Food Deserts and the Causes of Nutritional Inequality

Hunt Allcott, Rebecca Diamond, Jean-Pierre Dubé,
Jessie Handbury, Ilya Rahkovsky, and Molly Schnell*

November 9, 2018

Abstract

We study the causes of "nutritional inequality": why the wealthy eat more healthfully than the poor in the United States. Exploiting supermarket entry, household moves to healthier neighborhoods, and purchasing patterns among households with identical local supply, we reject that neighborhood environments contribute meaningfully to nutritional inequality. Using a structural demand model, we find that exposing low-income households to the same products and prices available to high-income households reduces nutritional inequality by only nine percent, while the remaining 91 percent is driven by differences in demand. These findings counter the common notion that policies to reduce supply inequities, such as "food deserts," could play an important role in reducing nutritional inequality. By contrast, the structural results predict that means-tested subsidies for healthy food could eliminate nutritional inequality at a fiscal cost of about 15 percent of the annual budget for the U.S. Supplemental Nutrition Assistance Program.

Keywords: Inequality, food deserts, grocery demand estimation, nutrition policy.

JEL codes: D12, I12, I14, L81, R20.

---

*Allcott: New York University and NBER, hunt.allcott@nyu.edu. Diamond: Stanford GSB and NBER, diamondr@stanford.edu. Dubé: Chicago Booth and NBER, jdube@chicagobooth.edu. Handbury: Wharton and NBER, handbury@wharton.upenn.edu. Rahkovsky: U.S. Department of Agriculture, irahkovsky@ers.usda.gov. Schnell: Stanford, mschnell@stanford.edu. We thank Prottoy Aman Akbar, Yue Cao, Hae Nim Lee, and Catherine Wright for exceptional research assistance; Charles Courtemanche, Sungho Park, and Andrea Carlson for sharing data; and Marianne Bitler, Anne Case, David Cuberes, Amanda Chuan, Janet Currie, Jan De Loecker, Gilles Duranton, Joe Gyourko, Jakub Kastl, Ephriam Liebtag, Ilyana Kuziemko, Todd Sinai, Diane Whitmore Schanzenbach, Jesse Shapiro, Tom Vogl, and David Weinstein for helpful comments. We also thank seminar participants at Amazon, the 2015 and 2017 ASSA Meetings, the Becker Friedman Institute at the University of Chicago, Brown, Columbia, Duke, the Federal Reserve Bank of Kansas City, the Federal Trade Commission, Microsoft Research, the 2015 and 2016 NBER Summer Institutes, New York University, the Paris School of Economics, Penn State, Princeton, the Pritzker School of Medicine, the Robert Wood Johnson Foundation, Stanford, Temple, Toronto, the Tilburg Christmas Research Camp, the University of New South Wales, the University of Pennsylvania, USC Marshall, the USDA, the University of Sydney, the 2014 Urban Economics Association Meeting, Warwick, Wharton, and Yale SOM. We are grateful for funding from the Chicago Booth Initiative on Global Markets, the Wharton Social Impact Initiative, the Research Sponsors' Program of the Wharton Zell-Lurie Real Estate Center, and the USDA Economic Research Service. This paper reflects the researchers' own analyses calculated based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. The findings and conclusions in this preliminary publication have not been formally disseminated by the USDA and should not be construed to represent any agency determination or policy. This paper subsumes and replaces our previous work, Handbury, Rahkovsky and Schnell (2015) and Allcott, Diamond and Dubé (2017).

I Introduction

A wave of recent studies has drawn increased attention to the causes and consequences of socioeconomic inequality (Aizer and Currie, 2014; Case and Deaton, 2015; Chetty et al., 2014; Chetty et al., 2016; Saez and Piketty, 2003). This inequality plays out in many different ways, including educational opportunities, health outcomes, mass incarceration, and social networks. In this paper, we study one additional correlate of socioeconomic status—what we eat and drink—and quantify the economic mechanisms that drive nutritional inequality. As nutritional intake is a key driver of obesity and other health outcomes, understanding why nutritional inequality exists is crucial for designing policies to address socioeconomic disparities in health outcomes.

A large body of literature has documented that low-income neighborhoods are more likely to be "food deserts"—that is, areas with low availability or high prices of healthy foods.[1] Many public health researchers, policymakers, and advocates further argue that food deserts are an important *cause* of unhealthy eating.[2] Despite limited evidence supporting this causal claim, both the federal government and local municipalities spend millions of dollars each year on supply-side policies that subsidize and assist grocers in underserved areas.[3]

It is certainly possible that differential access to healthy foods is at least partially to blame for nutritional inequality. For example, Bitler and Haider (2011) discuss how zoning restrictions, crime, and other factors could raise costs for grocers in low-income areas, and results from the Moving To Opportunity experiment suggest that moving to lower-poverty neighborhoods reduces obesity (Kling, Liebman and Katz, 2007; Ludwig et al., 2011). In the presence of market failures such as health cost externalities and imperfect information about how diet affects health, supply-side policies could improve economic efficiency in addition to public health and social equity. On the other hand, it is natural to imagine that the observed supply differences across neighborhoods are simply equilibrium responses to differences in consumer preferences. Thus, teasing apart supply-side versus demand-side explanations is crucial for understanding whether and how policymakers should intervene.

[1]See, for example, Algert et al. (2006); Alwitt and Donley (1997); Baker et al. (2006); Horowitz et al. (2004); Jetter and Cassady (2005); Larson et al. (2009); Powell et al. (2007); Sharkey et al. (2010).

[2]In an influential review article, Bitler and Haider (2011) write that "it appears that much of the existing research implicitly assumes that supply-side factors cause any food deserts that exist." For example, former First Lady Michelle Obama argues that "it's not that people don't know or don't want to do the right thing; they just have to have access to the foods that they know will make their families healthier" (Curtis, 2011), and Hilmers, Hilmers and Dave (2012, page 1652) write that "the disproportionate distribution of food sources that contributes to the development of unhealthy behaviors among these communities and the consequent disease burden deeply affect not only individuals and families, but also society as a whole." Economists, meanwhile, have largely ignored the relationship between disparities in supply and the nutrition-income relationship, focusing instead on the role of differential price elasticities in generating nutritional inequality (e.g., Bertail and Caillavet, 2008; Beydoun et al., 2008; Jones, 1997; Park et al., 1996).

[3]The Healthy Food Financing Initiative has awarded $220 million in subsidies and technical assistance since 2011 to grocery stores, farmers markets, and other suppliers of healthy foods in underserved areas (TRF 2017); the Agricultural Act of 2014 appropriated $125 million in federal funds to be spent annually to promote access to healthy foods in disadvantaged communities (Aussenberg, 2014); and projects aimed at "eliminating food deserts" were eligible for $100 million in Community Transformation Grants under the Affordable Care Act. Many state and local governments have also introduced programs to improve access to nutritious foods by providing loans, grants, and tax credits to qualifying businesses operating in underserved neighborhoods (CDC, 2011). In the United Kingdom, the 2001 Food Poverty Eradication Bill required local and national governments to document and take actions to eliminate food deserts.

This paper combines reduced-form analyses with a structural demand model to quantify the relative importance of local supply and demand factors in generating nutritional inequality. We exploit a rich combination of datasets, including Nielsen Homescan, a 61,000-household, nationally representative panel survey of grocery purchases, and Nielsen's Retail Measurement Services (RMS), a 35,000-store, national panel of UPC-level sales data covering about 40 percent of all U.S. grocery purchases. We match the Homescan data to surveys of panelists' nutrition knowledge, preferences, and health outcomes gathered by Nielsen for Allcott, Lockwood and Taubinsky (2018). Finally, we gather data on the entry dates and exact locations of all 6,721 new supermarkets that opened in the United States from 2004–2016, along with annual data on retail establishments in each zip code. We thus have an extraordinarily rich window into households' choice sets, information sets, local environments, and resulting consumption and health.

We begin by laying out the two basic facts that motivate the debate on food deserts and the causes of nutritional inequality. First, there is a meaningful nutrition-income relationship: households in the top income quartile buy groceries that are 0.56 standard deviations more healthful than the bottom income quartile, as measured by our version of the Healthy Eating Index, the standard measure of dietary quality. Second, RMS stores in low-income neighborhoods offer less healthy groceries than stores in high-income neighborhoods, and low-income neighborhoods have more drug and convenience stores as well as fewer large supermarkets, which offer wider variety of healthy options.

We then use three reduced-form analyses to test whether the local environment has an economically significant effect on healthy eating. The first analysis is an event study that looks within households, before versus after the entry of a new supermarket nearby. In both the full sample and the sample of households in food deserts, we find effects of supermarket entry on healthy eating that are occasionally statistically significant but always economically small. We document an intuitive reason for these small effects: while consumers shift their purchases toward the new entrants, these purchases are primarily substituted away from other supermarkets, not away from drug stores and convenience stores that offer less healthy choice sets. Indeed, even households living in zip codes with no supermarkets still buy 85 percent of their groceries from supermarkets. We can bound the short-run effect of differential access to supermarkets at no more than about 1.5 percent of the nutrition-income relationship.

Our second analysis tests the hypothesis that a broader set of place-related factors, including peer effects or supply differences other than supermarket density, contribute to nutritional inequality. To do this, we exploit the fact that thousands of households move between zip codes or counties while in the Homescan panel. Before a move, households exhibit no trend in healthy eating. After a move, households converge toward eating patterns in the new location by an amount that is statistically significant but economically small. Any endogeneity in moving decisions likely biases these estimates upward. While the panel is not long enough to study households for more than a few years after a move, we can bound the medium-term, partial equilibrium effects of place as contributing no more than two to three percent of the nutrition-income relationship.

Our third reduced-form analysis tests the possibility that improved access to healthy foods could have larger effects in the long term, as consumers' tastes and knowledge of healthy eating slowly evolve.

To do this, we study people across different income levels who consistently have the same access to healthy foods: households living in the same locations and shopping in the same stores. If differential access is entirely to blame for nutritional inequality, then the nutrition-income relationship should disappear when we compare low- and high-income households facing the same retail environment. In reality, less than 30 percent of the nutrition-income relationship is explained by fixed effects for zip code or retail chain. Any endogeneity in where to live and where to shop likely biases this number upward. Thus, while our long-term bound is looser than our short-term bounds, we can still conclude that differential local access to healthy foods explains no more than a third of the nutrition-income relationship.

These reduced form analyses cannot precisely tell us what low-income households would consume if they had the same supply conditions as high-income households. For example, simply adding a supermarket changes availability and prices of both healthy and unhealthy foods in ways that may not replicate conditions in high-income neighborhoods, and moving to a healthy area affects people's lives in many ways other than just changing grocery supply. Meanwhile, the within-location and within-store disparity can only provide a loose upper bound on the counterfactual due to the endogeneity of residential and store choice. To precisely answer the question of whether unequal supply of healthy groceries contributes to nutritional inequality, the second half of the paper uses a structural model to separate supply and demand and then simulates counterfactuals with equal supply. Building on Dubois, Griffith and Nevo (2014), we specify a utility function with Constant Elasticity of Substitution (CES) preferences over individual products, Cobb-Douglas preferences for product groups (milk, bread, candy, vegetables, etc.), and linear preferences for specific dietary quality characteristics (added sugar, salt, saturated fat, etc.). We derive a novel empirical specification that relates calorie consumption to prices and nutrients while accommodating potential unobserved product characteristics.

To separate supply and demand, we introduce a new instrument for prices. The instrument uses the variation in prices generated by grocery retail chains' differing comparative advantages in supplying different product groups combined with chains' differing geographic presence across markets. To illustrate, suppose there are two types of foods, apples and pizza, and two grocery chains, Safeway and Shaw's. Suppose Safeway is able to supply pizza cheaply, while Shaw's can supply apples cheaply. Then, cities dominated by Safeway will have relatively low prices for pizza, while cities dominated by Shaw's will have relatively low prices for apples. Our key identifying assumption is that geographic variation in prices due to the presence of specific chains is independent of geographic variation in unobserved preferences. Consistent with this assumption, we find that retail chains that offer lower prices on healthy foods are no more or less likely to be present in higher-income counties. The instrument has a very strong first stage, and it may be useful for other researchers in similar settings.

