; Jaravel, 2019, 2024). In practice, we find significant differences in the price indexes for high- and low-income households and their spatial distribution. We also find that if we aggregate our price indexes by combining the income-specific indexes into one average index for all income groups, such index is highly correlated with the BEA index, a finding that we interpret as supportive of our approach.

4.1 Main Price Indexes

In our main analysis, we use two baseline indexes—Laspeyres and GEKS-Fisher—while in the robustness analysis, we use four alternative indexes based on different assumptions.

(A) Laspeyres 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). 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 Laspeyres price index for commuting zone j and income group k is defined as:

$$P_{j,k}^{\text{Laspeyres}} = \sum_{i \in I} \frac{p_{i,j}}{\bar{p}_i} \cdot s_{i,k} \quad (4)$$

where $p_{i,j}$ is the price of good i in commuting zone j ; $\bar{p}_i$ 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.

The choice of using a Laspeyres index as one of our baseline indexes is motivated by the fact that it is the index used by the BLS to compute the official price index. The index in Equation 4 is a useful and transparent starting point. 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 is that the Laspeyres index is subject to substitution bias, since it does not account for the second-order utility benefits of allowing consumers to substitute away from high price goods. In addition, the Laspeyres index does not allow for variation in variety and supply of goods and services across space, which has been shown to be quantitatively important for measuring local prices (Handbury and Weinstein, 2015; Handbury, 2019).¹⁵

(B) GEKS-Fisher Index. This index is based on methods developed to measure purchasing price parities (PPP) across countries (Deaton and Heston, 2010). In our context, the PPP is the exchange rate at which one CZ’s income would be converted into another CZ’s income to achieve a given utility level. A Fisher price index is the Geometric mean of a Laspeyres and Paasche indexes for a given pair of cities. A desirable property of the Fisher index is that it is a second-order approximation for the true 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 is a generalization of the Fisher index that imposes transitivity. Specifically, a GEKS-Fisher index for a commuting zone j takes a geometric mean of pairwise Fisher index products based on all possible paths to j from an arbitrarily chosen commuting zone $j_0 \neq j$:

$$P_{j,k}^{\text{GEKS-Fisher}} = \left(\prod_{j' \in J} P_{j_0,j',k}^{\text{Fisher}} P_{j',j,k}^{\text{Fisher}} \right)^{\frac{1}{|J|}} \quad (5)$$

where $P_{j_1,j_2,k}^{\text{Fisher}} = \sqrt{P_{j_1,j_2,k}^{\text{Laspeyres}} / P_{j_2,j_1,k}^{\text{Laspeyres}}}$ and $P_{j_1,j_2,k}^{\text{Laspeyres}} = \sum_{i \in I} (s_{i,j_1,k} \cdot \frac{p_{i,j_2}}{p_{i,j_1}})$. j' is a commuting zone from the set of commuting zones J , and j_1 and j_2 refer to specific commuting zones in that set. $P_{i,j}$ is the price of good i in commuting zone j , and I refers to the set of goods.

¹⁵While 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.

4.2 Alternative Price Indexes

The baseline Laspeyres and GEKS-Fisher indexes defined above restrict all income groups living in the same commuting zone to face the same set of prices. This choice of baseline indexes is motivated by the notion that it is conceptually correct to think of consumers in a city as facing similar prices irrespective of their income and what bundle of goods they choose. We recognize, however, that in practice, high income and low income households may be living in different parts of a commuting zone and this, combined with spatial frictions—travel costs, for example—may result in high- and low-income consumers facing different prices. For example, while in theory high-income consumers in the Upper East side of New York city face the same prices as low-income consumers in the Bronx, in practice this may not apply to many consumption items. High-income consumers in the Upper East side are unlikely to travel to the Bronx to take advantage of a cheaper espresso and similarly, low-income consumers in the Bronx are unlikely to travel to the Upper East Side to take advantage of a more affordable burger (if it exists). This is consistent with the literature on sources of inflation heterogeneity which suggests that a lot of dispersion is caused by different prices being paid for the same product (Kaplan and Schulhofer-Wohl, 2017). (On the other hand, the notion that consumers in a location face similar prices irrespective of their income likely applies to housing and big ticket items, like cars and appliances.) To see if the assumption that all income groups living in the same commuting zone face the same set of prices makes a difference, in the robustness analysis we build Laspeyres and GEKS-Fisher indexes that allow prices to vary across income groups within the same commuting zone.

To further assess the robustness of our baseline estimates, we present additional estimates based on three alternative indices. We discuss the conceptual differences here, and refer the interested reader to Appendix C for the full details of their construction.

a) 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.

b) The Nested CES allows for more complex substitution patterns between products. It can also account for differences in the choice set across commuting zones by allowing for differences in product variety. There is an elasticity of substitution between the 21 high-level expenditure categories, and then expenditure category-specific substitution elasticities across product groups. Within each product group, there is a product-group specific substitution elasticity between unique varieties of products. We follow Handbury and Weinstein (2015) and Broda and Weinstein (2010) in building the nested CES index and correcting for variety.¹⁶

c) The Geary-Khamis index is based on PPP methods, like the GEKS-Fisher index. It is a

¹⁶To 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).

Paasche index that compares the local prices in a given CZ to nationwide average prices. The weights on the relative prices differences between CZs come from each local CZ. This is the method used by the BEA to estimate local price indices.

In practice, we will find below that our main findings are not sensitive to choice a particular index. The reason is that the alternative indexes appear highly correlated with the main ones and with one another. Appendix Table A4 shows the pairwise correlations. For parsimony, the table focuses on the version of the indexes for all income groups combined. (The income-specific versions is available upon request). Despite the fact that the alternative indexes are based on different assumptions and functional forms, they tend to be generally correlated with the two main indexes. The table also report correlations for the version of the Laspeyres and GEKS-Fisher indexes where the prices are allowed to vary by CZ and income group (not just CZ). By allowing for local prices to also vary by income group, we allow for the possibility that local prices differ by income level within CZs (allowing for price differences between low- and high-income neighborhoods).

The last row is particularly informative because it focuses on the price index produced by the BEA. Recall that the main limitation of the BEA index is that it is not available by income group. Aggregating our indexes by combining the income-specific indexes into one, and comparing it to the BEA offers a useful diagnostic to assess the validity of our indexes. While there is some overlap, the sources of our price data are mostly independent of the BEA’s sources. If measurement error in the BEA index is uncorrelated with measurement error in our data, the correlation with the BEA index can be interpreted as a reliability ratio for our indexes. Entries in the table indicate that the correlations between the BEA index and our Laspeyres and GEKS-Fisher indexes are 0.93 and 0.92, respectively. We interpret this finding as a validation of our indexes.

5 Geographical Differences in Prices and Consumption, by Income Group

5.1 Cost of Living

Columns 1 to 3 of 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 the Laspeyres index. 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 1.656, 1.553, 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). Examples of other expensive commuting zones include New York, NY; Washington, DC; and Anchorage AK (where most consumption goods need to be imported from afar). The least expensive commuting zones for low-income residents are Waycross, GA; Jonesboro, AR; and Batesville, AR with price indexes equal to 0.852, 0.849, and 0.844, respectively. In the last two columns we compare the Laspeyres and GEKS-Fisher price indexes. For parsimony, we report the version of the index for all income groups combined. The ranking and the estimates are similar.

The geographical price differences revealed by our indexes are economically large, and this is particularly true for low income families. Based on the Laspeyres index, the cost of living in San Jose is estimated to be 96% and 46% higher than the cost in Batesville 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. The corresponding numbers for the GEKS-Fisher index are 90% and 41%.

Figure 3 displays the spatial dispersion of the our price indexes across all 443 commuting zones. It confirms that the cost of living index for low-income households exhibits significantly higher spatial variation than the index for high-income households. The standard deviation of the Laspeyres index, for example, equals 0.107 and 0.056 for the low- and high-income group, respectively. Similarly, the 75-25 and 90-10 percentile differences for low-income households are much larger than the corresponding differences for high-income households. This finding is important, as it indicates that cities with a high cost of living index are expensive both for high and low income households, but they are significantly more expensive for low-income households. This reflects the fact that low-income households put a higher weight on housing expenditure, which is the item in the consumption basket whose price varies the most across commuting zones. The share of housing of high- and low-income households is 0.079 and 0.279, respectively. The alternative price indexes provide a similar picture, as shown in Appendix Table A5. In all cases, the low-income indexes exhibit highest spatial variation, followed by the middle-income indexes and the high-income indexes.

Another interesting feature of Figure 3 is that the spatial distribution of cost of living 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 example, the Laspeyres index for low-income families indicates that 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 and for other indexes.

The Importance of Housing Costs. The consumption item that is most responsible for the spatial variation in the cost of living is housing. This is due to the fact that its share of consumption is the largest and its price varies over space more than the price of any other goods (Moretti, 2013). 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, both because they have a lower share and their price varies less over space. Moreover, 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 Francisco than in Cleveland, holding quality constant, because retail space and

labor are more expensive in San Francisco.¹⁷ We regressed the price of non housing consumption categories on rent and confirmed that the price of most categories is positively correlated with rent. The only exception was the price of telecommunication, which is negatively correlated with rent, likely reflecting the lack of competition in the cable market in more sparsely populated low rent areas compared to more densely populated high rent areas.

Quantitatively, it is important to establish how much of the spatial variation in overall cost of living reflects variation in cost of housing vs. variation in the price of non-housing goods and services. The R-squared of the regression of our low- and high-income price indexes on housing rent in Appendix Table A6 are 0.95 and 0.90 for Laspeyres and 0.95 and 0.86 for GEKS-Fisher. This finding indicates that most of the spatial variation in cost of living stems from variation of housing costs. An implication of this finding is that local price indices can be well-approximated by using data only on local housing costs, especially for the low-income group. This finding may be useful for researchers who are interested in studying cost of living but don't have access to all the data necessary to build a complete price index. In this case, data on housing costs, weighted appropriately by income group, would capture most of the geographical differences across CZs in the price index.¹⁸

Accounting for Variety. Differences across commuting zones in the variety of products that are locally available are potentially important. Using data on grocery products from NielsenIQ, Handbury (2019) and Handbury and Weinstein (2015) have shown that correcting for differences across cities in product variety has a large impact on measured prices. Handbury (2019) shows that the correlation of the variety-corrected price index and city income is negative — so that richer cities have lower effective prices — while the price indices that don't include the variety adjustment show a positive correlation of city income. We find similar results with our variety-corrected nested CES index applied to Automotive Expenses, Child/Dependent Expenses, Clothing/Shoes/Jewelry, Electronics, Gasoline/Fuel, General Merchandise, Groceries, Healthcare/Medical, Hobbies/Entertainment, Office Supplies, Personal Care, Restaurants/Dining, Telecommunications, and Utilities, as shown in Appendix Table A7. The table also shows that this negative correlation holds for the overall price index. However, even with the variety correction, we find the overall price index is higher in high-income CZs, even though many sub-components of the price index, such as groceries, electronics, and general merchandise are lower in high-income CZs. Below, we find that our empirical findings based on the two baseline indexes are not sensitive to using the Nested CES index.

Changes in the Price Indexes Over Time. Our empirical analysis focuses on 2014. To

¹⁷Using NielsenIQ data, DellaVigna and Gentzkow (2019) find limited variation in grocery prices across locations.

¹⁸Appendix Table A6 also shows that the coefficients on rent are 0.244 (0.008) and 0.440 (0.011) for high- and low-income households, respectively. Recall that the share of housing in the indexes of high- and low-income households is 0.079 and 0.279, 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 other goods tend to be higher in areas with more expensive land.

document how spatial variation in local price has changed over the long run, we estimated our Laspeyres indexes for years after and before 2014. For 2020, we collected updated price data for most consumption categories and recomputed the indexes.¹⁹ Estimating price index for early years proved more difficult, since reliable price data at the CZ level are not available for most consumption categories. Thus, for the years 1980, 1990, 2000, 2010 we imputed the Laspeyres indexes: for each of the 145 detailed product groups, we estimated a CZ-level bivariate regression of log price vs. log rent with hedonic adjustment using the 2014 data, and used the estimated coefficients and rent data for each year to impute non-housing prices.