The estimates show a striking and systematic relationship between household income and preferences for healthy dietary characteristics. Higher-income households have stronger preferences for six out of the Healthy Eating Index's eight "healthy" dietary components (whole fruit, other fruit, whole grains, green vegetables and beans, other vegetables, and dairy) and two out of four "unhealthy" components (sodium and added sugar). These preference differences are economically significant:

households in the bottom income quartile are willing to pay \$0.43 per 1,000 calories to consume the bundle of dietary components that would receive the maximum Healthy Eating Index score instead of the minimum, while households in the top income quartile are willing to pay \$1.14—almost three times as much. We find that about 20 percent of the income-related preference differences are econometrically explained by education and another 14 percent are explained by nutrition knowledge.

We use our demand model to simulate policies in which bottom income quartile households are exposed to the prices and product availability experienced by top income quartile households. Consistent with our reduced-form results, we find only small effects of these supply-side factors on purchases. Only nine percent of the nutrition-income relationship is driven by differences in supply, while 91 percent of the relationship is driven by differences in demand. Overall, these findings suggest that supply-side policy initiatives aimed at eliminating food deserts will have limited effects on healthy eating in disadvantaged neighborhoods.

What policies might work? We use the demand estimates to simulate the effects of means-tested subsidies for healthy foods, inspired by an active debate about whether the Supplemental Nutrition Assistance Program (SNAP) should be modified to include such subsidies (Richards and Sindelar, 2013; Shenkin and Jacobson, 2010). We find that the fiscal cost of inducing low-income households to purchase groceries that are as healthy as high-income households would be only 15 percent of the current average SNAP benefit.

Our paper's main contribution to the literature is to bring together a unified set of insights about why the wealthy and the poor eat differently in the United States. Our results add to findings on the impacts of proximity to supercenters (Courtemanche and Carden, 2011; Volpe et al., 2013) and fast food restaurants (Anderson and Matsa, 2011; Currie et al., 2010; Davis and Carpenter, 2009; Dunn, 2010), along with case studies in the public health literature of individual grocery store entry (e.g. Cummins et al., 2005; Elbel et al., 2015; Song et al., 2009; Weatherspoon et al., 2012; Wrigley et al., 2003).[4] Our household migration event study adds a nutritional aspect to recent work which uses migration to understand the evolution of brand preferences (Bronnenberg, Dubé and Gentzkow, 2012), the caloric costs of culture (Atkin, 2016), the effects of urban sprawl on obesity (Eid et al., 2008), and the drivers of geographic variation in health and health care utilization (Finkelstein, Gentzkow, and Williams 2016; 2018a; 2018b; Molitor, 2018).[5] Recent work by Hut (2018) finds similarly small short-run effects of migration on healthy eating, but larger associations with local consumption patterns in the long run. Our structural demand analysis builds on the framework introduced by Dubois, Griffith and Nevo (2014), but adds a novel identification strategy and price instrument. Finally, the decomposition of our preference estimates builds on other work measuring correlates of health behaviors (Cutler and Lleras-Muney, 2010; Furnee et al., 2008; Grossman, 2015),

[4]Although not solely focused on retail environments, recent work further examines how changes in a range of socio-environmental factors, including restaurant and food store availability and food prices, have contributed to rising obesity over the last 50 years (Baum and Chou, 2016; Chou et al., 2004; Courtemanche et al., 2016; Cutler et al., 2003; Lakdawalla et al., 2005; Rashad, 2006; Rashad et al., 2006).

[5]Indeed, the debate over health care utilization has many parallels to our current setting, with a public health literature arguing that geographic disparities are primarily due to differences in providers and insurance coverage. This has provided support for supply-side policies such as insurance expansions and subsidies for doctors to work in underserved areas.

but the new Homescan add-on survey provides a remarkable opportunity to connect large-sample scanner data to measures of health preferences and nutrition knowledge.

Sections II through VIII, respectively, present data, stylized facts, reduced-form empirical analyses, demand model setup, demand model estimation and results, counterfactual estimates, and the conclusion.

II Data

II.A Nielsen Homescan and Retail Scanner Data

We use the Nielsen Homescan Panel for 2004–2016 to measure household grocery purchases.[6] Homescan includes about 169,000 unique households, of which we observe about 39,000 in each year for 2004–2006 and about 61,000 in each year for 2007–2016. Homescan households record UPCs of all consumer packaged goods they purchase from any outlet. We consider only food and drink purchases, and we further exclude alcohol and health and beauty products such as vitamins.

We focus on explaining income-related differences in the take-home market (i.e., grocery purchases) instead of overall diets because Homescan does not include data on away-from-home food purchased in establishments like restaurants.[7] One additional limitation of Homescan is that most households only record purchases of packaged items with UPCs, not non-packaged groceries such as bulk produce and grains. For 2004–2006, however, the data also include an 8,000-household "magnet" subsample that also recorded prices and weights of non-packaged groceries. We use the magnet data for robustness checks. Appendix Figure A1 shows that about 60 percent of magnet households' produce calories are from packaged goods that are observed in the full Homescan sample, and this proportion does not vary statistically by income.[8] This suggests that the focus on packaged groceries does not significantly detract from our results, both because produce represents a small share of overall grocery purchases and because packaged produce is a significant and reasonably representative portion of produce purchases.

Homescan households report demographic variables such as household income (in 16 bins), household composition, race, and the age, educational attainment, employment status, and weekly work hours for male and female household heads. For households with two household heads, we use the mean of age, education, and employment variables observed for both heads. We combine calorie needs by age and gender as reported in the U.S. government's Dietary Guidelines with Homescan household composition to get each household's daily calorie need. Our household size variable measures the number of adult "equivalents" in the household, where children are scaled into adults by their calorie needs. In addition to the standard Homescan data, we observe self-reports of the importance of

[6] See Einav, Leibtag and Nevo (2010) for a validation study of Homescan.

[7] The National Health and Nutrition Examination Survey (NHANES) finds that Americans consume 34 percent of calories away from home, including 25 percent in restaurants (USDA 2014b). For all income groups, the share of healthy and unhealthy macronutrients (protein, carbohydrates, saturated fat, etc.) consumed away from home is about the same as the share of calories consumed away from home, so grocery purchases are not a systematically biased measure of overall diet healthfulness.

[8] After excluding canned, frozen, and dried produce, about 40 percent of magnet households' fresh produce calories are from packaged items.

staying healthy, a detailed nutrition knowledge quiz, Body Mass Index (weight/height^2), and diabetes status from a Homescan PanelViews add-on survey carried out by Nielsen for Allcott, Lockwood and Taubinsky (2018) in 2017. Panel A of Table 1 presents descriptive statistics for Homescan households. Unless otherwise stated, all Homescan results are weighted for national representativeness.

The Nielsen Retail Measurement Services (RMS) data consist of weekly prices and sales volumes for each UPC sold at approximately 42,000 unique stores from 160 retail chains for 2006–2016, of which we observe about 35,000 in each year. We exclude liquor stores. RMS includes 53, 32, 55, and 2 percent of sales in the grocery, mass merchandiser, drug, and convenience store channels, respectively. As with Homescan, RMS does not include sales of bulk produce and other non-packaged items.

We deflate prices and incomes to 2010 dollars using the consumer price index for urban consumers for all items.

II.B Grocery Retail Establishments

Studying the effects of retailer entry requires reliable data on store opening dates to avoid attenuation bias. Some datasets, such as InfoUSA and the National Establishment Time Series, might be useful for cross-sectional analyses, but they do not record the opening dates of new establishments with sufficient precision (see Bitler and Haider, 2011, page 162). We measure supermarket entry using Nielsen's TDLinx dataset, a geocoded census of all food retailers in the United States, including the month each store opened and its exact latitude and longitude. In validation checks, we found that TDLinx data accurately match entry dates and locations for 1,914 supermarkets for which we have the exact address and opening date from four retailers' administrative records. We include only club stores, supercenters, and grocery stores (excluding "superettes"), and we further restrict to entries for which the retailer could be matched to a retailer code in the Homescan data, which excludes small "mom and pop" groceries.[9] For simplicity, we call the set of included stores "supermarkets." There were 6,721 entries of supermarkets in the United States between 2004 and 2016.

We also use Zip Code Business Patterns (ZBP), which gives a count of establishments by NAICS code and employment size category for every zip code as of March 10th of each year. The ZBP data are drawn from tax records, the U.S. Census Company Organization Survey, and other administrative data. Panel B of Table 1 presents ZBP descriptive statistics.

II.C Nutrition Facts and the Healthy Eating Index

Our nutrition facts are from databases available from the U.S. Department of Agriculture (USDA, 2016; 2018), which we match to UPCs using crosswalks developed by the USDA. The UPC-level USDA nutrition facts closely match a second database from a marketing data provider that we used in a previous version of this paper.

[9] To measure true changes in grocery availability experienced by consumers, we must use new physical establishments or stores that significantly improve their grocery selection, such as conversions from standard mass merchants to supercenters selling a full line of groceries. To implement this, we use a list of specific TDLinx stores transferred through mergers and acquisitions to exclude spurious "entrants" that were in fact in continuous operation, and we further drop potentially spurious entries where TDLinx shows a store of the same subtype in the same census block in the previous year.

To characterize goods and preferences on a single index of healthfulness per calorie, we use a slightly modified version of the standard dietary quality measure in the United States: the USDA's Healthy Eating Index (Guenther, Reedy and Krebs-Smith, 2008; Volpe and Okrent, 2012). The HEI scores diets on a scale from 0–100, adding points for consumption of "healthy" components (fruits, vegetables, whole grains, dairy, and proteins) and subtracting points for consumption of "unhealthy" components (refined grains, sodium, saturated fats, and added sugars). The HEI is usually used to score one individual's full diet, and it is non-linear in its components. For example, added sugar consumption reduces the HEI linearly at a prescribed "slope" until added sugars reach a threshold of 26 percent of calories, after which the HEI is unchanged. These non-linearities are debated by nutritionists and are inappropriate for many of our analyses, in which we observe partial diets for entire households (in Homescan) or purchases by many consumers (in store-level RMS data). We thus construct a linearized version of the HEI that continues scoring nutrients with the prescribed slope, regardless of whether consumption is beyond a minimum or maximum threshold.[10] We find that the linearization makes little difference: the correlation between "true" HEI and linearized HEI is 0.91 in our household-by-year Homescan data. For ease of interpretation, we normalize our linearized HEI to have a mean of zero and a standard deviation of one across households. We call this linearized, normalized HEI the "Health Index."

Appendix Table A1 shows that the strongest correlates of the Health Index in household-by-year Homescan data are purchases of fruits and vegetables (with correlation coefficients of $\rho \approx 0.4$ to 0.6), whole grains ($\rho \approx 0.50$), sea and plant protein ($\rho \approx 0.64$), added sugar ($\rho \approx -0.41$), and solid fats ($\rho \approx -0.44$). Appendix Table A2 shows that both the "true" HEI and our linearized HEI are highly correlated with Homescan panelists' Body Mass Index and diabetes status.

III Stylized Facts: Purchases and Supply of Healthful Foods

III.A Purchase Disparities: The Nutrition-Income Relationship

We begin by using the Homescan data to provide basic facts on socioeconomic disparities in dietary quality. Even basic stylized facts from these data are important: although a large body of literature exists on social class and diet quality, most previous work uses datasets that only cover a few states or municipalities, or national datasets such as NHANES that are an order of magnitude smaller than Homescan (Darmon and Drewnoski, 2008). For comparability to other work on inequality, we proxy for socioeconomic status (SES) using household income. Because household income varies with household size, age, and survey year for reasons unrelated to household SES, we control for these factors in our analyses.

Figure 1 presents binned scatterplots of how dietary quality varies with household income, residual of household size and age and year indicators. We use three illustrative measures—grams of added sugar per 1,000 calories purchased, share of bread calories from whole-grain breads, and share of total

[10] We omit the fatty acid ratio from our linearized HEI as it is less obvious how to linearize this ratio and saturated fats are already counted directly as a moderation component.

calories from packaged produce—and a summary measure, the normalized Health Index across all grocery purchases. These four measures paint a consistent picture: low-income households purchase less healthful foods. Appendix Figure A3 shows that these results are similar when using the magnet subsample, which includes bulk purchases as well as packaged items.[11]

To quantify this nutrition-income relationship, we calculate the difference in overall dietary quality between households in the highest and the lowest income quartiles. Conditional on household size and age and year indicators, the top income quartile (households with sample average income above about $84,000) buy groceries with a Health Index that is 0.56 standard deviations higher than the bottom income quartile (households with sample average income below about $28,000).[12] A key objective in the rest of the paper is to explain this 0.56 standard deviation difference.[13]

We can benchmark the potential impacts of these differences in dietary quality using relationships between dietary quality and health outcomes, with the important caveat that these relationships are from correlation studies and thus may not be causal. For example, Figure 1 illustrates that households in the bottom income quartile purchase approximately 9.7 more grams of added sugar per 1,000 calories than households in the top quartile. The estimates in Yang et al. (2014) imply that 9.7 fewer grams of added sugar per 1,000 calories is conditionally associated with a 26 percent decrease in death rates from cardiovascular disease. Similarly, the conditional correlations in our Appendix Table A2 imply that the 0.56 standard deviation difference in the Health Index is conditionally associated with a 0.67 difference in Body Mass Index (0.09 standard deviations) and a 1.9 percentage point difference in diabetes (11 percent).

III.B Supply Disparities by Neighborhood Income

Having documented socioeconomic disparities in consumption, we now provide basic facts on how supply of healthy groceries varies between high- and low-income neighborhoods. The four panels of Figure 2 present the relationship between zip code median income and the healthfulness of items offered in RMS stores, again measured by added sugar content, the share of bread UPCs that are

[11]Appendix Figure A2 presents analogues to Figure 1 considering each individual component of the Healthy Eating Index. Consistent with NHANES data, the most economically and statistically significant relationships in Homescan are that higher-income households tend to get a larger share of calories from fruits, vegetables, and whole grains and a smaller share from added sugar. Higher-income households also tend to consume more saturated fat and sodium, although these differences are considerably smaller in percentage terms. Thus, although the Healthy Eating Index imposes specific weights on the different dietary components, higher-income diets would tend to be classified as "more healthy" unless the weights change substantially.

[12]We have also estimated the short-term association between the Health Index and household income by conditioning on household-by-location fixed effects and other time-varying household demographics; see Appendix Table A3. The within-household association is about nine percent the size of the pooled OLS coefficient. Unless measurement error generates significant attenuation bias in the within-household estimate, this suggests that the bulk of the nutrition-income relationship is explained by long-term effects of income and correlates of income, not by the short-term effects of a tighter budget constraint. This is consonant with the finding of Hastings, Kessler and Shapiro (2018) that SNAP enrollment has no effect on healthy purchases.

[13]This difference is growing over time, increasing from 0.54 in 2004–2007 to 0.61 in 2012–2016. This finding of growing nutritional inequality is consistent with results from the NHANES data (Rehm et al., 2016; Wang et al., 2014), although there is less precision in NHANES due to the smaller sample size. This trend underscores the increasing relevance of the issues we study.

whole grain, the share of all UPCs that are packaged produce, and the mean Health Index of UPCs offered. The measures used in this figure weight UPCs only by calories in the package and not by quantity sold so that this figure reflects choice sets, not consumption. All four panels show the same qualitative result: stores in higher-income zip codes offer healthier items.[14]

The store types we call "supermarkets"—large grocery stores, supercenters, and club stores—generally offer a wider variety of healthy items and packaged and bulk produce compared to small grocery stores, convenience stores, and drug stores.[15] Using the Zip Code Business Patterns data, Figure 3 shows that lower-income neighborhoods tend to have fewer supermarkets and more drug and convenience stores per capita. While only 24 percent of zip codes (weighted by population) have no supermarket, 55 percent of zip codes with median income below $25,000 (weighted by population) have no supermarket.[16]

IV Reduced-Form Analysis: Effect of Access on Consumption

In the previous section, we documented both that low-income households consume less healthy groceries and that low-income neighborhoods have less local supply of healthy foods. As described in the introduction, these two stylized facts have raised questions in policy circles about the extent to which reduced availability of healthy foods causes nutritional inequality. Of course, supply and demand are determined in equilibrium. Neighborhood supply could also be correlated with demand due to simultaneity (where supply responds to demand, in addition to demand responding to supply) or due to other unobserved factors that systematically affect both supply and demand in low-income neighborhoods. To measure the elasticity of healthy grocery demand with respect to local supply, we need quasi-exogenous variation in local supply.

In this section, we use three reduced-form strategies. We begin with two event studies that look within households as their local retail environments change in the short and medium terms. In our third strategy, we look across households within a retail environment to bound the longer-term impact of retail access on the nutrition-income relationship.

IV.A Effects of Supermarket Entry

We begin by using an event study framework to measure the effects of supermarket entry on grocery purchases. Over our sample period, we observe 6,721 supermarket entries. While the ideal experiment

[14] This pattern of less-healthy choice sets and store types in low-income neighborhood stores is broadly consistent with a large body of public health literature (e.g., Larson et al., 2009; Sharkey et al., 2010; Powell et al., 2007). To our knowledge, however, we are the first to document this in a dataset as comprehensive as RMS.

[15] Appendix Table A4 demonstrates this in the RMS data. By definition, supercenters carry a full line of groceries, including produce. Club stores such as Sam's Club, BJ's, and Costco typically also carry a variety of grocery and produce items.

[16] In the working paper versions of this paper (Handbury, Rahkovsky and Schnell 2015; Allcott, Diamond and Dubé 2017), we showed that while healthy food costs more per calorie than unhealthy food, there is essentially no price difference for categories other than fresh produce. Furthermore, the relative price of healthy versus unhealthy food is actually slightly *lower* in low-income areas. Therefore, if price plays a role in the nutrition-income relationship, it would have to do so through a preference to reduce produce consumption in order to economize on calories.

to measure the effects of store entry would randomly assign new supermarkets to different neighborhoods, stores enter for reasons that may create endogeneity concerns. In the analysis that follows, we include household fixed effects and measure how grocery consumption changes after a supermarket opens nearby. While supermarkets often open and close in response to long-term changes in neighborhood composition, it seems implausible that supermarkets plan their openings for the exact time at which households will suddenly change their demand patterns.

We consider the impact of supermarket entries that occur within a 0–10 or 10–15 minute drive of households' census tract centroids.[17] We compute driving times between each census tract centroid and the address of each entering supermarket using the Google Maps application program interface (API) and assuming no congestion delay. In our data, 66 percent of households experience zero entries and 19 percent of households experience only one entry within a 10 minute drive.

Let S_{bct} be the count of supermarket entries that have occurred within driving distance band b (where $b = [0, 10)$ or $[10, 15)$) of census tract c as of quarter t, and let $\boldsymbol{X}_{it}$ denote the vector of potentially time-varying household covariates presented in Table 1.[18] Letting Y_{ict} denote an outcome for household i in census tract c in quarter t, we run the following regression in household-by-quarter Homescan data:

$$Y_{ict} = \sum_{b \in \{[0,10),[10,15)\}} \tau_b S_{bct} + \gamma \boldsymbol{X}_{it} + \mu_{d(c)t} + \phi_{ic} + \varepsilon_{ict}, \tag{1}$$

where $\mu_{d(c)t}$ is a vector of census division-by-quarter indicators, and ϕ_{ic} is a household-by-census tract fixed effect. As we study in Section IV.B, some Homescan households move while in the sample. Conditioning on ϕ_{ic} isolates variation in supply due to entry, not relocation. Because the set of households exposed to local entry are not nationally representative, we do not use the Homescan sample weights for this analysis. When estimating Equations (1) and (2), we use robust standard errors with two-way clustering by household and census tract.

Before estimating Equation (1), we first show graphical results of the event study. We define E_{bcqt} as an indicator variable denoting whether one supermarket had entered in distance band b of census tract c by q quarters after quarter t. B_{bit} is an indicator variable for whether observation it is part of a balanced panel around one supermarket entry in distance band b, meaning that the household is observed in the same census tract continuously for all four quarters before and all eight quarters after one supermarket entry and experienced only one entry in distance band b during that window. We run the following regression in household-by-quarter Homescan data:

$$Y_{ict} = \sum_{b \in \{[0,10),[10,15)\}} \left[B_{bit} \left(\sum_{q} \tau_{bq} E_{bcqt} + \alpha_b \right) \right] + \gamma \boldsymbol{X}_{it} + \mu_{d(c)t} + \phi_{ic} + \varepsilon_{ict}. \tag{2}$$

The interaction with B_{bit} ensures that we identify τ_{bq} and α_b using only the households in the balanced

[17] According to the 2009 National Household Travel Survey, the median and 75th percentile of shopping travel times are 10 and 15 minutes, respectively.

[18] Specifically, $\boldsymbol{X}_{it}$ includes the natural log of income, natural log of years of education, indicators for each integer age from 23–90, household size, race indicators, an indicator for whether the household heads are married, employment status, and weekly work hours.

panel, although we include the full sample in the regression to improve the precision on the household covariates and fixed effects. The omitted category is $q = -1$, so all coefficients are relative to the outcome in the last quarter before entry.

Figure 4 presents the $\tau_{[0,10)q}$ coefficients and 95 percent confidence intervals. The panel on the left includes the full sample, while the panel on the right includes only the 23 percent of the sample living in zip codes with no supermarkets, which we call "food deserts."[19]

The top left panel shows that the share of a household's grocery expenditures that are spent at any of an entrant retailers' locations increases by about two percentage points one year after entry. These effects on total expenditures across all stores owned by the entrant retailer likely understate expenditures at the specific entering store, as some expenditures may be diverted from the entrant retailer's other stores.[20] The top right panel shows that the point estimates are larger in food deserts, rising closer to three percent.[21]

The middle panels in Figure 4 show households' expenditure shares at all grocery stores, supercenters, and club stores. We keep the y-axis on the same scale between the top and middle panels so that the magnitudes can be easily compared. To the extent that supermarket entry simply diverts sales from other supermarkets that offer a similar variety of healthy groceries, the changes in healthy grocery availability—and thus the possible effects on healthful purchases—will be limited. Indeed, total expenditure shares across grocery stores, supercenters, and club stores increase by only a fraction of a percentage point in the full sample, with no statistically detectable effect in the food desert subsample. Thus, the primary effect of supermarket entry is to divert sales from other supermarkets.[22]

The bottom panels of Figure 4 presents results with the Health Index of purchased groceries as the dependent variable. Both panels show no detectable increase in healthy purchases after supermarket entry. In any given quarter, we can reject Health Index increases of more than about 0.02 (0.05) standard deviations in the full (food desert) sample.

Table 2 presents estimates of Equation (1) using the same dependent variables as in Figure 4. Panel (a) considers the effects on expenditure shares, first at entrant retailers and then across all grocery stores, supercenters, and club stores. Unsurprisingly, all effects are significantly larger for stores entering within 10 minutes of a household's census tract centroid than for stores entering 10–15 minutes away. Columns 1 and 2 consider the full Homescan sample, columns 3 and 4 limit the sample to households in the bottom income quartile, and columns 5 and 6 limit to households in food deserts.[23] Perhaps unsurprisingly, the expenditure share changes are generally larger for low-income

[19] Appendix Figure A4 presents the analogous figures for the $\tau_{[10,15)q}$ coefficients. The effects are smaller, as would be expected given that the entering stores are 10–15 minutes away instead of 0–10 minutes away. Appendix Figure A5 presents analogous figures for balanced panel windows including eight (instead of four) quarters before entry. The results are similar.

[20] We cannot look at expenditures at the specific entering establishment, because most Homescan panelists record only the retail chain where they shop, not the specific store.

[21] This gradual adjustment of purchases over the first few quarters after entry is consistent with results in Atkin, Faber and Gonzalez-Navarro (2016), who study retail expansion in Mexico.

[22] These results are consistent with those of Hwang and Park (2016), who look at a subset of Walmart Supercenters that opened between 2003 and 2006.

[23] To match the high- versus low-income quartile differences calculated in Section III.A, we define the "bottom income quartile" as households in the lowest quartile of residuals from a regression of household average income across all years

households and households in food deserts. However, consistent with Figure 4, most or all of entrant chains' expenditure share gains consist of diverted sales from other supermarkets.[24] Appendix Table A5 shows that supermarket entry reduces expenditure shares at drug and convenience stores, which offer fewer healthy UPCs, by only a small fraction of a percentage point—even among households in food deserts.