Panel of A of Appendix Table A8 reports spatial variation of the overall-income indexes across the 443 CZs. For year 2020, we report two versions. The preferred one is “2020*”, which makes use of actual non-housing prices. The one labeled “2000” follows the imputation used for 1980-2010 and is reported only for comparison with earlier years. The key finding is that prices become spatially dispersed over time. For example, the standard deviation of the index increases from 0.049 in 1980 to 0.97 in 2020. The 75-25, 90-10, and max-min are all larger today than they were in 1980. For the 2020 indexes, using direct variation in non-housing prices generate similar, albeit slightly higher spatial dispersion compared with using imputed variation: 0.104 vs. 0.097, respectively. Estimates by income group in Panel B show that spatial dispersion is larger for low-income households and smaller for high income households.

5.2 Consumption

We now study how market consumption varies across commuting zones as a function of local cost of living. Here we hold household income fixed. The goal is to compare the consumption experienced by households in expensive commuting zones to the consumption experienced by households with the same income in affordable commuting zones. In the next section, we study how consumption varies across commuting zones when we allow household income to vary across areas.

We use our price indexes to deflate expenditures and quantify “real” consumption for each commuting zone and income group. Specifically, 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. We then estimate the following model:

$$\ln(C_{h,j,k}/P_{j,k}) = \delta_{j,k} + \beta_k \ln Y_{h,j,k} + \varepsilon_{h,j,k} \quad (6)$$

where $P_{j,k}$ is either $P_{j,k}^{\text{Laspeyres}}$ or $P_{j,k}^{\text{GEKS-Fischer}}$; $\delta_{j,k}$ is a vector of commuting zone-income group fixed effects; 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. We are interested in $\delta_{j,k}$, which measures the conditional mean log consumption in commuting zone j of income group k . We report estimates where consumption

¹⁹We use the 2020 NielsenIQ Consumer Panel data for *Child/Dependent Expenses*, *Electronics*, *General Merchandise*, *Groceries*, *Hobbies/Entertainment*, *Office Supplies*, and *Personal Care*. We update our price data to the year 2024 for *Telecommunications*, 2021 for *Healthcare/Medical*, 2023 for electricity and 2021 for gas in *Utilities*, 2024 for *Clothing/Shoes/Jewelry*, and 2019 for *Restaurants/Dining*.

is evaluated at post-tax income equal to \$30,000, \$80,000, and \$285,000 for low-, middle-, and high-income consumers, respectively.

Appendix Table A9 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 households, ranked based on consumption estimated using the Laspeyres index.²⁰ 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 table shows that households in expensive areas can afford a much lower level of consumption compared to households with the same income in cheaper areas. The three commuting zones with the highest Laspeyres consumption of low income families (column 2) are Elizabeth City, NC; Traverse City, MI; Champaign, IL. They have level of consumption measured in real terms equal to \$47,939; \$43,525; and \$42,721. Three cities with the lowest consumption are Honolulu, HI; San Francisco, CA; and San Jose, CA. The corresponding values are \$26,453; \$25,771; and \$24,257. Other examples of commuting zones with low consumption are New York, NY (\$28,278); Washington, DC (\$28,049); Seattle, WA (\$30,139); and Los Angeles, CA (\$28,564). The GEKS-Fischer consumption is similar (column 4).

These geographical differences are economically large, and this is especially true for low-income households. Based on Laspeyres consumption, low-income households who live in the top commuting zone in the top group enjoy a level of consumption that is 98% higher than households with the same income who live in the bottom commuting zone in the bottom group. Put differently, moving from the bottom CZ to the top CZ and holding household income constant would imply almost a doubling of market consumption measured in real terms. The corresponding number for high-income households is 66%, as shown in Appendix Table A10. The fact that the geographical differences in standard of living are larger for low-income households reflects the finding in the previous section that the cost of living index for low-income households exhibits higher spatial variation than the index for high-income households.

To see more systematically the relationship between consumption and cost of living, Figure 4 plots log consumption against log income-group-specific price index across all 443 commuting zones in our data. 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. Since we are holding income fixed, it is probably not too surprising to find that consumption and local prices are negatively correlated. We are interested in the quantitative differences in consumption across CZs as a function of local prices and how these differences vary by income. The elasticities of consumption with respect to income-group-specific local prices are economically large. Specifically, the Laspeyres elasticities are -0.899 (0.009), -0.984 (0.021), and -1.066 (0.038) for low-, middle-, and high-income households, respectively. This means that a 10% higher cost of living index

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

is associated in an almost exactly proportional decline (10.7%) in consumption for high income households, and a less than proportional (9.0%) decline in consumption for low income households. The GEKS-Fisher elasticities are similar.²¹

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). The permanent income hypothesis predicts that the elasticity of consumption with respect to the (permanent) 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 (Laspeyres p-value = 0.0001).

The finding that the elasticity is larger for high-income than low-income households could be explained by the fact that the latter are closer to a minimum subsistence level, so that small consumption cuts cost more in terms of utility. Alternatively, it is in principle 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); or 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.

Since the consumption of low-income households declines less than proportionally with the price index, and we are holding income fixed, we should also see that low income households in expensive commuting zones tend to have negative savings. Panel A of Figure 5 shows the correlation between 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 our two price indexes. The figure confirms that low-income households in expensive commuting zones have a higher probability of negative savings than low-income households in cheap commuting zones.²² Panels B and C in the figure show a positive relationship between the price index two indicators of financial distress: 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”).²³ However, the evidence in Panel B and C is indirect, noisy and should be considered suggestive at best.

Overall, we draw two main conclusions. First, the difference in the amount of consumption that high-income households can afford in cheap and expensive cities is quantitatively very large. Second, while the consumption of high-income households declines proportionally to the price index, the consumption of low-income households declines less than proportionally and, as a consequence, a significant number of low income households in expensive cities has negative savings in a given year.

²¹The finding that the elasticity is larger for high- and middle-income consumers than low-income consumers is robust to using alternative price indexes although the magnitudes vary (Appendix Table A11).

²²For 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.

²³This is consistent with Keys et al. (2020), who find a sizable effect of location on personal bankruptcy.

This last finding could indicate that low-income households in expensive areas tend to have trouble in making ends meet, and there is some indirect evidence from overdraft fees and bankruptcy, but we can't draw definitive conclusions on this point: 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, so we leave this question for future researchers.

Consumption in Physical Units. We 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 analysis, we don't need any deflation, because we observe physical quantity of consumption *directly from the raw data*. Thus, this provides direct, "model free" validation of the main evidence.

We regress the mean-adjusted quantity of a product purchased by a household in a year on the log of the price index controlling for household income and other characteristics. A product is defined as a twelve-digit UPC. This definition is very detailed, as there are 785,609 UPC codes in the data. In Appendix Table A12 and Appendix Figure A6, we find that the majority of the coefficients are negative, confirming that households in more expensive commuting zones tend to have lower levels of consumption of grocery items and some non-grocery items than households with the same income level in less expensive commuting zones. In addition, 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 elasticities for all three groups are much smaller than the elasticities estimated for overall consumption above. This reflects the fact that most of the consumption items in the NielsenIQ data are grocery items and many grocery items are necessities, which are harder to cut. It also reflects the fact that groceries exhibit less geographic price variation, making them a relative bargain in expensive cities.

6 Where is Standard of Living the Highest? Geographical Differences in 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. For some individuals, income is completely independent of location. This would be the case, for example, for individuals whose only source of income is social security. In this case, estimates in the previous section would be informative of the level of consumption that this person may expect in each commuting zone. In most cases, however, the income level that a specific worker can attain depends on their location: for a given level of human capital, some cities offer high labor earnings (and therefore high income), while others offer low labor earnings (and therefore low income). Empirically, 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. Ultimately, the amount of market consumption that a given household can afford 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 this section we seek to measure the market consumption 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 evidence in this section allows us to answer two questions: (a) Is expected consumption 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 expected consumption and local cost of living the same for high- and low-skill households?

We use 2012–2016 ACS data in combination with consumption estimates from the previous Section to estimate the income and consumption that a given household may expect in each commuting zone as a function of education and demographics. We then relate the expected consumption in each CZ to the relevant price index, by income group. In particular, 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 consumption. Next, for each commuting zone and skill level, we predict mean pre-tax income, post-tax income (computed using Taxsim) 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.

The analysis in the previous section 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 assume 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 exam-

ple, 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.

6.1 Consumption in the Largest 20 Commuting Zones

We focus on three skill groups, based on the schooling level of the household head: 4-year college or more; high school or some college; and less than high school. The maps in Appendix Figure A7 show the geographical distribution of consumption by skill level. Since the maps are not easy to read, Tables 2, 3, and 4 present our findings for the largest thirty commuting zones, ordered by pre-tax nominal income for each skill group. The estimates hold constant the combination of household characteristics that define a type (education, age, gender, race, Hispanic origin, education, marital status, and children).

Table 2 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). New York, NY, Newark, NJ, and Boston, MA follow closely. Column 3 reports the corresponding after-tax income obtained by subtracting personal federal and state taxes from column 1. Columns 5 and 7 show our estimates of the levels of expected consumption based on the Laspeyres and GEKS-Fisher indexes. They quantify the standard of living that a family with this level of schooling can expect in each commuting zone. For San Francisco and New York, the entries in columns 5 and 7 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 two cities is modest. For example, in terms of GEKS-Fisher consumption, San Francisco and New York are at the 91st and 95th percentile respectively (column 8). Thus, despite some of the highest costs of living in the US, these cities remain near the top 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.²⁴

²⁴Los Angeles and San Diego are examples of cities that experience large 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 they drop by 41 and 79 percentiles, respectively. Other examples of cities with large negative percentile changes are Seattle (-42), Minneapolis (-54), and Denver (-58).

Table 3 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 many expensive cities appear to offer significantly lower consumption than affordable cities. For example, while in terms of adjusted pre-tax income, San Francisco and New York remain near the top, in terms of GEKS-Fisher consumption they are below the median, dropping to the 31st and 42nd percentile, respectively. High school graduates in these cities can expect to earn some of the highest nominal pre-tax incomes in the nation, but the nominal pre-tax incomes there are simply not high enough to offset the high costs of living.²⁵

Table 4 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 GEKS-Fisher consumption of high school drop-outs in San Francisco and New York is near the bottom of the distribution, at the 3rd and 10th percentiles, respectively. 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.

An interesting feature of our estimates that is worth noting is that spatial variation in consumption is much smaller than spatial variation in pre-tax income. Across all 443 commuting zones, the standard deviation in adjusted pre-tax income for high skill households is 10,765, while the standard deviation in the two consumption measure is only 4,643 and 4,692—or less than half. For low- and middle-skill households, the corresponding numbers are 4,692 vs. 2,800 and 2,712; and 6,140 vs 2,780 and 2,723, respectively. 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.

6.2 Consumption vs. Local Prices

To understand more systematically how standard of living of each skill group varies as a function of local prices, Figure 6 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.

The top panel is based on the Laspeyres index. We observe a positive correlation between pre-tax income and local cost of living for all three skill groups. This is hardly surprising, as more expensive cities have long been known to offer higher 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 1.075 (0.060), 0.907 (0.038), and 0.733 (0.046) 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 (an increase of 10.8%) and a less than proportional increase in expected pre-tax income for the low skill group (an increase of only 7.3%).

²⁵Los Angeles, Chicago, and Denver are other examples of cities in the bottom tercile of the GEKS-Fisher consumption distribution. Boston, on the other hand, is an exception: it is a high cost city that offers high consumption to its middle-skill residents.