Panel (b) of Table 2 presents effects on the Health Index, which is normalized to have a standard deviation of one across households. All of the six point estimates are positive, and in two cases they are statistically significant, but the effect sizes are economically small. Appendix Table A5 repeats these estimates using three alternative definitions of food deserts. Again, we find economically small, although in one case statistically significant, effects of supermarket entry on healthy eating. The insignificant effects are not due to limited power: with 60,000 households and all supermarket entries in the entire United States over a 13-year period, we are able to detect very small effects.

We can use these estimates to determine the share of the nutrition-income relationship that can be explained by having more local supermarkets. Recall from Section III.A that households in the top income quartile buy groceries with a Health Index that is 0.56 standard deviations higher than those in the bottom quartile (conditional on household size and age and year indicators). The upper bound of the 95 percent confidence interval from column 2 of Panel (b) implies that one supermarket entry increases a household's Health Index by no more than about 0.02 standard deviations for bottom income quartile households. Using ZBP data, we calculate that high-income (low-income) Homescan households have an average of 2.47 (2.03) supermarkets in their zip code, which implies an average difference of 0.44 supermarkets. Thus, we can conclude that local access to supermarkets explains no more than $0.02 \times 0.44/0.56 \approx 1.5\%$ of the Health Index difference between high- and low-income households. In short, differences in local access to supermarkets do not appear to be driving the nutrition-income relationship.

Given the academic and policy attention directed towards local access to healthy grocery options, it is remarkable that supermarket entry seems to matter so little for consumption. However, the limited impact of supermarket entry is less surprising in light of two key facts. First, using data from the 2009 National Household Travel Survey (NHTS), a nationally representative survey that gathers demographics, vehicle ownership, and "trip diaries" from 150,000 households, we find that the average American travels 5.2 miles to shop, with 90 percent of shopping trips being made by car. (See Appendix Figures A6 and A7.) While average distances are slightly shorter among low-income households (4.8 miles) and slightly longer among households living in food deserts (nearly 7 miles), the take-away remains the same: households are willing to travel long distances to purchase their groceries. Second, Appendix Figure A8 shows that as a result of this travel, households in food deserts spend

observed in the sample on household size and age and year indicators.

[24]Appendix Table A6 presents additional estimates considering only entry by supercenters, i.e. excluding grocery and club stores. As might be expected from looking at larger stores, the expenditure share effects are larger, but the effects on the Health Index are again economically small and mostly statistically insignificant. There is one marginally significant unexpected result at the top right of Table 2, suggesting that entry reduces grocery, supercenter, and club store expenditures for households in food deserts. We think of this as an anomaly, as there are no similar results in our prior working papers nor in the supercenter entry study in Appendix Table A6.

only about one percent less of their grocery budgets at grocery stores, supercenters, and club stores than households that are not in food deserts.

Of course, households in food deserts do benefit from more variety and reduced travel costs when a new supermarket enters nearby. However, because most consumers already travel to shop in supermarkets, local supermarket entry does not significantly change choice sets and thus should not be expected to affect healthy eating.

IV.B "Place Effects" Identified by Movers

Although we have shown that supermarket entry has little effect on healthy eating, a related hypothesis is that a broader class of "place effects" could drive grocery purchases. For instance, peer effects from the eating habits of friends and neighbors as well as general local knowledge and image concerns related to healthy eating could drive a household's choices. To illustrate the possibility of place effects, Figure 5 presents a map of the estimated Health Index of packaged grocery purchases by county using the 2006–2016 RMS data.[25] The county-level Health Index is highly correlated with county mean income (correlation coefficient $\rho \approx 0.42$) and with Chetty et al. (2016)'s county-level life expectancy measure ($\rho \approx 0.61$), underscoring both the inequities and the potential implications of what Americans eat and drink. Of course, this geographic variation could reflect any combination of causal place effects and geographic sorting of people with similar preferences.

To test for place effects, we measure within-household changes in grocery purchases after the 20,031 cross-zip code moves and $11,728$ cross-county moves that Homescan households made during our sample period. While the ideal experiment to measure place effects would randomly assign households to different neighborhoods, households in our data move for reasons that may create endogeneity concerns. For example, Appendix Figure A11 and Appendix Table A7 show that moves to healthier counties (although not moves to healthier zip codes) are associated with increased household income. In what follows, we study how the estimates are affected by including controls for observed changes in income, job responsibilities, household composition, and marriage status that could generate endogeneity. Any remaining endogeneity would bias us *toward* finding place effects, because unobserved lifestyle changes and unreported salary increases that cause people to move to healthier places should also cause healthier eating. We therefore interpret the results in this section as upper bounds on true place effects.

Define H_m as the estimated local Health Index of packaged groceries purchased in geographic area m, where m is either a zip code or a county. For county-level H_m, we use the estimates mapped in Figure 5, and we use the same approach and RMS data to calculate H_m by zip code.[26] We estimate

[25] Because the RMS data do not contain the complete census of stores, the distribution of store types in the RMS sample may not match a county's true distribution. For example, the RMS sample might include most of the grocery stores in county A, but few of the grocery stores and most of the drug stores in county B. To estimate the county average Health Index, we thus take the calorie-weighted average Health Index of groceries sold in RMS stores and regression-adjust for the difference between the distribution of store channel types in the RMS data versus the true distribution of store channel types observed in ZBP data.

[26] The average Homescan panelist lives in a zip code with 4.3 RMS stores and a county with 104 RMS stores.

the following regression in household-by-year Homescan data:

$$Y_{imt} = \tau H_m + \gamma \boldsymbol{X}_{it} + \mu_t + \phi_i + \varepsilon_{imt}, \tag{3}$$

where μ_t denotes year indicators, ϕ_i is a household fixed effect, and $\boldsymbol{X}_{it}$ is again a vector of potentially time-varying household covariates described in Table 1. By conditioning on household fixed effects, we isolate changes in grocery purchases associated with changes in neighborhood variables generated by moves. Because most Homescan households are in the sample for only a few years, this within-household design only allows us to estimate place effects over the medium term—that is, a few years after a move.[27] As the set of movers is not nationally representative, we again do not use the Homescan sample weights for this analysis. When estimating Equations (3) and (4), we use robust standard errors with two-way clustering by household and geographic area (zip or county).

Before estimating Equation (3), we first show graphical results of the event study. As before, let B_{it} be an indicator for whether observation it is part of a balanced panel around a move, meaning that the household is observed continuously from one year before the move to two years after and at least 50 percent of the household's trips to RMS stores are to stores located in the household's end-of-year county of residence in all three years other than the year of the move.[28] These restrictions result in balanced panels that include 2,869 cross-zip code moves and 2,277 cross-county moves.

Letting Δ_i denote the change in the average Health Index between a household's final and original location, we estimate the following regression in household-by-year Homescan data:

$$Y_{it} = B_{it} \cdot \left(\alpha \Delta_i + \sum_y (\tau_y \Delta_i + \omega_y) \right) + \gamma \boldsymbol{X}_{it} + \mu_t + \phi_i + \varepsilon_{it}, \tag{4}$$

where y indexes years around the move, with the pre-move year $y = -1$ as the omitted category. The ω_y coefficients are intercepts for each year, α measures the association between the household Health Index and the change in local environment in the year before the move ($y = -1$), and τ_y measures the difference in that association between $y = -1$ and each other year in the event study window. As in Section IV.A, the interaction with B_{it} means that we identify the coefficients of interest using only households in the balanced panel, although we include the full sample in the regression to improve the precision on the demographic associations γ, year effects μ_t, and household fixed effects ϕ_i.

Figure 6 presents this event study analysis for cross-county moves.[29] The top left panel shows

[27] One way to study longer-term trends is to exploit retrospective data on where Homescan panelists previously lived, as in Bronnenberg, Dubé and Gentzkow (2012) and Hut (2018). We have collected Homescan panelists' lifetime histories of census place of residence, but we did not include this analysis in our paper because of endogeneity concerns that might be more relevant here than in the Bronnenberg, Dubé and Gentzkow (2012) study of brand choice.

[28] Although Homescan panelists may report their moves to Nielsen within a few days or weeks of moving, Nielsen only reports the household's county of residence at the end of each year. The 50 percent local shopping restriction aims to eliminate households that move multiple times within a year or otherwise are less strongly exposed to the retail environment in the county where they report living.

[29] Appendix Figures A9 and A10, respectively, present an analogue of Figure 6 for cross-zip code moves and results for balanced panel windows that include more years before and after moves. Some of the post-move point estimates are again positive, but there is no statistically significant evidence that the average household's Health Index converges toward the Health Index of the new area after a move, nor is there evidence of potentially problematic pre-move trends.

the share of shopping trips to RMS stores that are in the new versus the old county for households with $B_{it} = 1$. For these households, almost all of the trips are in the old county before the move, and almost all of the trips are in the new county after. The top right panel presents the distribution across households of the local environment change Δ_i in units of the (normalized) Health Index. The median cross-county mover experiences a local Health Index change of 0.13 standard deviations; this variation in Δ_i is what identifies τ_y.

The bottom panels of Figure 6 shows the estimated τ_y coefficients and 95 percent confidence intervals. The bottom left panel shows results excluding household demographics $\boldsymbol{X}_{it}$, while the bottom right panel includes $\boldsymbol{X}_{it}$. In both cases, there is no statistically significant post-move Health Index change associated with Δ_i, although the point estimates are positive. In other words, the figures show suggestive but insignificant evidence that households purchase more (less) healthy groceries when they move to counties where other households purchase more (less) healthy groceries.

Table 3 presents estimates of Equation (3). Columns 1 and 2 consider cross-zip code moves, while columns 3 and 4 consider cross-county moves. Sample sizes are slightly smaller in columns 1 and 2 because H_m is missing for zip codes with no RMS stores. In all four columns, $\hat{\tau}$ is positive and statistically significant.[30] Columns 2 and 4 include controls for household demographics $\boldsymbol{X}_{it}$; including these demographics has very little impact on the results. However, adding demographic controls only slightly increases the regression R^2, suggesting that unobserved within-household changes could be relevant (Oster, 2016).

Using the above results, we can bound the extent to which location explains the nutrition-income relationship. We consider a partial equilibrium thought experiment in which an individual household moves from a low- to high-income retail environment, leaving aside general equilibrium effects that would occur if this happened on a large scale. The average household in the top quartile of income (residual of household size and age and year indicators) lives in a zip code (county) with a Health Index 0.11 (0.08) higher than households in the lowest income quartile. The upper bound of the confidence intervals on τ for zip codes (counties) from Table 3 is about 0.10 (0.22), and the difference between the high- and low-income Health Index is 0.56 standard deviations. Thus, in combination with the assumption that any endogeneity would bias τ upward relative to the causal effect of place, we conclude that medium-term place effects can explain no more than $0.11 \times 0.10/0.56 \approx 2.0\%$ of the high- versus low-income difference in the Health Index using cross-zip code moves, and no more than $0.08 \times 0.22/0.56 \approx 3.2\%$ using cross-county moves. While we have the power to detect statistically significant associations, it is clear that place effects do not explain a large share of nutritional inequality.

[30] As a point of comparison, Appendix Table A8 uses this same strategy to replicate Bronnenberg, Dubé and Gentzkow (2012)'s immediate brand choice effect, focusing on Coke versus Pepsi. Specifically, we estimate Equation (3) using county-level Coke market shares for H_m, where Coke market shares are defined as [Coke calories purchased/(Coke and Pepsi calories purchased)]. We estimate a highly statistically significant $\hat{\tau} \approx 0.16$, which implies that moving to a county with a 10 percentage point higher Coke market share is associated with a 1.6 percentage point increase in the share of a household's Coke and Pepsi purchases that are of Coke.

IV.C The Nutrition-Income Relationship Within Locations and Stores

We showed in Sections IV.A and IV.B that household consumption responds minimally to improvements in access in the short and medium run. It is possible, however, that improved access to healthy foods has larger effects in the long run, as consumers' tastes and knowledge of healthy eating slowly evolve. In the analysis that follows, we bound the long-run impacts of equalizing local supply by studying households with different incomes but the same retail environment. If differential local supply is entirely to blame for nutritional inequality, then the nutrition-income relationship should disappear when we compare the purchases of households living in the same neighborhood or shopping in the same store. It follows that the component of the nutrition-income relationship that *persists* within retail environments provides us with an estimate of the fraction of nutritional inequality that *cannot* be explained by access.

The ideal experiment to measure the long-run effects of supply would be to randomly assign households to retail environments and observe them over time. In our data, however, households select into neighborhoods and stores based on factors that may be correlated with their demand for healthy groceries. In what follows, we again control for observable household characteristics, but we expect that households observed in the same retail environments are unobservably more similar than households in the full cross-section. Thus, our measure of within-location nutritional inequality likely underestimates the nutrition-income relationship that would persist across the United States if supply were equalized. We therefore interpret the difference between the nutrition-income relationship in the full cross-section and the nutrition-income relationship that persists conditional on access as an upper bound for the share of the nutrition-income relationship that can be explained by supply.

In this section, we measure the nutrition-income relationship within two types of markets: residential zip codes and retail chains. As households living in the same location have access to the same stores, we first measure the nutrition-income relationship that remains among households who live in the same zip code in a given year. Even within a zip code, however, distance to retail outlets varies depending on the location of the household, and factors such as car ownership or proximity to public transportation may yield differences in the ability of households to travel to stores. To control for the possibility that households living in the same neighborhood may still have differential access, we then measure the nutrition-income relationship that persists among households that shop in the same retail chain in a given year.[31]

We estimate the within-market nutrition-income relationship with the following regression specification in household-by-year Homescan data:

$$Y_{imt} = \sum_b \tau_b I_{bi} + \gamma \boldsymbol{X}_{it} + \mu_t \left(+\phi_{mt}\right) + \varepsilon_{imt}, \tag{5}$$

where $\{I_{bi} \mid b \in [1, 15]\}$ is a set of indicators denoting into which of 15 bins household i's sample

[31] The results are very similar when looking at households that shop in the same retail chain in a given year and reside in the same county. This is consistent with findings that product pricing is fairly uniform across establishments within the same chain (DellaVigna and Gentzkow, 2017).

average real income falls, $\boldsymbol{X}_{it}$ is either household size and age indicators or the full vector of household covariates presented in Table 1, and μ_t are year fixed effects. For the analysis examining the nutrition-income relationship within zip codes, Y_{imt} is the Health Index of all purchases made by household i living in zip code m in year t, and ϕ_{mt} are residential zip code-by-year fixed effects. For the analysis examining the nutrition-income relationship within retail chains, Y_{imt} is the Health Index of purchases made by household i in retail chain m in year t, ϕ_{mt} are retail chain-by-year fixed effects, and we limit the sample to bundles accounting for over 50 percent of a household's annual calories purchased, which implies that there is no more than one observation per household-year.[32]

We estimate Equation (5) both with and without market-by-year fixed effects and plot the resulting estimates of τ_b against mean bin income to examine how the nutrition-income relationship changes when we compare households in the same retail environment. The within-market nutrition-income relationship is only identified in markets composed of multiple income bins. To compare the estimates with and without the market fixed effects, we thus restrict the sample to observations in market-year cells in which more than one income bin is represented. [33]

Figure 7 presents the results. The panels on the left present estimates with $\boldsymbol{X}_{it}$ including only household size and age controls, while the panels on the right present estimates controlling for the full vector of household demographics $\boldsymbol{X}_{it}$. The top panels present results with ϕ_{mt} representing zip code-by-year fixed effects, while the bottom panels present results with ϕ_{mt} representing retail chain-by-county-by-year fixed effects. In both cases, the majority of the nutrition-income relationship persists even when looking within a local market. In the sample used for the top panels, the Health Index difference between the top and bottom income quartiles drops from 0.55 to 0.45 standard deviations when adding zip code-by-year fixed effects. In the sample used for the bottom panels, that Health Index difference drops from 0.59 to 0.41 standard deviations when looking within retail chain and year. In combination with the assumption that any endogeneity in household location or shopping decisions would bias our estimates upward, these results demonstrate that differences in local supply can explain no more than a third of nutritional inequality in the long run.

V A Model of Grocery Demand

From the model-free analyses in the previous section, we know that supermarket entry and moving to areas with healthier eating patterns both have limited effects on healthy eating. While compelling, these event studies are conceptually imprecise, because store entry and migration have many effects. For example, supermarket entry changes both pricing and access to healthy and unhealthy foods and may not replicate the grocery supply conditions of a wealthy area. Similarly, moving to different neighborhoods can affect healthy eating in many ways beyond changing the retail environment.

[32] We limit the sample to bundles reflecting over 50 percent of a household's annual calories purchased to mitigate the concern that our within-chain nutrition-income relationship is biased by the presence of small, chain-specific bundles that fail to reflect a households' overall consumption.

[33] These within-market identifying samples contain 91 percent of household-year observations that were used to estimate the nutrition-income gradient in Figure 1 and nearly 100 percent of the household-chain-year observations that account for over 50 percent of the household's annual calories purchased.

The only way to isolate precisely the effect of changing *only* grocery supply conditions from conditions in poor neighborhoods to conditions in wealthy neighborhoods, holding other factors constant, is to estimate a structural model and simulate specific counterfactuals. In addition, the model will allow us to explore the effectiveness of means-tested healthy food subsidies as an alternative policy to reduce nutritional inequality.

V.A Model Setup

We build our structural approach on the framework introduced by Dubois, Griffith and Nevo (2014), though some of our notational conventions differ to accommodate differences in our estimation strategy. Let $\mathcal{S}$ denote the set of stores in a household's choice set. We assume the household has full information about the prices and availability of products across all the stores in $\mathcal{S}$. In a given week, the household visits a subset $s \in \mathbb{P}(\mathcal{S})$ of the stores and incurs shopping costs $\alpha d(s)$, where α is a travel cost per hour and $d(s)$ is the total travel time.

Each week, the household jointly decides which stores to visit and what bundle of goods to purchase to determine its total consumption of calories and dietary characteristics. Let $\boldsymbol{y} = (y_1, ..., y_{\mathcal{N}})$ denote the quantities purchased (measured in calories) of each of the $\mathcal{N}$ food products (UPCs) available across all the stores, and let $\boldsymbol{p} = (p_1, ..., p_{\mathcal{N}})$ denote the corresponding prices paid per calorie. Let $\boldsymbol{\Psi} = (\Psi_1, ..., \Psi_{\mathcal{N}})$ denote the perceived qualities of each of the goods. Finally, let x denote the composite good capturing all the other weekly expenditures, with price normalized to $p_x = 1$ and perceived quality $\Psi_x = 1$.

Each of the $n = 1, ..., \mathcal{N}$ products is characterized by C characteristics $\{a_{n1}, ..., a_{nC}\}$, which could include taste, health implications, shelf life, etc. Define the $\mathcal{N} \times C$ matrix $\boldsymbol{A} = \left\{ \begin{array}{c} a_{11}, ..., a_{1C} \\ ... \\ a_{n1}, ..., a_{nC} \end{array} \right\}$, which measures the content per calorie of each of the C characteristics in each of the $\mathcal{N}$ different goods. The $C \times 1$ vector $\boldsymbol{z} = \boldsymbol{A}'\boldsymbol{y}$ denotes the total characteristic consumption associated with the household's bundle of purchases.

Each week, the household optimizes its calories and characteristics purchased subject to its budget constraint w and the opportunity cost of time spent shopping $\alpha d(s)$. Let $\boldsymbol{\Theta}$ denote a vector of preference parameters. The household's weekly utility maximization problem is:

$$\begin{aligned} & \max_{s \epsilon \mathbb{P}(\mathcal{S}), x, y} U(x, \boldsymbol{z}, \boldsymbol{y}; \boldsymbol{\Theta}, \boldsymbol{\Psi}) - \alpha d(s) \\ s.t. & \\ & \sum_{n=1}^{\mathcal{N}} y_n p_n + x \leq w. \end{aligned} \tag{6}$$

We assume that the utility function $U(x, \boldsymbol{z}, \boldsymbol{y}; \boldsymbol{\Theta}, \boldsymbol{\Psi})$ is continuous, increasing, and strictly quasi-concave. Since U is increasing, the household will spend its entire budget (the budget constraint will bind), and at least one good will always be consumed. We assume that an interior quantity of the composite good is always consumed.

V.B A CES Model of Utility

To formulate a tractable model, we use the Dubois, Griffith and Nevo (2014) utility framework with constant elasticity of substitution (CES) preferences over calories from each of $\mathcal{K}_j$ products within each product group j, Cobb-Douglas preferences over J product groups, and linear preferences over characteristics:

$$U\left(x, \boldsymbol{z}, \boldsymbol{y}; \Theta, \Psi\right) = \sum_{j=1}^{J} \mu_j \ln \left(\sum_{k=1}^{\mathcal{K}_j} \Psi_{kj} y_{kj}^{\theta_j} \right) + \sum_{c=1}^{C} \beta_c z_c + \lambda x. \tag{7}$$

The J product groups, such as carbonated soft drinks, bread, and milk, include all $\mathcal{N}$ products available across all $\mathcal{S}$ stores. Each group contains $k = 1, ..., \mathcal{K}_j$ products. μ_j captures a household's satiation rate over calories consumed in group j. θ_j determines the household's satiation rate over calories consumed through product k in group j. Ψ_{kj} allows for perceived product differentiation, so that the household's marginal benefit of calories can differ across products within a group. λ represents the marginal utility of consuming the outside good. Finally, β_c represents the marginal utility of consuming characteristic c.

Let K_j denote the number of products that the household actually purchases in group j. The first-order conditions from maximizing Equation (7) can be summed over all products purchased in group j and re-written as follows:

$$\sum_{k=1}^{K_j} p_{kj} y_{kj}^{*} = \sum_{c=1}^{C} \frac{\beta_c}{\lambda} \sum_{k=1}^{K_j} a_{kjc} y_{kj}^{*} + \frac{\mu_j \theta_j}{\lambda}, \ j = 1, ..., J. \tag{8}$$

The term $\frac{\mu_j \theta_j}{\lambda}$ represents what the household would spend on product group j if products in that group had zero characteristics. The term $\sum_{c=1}^{C} \frac{\beta_c}{\lambda} \sum_{k=1}^{K_j} a_{kjc} y_{jk}^{*}$ captures the household's additional expenditures in group j due to the products' characteristics. A household will spend more in group j (the left-hand side will be larger) if it satiates more slowly in that group (μ_j and θ_j are larger) or if it gets more characteristics that it values from that group (a_{kjc} is larger for characteristics with more positive β_c). Higher marginal utility of outside good consumption λ reduces grocery expenditures.

VI Estimation and Results

VI.A Empirical Model

To apply the model to data, let $i = 1, ..., I$ index households, let $t = 1, ..., T$ index years, and let τ index subperiods (such as weeks or shopping trips) within year t. The first-order conditions from Equation (8) can be aggregated over τ to the household-by-year level:

$$\sum_{\tau\epsilon t}\sum_{k=1}^{K_j} p_{kj\tau} y_{ikj\tau} = \sum_{c=1}^{C} \frac{\beta_c}{\lambda_i} \sum_{k=1}^{K_j} \sum_{\tau\epsilon t} a_{kjc} y_{ikj\tau} + \sum_{\tau\epsilon t} \frac{\mu_{ij\tau}\theta_{ij\tau}}{\lambda_i}. \tag{9}$$

Equation (9) illustrates the appeal of the model, in that it allows for estimation of characteristic and product group preferences from data aggregated to the level of household-by-product group-by-year. While the model is precisely microfounded, this aggregation allows us to avoid dealing with parameters driving UPC-level preferences and weekly dynamics such as stockpiling. The model potentially allows for considerable heterogeneity across households and time. As described below, we will estimate separate parameters for each of four household income groups, assuming homogeneous β_c parameters within each group.

To economize on notation, define total calories purchased by household i in product group j in year t as $Y_{ijt} = \sum_{\tau\epsilon t}\sum_{k\epsilon J} y_{ikj\tau}$. Define $\tilde{p}_{ijt}$ and $\tilde{a}_{ijct}$, respectively, as the calorie-weighted average price paid and average amount of characteristic c for household i's purchases in group j in year t. Define $\tilde{\delta}_{ijt} = \sum_{\tau\epsilon t} \frac{\mu_{ij\tau}\theta_{ij\tau}}{\lambda_i}$ as the strength of household i's preferences for group j in year t. Finally, define $\tilde{\beta}_c = \frac{\beta_c}{\lambda_i}$, the money-metric marginal utility of each characteristic. Equation (9) can now be written more compactly as:

$$\tilde{p}_{ijt} Y_{ijt} = \sum_{c=1}^{C} \tilde{\beta}_c \tilde{a}_{ijct} Y_{ijt} + \tilde{\delta}_{ijt}. \tag{10}$$

As in the existing literature that uses the characteristics approach to demand, we allow for a product characteristic that is unobserved to the econometrician (Berry, 1994). Let characteristic $c = 1$ be unobserved, and let characteristics $c = 2, ..., C$ be observed. We denote $\xi = \tilde{\beta}_1 \tilde{a}_{ij1t}$ as the unobserved characteristic, which again is assumed to be constant within each income group.