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

The findings for consumption are the most important part of the evidence. For high-skill households, there is essentially no relationship between consumption and cost of living. The coefficient is 0.014 (0.056)—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. The reason is that, compared with affordable cities, expensive cities offer incomes high enough to exactly offset the difference in cost of living and personal taxes, so that consumption ends up being independent of local prices.

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 -0.195 (0.028), 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 8.7%. Moving from the commuting zone with the lowest cost of living index (Batesville) to the commuting zone with the highest index (San Jose) would imply a decline in the standard of living by 11.3%.

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 -0.373 (0.033), implying vast geographical differences in consumption. Moving from Cleveland to San Jose implies a 16.6% decline in the standard of living. Moving from Batesville to San Jose implies a 21.7% 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.

The bottom panel is based on the GEKS-Fisher index. The picture that emerges is qualitatively similar. Quantitatively, the elasticity of consumption for high skill households is 0.117 (0.065) and it is not statistically different from zero. The corresponding estimates for the middle- and low-skill groups are -0.121 (0.033) and -0.328 (0.040)—negative and statistically significant like in the top panel.

Thus, we conclude that the consumption of less skilled households declines significantly with local prices, while the consumption of college graduates does not. The existence of large differences in the consumption of less skilled households between expensive and affordable CZs raises the question of how such differences can persist in equilibrium. First, it is possible that less skilled households are less mobile and have a lower ability to arbitrage away geographical differences in consumption. In turn, differences in mobility may reflect higher mobility frictions faced by less educated households (credit constraints or lack of information) or an optimal choice driven by idiosyncratic ties to local areas, such as proximity to one’s friends and family. Moreover, going to college is often associated with relocation, and that this in itself could contribute to a higher propensity to move later in life to areas with higher standard of living.

Second, it is possible that expensive CZs have non-market amenities that offsets differences in market consumption for less skilled households. A third possibility is that expensive cities offer more upward mobility to the children of their less skilled residents. In this scenario, less skilled residents are willing to accept a current lower consumption to offer their children a future higher income. To test this hypothesis, we relate two measures of upward mobility from Chetty et al. (2014) to our price index. The first measure is absolute upward mobility, defined as the expected rank of children whose parents are at the 25th percentile of the national income distribution. The second measure is the probability that parents with family income in bottom quintile of national parent income distribution have a child with family income in top quintile of national child income distribution. In Appendix Figure A8, we plot these two measures against our price index. The figure shows that high-cost-of-living cities offer higher upward mobility for low-income households. The slopes indicate that the spatial difference in mobility between high and low cost cities is both statistically significant and economically large.

6.3 Within-City Inequality in Consumption

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, Laspeyres consumption of college graduates in San Francisco, San Jose, and New York is 1.97, 1.94, and 1.94 times higher than consumption of high school drop-outs. The corresponding ratios in the three cheapest commuting zones, Batesville, AR, Jonesboro, AR; and Waycross, GA are 1.52, 1.57, and 1.64.

The top panel of Figure 7 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 Laspeyres slopes are 0.208 (0.033) and 0.386 (0.038), and the GEKS-Fisher slopes are 0.238 (0.037) and 0.445 (0.041), 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.²⁶

This finding 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.

6.4 City Size and College Share

We conclude by investigating how standard of living varies across cities as a function of two other city characteristics that are prominent in the literature on spatial wage differences: city size and share of residents with a college degree. The top panel in Figure 8 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. What is more interesting and novel 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, indicating that residents of large cities enjoy lower standard of living than residents of 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. While the elasticity is positive for high-skill households, it is negative for middle- and low-skill households.

Since cities' prices, population, and college shares are positively correlated, in Appendix Table A13 we investigate a multivariate regression where we include all three, we allow the elasticities to vary by skill group and we condition on the same set of controls used above. 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 and low-skill groups. The correlation of consumption and college share appears close to zero for all three groups, but the effects are somewhat 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.

²⁶Bertrand 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)

6.5 Robustness

One concern with our data is that our income variable 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, marital 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 A14 shows that our main empirical results are not particularly sensitive. The regression coefficients of log Laspeyres consumption (inclusive of government transfers) and log price index change from 0.014 to 0.004 for high skill, from -0.195 to -0.203 for middle skill, and remains -0.373 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.

Among means-tested transfer program that differs from state-to-state in the United States, Medicaid is the most expensive by far (Schmidt et al., 2025). Since passage of the Affordable Care Act (ACA) allowed states to expand the program’s eligibility to include low income able-bodied prime-age childless adults, which were originally not eligible. In 2014, 20 states expanded Medicaid eligibility accordingly, creating substantial state-level differences in the program’s scope (McPherson, 2024). To see if so that low income residents in expensive areas receive more Medicaid transfers than low income residents in affordable areas, we regressed the share of residents enrolled in Medicaid in 2014 on our price indexes. Our estimates fail to reject that the coefficient is equal to zero, indicating that CZs in states that expanded Medicaid eligibility do not have higher average local prices.

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 do not significantly affect our findings.

Finally, Appendix Table A15 shows estimates based on the alternative price indexes in Panel A. Quantitatively, the coefficients vary in magnitude, as one might expect. But in all cases, the estimated elasticities are most negative for the low skill group, and either not statistically significant or slightly positive for the high skill group, with the middle skill group in the middle. Finally, we find in Panels B and C that classifying the middle skill group further into some-college and high-school graduates or excluding retired population yield similar estimates as our main results.

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 than 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 21.7% decline in market consumption.

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

References


- AGUIAR, M. AND M. BILS (2015): “Has Consumption Inequality Mirrored Income Inequality?” *American Economic Review*, 105, 2725–2756.
- ALBOUY, D. (2008): “Are big cities bad places to live? Estimating quality of life across metropolitan areas,” Tech. rep., National Bureau of Economic Research.
- AMERICAN WATER WORKS ASSOCIATION AND RAFTELIS FINANCIAL CONSULTANTS, INC. (2022): “2014 Water and Wastewater Rate Survey,” .
- 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>.

- DEATON, A. AND A. HESTON (2010): “Understanding PPPs and PPP-based national accounts,” *American Economic Journal: Macroeconomics*, 2, 1–35.
- 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. AND W. F. DIAMOND (2024): “Racial differences in the total rate of return on owner-occupied housing,” Tech. rep., National Bureau of Economic Research.
- 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%20service%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.
- GANONG, P. AND D. SHOAG (2017): “Why has regional income convergence in the US declined?” *Journal of Urban Economics*, 102, 76–90.
- 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.
- (2024): “Distributional Consumer Price Indices,” *LSE*.
- KAPLAN, G. AND S. SCHULHOFER-WOHL (2017): “Inflation at the Household Level,” *Journal of Monetary Economics*.
- 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*.
- MCPHERSON, C. (2024): “How Magnetic can Welfare be?” .
- 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.
- ROBACK, J. (1982): “Wages, Rents, and the Quality of Life,” *Journal of Political Economy*.
- ROSSI-HANSBERG, E., P.-D. SARTE, AND F. SCHWARTZMAN (2019): “Cognitive hubs and spatial redistribution,” Tech. rep., National Bureau of Economic Research.
- 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.
- SCHMIDT, L., L. SHORE-SHEPPARD, AND T. WATSON (2025): “Did Welfare Reform End the Safety Net as We Knew It? The Record since 1996,” *The Journal of Economic Perspectives*, 39, 101–128.
- 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. CAMMENG (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 (\$67,623), 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 and housing expenditures more comparable. For our bank data, we add healthcare costs paid by the government and insurers; and we adjust housing costs for homeowners. For the CEX: we adjust housing costs for homeowners. See text for details.

Figure 3: Spatial Distribution of Price Indexes

(a) Laspeyres Price Index




(b) GEKS-Fischer Price Index




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

Figure 4: **Consumption vs. Price Index**



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


Figure 5: **Negative Savings, Overdraft, and 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. Values on x-axis are measured in log scale but we label actual values for easier interpretation. See text for details.

Figure 6: Pre-Tax Income, Post-Tax Income, and Consumption vs. Price Index by Skill

(a) Laspeyres price index




(b) GEKS-Fischer price index




Notes: We plot 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 CZs. Values on both x-axis and y-axis are measured in log scale but we label actual value.

Figure 7: 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 the price index across all commuting zones in the sample.

Figure 8: 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.

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

City	Laspeyres Low Income Price Index	Laspeyres Middle Income Price Index	Laspeyres High Income Price Index	Laspeyres Overall Price Index	GEKS-Fischer Overall Price Index
Most Expensive					
San Jose, CA	1.656	1.482	1.349	1.562	1.479
San Francisco, CA	1.553	1.416	1.308	1.478	1.423
Honolulu, HI	1.506	1.439	1.375	1.477	1.391
San Diego, CA	1.462	1.372	1.284	1.418	1.367
Santa Barbara, CA	1.457	1.359	1.280	1.406	1.371
New York, NY	1.413	1.316	1.243	1.359	1.314
Washington, DC	1.400	1.307	1.211	1.352	1.283
Los Angeles, CA	1.381	1.295	1.218	1.336	1.301
Anchorage, AK	1.371	1.294	1.245	1.336	1.267
White Plains, NY	1.343	1.267	1.198	1.304	1.247
Kapaa, HI	1.343	1.319	1.329	1.338	1.290
Newark, NJ	1.322	1.251	1.181	1.286	1.228
Edison, NJ	1.314	1.230	1.156	1.269	1.229
Ketchikan, AK	1.303	1.230	1.216	1.268	1.209
Miami, FL	1.282	1.233	1.169	1.261	1.217
Median					
Chincoteague, VA	1.002	1.009	1.010	1.005	1.000
Wichita, KS	1.000	1.026	1.042	1.019	1.019
Cleveland, OH	1.000	1.000	1.000	1.000	1.000
Lewiston, ID	1.000	1.007	1.005	1.004	0.995
Columbia, SC	1.000	1.013	1.023	1.012	1.006
Least Expensive					
Cape Girardeau, MO	0.872	0.923	0.959	0.904	0.910
Russellville, AR	0.871	0.911	0.938	0.893	0.899
Marquette, MI	0.870	0.903	0.960	0.888	0.866
Presque Isle, ME	0.869	0.904	0.957	0.890	0.874
Harrison, AR	0.868	0.922	0.957	0.901	0.900
Somerset, KY	0.868	0.917	0.948	0.898	0.901
Greenville, MS	0.866	0.921	0.959	0.898	0.897
Pikeville, KY	0.864	0.907	0.934	0.886	0.887
Paducah, KY	0.863	0.892	0.914	0.879	0.881
Jasper, AL	0.862	0.917	0.953	0.894	0.890
London, KY	0.861	0.910	0.950	0.889	0.885
Natchez, MS	0.861	0.929	0.962	0.903	0.909
Waycross, GA	0.852	0.889	0.924	0.873	0.867
Jonesboro, AR	0.849	0.898	0.919	0.878	0.882
Batesville, AR	0.844	0.895	0.925	0.873	0.873

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

City	<u>Adj Pre-tax Income</u>		<u>Adj Post-tax Income</u>		<u>Laspeyres Consumption</u>		<u>GEKS-Fischer Consumption</u>	
	value (1)	pct (2)	value (3)	pct (4)	value (5)	pct (6)	value (7)	pct (8)
1. San Jose, CA	143,935	100	111,323	100	65,846	63	69,179	86
2. San Francisco, CA	139,465	100	108,985	100	68,142	80	70,656	91
3. Washington, DC	138,555	100	105,959	100	70,690	93	74,332	98
4. New York, NY	130,567	99	102,409	99	69,958	90	72,322	95
5. Newark, NJ	127,971	99	101,405	99	72,232	96	75,572	99
6. Boston, MA	125,160	99	99,353	99	74,188	99	76,796	100
7. Hartford, CT	124,391	98	98,270	98	73,383	98	75,514	99
8. Philadelphia, PA	117,250	98	93,887	98	73,317	98	75,654	99
9. Los Angeles, CA	116,548	98	92,502	97	63,549	44	65,392	57
10. Baltimore, MD	116,400	97	91,300	97	70,597	92	72,735	95
11. Houston, TX	115,878	97	94,099	98	74,048	98	75,303	98
12. San Diego, CA	114,430	97	90,781	96	58,021	9	60,243	18
13. Chicago, IL	112,847	96	88,814	95	66,522	69	68,024	78
14. Dallas, TX	110,199	95	90,026	96	70,869	93	72,101	94
15. Seattle, WA	110,021	95	90,254	96	63,501	43	64,928	53
16. Sacramento, CA	107,682	94	85,967	93	61,827	33	63,026	36
17. Denver, CO	106,517	93	85,254	92	61,209	28	62,705	35
18. Atlanta, GA	104,703	92	82,065	90	66,005	65	66,953	70
19. Minneapolis, MN	103,540	91	82,160	90	61,639	30	63,029	37
20. Fort Worth, TX	103,097	91	84,586	91	69,322	87	70,643	91

Notes: Entries are average adjusted household pre-tax income, adjusted post-tax income, and consumption across the largest 40 commuting zones. For each variable, we report its corresponding unweighted percentile among all 443 CZs in our data.

Table 3: **Pre-tax Income, Post-tax Income, and Consumption — Middle Skill**

City	<u>Adj Pre-tax Income</u>		<u>Adj Post-tax Income</u>		<u>Laspeyres Consumption</u>		<u>GEKS-Fischer Consumption</u>	
	value (1)	pct (2)	value (3)	pct (4)	value (5)	pct (6)	value (7)	pct (8)
1. San Jose, CA	85,629	100	71,716	100	42,738	7	44,589	20
2. Washington, DC	84,271	100	68,508	99	46,474	49	48,725	78
3. San Francisco, CA	82,953	100	69,442	100	44,090	18	45,484	31
4. New York, NY	76,419	98	64,267	98	44,733	25	46,115	42
5. Newark, NJ	75,980	98	64,535	98	46,790	53	48,794	79
6. Boston, MA	74,897	97	63,377	98	48,806	81	50,361	92
7. Hartford, CT	73,803	96	62,655	97	47,561	64	48,802	79
8. San Diego, CA	73,711	96	62,379	96	40,993	1	42,356	3
9. Los Angeles, CA	73,176	96	61,955	96	43,509	12	44,588	20
10. Baltimore, MD	72,310	95	60,126	95	48,134	73	49,500	86
11. Seattle, WA	71,095	95	61,199	95	44,543	23	45,478	31
12. Philadelphia, PA	69,680	94	59,233	93	46,999	56	48,302	72
13. Sacramento, CA	68,895	93	58,806	93	43,796	15	44,533	19
14. Denver, CO	68,816	93	58,198	92	42,924	7	43,831	11
15. Chicago, IL	68,122	92	57,008	91	43,806	15	44,706	22
16. Minneapolis, MN	67,367	91	56,850	90	43,956	17	44,918	24
17. Houston, TX	67,335	91	57,978	92	47,756	67	48,478	75
18. Dallas, TX	65,811	89	56,882	90	46,285	46	47,017	54
19. Fort Worth, TX	65,370	88	56,438	88	47,761	67	48,603	75
20. Portland, OR	65,156	87	54,797	82	42,604	6	43,395	8

Notes: Entries are average adjusted household pre-tax income, adjusted post-tax income, and consumption across the largest 40 commuting zones. For each variable, we report its corresponding unweighted percentile among all 443 CZs in our data.

Table 4: **Pre-tax Income, Post-tax Income, and Consumption — Low Skill**

City	<u>Adj Pre-tax Income</u>		<u>Adj Post-tax Income</u>		<u>Laspeyres Consumption</u>		<u>GEKS-Fischer Consumption</u>	
	value (1)	pct (2)	value (3)	pct (4)	value (5)	pct (6)	value (7)	pct (8)
1. San Jose, CA	63,533	99	56,160	99	33,879	1	35,197	1
2. Washington, DC	62,939	99	54,147	99	37,593	17	39,352	37
3. San Francisco, CA	61,246	99	53,981	98	34,644	1	35,628	3
4. Newark, NJ	59,284	98	52,713	98	38,803	32	40,391	54
5. Boston, MA	57,837	97	51,418	97	41,023	66	42,267	80
6. New York, NY	57,312	96	51,175	97	36,101	5	37,162	10
7. Hartford, CT	55,984	95	50,251	95	39,343	39	40,303	53
8. Seattle, WA	55,750	94	49,863	95	37,281	14	38,030	19
9. San Diego, CA	54,723	94	48,998	94	32,940	0	33,903	0
10. Los Angeles, CA	54,138	93	48,470	93	34,751	2	35,510	2
11. Denver, CO	54,040	93	47,762	93	35,578	4	36,252	5
12. Baltimore, MD	53,382	93	46,881	90	38,868	33	39,925	47
13. Chicago, IL	53,352	92	46,670	89	36,621	9	37,330	12
14. Portland, OR	52,918	92	46,480	89	36,977	12	37,613	14
15. Philadelphia, PA	52,662	91	46,960	91	37,780	19	38,738	28
16. Minneapolis, MN	52,339	90	46,676	90	36,415	7	37,197	10
17. Sacramento, CA	51,454	88	46,386	88	35,364	3	35,891	3
18. Las Vegas, NV	50,948	86	45,687	84	38,805	32	39,113	33
19. Houston, TX	49,516	78	44,638	79	38,147	23	38,673	26
20. Dallas, TX	49,199	77	44,505	77	37,325	14	37,874	16

Notes: Entries are average adjusted household pre-tax income, adjusted post-tax income, and consumption across the largest 40 commuting zones. For each variable, we report its corresponding unweighted percentile among all 443 CZs in our data.

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 recalculate our “Healthcare/Medical” spending as a sum of total healthcare charges, total pharmacy charges, and the original non-recreation health

spending.²⁷ 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 scaling 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 5.23 (0.52), 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 3.63 (1.18).

Second, the healthcare spending category in the raw bank data consists of OOP non-drug health spending, OOP prescription spending, and health-related recreational spending that is not covered by insurance. Diamond et al. (2018) use various keywords contained in the transaction descriptions, train a machine learning model, and classify all raw healthcare transactions into these three types. We use their data to calculate the imputation relationships. Precisely, we regress the classified OOP non-drug health spending on raw healthcare spending interacted with income group indicators. We do the same for OOP prescription spending. We then use these regressions to impute OOP non-drug health spending and OOP prescription spending for all households in 2014.

Finally, for each household in the bank data, we use our estimated coefficients from the first step to impute total non-drug health and total prescription expenditures for a given level of observed OOP non-drug health and OOP prescription expenditures, respectively, from the second step. The mean and median for non-drug health expenditures are \$889 and \$259 for OOP spending; and \$6,021 and \$3,992 for total spending. The corresponding numbers for prescription expenditures is \$1,324 and \$485 for OOP spending; and \$6,109 and \$2,689 for total spending.

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,

²⁷Our 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).

centering around 2014. We define two measures of housing costs in the CEX data. The first measure is “housing costs to be subtracted” or owner costs, including purchase costs, closing costs, mortgage payments, and down payments. The second measure is “housing costs to be added” or 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. These regressions are reported in Appendix Table A17.

Second, since some of our price indices require expenditure shares at the CZ level, we need to translate our findings in the CEX to each CZ, even though the CEX does not report CZs in the data. To do this, we regress our defined housing costs from the CEX (both measures, separately) on property value, its squared term, post-tax income, and number of rooms separately by region $\times$ income group. These covariates are all variables available in both the CEX and ACS. 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 and how they vary across CZs. 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 $\times$ 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.

These three steps measure “housing expenditures” for all households in our data. Together with households’ total expenditures, this ultimately allows us to compute the housing shares by income group displayed in Appendix Table A2. In addition, we explore a “housing calculation alternative” (as labeled in our robustness price index). Instead of using the imputation relationship in CEX for owner-occupied housing, we directly impute annual rent-equivalent values for homeowners using $0.078 \times$ observed housing value, as estimated in (Diamond and Diamond, 2024). These two methods yield quantitatively similar expenditure shares on housing.

To calculate our price indexes, we need both the “shares” and “prices”. Finally, turning to the price component, we now describe how we measure housing prices across CZs in the pooled 2012-2016 ACS data.

- First, to ensure that the spatial differences in housing prices are not driven by differences in quality, we define granular housing types based on five housing characteristics: year built, unit structure, number of rooms, number of bedrooms, and presence of amenities (e.g., kitchen, stove, sink, shower, toilet, hotwater, and plumbing). There are 381 housing types (h) in total.
- Second, for a given housing type and CZ, we take the *mean of contract rents* $\overline{Rent}_{hj}$ observed among all renters who live in this CZ of this type. Note that rent is missing if there are no

renters. In all such cases, we regress type-specific rent (in logs) on average rents and these 5 types of property characteristics across CZs and then impute all missing values. Evidently, this imputation matters very little in the final calculation of the final rent since the types that need imputation are the less-common ones carrying lower weights.

- As described above, there are two ways to collapse the 381 type-specific rents to obtain the average-type rent for a given CZ. Both approaches use the same housing-type-specific CZ-specific rents $\overline{Rent}_{hj}$ from the second step as prices. However, the two approaches differ in type-specific housing expenditures that go into the housing share calculation, specifically:

$$\begin{aligned} HousingExpenditure_{hjk}^{MAIN} &= (N_{hjk}^{renter} + N_{hjk}^{owners}) \cdot \overline{Rent}_{hj} \\ HousingExpenditure_{hjk}^{SUPP} &= N_{hjk}^{renter} \cdot \overline{Rent}_{hj} + N_{hjk}^{owners} \cdot \left(\frac{0.078 \cdot \overline{PropVal}_{hj}}{12} \right) \end{aligned}$$

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 NielsenIQ-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, marital 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 NielsenIQ 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 prod-

uct 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 $k(h) \in \{\text{overall, low, middle, high}\}$ 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 estimate mean rents in each CZ for each of our 381 types of housing. In particular, we interact the following five housing characteristics to define “housing types” (n): (1) Year the structure was built (before 1950, 1950–1969, 1970–1989, and from 1990 onward); (2) Unit structure (one-family house, multiple family building, and other remaining structures); (3) Number of rooms (at most three rooms, four rooms, five rooms, six to seven rooms, and eight rooms or more); (4) Number of bedrooms (at most one bedroom, two bedrooms, three bedrooms, and four bedrooms or more); and (5) Presence of facilities (having all of the following facilities: kitchen, sink, shower, stove, hot water, and plumbing), ; and lacking at least one facility). There are $N = 381$ 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. We have also tried imputed rental prices for owner-occupied units using $0.078 \times$ observed housing value (Diamond and Diamond, 2024). These are listed in the tables as housing alternatives indices. This alternative method for imputed rents for owner occupied housing is used in both the rent differences across space (the local prices) and the expenditure shares on housing. Both approaches yield similar results.

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 NielsenIQ.

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 Millennials; baby boomers favored the Toyota Camry, which we also include in our 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 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 NielsenIQ data to get the costs of car maintenance in five UPCs. Refer to the NielsenIQ 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/Jewelry. 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 jewelry: 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 Pricelist, 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 missing restaurant prices in Alaska, we use our imputed Washington price.)

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

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 3).