We now depart from the empirical strategy in Dubois, Griffith and Nevo (2014).[34] Instead of estimating Equation (10) directly, we solve for total calories Y_{ijt} and take logs of both sides. We also separate the household's product group preferences $\ln\tilde{\delta}_{ijt}$ into a product group fixed effect δ_j, a geographic market fixed effect ϕ_m (which in practice will be a county), a year fixed effect ϕ_t, and the household-specific deviation ε_{ijt}, so $\ln\tilde{\delta}_{ijt} = \delta_j + \phi_m + \phi_t + \varepsilon_{ijt}$. Our final estimating equation for each income group is thus

$$\ln Y_{ijt} = -\ln\left(\tilde{p}_{ijt} - \sum_{c=2}^{C} \tilde{\beta}_c \tilde{a}_{ijct} - \xi\right) + \delta_j + \phi_m + \phi_t + \varepsilon_{ijt}. \tag{12}$$

Equation (12) uses intuitive variation to estimate the preference parameters. The independent

[34] Dubois, Griffith and Nevo (2014) directly estimate the average nutrient preference parameters β_c using a version of Equation (10):

$$\tilde{p}_{ijt} Y_{ijt} = \sum_{c=2}^{C} \tilde{\beta}_c \tilde{a}_{ijct} Y_{ijt} + \epsilon_{ijt}, \tag{11}$$

including additional fixed effects and instrumenting for $\tilde{a}_{ijct} Y_{ijt}$ using variation in local product availability. Unfortunately, the error term in this regression contains purchases of the unobserved nutrient: $\epsilon_{ijt} = \xi Y_{ijt} + \tilde{\delta}_{ijt}$. Thus, if the instrument affects consumption Y_{ijt}, it *mechanically* is also correlated with the error term. No instrument can address this mechanical endogeneity problem.

variable inside the parentheses is an "implicit price": the actual price adjusted for the utility value of the characteristics in product group j. As this implicit price increases, quantity purchased decreases. Despite using a Cobb-Douglas function form that typically restricts product group price elasticities to one, variation in the level of the unobserved characteristic will lead to variation in price elasticities. Specifically, an income group's price elasticity is determined by the absolute magnitudes of $\tilde{\beta}_c$ and ξ: larger (smaller) $\tilde{\beta}_c$ and ξ parameters scale down (up) the importance of price variation in determining quantity purchased Y_{ijt}. By allowing these preference parameters to vary by income group, we allow different income groups to differentially value both characteristics and price.

VI.B Identifying Characteristic Preferences

In Section VI.D, we detail a method of moments estimator of Equation (12). A key moment condition for identifying characteristic preferences $\tilde{\beta}_c$ will be

$$\mathbb{E}\left((\delta_j + \varepsilon_{ijt})\tilde{a}_{ijct}\right) = 0. \tag{13}$$

This moment condition generates two types of variation with which to identify characteristic preferences $\tilde{\beta}_c$. First, we use the variation *between* product groups: we infer that $\tilde{\beta}_c$ for sodium is high if consumers spend more on product groups with high sodium. Second, we use variation across households *within* product groups: we also infer that $\tilde{\beta}_c$ for sodium is high if consumers who purchase especially salty products within a group purchase more calories in that group.

We face a standard econometric challenge with the characteristics approach to demand: characteristics are not randomly assigned to products, and valid instruments for characteristics are unavailable.[35] Because unobserved product characteristics could be correlated with $\tilde{a}_{ijct}$, we should not necessarily interpret $\tilde{\beta}_c$ as the causal effect of characteristic c on demand. For example, because salt acts as a preservative, preferences for salt could also capture preferences for longer shelf life. This could be problematic if we were, for example, predicting how demand would respond to product reformulations that change specific characteristics in isolation. This issue is less problematic for our analyses, in which we simply want to interpret how preferences for groceries including healthy or unhealthy characteristics vary by income and then simulate how demand would respond to changes in prices or overall bundles of products available.

[35] To interpret $\tilde{\beta}_c$ as the causal effect of characteristics on demand, the identifying assumption in Equation (13) has two economic implications, which parallel these two types of identifying variation. First, we must assume that product group preferences δ_j are uncorrelated with characteristic contents $\tilde{a}_{ijct}$. Second, individual households' idiosyncratic tastes ε_{ijt} must be uncorrelated with the variation in $\tilde{a}_{ijct}$ across households but within product groups. In the context of the model, households' idiosyncratic perceptions of product kj's quality Ψ_{ikj} influence y^*_{ikjt}, the quantities of individual products purchased within the group, and hence affect calorie-weighted average characteristics $\tilde{a}_{ijct}$. Thus, for the identifying assumption to hold, Ψ_{ikj} must be independent of the terms that make up ε_{ijt}: households' idiosyncratic preference parameters μ_{ijt} and θ_{ijt}.

VI.C Price Endogeneity

A more worrisome source of endogeneity arises from a potential correlation between a household's idiosyncratic product group preferences ε_{ijt} and the household's average price paid $\tilde{p}_{ijt}$:

$$\mathbb{E}\left(\varepsilon_{ijt}\tilde{p}_{ijt}\right) \neq 0. \tag{14}$$

Such endogeneity could arise if households endogenously shop at stores offering lower prices for product groups on which they spend more, if households seek out systematically different quality levels (and thus price levels) in groups on which they spend more, or through standard simultaneity bias, in which retailers set higher markups in response to higher demand.

To address the possibility of price endogeneity, we develop a new instrument for prices. The underlying intuition for our instrument is that retail chains differ in their sourcing and distribution costs across products, giving different chains heterogeneous comparative advantages in supplying different products. Then, because different chains are present in different geographic areas, the relative prices of different products also vary across areas. To illustrate, consider a simple example in which there are two types of foods, apples and pizza, and two grocery chains, Safeway and Shaw's. Suppose Safeway is able to supply pizza cheaply, while Shaw's can supply apples cheaply. Then, areas dominated by Safeway will have relatively low pizza prices, while areas dominated by Shaw's will have relatively low apple prices.

We construct our instrument as follows. For retail chain r in market (i.e. county) m during time period t, let $\ln(p_{krt,-m})$ denote the average log price of UPC k in stores from the same chain but in all markets excluding market m, denoted by $-m$. Let $\ln(p_{kt,-m})$ denote the national average log price of UPC k in period t in all markets excluding m. We exclude market m to ensure that the IV reflects a chain's comparative advantages in supplying product k based on other markets, not local demand conditions in market m. Retail chain r's cost advantage in supplying UPC k relative to the national average is thus $\Delta \ln(p_{krt,-m}) = \ln(p_{krt,-m}) - \ln(p_{kt,-m})$.

Let N_{rmt} denote retailer r's number of establishments in market m, let N_{jrt} denote retailer r's nationwide unit sales in product group j, and let N_{kt} denote the total nationwide unit sales of product k in year t. The price instrument P_{jmt} is the weighted average cost advantage that chains in market m have for UPCs in product group j:

$$P_{jmt} = \frac{\sum_{r\in m} N_{rmt}N_{jrt} \cdot \sum_{k=1}^{\mathcal{K}_j} N_{kt}\Delta \ln(p_{krt,-m})}{\sum_{r\in m} N_{rmt}N_{jrt} \cdot \sum_{k=1}^{\mathcal{K}_j} N_{kt}}. \tag{15}$$

Our identifying assumption is that household i's idiosyncratic preferences for product group j are uncorrelated with the price instrument P_{jmt} for group j in household i's market:

$$\mathbb{E}\left(\varepsilon_{ijt}P_{jmt}\right) = 0. \tag{16}$$

The variation in our instrument reflects geographic variation in different retailers' presence across markets. Because our model includes product group and market fixed effects, our instrument relies on variation in the relative prices across product groups within a market.[36] The key economic content of our identifying assumption is that chains do not have comparative pricing advantages in product groups where their customers have stronger tastes.

A suggestive test of our identifying assumption is whether there is a correlation between county income and the instrument's predicted relative price for healthy product categories. Although not fully dispositive in either direction, such a correlation would suggest that chains located in places with higher demand for healthy foods tend to have a comparative advantage in supplying healthy food, which would violate the identifying assumption. To test this, we regress our price instrument on the interaction of product group j's average Health Index H_j with a dummy for whether the county average income in 2010 is above the nationwide median, conditional on the same product group and county fixed effects from Equation (12):

$$P_{jmt} = \gamma \cdot \mathbf{1}(Income_m > National\ median\ income) * H_j + \delta_j + \phi_m + \epsilon_{jmt}. \tag{17}$$

γ measures whether healthier product groups are differentially priced in higher-income counties than lower-income counties. Figure 8 presents a binned scatterplot of this regression after partialling out the fixed effects. We see zero correlation between a product group's Health Index and our price instrument in higher-income counties versus lower-income counties, reducing concerns that our price IV is driven by demand-side variation.

Appendix Figure A12 presents maps illustrating the geographic presence of several large retail chains in RMS. There is substantial variation within and between regions of the country. Appendix Figure A13 presents maps illustrating the resulting variation in P_{jmt} across U.S. counties for example product groups. Produce is predicted to be cheap on the west coast and expensive on the east coast, likely because so much produce is grown in California. Yogurt is predicted to be cheap for a dispersed set of counties in the midwest. Cookies are predicted to be relatively cheap in Massachusetts and West Texas. The figures make clear that there is substantial national and within-region variation, that this variation is not closely related to county income, and that there are substantial differences in the spatial patterns across product groups.

Our conversations with insiders from the grocery industry suggest that several factors contribute to differences in comparative advantage across retail chains. First, different products are produced in different parts of the country, generating transportation costs that vary across retail chains in different regions. Second, some retailers are larger than others, and economies of scale vary across product groups. Third, although the Robinson-Patman Act prohibits wholesale price discrimination, the Act is increasingly unenforced (Lipman, 2012), and producers often offer more subtle contractual incentives that generate variation in effective marginal cost across retailers.

Appendix Figure A14 plots the standard deviation in our instrument for the six product "depart-

[36] DellaVigna and Gentzkow (2017) show that retail chains whose stores are in markets with more inelastic demand tend to charge higher prices than other retail chains whose stores are in markets with more elastic demand. Our market fixed effects are designed to address this form of endogeneity in response to overall market demand patterns.

ments" (broadly aggregated product categories as defined by Nielsen), after residualizing against year, market, and product group fixed effects. Fresh produce has the most variation, followed in order by dairy, frozen foods, packaged meats and deli items, and dry grocery. This ordering is consistent with how costs vary across chains. Because produce is grown more in certain areas of the country, transportation costs differ considerably across chains in different regions. Furthermore, fresh produce and dairy require refrigeration and are highly perishable, making their cost quite sensitive to a chain's distribution network. By contrast, dry grocery items require no refrigeration and have long shelf lives, and are thus easy to transport.

The instrument is very powerful. Appendix Figure A15 shows that there is a robust linear relationship between log prices and the instrument, controlling for market and product group fixed effects. A linear version of our IV procedure has first stage F-statistics ranging from 243 to 260 in the four income groups.

This instrument is novel in the literature, and it can be used in situations in which other instruments do not generate identification or fail the exclusion restriction. DellaVigna and Gentzkow (2017), for example, use price variation from individual stores' short-term promotions. This variation is useful in identifying a store's residual demand elasticity; it may be less useful for identifying households' preferences, especially if households substitute across stores or stockpile goods bought on sale. Hausman (1996) uses variation in prices over time in other markets, which is valid under the assumption that demand shocks are uncorrelated across markets. By contrast, our instrument generates cross-sectional identification, while relying on an exclusion restriction that could be more plausible in many applications.