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 NielsenIQ 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 $i \in \{\text{General Merchandise, Groceries, Hobbies/Entertainment, Personal Care}\}$. 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 NielsenIQ 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

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:

$$P_{j,k}^{\text{CES}} = \prod_{i \in I} \left(\frac{p_{i,j}}{\bar{p}_i} \right)^{\omega_{i,j,k}} \quad (7)$$

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 NielsenIQ 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 NielsenIQ. For high-level product categories not covered by NielsenIQ, there 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 NielsenIQ groups, we use UPC codes to identify unique varieties. For the rest of the product categories not covered by NielsenIQ, 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:

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

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

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 NielsenIQ 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_{img} 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:

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

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

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

$CEPI_{ij}$ 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:

$$\begin{aligned}
w_{ij} &= \frac{\frac{sh_{ij}-sh_i}{\ln sh_{ij}-\ln sh_i}}{\sum_{i' \in I} \left(\frac{sh_{i'j}-sh_{i'}}{\ln sh_{i'j}-\ln sh_{i'}} \right)}, & w_{gj} &= \frac{\frac{sh_{gj}-sh_g}{\ln sh_{gj}-\ln sh_g}}{\sum_{m \in M_i} \sum_{g' \in G_{im}} \left(\frac{sh_{g'j}-sh_{g'}}{\ln sh_{g'j}-\ln sh_{g'}} \right)}, \\
w_{imj} &= \frac{\frac{sh_{mj}-sh_m}{\ln sh_{mj}-\ln sh_m}}{\sum_{m' \in M_i} \left(\frac{sh_{m'j}-sh_{m'}}{\ln sh_{m'j}-\ln sh_{m'}} \right)}, & w_{imj} &= 1 \text{ for non-nielsenIQ categories.} \\
sh_{ij} &= \frac{\sum_{m \in M_i, g \in \{G_i, G_{im}\}} E_{gj}}{\sum_{i \in I} \sum_{m' \in M_i, g' \in \{G_i, G_{im}\}} E_{g'j}}, & sh_i &= \frac{\sum_{m \in M_i, g \in \{G_i, G_{im}\}} E_g}{\sum_{i \in I} \sum_{m' \in M_i, g' \in \{G_i, G_{im}\}} E_{g'}}, \\
sh_{gj} &= \frac{E_{gj}}{\sum_{g' \in G_i} E_{g'j}}, & sh_g &= \frac{E_g}{\sum_{g' \in G_i} E_{g'}}, \\
sh_{mj} &= \frac{\sum_{g \in G_{im}} E_{gj}}{\sum_{m' \in M_i, g' \in G_{im}} E_{g'j}}, & sh_m &= \frac{\sum_{g \in G_{im}} E_g}{\sum_{m' \in M_i, g' \in G_{im}} E_{g'}}.
\end{aligned}$$

E_g is national total expenditure on variety g . To measure the nested-CES price index, the only parameter that is not directly observed in the data is σ_{im} , the elasticity of substitution between varieties within the lower-level nests, which we calibrate. Indeed, if we were to assume that there were no variety differences across space, even this parameter would be directly inverted from the data, using the same Sato-Vartia weighting method.

For housing, we assume there is only one variety and it's available everywhere. For products with price data not from NielsenIQ, 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:

$$P_{j,k}^{\text{Geary-Khamis}} = \frac{\sum_{i \in I} (p_{i,j} \cdot q_{i,j,k})}{\sum_{i \in I} (\pi_{i,k} \cdot q_{i,j,k})} \quad (8)$$

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

D Consumption in Physical Units in NielsenIQ 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.

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.

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.

For each product group, we first convert ounce, pound, milliliter, liter, and quart to kilogram, assuming density of water ($1,000 \text{ kg/m}^3$). 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.

We regress the mean-adjusted quantity of a product 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.²⁸ We run one regression for each product, obtaining 785,609 elasticities. To summarize the results, we compute the average of the estimated elasticities for each of the 116 product groups in our data.

Table A12 show some examples. The second 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. The elasticity of consumption of carbonated beverages with respect to the Laspeyres cost of living index is -1.002 (0.107) and -1.510 (0.299) 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.

Since it is difficult to draw conclusions based on selected examples, Figure A6 plots the distribution of all the 116 estimated mean 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 72% are negative for low-, middle-, and high-income households. 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 -0.104, -0.139, and -0.477. 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

²⁸In 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.

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.²⁹

E Estimating Consumption by Education Group

Here, we describe in detail the data and the methodology used in Sections 6 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, marital status, age of household head and spouse, and wages of household head and spouse (if exists). We always use joint filing for households with the spouse present and use single filing otherwise. We subtract state taxes, federal taxes, and social securities (these are outputs from the software) from household pre-tax income to obtain household post-tax income. To make households in the ACS data consistent with those in our bank account data, we drop households with missing post-tax income, households with post-tax income less than \$10,000, and households not belonging to the 443 commuting zones identified in our data. These restrictions together leave 5,302,154 households in the ACS data.

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

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

1. Age — based on mean age of household head and spouse (if exists):
 - Less than 30 years old 446,250 (8.42%)
 - From 30 to less than 45 years old 1,249,376 (23.56%)
 - From 45 to less than 65 years old 1,647,023 (31.06%)
 - At least 65 years old 1,959,505 (36.96%)
2. Gender — based on a composition of household head and spouse (if exists):
 - Household head is male OR both are male 959,606 (18.10%)
 - Household head is female OR both are female 1,486,558 (28.04%)
 - One person is male and the other person is female 2,855,990 (53.86%)
3. Race — based on a composition of household head and spouse (if exists):
 - Household head is white OR both are white 4,190,909 (79.04%)
 - At least one person is nonwhite 1,111,245 (20.96%)
4. Hispanic Origin — based on a composition of household head and spouse (if exists):
 - At least one person has Hispanic origin 381,620 (7.20%)

²⁹We have run regressions relating the price of each goods to the cost of living index, conditioning on the same controls. We find that goods are more costly in expensive areas than in cheap areas. This is despite the fact that grocery goods can be considered traded. The median estimated price coefficients on the index for the low-, middle-, and high-income groups are 0.091, 0.133, and 0.230.

- 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 $\geq$ college OR household head is $\geq$ college 1,455,299 (27.45%)
 - One is $\geq$ college AND the other is $<$ college 683,094 (12.88%)
 - Both are $\geq$ highsch $<$ college OR head is $\geq$ highsch $<$ college 2,527,382 (47.67%)
 - One is $\geq$ 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 $v = 1, 2, \dots, 20$ 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 .

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: Consumption Expenditure by Mean of Payment: Bank and Fed




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.

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: **Households by CZ: 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). $R^2=0.8147$

Appendix Figure A5: **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 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 expenditure types defined consistently in each set and we aggregate some categories with definitions that do not 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.


Appendix Figure A6: **Distribution of Estimated Elasticities — NielsenIQ Data**




Notes: Each panel plots the distribution of estimated elasticities by product group. We weight by average household expenditure for each product group. Vertical lines denote the median. Elasticities are recovered from regressions of mean-adjusted quantity of consumption on the local price index controlling for household characteristics: income, size, age and presence of children, type of residence, composition, 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.

Appendix Figure A7: 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 A8: Upward Mobility Measures vs. Price Index



Notes: *Panel (a)* Absolute upward mobility is the expected rank of children whose parents are at the 25th percentile of the national income distribution. *Panel (b)* Probability child has family income in top quintile of national child income distribution conditional on having parents with family income in bottom quintile of national parent income distribution. For all panels, observations are weighted by commuting zone populations. Slopes and robust standard errors are reported on all plots. $N = 443$.

Appendix Table A1: **Summary Statistics**

	Overall Income (1)	Low Income (2)	Middle Income (3)	High Income (4)
Panel A. Raw Measures of Income and Expenditure				
Post-tax Income				
Mean	81,010.77	29,638.88	91,121.28	448,699.56
Median	52,955.79	29,495.21	81,021.33	288,099.91
Expenditure				
Mean	74,631.26	29,902.46	82,135.14	406,517.88
Median	47,750.03	27,652.95	71,200.55	251,517.69
Panel B. Adjusted Measures of Income and Expenditure				
Post-tax Income				
Mean	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 (1)	Low Income (2)	Middle Income (3)	High Income (4)	Standard Deviation (5)
Automotive Expenses	2.65%	2.06%	3.24%	5.10%	0.482
Charitable Giving	0.30%	0.22%	0.40%	0.49%	0
Child/Dependent Expenses	0.38%	0.24%	0.56%	0.54%	0.029
Clothing/Shoes/Jewellery	4.75%	4.00%	5.45%	8.66%	0.251
Electronics	1.13%	0.98%	1.29%	1.46%	0.135
Education	0.69%	0.56%	0.78%	1.86%	0
Financial Fees	0.74%	0.59%	0.86%	1.66%	0
Gasoline/Fuel	5.15%	5.14%	5.27%	3.46%	0.053
General Merchandise	10.53%	9.01%	12.45%	10.97%	0.237
Groceries	6.48%	5.46%	7.68%	7.94%	0.156
Healthcare/Medical	17.55%	21.01%	13.36%	13.87%	0.137
Hobbies/Entertainment	3.27%	2.88%	3.67%	4.65%	0.169
Housing	23.02%	27.89%	17.80%	7.87%	0.305
Insurance	3.92%	3.09%	4.80%	6.61%	0
Office Supplies	0.19%	0.16%	0.21%	0.37%	0.055
Personal Care	1.56%	1.35%	1.76%	2.56%	0.197
Printing and Postage	0.18%	0.16%	0.18%	0.41%	0
Restaurants/Dining	7.82%	7.69%	8.06%	6.71%	0.232
Telecommunications	4.96%	4.09%	6.01%	5.80%	0.358
Travel	2.17%	1.74%	2.57%	4.31%	0
Utilities	2.58%	1.69%	3.59%	4.70%	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 NielsenIQ Product Groups**

	Low Income	Medium Income	High Income
Groceries			
Baby Food	0.08%	0.15%	0.65%
Baked Goods - Frozen	0.45%	0.43%	0.31%
Baking Mixes	0.41%	0.41%	0.37%
Baking Supplies	0.53%	0.59%	0.54%
Beer	1.67%	1.21%	1.72%
Bread, Baked Goods	4.52%	4.13%	3.72%
Breakfast Food	0.84%	1.08%	1.29%
Breakfast Food, Frozen	0.59%	0.62%	0.61%
Butter, Margarine	0.98%	0.90%	0.78%
Candy	2.67%	2.91%	2.19%
Carbonated Beverages	3.49%	2.90%	2.75%
Cereal	1.66%	1.87%	1.91%
Charcoal, Logs	0.07%	0.09%	0.06%
Cheese	3.36%	3.54%	3.71%
Coffee	1.88%	1.99%	2.42%
Condiments, Gravies, Sauces	1.49%	1.45%	1.72%
Cookies	1.14%	1.14%	0.83%
Cottage Cheese, Sour Cream	0.66%	0.67%	0.60%
Crackers	0.77%	0.84%	0.57%
Desserts, Fruits, Toppings	0.23%	0.25%	0.24%
Desserts, Gelatins, Syrup	0.49%	0.54%	0.49%
Detergents	1.20%	1.21%	1.47%
Disposable Diapers	0.17%	0.31%	0.65%
Dough Products	0.35%	0.35%	0.34%
Dressings, Salads, Prepared Foods	5.90%	6.56%	4.75%
Eggs	0.85%	0.76%	0.65%
Flour	0.13%	0.14%	0.12%
Fresh Meat	0.39%	0.42%	0.53%
Fresh Produce	5.60%	6.43%	7.29%
Fresheners, Deodorizers	0.34%	0.36%	0.29%
Fruit - Canned	0.43%	0.30%	0.40%
Fruit, Dried	0.29%	0.36%	0.38%
Gum	0.13%	0.21%	0.22%
Household Cleaners	0.54%	0.58%	0.74%
Household Supplies	0.71%	0.65%	0.56%
Ice	0.01%	0.01%	0.00%
Ice Cream	1.45%	1.31%	1.31%
Jams, Jellies, Spreads	0.61%	0.58%	0.65%
Juice, Drinks - Canned-Bottled	1.75%	1.80%	1.93%
Juice, Drinks - Frozen	0.07%	0.05%	0.03%
Laundry Supplies	0.60%	0.66%	0.82%
Liquor	0.73%	1.55%	1.39%
Milk	2.16%	2.39%	2.50%
Nuts	0.87%	1.15%	1.24%
Packaged Meats - Deli	4.08%	3.83%	3.18%
Packaged Milk, Modifiers	0.68%	0.70%	0.47%
Paper Products	3.47%	3.43%	2.99%
Pasta	0.41%	0.51%	0.50%
Pet Food	5.55%	4.03%	3.96%
Pickles, Olives, Relish	0.37%	0.36%	0.34%
Pizza, Snacks - Frozen	0.94%	1.19%	1.13%
Prepared Food - Dry Mixes	0.92%	0.94%	0.88%
Prepared Food - Ready-to-Serve	0.77%	0.75%	0.91%
Prepared Foods - Frozen	2.80%	2.96%	3.44%
Puddings, Dessert - Dairy	0.06%	0.07%	0.01%
Salad Dressings, Mayo, Toppings	0.83%	0.72%	0.72%
Seafood, Canned	0.38%	0.39%	0.31%
Shortening, Oil	0.47%	0.45%	0.42%
Snacks	3.61%	4.09%	3.89%
Snacks, Spreads, Dips - Dairy	0.28%	0.40%	0.38%
Soap, Bath Additives	0.69%	0.63%	0.73%
Soft Drinks - Non-Carbonated	0.99%	0.94%	0.90%

Soup	1.27%	1.12%	1.12%
Spices, Seasoning, Extracts	0.53%	0.48%	0.51%
Sugar, Sweeteners	0.43%	0.41%	0.31%
Table Syrups, Molasses	0.12%	0.12%	0.23%
Tea	0.78%	0.73%	0.93%
Tobacco	3.15%	2.16%	0.43%
Unprepared Meat, Poultry, Seafood	7.15%	6.90%	7.72%
Vegetables - Canned	1.00%	0.91%	0.92%
Vegetables - Frozen	1.18%	1.05%	1.11%
Vegetables, Grains - Dried	0.39%	0.49%	0.32%
Wine	1.56%	1.18%	2.50%
Wrapping Materials, Bags	0.79%	0.88%	0.90%
Yeast	0.00%	0.00%	0.00%
Yogurt	1.08%	1.33%	2.04%
General Merchandise			
Batteries, Flashlights	12.90%	13.77%	9.03%
Canning, Freezing Supplies	0.99%	1.00%	0.91%
Cookware	3.71%	2.97%	3.86%
Floral, Gardening	9.91%	13.64%	14.40%
Glassware, Tableware	4.62%	4.04%	4.05%
Hardware, Tools	5.23%	6.44%	5.31%
Housewares, Appliances	29.84%	26.00%	33.38%
Insecticides, Pesticides, Rodenticides	4.20%	4.42%	3.15%
Kitchen Gadgets	8.66%	9.24%	9.11%
Light Bulbs, Electric Goods	14.01%	11.53%	12.06%
Party Needs	0.37%	0.22%	0.17%
Photographic Supplies	2.63%	4.14%	2.60%
Seasonal	0.98%	0.98%	0.38%
Sewing Notions	0.41%	0.48%	0.33%
Shoe Care	0.25%	0.22%	0.28%
Soft Goods	1.29%	0.91%	0.98%
Hobbies/Entertainment			
Books, Magazines	9.30%	8.45%	5.06%
Pet Care	90.03%	90.77%	94.01%
Toys, Sporting Goods	0.67%	0.78%	0.93%
Personal Care			
Cosmetics	2.88%	4.30%	3.15%
Cough and Cold Remedies	5.64%	6.33%	6.14%
Deodorant	1.94%	2.04%	2.26%
Diet Aids	0.55%	0.58%	1.72%
Ethnic Haba	0.10%	0.04%	0.06%
Feminine Hygiene Products	0.36%	0.28%	0.20%
First Aid	2.29%	1.98%	2.02%
Fragrances - Women	0.67%	0.48%	0.92%
Grooming Aids	1.13%	0.99%	0.93%
Hair Care	5.59%	6.06%	7.11%
Medications, Remedies, Health Aids	48.94%	42.78%	36.66%
Men's Toiletries	0.33%	0.43%	0.41%
Oral Hygiene	5.64%	6.31%	7.85%
Sanitary Protection	1.93%	2.08%	1.67%
Shaving Needs	2.11%	2.09%	4.47%
Skin Care Preparations	4.92%	5.43%	6.84%
Vitamins	14.99%	17.80%	17.59%

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

Appendix Table A4: Price Index Correlations

	Laspeyres	Laspeyres, housing alternatives	GEKS-Fischer	GEKS-Fischer, housing alternatives	CES	Nested CES ($\sigma=11.5$)	Nested CES ($\sigma=7$)	Geary-Khamis	BEA
Laspeyres	1								
Laspeyres, housing alternatives	.99	1							
GEKS-Fischer	.99	.99	1						
GEKS-Fischer, housing alternatives	.99	.99	.99	1					
CES	.99	.99	.99	.99	1				
Nested CES ($\sigma=11.5$)	.92	.92	.93	.93	.92	1			
Nested CES ($\sigma=7$)	.80	.80	.81	.81	.80	.96	1		
Geary-Khamis	.99	.98	.99	.99	.99	.93	.81	1	
BEA	.93	.93	.93	.93	.93	.84	.72	.92	1

Notes: The numbers of observations are 302 for BEA and Nested CES indexes, 305 for all other correlations with BEA index, 412 for all other correlations with Nested CES indexes, and 443 otherwise. Correlation matrix of all alternative price indices. See text for detailed definition of each index.

Appendix Table A5: 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
Laspeyres	1.027	0.101	0.873	1.562	1.018	0.115	0.844	1.656	1.029	0.089	0.889	1.482	1.033	0.069	0.912	1.375
Laspeyres, housing alternatives	1.022	0.096	0.872	1.491	1.024	0.114	0.844	1.586	1.019	0.082	0.892	1.402	1.022	0.063	0.915	1.330
GEKS-Fischer	1.020	0.090	0.866	1.479	1.014	0.107	0.849	1.609	1.021	0.077	0.885	1.406	1.022	0.056	0.905	1.272
GEKS-Fischer, housing alternatives	1.016	0.087	0.872	1.431	1.019	0.105	0.849	1.552	1.013	0.073	0.891	1.356	1.016	0.052	0.913	1.2602
CES	1.018	0.094	0.859	1.495	1.011	0.111	0.840	1.625	1.019	0.081	0.876	1.426	1.023	0.061	0.903	1.292
Nested CES ($\sigma = 11.5$)	1.071	0.067	0.914	1.335	1.065	0.087	0.903	1.524	1.072	0.067	0.894	1.344	1.063	0.054	0.943	1.243
Nested CES ($\sigma = 7$)	1.119	0.088	0.932	1.383	1.104	0.093	0.929	1.527	1.124	0.089	0.911	1.399	1.103	0.078	0.958	1.353
Geary-Khamis	1.016	0.087	0.839	1.440	1.012	0.105	0.827	1.607	1.018	0.075	0.862	1.382	1.018	0.054	0.889	1.253
BEA	1.040	0.073	0.874	1.359	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

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

Appendix Table A6: Price Index vs. Rent

	Overall Income (1)	Low Income (2)	Middle Income (3)	High Income (4)
Laspeyres	0.381*** (0.010)	0.440*** (0.011)	0.333*** (0.009)	0.244*** (0.008)
R^2	0.953	0.954	0.944	0.897
PPP, GEKS-Fischer	0.340*** (0.010)	0.410*** (0.013)	0.290*** (0.010)	0.196*** (0.009)
R^2	0.947	0.948	0.933	0.858

Notes: All columns use a log-log specification. We report the coefficient on log rent. 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 A7: Variety Effect Decomposition

Dependent variable:	Log Nested CES Price Index				
Effect decomposition:	price only	variety only		price and variety	
Elasticity parameter (σ):	— (1)	11.5 (2)	7 (3)	11.5 (4)	7 (5)
A. Automotive Expenses					
Log Median CZ Income	0.046 (0.026)	-0.068*** (0.007)	-0.119*** (0.013)	-0.021 (0.024)	-0.071** (0.024)
B. Child/Dependent Expenses					
Log Median CZ Income	0.058** (0.019)	-0.048*** (0.004)	-0.084*** (0.008)	0.012 (0.018)	-0.023 (0.018)
C. Clothing/Shoes/Jewelry					
Log Median CZ Income	0.161* (0.080)	-0.040*** (0.004)	-0.069*** (0.007)	0.121 (0.079)	0.092 (0.078)
D. Electronics					
Log Median CZ Income	-0.001 (0.012)	-0.329*** (0.030)	-0.576*** (0.053)	-0.328*** (0.034)	-0.573*** (0.056)
E. Gasoline/Fuel					
Log Median CZ Income	0.157*** (0.032)	-0.029*** (0.003)	-0.050*** (0.005)	0.128*** (0.031)	0.106*** (0.031)
F. General Merchandise					
Log Median CZ Income	0.099** (0.033)	-0.341*** (0.039)	-0.596*** (0.069)	-0.243*** (0.050)	-0.498*** (0.075)
G. Groceries					
Log Median CZ Income	0.056*** (0.015)	-0.101*** (0.012)	-0.177*** (0.020)	-0.046* (0.021)	-0.122*** (0.028)
H. Healthcare/Medical					
Log Median CZ Income	0.231** (0.079)	-0.039*** (0.004)	-0.069*** (0.007)	0.191* (0.078)	0.162* (0.078)
I. Hobbies/Entertainment					
Log Median CZ Income	-0.002 (0.023)	-0.297*** (0.035)	-0.519*** (0.062)	-0.301*** (0.041)	-0.525*** (0.065)
J. Office Supplies					
Log Median CZ Income	0.039** (0.015)	-0.026*** (0.003)	-0.046*** (0.005)	0.013 (0.016)	-0.007 (0.017)
K. Personal Care					
Log Median CZ Income	0.076*** (0.016)	-0.322*** (0.040)	-0.564*** (0.070)	-0.247*** (0.038)	-0.490*** (0.067)
L. Restaurants/Dining					
Log Median CZ Income	0.128*** (0.027)	-0.075*** (0.006)	-0.131*** (0.011)	0.053* (0.024)	-0.003 (0.023)
M. Telecommunications					
Log Median CZ Income	-0.354** (0.126)	-0.020*** (0.002)	-0.035*** (0.003)	-0.374** (0.126)	-0.388** (0.126)
N. Utilities					
Log Median CZ Income	0.495*** (0.118)	-0.252*** (0.027)	-0.442*** (0.048)	0.243* (0.110)	0.054 (0.107)
Overall					
Log Median CZ Income	0.379*** (0.038)	-0.092*** (0.009)	-0.160*** (0.016)	0.290*** (0.035)	0.224*** (0.034)

Notes: This table decomposes the impact of the price effect versus the supply of variety in the nested CES price index across each sub-component of the price index. Standard errors are clustered at the commuting zone level. Each number represents the regression coefficient of the respective price index on the CZ mean household income. See text for details.