VI.D Method of Moments Estimation

For estimation, we construct separate datasets for each of four household income quartiles, where income is residual of household size and age and year indicators to be consistent with the earlier parts of the paper. Data are at the household-by-product group-by-year level. We define $J = 45$ product groups using a slight modification of Nielsen's original "product group" variable, combining a handful of groups with infrequent purchases so as to minimize observations with zero purchases. We drop any remaining observations with zero purchases, as the first-order condition does not hold for these observations.[37] We define $C - 1 = 12$ observed product characteristics for the 12 dietary components that enter the Health Index: grams of sodium per 1,000 calories, ounces of whole grains per 1,000 calories, etc. For each income group, we estimate four parameter vectors: the $(C-1) \times 1$ vector $\tilde{\boldsymbol{\beta}}$ of preferences for observed characteristics, the scalar ξ representing the unobserved characteristic, the $J \times 1$ vector $\boldsymbol{\delta}$ of product group fixed effects, and the $M \times 1$ vector $\boldsymbol{\phi}$ of market fixed effects.

To specify the moment conditions, let $\boldsymbol{D}_j$ be a $J \times 1$ vector of dummy variables for whether the observation is in product group j, and let $\boldsymbol{D}_m$ be an $M \times 1$ vector of dummy variables for whether

[37] We drop 10.6% of observations at the household-by-product group-by-year level because they have zero purchases. "Baby food" is the product group with the most missing observations. The differences across income groups in a product group's share of missing observations are not correlated with the product group's average characteristic contents, suggesting that these dropped observations do not affect our estimated preference heterogeneity across income groups.

household i is in market m. The model estimation relies on the following set of $(C + J + M)$ identifying moments that just identify our model parameters:

$$\begin{aligned}
&\mathbb{E}\left((\delta_j + \varepsilon_{ijt})\tilde{a}_{ijct}\right) &&= 0 \quad , \ c = 1, ..., C \\
&\mathbb{E}\left(\varepsilon_{ijt} P_{jmt}\right) &&= 0 \\
&\mathbb{E}\left(\varepsilon_{ijt} D_{ijt}\right) &&= 0 \quad , \ j = 1, ..., J \\
&\mathbb{E}\left(\varepsilon_{ijt} D_{im}\right) &&= 0 \quad , \ m = 1, ..., M
\end{aligned} \quad . \tag{18}$$

Loosely, the first set of moments identifies the $\tilde{\boldsymbol{\beta}}$ parameters, the second identifies ξ, the third set identifies $\boldsymbol{\delta}$, and the fourth set identifies $\boldsymbol{\phi}$. We construct the sample analogue of the moment conditions from Equation (18) in the following $(C + J + M) \times 1$ vector: $\boldsymbol{g}_{ijt} = \boldsymbol{g}_{ijt}\left(\boldsymbol{\delta}, \boldsymbol{\phi}, \tilde{\boldsymbol{\beta}}, \xi\right)$. We can then write the method of moments estimator as:

$$\left(\hat{\boldsymbol{\delta}}, \hat{\boldsymbol{\phi}}, \hat{\tilde{\boldsymbol{\beta}}}, \hat{\xi}\right) = \underset{(\boldsymbol{\delta}, \boldsymbol{\phi}, \tilde{\boldsymbol{\beta}}, \xi)}{\arg\min} \left(\frac{1}{IJT}\sum_i \sum_j \sum_t \boldsymbol{g}_{ijt}\right)' \left(\frac{1}{IJT}\sum_i \sum_j \sum_t \boldsymbol{g}_{ijt}\right). \tag{19}$$

Appendix D.A presents details of the estimator and its standard errors.

VI.E Estimation Results

Table 4 reports the estimated characteristic preference parameters $\hat{\tilde{\beta}}_c$ for each of the four income quartiles. The units on each parameter are those originally specified by the Healthy Eating Index. For example, produce is measured in cups, while protein is measured in ounces. The estimated $\hat{\tilde{\beta}}_c$ measures the average willingness to pay (WTP) to have a unit of characteristic c instead of a calorie containing no healthy or unhealthy characteristics—for example, a calorie of unsaturated fat. This normalization also removes differences in λ, the marginal utility of a dollar, across income groups.

Among the eight "healthy" dietary components, six of them (whole fruit, other fruit, whole grains, greens and beans, vegetables other than greens and beans, and dairy) display a strong and almost uniformly monotonic increase in willingness to pay as household income increases. The only healthy characteristic not fitting this pattern is protein from fish and plants. All households dislike this type of protein, especially high-income households. Meat protein is valued similarly across the income distribution. Among the four "unhealthy" dietary components, two of them (sodium and added sugar) are are more strongly disliked by high-income households. The estimated WTP for added sugar is especially striking: bottom income quartile households are willing to pay \$0.0002 to consume a gram of added sugar, while top income quartile households are willing to pay \$0.004 to avoid consuming a gram of added sugar. Added sugar is the only component for which high- and low-income households have opposite-signed preferences, highlighting substantial differences in preferences for added sugar across the income distribution. High-income households have stronger preferences for the remaining two unhealthy components, refined grains and solid fats. The magnitudes of some preference differences are large: the highest-income quartile dislikes sodium nearly three times as much as the lowest-income quartile, and likes dairy about twice as much. We also find that higher-income households have the lowest estimated unobserved characteristic. Accordingly, higher-income households are less sensitive to prices.

To derive a summary statistic for overall preferences for healthy eating, we sum the $\hat{\hat{\beta}}_c$, weighting by the healthfulness of a unit of each dietary component. To implement this, recall that the Healthy Eating Index grants points for consuming more of "healthy" components and less of "unhealthy" components, with a minimum component score of zero and a maximum score of five or ten achieved at minimum and maximum thresholds. If $s_c \in \{5, 10\}$ is the maximum HEI score for dietary component c and r_c is the difference in consumption of component c to receive the maximum instead of the minimum score (with $r_c > 0$ for "healthy" components and $r_c < 0$ for "unhealthy" components), column 14 reports $\sum_c \hat{\hat{\beta}}_c s_c r_c$. All income groups value healthy groceries, but the highest-income group is willing to pay the most, making healthy eating a normal good. The lowest-income quartile is willing to pay \$0.43 per 1,000 calories to consume the maximum-scoring bundle, while the highest-income quartile is willing to pay \$1.14 per 1,000 calories.

As discussed in Section VI.B, the $\hat{\hat{\beta}}_c$ parameters are most safely interpreted as preferences for dietary component c and any unmeasured correlates. For example, foods differ in shelf life and preparation time, and these characteristics could both be correlated with dietary components and be valued differently by low- versus high-income households. To consider this possibility, we collected data on each UPC's shelf life (from the U.S. government's FoodKeeper app (HHS 2015)) and convenience of preparation (from Okrent and Kumcu (2016)). Appendix Table A9 reports the model estimates including these two additional characteristics. Higher-income households have stronger preferences for convenience and fresher foods with shorter shelf lives, but the patterns of preferences for dietary components across the income distribution are similar to the primary estimates. Willingness to pay to consume the maximum-scoring bundle of dietary components is now lower for all income quartiles: \$0.20 per 1000 calories for the lowest-income quartile, and \$0.63 for the highest-income quartile. Even though these dollar values are lower, the ratio of the high-income WTP to the low-income WTP is now larger: $0.63/0.20 \approx 3.15$, compared to $\$1.14/\$0.43 \approx 2.65$ in the primary specification.

VII Explaining and Reducing Nutritional Inequality

VII.A Decomposing Consumption Differences into Supply versus Demand Factors

Using the model estimates from Section VI.E, we decompose the observed differences in healthy eating across income groups into underlying supply-side factors (prices and availability) and demand-side factors (preferences for product groups and characteristics). Because our model is estimated at the product group level, our counterfactuals only allow households to re-optimize their calorie demand across product groups. We do not analyze how households would change their relative quantities across UPCs within product groups.

For this section, we index parameters for each of the four household income groups (residual of household size and age and year indicators) by $g \in \{1,2,3,4\}$. For a given set of prices $\tilde{p}_{gj}$, observed and unobserved characteristics $\tilde{a}_{gjc}$ and ξ_g, product group preferences δ_{gj}, and characteristic preferences $\tilde{\beta}_{gc}$, we can predict $\hat{Y}_{gj}$, the total calories that income group g consumes within product group j:

$$\hat{Y}_{gj} = \frac{\exp(\delta_{gj})}{\tilde{p}_{gj} - \sum_{c=2}^{C} \tilde{\beta}_{gc}\tilde{a}_{gjc} - \xi_g}. \tag{20}$$

For simplicity, this equation excludes the market fixed effect ϕ_m. This omission just scales consumption up or down for all product groups within an income group, which does not affect our calculation of the overall Health Index.

For each income group, we construct a representative product for each product group, and we calculate the resulting representative price $\tilde{p}_{gj}$, observed characteristics $\tilde{a}_{gjc}$, and Health Index H_{gj}. This representative product is the weighted average of products available in RMS stores where each income group shops, weighting stores by their share of nationwide trips and weighting UPCs by their share of nationwide calorie consumption. Specifically, let Q_{kj} denote the total nationwide quantity of calories sold of UPC k in product group j, let N_{ge} denote the number of trips made by Homescan households of income group g to store e, and let $\mathbf{1}(kj \in e)$ denote the indicator function for whether product kj is stocked in store e. The Health Index of the representative product is the weighted average of the Health Index for each UPC within the product group:

$$H_{gj} = \frac{\sum_k Q_{kj} \sum_s N_{ge}\mathbf{1}(kj \in e) H_{kj}}{\sum_k Q_{kj} \sum_s N_{ge}\mathbf{1}(kj \in e)}. \tag{21}$$

The representative price and observed characteristics are calculated analogously, substituting $\tilde{p}_{gj}$ and $\tilde{a}_{gjc}$ for H_{jk}.

To evaluate the bundle of goods characterized by predicted group-level consumption $\hat{Y}_{gj}$, we then calculate the overall Health Index:

$$\hat{H}_g = \sum_j \hat{Y}_{gj} H_{gj}. \tag{22}$$

For this subsection, we re-normalize the Health Index so that the initial difference between the highest and lowest income groups equals one. The left-most points on Figure 9 display each income group's initial Health Index level, calculated by substituting the predictions of Equation (20) into Equation (22).

We can now simulate the changes in healthy eating that would occur under different counterfactual scenarios for supply- and demand-side factors. We begin with the supply side, motivated by the arguments that food deserts are a key cause of the nutrition-income relationship. Our first counterfactual explores one aspect of the food desert policy discussion by assessing the role of product group prices in healthy eating differences across the income groups. For each income group, we set product group prices $\tilde{p}_{gj}$ to the levels observed in the highest income group, $g = 4$. Thus, for each income group g, we calculate the product group calorie demand as:

$$\hat{Y}_{gj} = \frac{\exp(\delta_{gj})}{\tilde{p}_{4j} - \sum_{c=2}^{C} \tilde{\beta}_{gc}\tilde{a}_{gjc} - \xi_g}. \tag{23}$$

We then compute the corresponding Health Index by substituting the quantities in Equation (23) into Equation (22). Figure 9 shows that prices do not appear to explain much of the difference in the Health Index across income groups.

Our second counterfactual explores the part of the food desert discussion related to availability, by setting the characteristic levels $\tilde{a}_{gjc}$ in each product group equal to those observed in the highest income group. We make this change in addition to equating the price level across income groups. We now recompute each income group's total calorie demand in product group j as follows:[38]

$$\hat{Y}_{gj} = \frac{\exp(\delta_{gj})}{\tilde{p}_{4j} - \sum_{c=2}^{C} \tilde{\beta}_{gc}\tilde{a}_{4jc} - \xi_g}. \tag{24}$$

Figure 9 shows that equalizing the availability of nutrients decreases the Health Index difference between the highest and lowest income groups. In combination, equalizing prices and availability decreases the difference by nine percent. The combined effect of equalizing prices and availability decreases the difference by nine percent. These first two counterfactuals confirm our findings from Section IV: differences in supply do not explain very much of the nutrition-income relationship. Contrary to the view of some advocates and policymakers, even fully equating availability and prices of healthy foods for all households would reduce nutritional inequality by only nine percent.