Appendix Table A8: **Decennial Laspeyres Price Indexes: 1980-2020**

Panel A. *Spatial Distribution of the Overall-income Index*

Year	Mean	S.D.	Min	P1	P10	P25	P50	P75	P90	P99	Max
Main	1.027	0.101	0.873	0.884	0.925	0.960	1.007	1.068	1.148	1.406	1.562
1980	0.982	0.049	0.859	0.888	0.925	0.949	0.977	1.008	1.035	1.159	1.220
1990	0.973	0.069	0.844	0.872	0.902	0.930	0.962	1.001	1.052	1.213	1.474
2000	0.998	0.065	0.804	0.881	0.932	0.958	0.988	1.025	1.076	1.232	1.415
2010	1.040	0.079	0.866	0.900	0.958	0.987	1.024	1.083	1.142	1.316	1.452
2014	1.001	0.099	0.840	0.846	0.896	0.934	0.983	1.049	1.120	1.381	1.477
2020	1.080	0.097	0.890	0.937	0.989	1.017	1.051	1.119	1.208	1.415	1.578
2020*	1.099	0.105	0.911	0.948	1.005	1.033	1.069	1.138	1.236	1.485	1.645

Panel B. *Income-group-specific Index Summary Statistics*

Year	Low Income				Middle Income				High Income			
	Mean	S.D.	Min	Max	Mean	S.D.	Min	Max	Mean	S.D.	Min	Max
Main	1.018	0.115	0.844	1.656	1.029	0.089	0.889	1.482	1.033	0.069	0.912	1.375
1980	0.978	0.058	0.833	1.264	0.984	0.042	0.880	1.188	0.989	0.030	0.913	1.132
1990	0.967	0.083	0.814	1.567	0.976	0.059	0.866	1.402	0.983	0.042	0.903	1.279
2000	0.998	0.078	0.767	1.497	0.998	0.055	0.832	1.353	0.999	0.039	0.877	1.245
2010	1.048	0.094	0.841	1.541	1.034	0.067	0.886	1.384	1.024	0.047	0.917	1.266
2014	1.002	0.118	0.809	1.571	1.001	0.084	0.863	1.405	1.000	0.059	0.900	1.281
2020	1.096	0.117	0.869	1.693	1.068	0.083	0.906	1.490	1.048	0.058	0.932	1.339
2020*	1.115	0.122	0.899	1.769	1.087	0.092	0.920	1.554	1.065	0.071	0.936	1.402

Panel C. *Correlation among Overall-income Indexes*

	Main	1980	1990	2000	2010	2014	2020	2020*
Main	1							
1980	0.446	1						
1990	0.678	0.651	1					
2000	0.740	0.543	0.848	1				
2010	0.754	0.476	0.790	0.849	1			
2014	0.955	0.387	0.666	0.707	0.749	1		
2020	0.793	0.443	0.746	0.861	0.924	0.764	1	
2020*	0.824	0.467	0.757	0.856	0.904	0.768	0.976	1

Notes: Panel A reports the spatial distribution of our main price index across commuting zones in row 1. In rows 2-6, we calculate price indexes for the years 1980, 1990, 2000, 2014, and 2020 using rents with hedonic adjustments in Table ?? and the imputed non-housing prices based on Figure A8. Note that variation in these indexes comes exclusively from variation in rents. The last row reports our post-pandemic price index. As opposed to using the imputed non-housing prices, we update some prices for which data are available for later years: 2020 NielsenIQ Consumer Panel data for *Child/Dependent Expenses*, *Electronics*, *General Merchandise*, *Groceries*, *Hobbies/Entertainment*, *Office Supplies*, and *Personal Care*, 2024 for *Telecommunications*, 2021 for *Healthcare/Medical*, 2023 for electricity and 2021 for gas in *Utilities*, 2024 for *Clothing/Shoes/Jewelry*, and 2019 for *Restaurants/Dining*. For housing prices, we use pooled 2020-2024 ACS data. We label the resulting price index “2020*”, although it includes prices for various years around 2020. Panel B summarizes income-group-specific indexes. Panel C reports pairwise correlations among these indexes.

Appendix Table A9: Mean Household Consumption by Commuting Zone – Low-Income

City Name (1)	Laspeyres		GEKS-Fischer	
	Consumption (2)	Index (3)	Consumption (4)	Index (5)
Highest Consumption				
Elizabeth City, NC	47,939	0.964	47,986	0.965
Traverse City, MI	43,525	0.959	43,458	0.962
Champaign, IL	42,721	1.017	42,944	1.013
Huntington, WV	42,458	0.920	42,208	0.927
Youngstown, OH	42,287	0.909	42,051	0.914
Beaumont, TX	42,108	0.982	42,099	0.983
Kingsport, TN	42,073	0.912	41,872	0.917
Johnstown, PA	42,038	0.915	42,103	0.914
State College, PA	42,027	0.962	41,876	0.966
Mobile, AL	41,962	1.004	41,689	1.011
Columbia, MO	41,805	0.926	41,806	0.926
Florence, SC	41,405	0.933	41,420	0.934
Morgantown, WV	41,166	0.930	41,153	0.931
Springfield, IL	40,992	0.931	41,068	0.930
Lafayette, LA	40,893	0.934	40,651	0.940
Median Consumption				
Kalispell, MT	37,146	1.023	37,320	1.020
El Paso, TX	37,034	1.025	37,061	1.024
Amarillo, TX	36,883	1.037	36,989	1.035
Deltona, FL	36,879	1.066	36,966	1.064
Lincoln, NE	36,873	1.016	37,314	1.005
Lowest Consumption				
Fairbanks, AK	30,171	1.239	31,615	1.182
Denver, CO	30,162	1.257	30,617	1.238
Seattle, WA	30,139	1.282	30,692	1.258
White Plains, NY	30,021	1.343	30,938	1.303
Edison, NJ	29,990	1.314	30,397	1.297
Newark, NJ	29,775	1.322	30,888	1.274
Los Angeles, CA	28,564	1.381	29,048	1.358
New York, NY	28,278	1.413	29,045	1.376
Anchorage, AK	28,166	1.371	29,511	1.308
Washington, DC	28,049	1.400	29,249	1.342
San Diego, CA	27,010	1.462	27,635	1.428
Santa Barbara, CA	27,008	1.457	27,404	1.435
Honolulu, HI	26,453	1.506	27,501	1.448
San Francisco, CA	25,771	1.553	26,324	1.520
San Jose, CA	24,257	1.656	24,949	1.609

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. Commuting zones ordered by Laspeyres index.

Appendix Table A10: Mean Household Consumption by Commuting Zone – High-Income

City Name (1)	Laspeyres		GEKS-Fischer	
	Consumption (2)	Index (3)	Consumption (4)	Index (5)
Highest Consumption				
Toledo, OH	290,451	0.976	301,314	0.976
Pittsburgh, PA	278,827	1.000	285,199	0.991
Erie, PA	278,371	0.997	289,720	1.000
Huntington, WV	278,356	0.951	282,164	0.955
Kalamazoo, MI	278,261	1.002	292,662	0.994
Canton, OH	277,418	0.973	285,134	0.980
Warsaw, IN	276,938	0.981	286,817	0.989
Lufkin, TX	274,663	0.978	285,746	0.983
Louisville, KY	274,053	0.987	280,826	0.980
Cleveland, OH	273,908	1.000	277,171	1.000
Cincinnati, OH	272,956	1.000	280,446	0.986
Sandusky, OH	272,863	0.999	283,300	1.001
South Bend, IN	272,297	1.000	277,302	1.000
Johnstown, PA	271,009	0.979	281,443	0.970
Elizabeth City, NC	270,201	1.006	289,625	1.005
Median Consumption				
Rochester, MN	239,492	1.026	241,420	1.028
Brunswick, GA	239,466	1.051	247,217	1.042
Tuscaloosa, AL	239,068	1.021	242,363	1.021
Denton, TX	239,048	1.088	243,916	1.071
Memphis, TN	238,943	1.004	240,774	1.001
Lowest Consumption				
Corpus Christi, TX	210,858	1.087	211,259	1.073
Chico, CA	210,275	1.121	211,351	1.106
Ogdensburg, NY	208,980	1.044	208,491	1.036
Yuma, AZ	208,527	1.032	205,809	1.026
Los Angeles, CA	208,467	1.218	215,516	1.174
New York, NY	208,165	1.243	216,405	1.192
Seattle, WA	208,119	1.151	212,974	1.120
Santa Barbara, CA	207,499	1.280	214,097	1.233
Anchorage, AK	200,313	1.245	207,225	1.177
Medford, OR	200,034	1.056	198,435	1.044
Olympia, WA	198,412	1.180	198,024	1.141
San Francisco, CA	195,801	1.308	203,694	1.250
San Diego, CA	189,616	1.284	196,863	1.226
San Jose, CA	186,800	1.349	196,650	1.271
Honolulu, HI	175,105	1.375	185,476	1.272

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. Commuting zones ordered by Laspeyres index.

Appendix Table A11: **Consumption vs. Price Index by Income Group– Robustness**

Index	Low-Income	Middle-Income	High-Income
Laspeyres	−0.899*** (0.009)	−0.984*** (0.021)	−1.066*** (0.038)
Laspeyres, housing alternatives	−0.924*** (0.013)	−1.012*** (0.035)	−1.067*** (0.063)
GEKS-Fischer	−0.892*** (0.009)	−0.982*** (0.024)	−1.090*** (0.047)
GEKS-Fischer, housing alternatives	−0.920*** (0.012)	−0.989*** (0.035)	−1.062*** (0.063)
CES	−0.895*** (0.009)	−0.981*** (0.023)	−1.082*** (0.043)
Nested CES ($\sigma=11.5$)	−0.864*** (0.011)	−0.963*** (0.030)	−1.099*** (0.068)
Nested CES ($\sigma=7$)	−0.857*** (0.012)	−0.966*** (0.029)	−1.116*** (0.061)
Geary-Khamis	−0.891*** (0.009)	−0.982*** (0.025)	−1.098*** (0.049)
BEA	−0.840*** (0.015)	−0.983*** (0.024)	−1.028*** (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 A12: Elasticity of Consumption wrt Price Index — NielsenIQ Data

Product Group	Unit	Low Income			Middle Income			High Income		
		$\hat{\beta}_{Laspeyres}$	$\hat{\beta}_{GEKS-Fischer}$	$\bar{Y}$	$\hat{\beta}_{Laspeyres}$	$\hat{\beta}_{GEKS-Fischer}$	$\bar{Y}$	$\hat{\beta}_{Laspeyres}$	$\hat{\beta}_{GEKS-Fischer}$	$\bar{Y}$
Beer	KG	0.535 (0.328)	0.588* (0.349)	21.2	0.350 (0.344)	0.458 (0.391)	22.5	-0.655 (0.548)	-0.555 (0.621)	18.7
Carbonated Beverages	KG	-0.870*** (0.092)	-0.921*** (0.102)	128.5	-0.936*** (0.142)	-1.057*** (0.165)	129.8	-1.225*** (0.223)	-1.411*** (0.263)	122.1
Cookies	KG	-0.308*** (0.094)	-0.340*** (0.099)	6.2	-0.479*** (0.150)	-0.491*** (0.174)	6.3	-0.280 (0.224)	-0.419 (0.272)	5.8
Deodorant	KG	-0.236** (0.099)	-0.257** (0.104)	0.3	-0.420*** (0.076)	-0.477*** (0.085)	0.4	-0.986*** (0.246)	-1.176*** (0.312)	0.4
Eggs	CT	-0.166** (0.068)	-0.175** (0.072)	182.8	-0.145 (0.104)	-0.169 (0.121)	195.8	-0.490** (0.208)	-0.647** (0.253)	186.4
Housewares, Appliances	CT	-0.880*** (0.094)	-0.937*** (0.099)	2.2	-1.052*** (0.169)	-1.170*** (0.189)	2.4	-1.592*** (0.236)	-1.821*** (0.268)	2.4
Kitchen Gadgets	CT	-0.499** (0.193)	-0.536*** (0.204)	39.0	-0.236 (0.294)	-0.231 (0.338)	55.2	1.030** (0.449)	1.319** (0.564)	64.4
Laundry Supplies	KG	-0.432*** (0.097)	-0.463*** (0.101)	11.0	-0.428*** (0.126)	-0.487*** (0.143)	12.2	-0.988*** (0.175)	-1.155*** (0.216)	12.0
Light Bulbs, Electric Goods	CT	-0.967*** (0.122)	-1.037*** (0.128)	6.7	-1.294*** (0.152)	-1.466*** (0.173)	7.3	-2.064*** (0.306)	-2.477*** (0.356)	7.7
Nuts	KG	0.138 (0.122)	0.151 (0.130)	3.1	0.272 (0.187)	0.319 (0.216)	4.2	-0.314 (0.433)	-0.421 (0.575)	4.8
Pet Food	KG	-0.706*** (0.126)	-0.744*** (0.137)	57.6	-0.782*** (0.203)	-0.871*** (0.233)	55.3	-1.229*** (0.474)	-1.332** (0.600)	47.7
Pizza, Snacks - Frozen	KG	-0.599*** (0.150)	-0.648*** (0.160)	5.8	-0.800*** (0.197)	-0.907*** (0.228)	6.0	-1.292*** (0.278)	-1.511*** (0.342)	5.6
Stationery, School Supplies	CT	-0.520** (0.207)	-0.551** (0.227)	228.4	-0.452 (0.281)	-0.481 (0.317)	278.5	-0.491 (0.520)	-0.366 (0.637)	275.6
Vegetables - Frozen	KG	-0.274** (0.131)	-0.311** (0.139)	10.0	-0.438*** (0.162)	-0.539*** (0.180)	11.2	-1.112*** (0.279)	-1.530*** (0.306)	10.0

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 cz and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.01.

Appendix Table A13: **Consumption vs. Price Index, City Size, and College Share**

	Log Consumption	Log Consumption
Log price index	-0.279*** (0.077)	-0.237*** (0.086)
Log price index × middle-skill	0.041 (0.087)	0.044 (0.098)
Log price index × low-skill	-0.055 (0.089)	-0.071 (0.101)
Log city size	0.031*** (0.008)	0.032*** (0.008)
Log population × middle-skill	-0.026*** (0.009)	-0.026*** (0.009)
Log population × low-skill	-0.038*** (0.009)	-0.039*** (0.009)
Log college share	0.040 (0.040)	0.052 (0.041)
Log college share × middle-skill	-0.035 (0.045)	-0.036 (0.047)
Log college share × low-skill	-0.027 (0.046)	-0.028 (0.049)
Middle-skill	-0.062 (0.150)	-0.060 (0.152)
Low-skill	-0.055 (0.150)	-0.056 (0.152)
Index	Laspeyres	GEKS-Fischer
N	1,329	1,329

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. *** p<0.01, ** p<0.05, * p<0.10.

Appendix Table A14: **Consumption Against Price Index – Robustness**

	College Graduates (1)	High School Graduates (2)	High School Dropouts (3)
A. Baseline			
Laspeyres price index	0.014 (0.056)	-0.195*** (0.028)	-0.373*** (0.033)
B. Income Includes Imputed Food Stamps, TANF, and Housing Assistance			
Laspeyres price index	0.004 (0.057)	-0.203*** (0.029)	-0.373*** (0.035)
C. Consumption Does Not Include Any Imputation			
Laspeyres price index	-0.052 (0.060)	-0.253*** (0.037)	-0.444*** (0.049)
D. Baseline			
GEKS-Fischer price index	0.117** (0.065)	-0.121*** (0.033)	-0.328*** (0.040)
E. Income Includes Imputed Food Stamps, TANF, and Housing Assistance			
GEKS-Fischer price index	0.107* (0.066)	-0.131*** (0.035)	-0.329*** (0.042)
F. Consumption Does Not Include Any Imputation			
GEKS-Fischer price index	0.042 (0.070)	-0.186*** (0.042)	-0.407*** (0.055)

Notes: N = 443. Panel A reports our main results of the relationship between consumption and 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.

Appendix Table A15: **Consumption vs. Price Index by Education Group–Robustness**

Index	College Graduate (1)	High School Graduate (2)	High School Dropout (3)
Laspeyres	0.014 (0.056)	−0.195 (0.028)	−0.373 (0.033)
Laspeyres, housing alternatives	0.052 (0.055)	−0.158 (0.027)	−0.339 (0.035)
GEKS-Fischer	0.117 (0.065)	−0.121 (0.033)	−0.328 (0.040)
GEKS-Fischer, housing alternatives	0.111 (0.066)	−0.124 (0.033)	−0.329 (0.041)
CES	0.078 (0.061)	−0.154 (0.031)	−0.353 (0.037)
Nested CES ($\sigma=11.5$)	0.226 (0.116)	0.002 (0.061)	−0.244 (0.058)
Nested CES ($\sigma=7$)	−0.082 (0.152)	−0.166 (0.085)	−0.331 (0.060)
Geary-Khamis	0.136 (0.070)	−0.110 (0.037)	−0.329 (0.044)
BEA	0.219 (0.060)	0.102 (0.032)	−0.050 (0.042)
Panel B. Some College vs. High School Graduates			
	Some College & High School (1)	Some College only (2)	High School only (3)
Log Laspeyres Index	−0.195 (0.028)	−0.179 (0.034)	−0.288 (0.031)
Log GEKS-Fischer Index	−0.121 (0.033)	−0.104 (0.041)	−0.229 (0.037)
Panel C. Excluding Retired Individuals			
	College Graduates (1)	High School Graduates (2)	High School Dropouts (3)
Log Laspeyres Index	0.015 (0.056)	−0.194 (0.028)	−0.366 (0.034)
Log GEKS-Fischer Index	0.118 (0.065)	−0.121 (0.034)	−0.321 (0.041)

Notes: N = 443. We report the relationship between consumption and local price indexes by skill group. Panel A reports estimates based on 9 different index definitions. Panel B reports estimates separately for households with some college and those with high school degrees. Panel C reports estimates based on the ACS household sample after we exclude 5.3% retired individuals.

Appendix Table A16: **Expenditure Shares of Thirty Most Common Housing Types**

No	Housing Characteristics					Income Group			
	Structure	Year Built	# Rooms	# Bedrooms	Amenities	All	Low	Middle	High
1.	One-Family	1990 Onward	Eight or More	Four or more	All	8.1%	2.1%	9.9%	23.0%
2.	One-Family	1950–1969	Six-Seven	Three	All	5.1%	4.6%	5.6%	3.6%
3.	One-Family	1970–1989	Eight or More	Four or more	All	4.6%	1.7%	5.7%	10.7%
4.	One-Family	1970–1989	Six-Seven	Three	All	4.1%	3.5%	4.8%	2.5%
5.	One-Family	1990 Onward	Six-Seven	Three	All	3.8%	2.7%	4.7%	3.2%
6.	One-Family	1950–1969	Eight or More	Four or more	All	3.5%	1.6%	4.1%	7.8%
7.	One-Family	1990 Onward	Six-Seven	Four or more	All	3.1%	1.7%	4.0%	3.6%
8.	One-Family	1950–1969	Five	Three	All	2.9%	3.2%	2.8%	1.1%
9.	One-Family	Before 1950	Six-Seven	Three	All	2.8%	2.5%	3.0%	3.0%
10.	One-Family	1950–1969	Six-Seven	Four or more	All	2.7%	1.8%	3.3%	2.6%
11.	Multiple-Family	1970–1989	Four	Two	All	2.6%	4.0%	1.9%	0.7%
12.	One-Family	1970–1989	Six-Seven	Four or more	All	2.5%	1.6%	3.2%	2.5%
13.	Multiple-Family	1970–1989	At Most Three	At most one	All	2.4%	4.6%	1.3%	0.6%
14.	Multiple-Family	Before 1950	At Most Three	At most one	All	2.4%	4.1%	1.5%	0.7%
15.	One-Family	1970–1989	Five	Three	All	2.2%	2.4%	2.3%	0.7%
16.	Multiple-Family	1950–1969	At Most Three	At most one	All	2.1%	3.9%	1.2%	0.4%
17.	Multiple-Family	1990 Onward	At Most Three	At most one	All	2.1%	3.2%	1.6%	1.0%
18.	One-Family	Before 1950	Eight or More	Four or more	All	2.0%	1.0%	2.2%	5.8%
19.	One-Family	1990 Onward	Five	Three	All	1.9%	1.9%	2.0%	0.9%
20.	Multiple-Family	1990 Onward	Four	Two	All	1.8%	2.4%	1.5%	1.0%
21.	Multiple-Family	1950–1969	Four	Two	All	1.7%	2.8%	1.2%	0.5%
22.	Multiple-Family	Before 1950	Four	Two	All	1.5%	2.3%	1.2%	0.7%
23.	One-Family	1990 Onward	Eight or More	Three	All	1.5%	0.7%	1.9%	1.9%
24.	One-Family	1950–1969	Eight or More	Three	All	1.3%	0.9%	1.6%	1.6%
25.	One-Family	1970–1989	Eight or More	Three	All	1.3%	0.8%	1.6%	1.6%
26.	One-Family	Before 1950	Five	Three	All	1.0%	1.2%	1.0%	0.5%
27.	One-Family	Before 1950	Six-Seven	Four or more	All	1.0%	0.8%	1.1%	1.1%
28.	One-Family	Before 1950	Eight or More	Three	All	0.9%	0.6%	1.0%	1.8%
29.	One-Family	Before 1950	Five	Two	All	0.9%	1.2%	0.7%	0.4%
30.	One-Family	1950–1969	Four	Two	All	0.9%	1.4%	0.6%	0.1%

Notes: The table reports mean expenditure shares within housing among the 30 most common housing types for the three income groups. Expenditure on each housing type is measured by the share of households consuming that type of housing, as measured in the ACS, multiplied by the rental price for that type of housing.

Appendix Table A17: **Housing Cost Regressions for Homeowner Housing Expenditure**

Dependent variable: (in levels)	Housing Costs to be Added (1)	Housing Costs to be Subtracted (2)
Constant	4580.1*** (190.9)	553.1* (286.3)
Mid-income	-838.4*** (287.2)	158.1 (392.1)
High-income	1622.7 (1589.8)	3158.7 (2646.8)
Income	0.0211*** (0.0030)	0.0923*** (0.0055)
Income $\times$ Mid-income	0.0020 (0.0032)	-0.0313*** (0.0059)
Income $\times$ High-income	-0.0227*** (0.0040)	-0.0881*** (0.0080)
Property value	0.0445*** (0.0006)	0.0214*** (0.0009)
Property value $\times$ Mid-income	0.0018* (0.0008)	0.0140*** (0.0011)
Property value $\times$ High-income	-0.0001 (0.0031)	0.0227*** (0.0059)
Property value squared	-0.0156*** (0.0007)	-0.0088*** (0.0008)
Property value squared $\times$ Mid-income	0.0016** (0.0008)	-0.0059*** (0.0009)
Property value squared $\times$ High-income	0.0037** (0.0017)	-0.0056* (0.0030)
Number of Bedrooms	329.8*** (30.1)	263.9*** (43.0)
Number of Bedrooms $\times$ Mid-income	37.9 (46.7)	-37.0 (56.8)
Number of Bedrooms $\times$ High-income	432.4*** (115.2)	341.3* (199.8)
State fixed effects	Yes	Yes
R^2	0.709	0.344
Number of households	70,288	70,288

Notes: The dependent variable in Column 1 is the household-reported estimated rental value of the housing unit, as reported by the homeowner in the 2012-2016 Consumer Expenditure Survey (CEX). All covariates are also measured in the CEX data. Column 2 is the reported household spending on mortgages, closing costs, and down payments. Column 1 is used to impute the spending of homeowners on housing services. Column 2 is used to impute the housing spending to be removed from homeowners expenditure, as spending on assets and debt does not reflect consumption.