We now explore the role of demand-side differences. In addition to the changes in prices and availability of characteristics, we now also set the characteristic preferences $\tilde{\beta}_{gc}$ in each income group to those of the highest-income group, generating calorie demand

$$\hat{Y}_{gj} = \frac{\exp(\delta_{gj})}{\tilde{p}_{4j} - \sum_{c=2}^{C} \tilde{\beta}_{4c}\tilde{a}_{4jc} - \xi_4}. \tag{25}$$

Figure 9 shows that equalizing the nutrient preferences closes most of the gap in healthy eating. While the bundles chosen by the highest-income group are still healthier than those of the lower income groups, the gap between the highest and lowest income groups has declined by 46 percent relative to the baseline.

Finally, in addition to prices, characteristic availability, and characteristic preferences, we also set the product group preferences δ_j equal to those of the highest income group, generating total calorie demand:

$$\hat{Y}_{gj} = \frac{\exp(\delta_{4j})}{\tilde{p}_{4j} - \sum_{c=2}^{C} \tilde{\beta}_{4c}\tilde{a}_{4jc} - \xi_4}. \tag{26}$$

[38] The unobserved characteristic ξ represents a combination of the amount of the unobserved characteristic and the income group's preference for it. While this is a mix of supply and demand forces, we will attribute it to demand because it primarily determines the consumer's demand elasticity with respect to price, a demand side force.

By construction, this last counterfactual mechanically equalizes the observed purchases across each of the income groups, as seen in Figure 9.

In summary, over 90 percent of the nutrition-income relationship is due to demand-side factors related to preferences, while less than 10 percent is explained by the supply side. Consistent with our model-free analyses in Section IV, this finding counters arguments that food deserts are important contributors to nutritional inequality.

VII.B Using Observables to Explain Demand for Healthy Groceries

The results of the previous section highlight that demand-side factors, not supply, are central to explaining the nutrition-income relationship. But what explains the differences in demand across income levels? In this section, we first show which household-level observable characteristics predict demand for healthy groceries. Then, we show which household observables explain the correlation between income and healthy grocery demand.

The structural model allows us to implement a key distinction: instead of analyzing equilibrium *purchases* of healthy groceries, we isolate each household's *demand* for healthy groceries, while holding supply conditions constant. To do this, we compute the sample average prices and observed characteristics for each product group, denoted $\bar{p}_j$ and $\bar{a}_{jc}$. As in the decomposition in Section VII.A, we interpret these two parameters as reflecting "supply." Using the parameter estimates, we also back out each observation's fitted error term, ε_{ijt}, which captures heterogeneity in preferences within an income group. For household i in year t, we then compute the total calorie demand in all product groups at the sample average supply parameters:

$$\hat{Y}_{ijt} = \frac{\exp(\delta_{gj} + \phi_{gm} + \varepsilon_{ijt})}{\bar{p}_j - \sum_{c=2}^{C} \tilde{\beta}_{gc}\bar{a}_{jc} - \xi_g}. \tag{27}$$

We then insert $\hat{Y}_{ijt}$ into an analogue of Equation (22) to calculate H^D_{it}, the Health Index of household i's grocery demand at sample average supply, and normalize H^D_{it} to have a standard deviation of one.

To determine the correlates of demand for healthy groceries, we run the following regression:

$$H^D_{it} = \alpha \ln w_i + \gamma^1 \boldsymbol{X}^1_{it} + \gamma^0 \boldsymbol{X}^0_{it} + \mu_t + \varepsilon_i, \tag{28}$$

where w_i denotes household i's sample average income, μ_t are year indicators, $\boldsymbol{X}^0_{it}$ denotes household size and age indicators (the same controls used in Figure 1), and $\boldsymbol{X}^1_{it}$ denotes the remaining demographics from Section IV (natural log of years of education, race indicators, an indicator for whether the household heads are married, employment status, and weekly work hours), plus two additional variables from the Homescan add-on survey carried out by Nielsen for Allcott, Lockwood and Taubinsky (2018): the self-reported importance of staying healthy and nutrition knowledge, both normalized to have a standard deviation of one.[39] As discussed in Section III.A, we control for $\boldsymbol{X}^0_{it}$ so as to inter-

[39] The health importance variable is from the question, "In general, how important is it to you to stay healthy, for example by maintaining a healthy weight, avoiding diabetes and heart disease, etc.?" Original responses were on a scale from 0 to 10. Nutrition knowledge is the score on 28 questions from the General Nutrition Knowledge Questionnaire,

pret w_i as a rough measure of the household's socioeconomic status, while we think of $\boldsymbol{X}_{it}^1$ as possible mediators of the relationship between socioeconomic status and healthy grocery demand. The sample sizes are smaller than in previous tables because we restrict the sample to only those households that responded to the survey. Standard errors are clustered by household, and observations are weighted by the Homescan sample weights.

Column 1 of Table 5 presents results. Education is strongly correlated with healthy grocery demand: the coefficient on the natural log of years of education is 0.685, meaning that at the sample mean education level, a one-year increase in education is associated with a 0.05 standard deviation increase in demand for healthy groceries. A one standard deviation increase in nutrition knowledge is associated with a 0.13 standard deviation increase in demand for healthy groceries. In separate regressions in which we exclude nutrition knowledge, the education coefficient rises by about 20 percent, suggesting that health knowledge is one mechanism through which education may influence healthy behavior.[40]

The second goal of this section is to determine which observables explain the relationship between income and demand for healthy groceries. We use the approach of Gelbach (2016), which is to think of covariates $\boldsymbol{X}_{it}^1$ as "omitted variables" in the relationship between healthy grocery demand H_{it}^D and household income $\ln w_i$, conditional on $\boldsymbol{X}_{it}^0$, and then calculate the "omitted variables bias" from excluding each specific covariate. Column 2 presents the relationship between H_{it}^D and $\ln w_i$, conditional on $\boldsymbol{X}_{it}^0$, giving a coefficient we denote as $\tilde{\alpha}$. The income coefficient in column 1 is $0.134/0.267 \approx 0.50$ of the coefficient in column 2. Therefore, these observables explain about one-half of the relationship between income and healthy grocery demand.

Column 3 of Table 5 presents estimates of eight separate auxiliary regressions, $\ln w_i = \Gamma X_{vit} + \gamma^0 \boldsymbol{X}_{it}^0 + \epsilon_i$, which give $\hat{\Gamma}$, the conditional covariance between the natural log of income and the individual variable X_v. Following Gelbach (2016), we can then estimate variable X_v's contribution to the relationship between income and demand for healthy groceries by using the omitted variable bias formula: $\hat{\pi}_v = \hat{\Gamma}_v \hat{\gamma}_v$. As with standard omitted variable bias, a covariate will explain more of the relationship if it is more strongly associated with healthy grocery demand or with income. Finally, dividing by $\tilde{\alpha}$ gives variable X_v's estimated contribution as a share of the unconditional relationship:

$$\tilde{\pi}_v = \frac{\hat{\Gamma}_v \hat{\gamma}_v}{\tilde{\alpha}}. \tag{29}$$

As shown in Gelbach (2016), these contributions can be aggregated to consider vectors of indicator variables jointly, and we do this for the vector of census division indicators. Figure 10 presents the estimated $\tilde{\pi}_v$ parameters and 95 percent confidence intervals. Education explains the largest share of the relationship between demand for healthy groceries and income, at about 20 percent. Nutri-

which is standard in the public health literature (Kliemann et al., 2016). One example question is, "If a person wanted to buy a yogurt at the supermarket, which would have the least sugar/sweetener?" The possible responses are "0% fat cherry yogurt," "Plain yogurt," "Creamy fruit yogurt," and "Not sure"; the correct answer is "Plain yogurt."

[40] This is consistent with literature showing that education is highly predictive of health behaviors and outcomes (e.g., Grossman (2015) and Furnee et al. (2008)), and with the Cutler and Lleras-Muney (2010) finding that knowledge explains 10–20 percent of the relationship between education and drinking and smoking.

tion knowledge explains the second-largest share, at about 14 percent. These results are correlations, so they do not reflect the causal effect of additional education or nutrition knowledge interventions. Notwithstanding, they are certainly suggestive of how improved educational opportunities and nutrition information could play a role in reducing nutritional inequality.

VII.C Using Subsidies to Reduce Nutritional Inequality

Our analyses suggest that supply-side policies such as encouraging supermarket entry will have limited effects on healthy eating. In this section, we study an alternative policy: subsidies for healthy foods. There are many types of taxes and subsidies that could affect healthy eating. To ease interpretation, we focus on a simple subsidy that scales in a product's healthfulness and is available only to the bottom quartile of income distribution, again conditional on household size and age and year indicators. While the exact setup of these counterfactuals is developed to fit with our previous analyses, we think of this as a stylized implementation of a healthy food subsidy within SNAP, the U.S. government's means-tested nutritional support program.

Specifically, we consider an ad valorem subsidy for each product that is proportional to its Health Index. In order to target only healthy foods, we set the subsidy to zero for all foods with a below-average Health Index. Because our demand estimates focus on between-product group demand, we continue analyzing composite products representing the calorie-weighted average product sold within product group j among the products purchased by income group g. The percent subsidy for product j for households in the bottom income quartile ($g = 1$) is

$$s_{1j} = \begin{cases} \min\{s(H_{1j} - \bar{H}), 0.95\} & \text{if } (H_{1j} - \bar{H}) > 0 \\ 0 & otherwise. \end{cases} \tag{30}$$

We limit the subsidy to a maximum of 95 percent of the product group's price, which binds in a few cases at the highest subsidy considered below.

Income quartile $g = 1$'s subsidized price for product group j is $\tilde{p}^s_{1j} = \tilde{p}_{1j}(1 - s_{1j})$, where $\tilde{p}_{1j}$ is the calorie-weighted average price paid among bottom income quartile households in the data. Bottom-quartile households would have the following subsidized demand for each product group:

$$\hat{Y}_{1j} = \frac{\exp(\delta_{1j})}{\tilde{p}^s_{1j} - \sum_{c=2}^{C} \tilde{\beta}_{1c}\tilde{a}_{1jc} - \xi_1}. \tag{31}$$

By construction, the subsidy is not available to the top three income quartiles, so their demand is unchanged. We numerically solve for the subsidy s that increases the bottom income quartile's Health Index by a given amount.

Table 6 presents results. Column 1 presents the subsidy that reduces the Health Index gap between top-quartile and bottom-quartile households by 0.9 percent—the point estimate of the impact of supermarket entry from Section IV.A.[41] Column 2 presents the subsidy that reduces the Health Index

[41] Our point estimate from Table 2, panel (b), column 2 is that one supermarket entry within 10 minutes' drive of a

gap by nine percent—the effect of equating supply estimated in Section VII.A. Column 3 presents the subsidy that brings bottom quartile households' Health Index to the level of the top quartile.

The first two rows present the subsidy parameter s and the mean percent subsidy for products with an above-median Health Index, which are the products receiving non-zero subsidies. The third row presents the average subsidy payment per bottom-quartile household. The bottom row presents the total subsidy payment, aggregating over all 31.45 million households in the bottom income quartile.

The results in column 1 show that an annual subsidy of $84 million would increase healthy eating by the same amount as one additional supermarket entry within a 10-minute drive of *all* bottom-quartile households. In comparison, the Healthy Food Financing Initiative has spent about $220 million of its $400 million budget on store subsidies (TRF 2017), and various state programs have spent tens of millions (CDC, 2011). Of course, government expenditures are far from a complete measure of social costs; the impacts of these supply-side programs on store entry are unclear. What we can conclude is that a supermarket entry subsidy of more than $2.62 per year per household within a 10-minute drive would be less cost-effective than this means-tested subsidy at increasing the Health Index of bottom-quartile households' grocery purchases.

Column 2 shows that an annual subsidy of $830 million would increase healthy eating by the same amount as providing bottom-quartile households the same supply conditions as top-quartile households. From this, we can similarly conclude that even a suite of supply-side policies that are somehow able to achieve this full equalization of supply will only be cost effective if they cost less than $830 million per year.

Column 3 shows that a subsidy of $11 billion per year could raise bottom-quartile households' Health Index all the way to the level of top-quartile households. In comparison, the annual SNAP budget in 2016 was $71 billion (CBPP 2018). Thus, a healthy food subsidy within SNAP could eliminate this measure of nutritional inequality at a cost of only about 15 percent of the SNAP budget. Of course, SNAP has many objectives, and there are many economic and political considerations around modifying the SNAP program to encourage healthy eating (Richards and Sindelar, 2013; Schanzenbach, 2017; Shenkin and Jacobson, 2010). But this result on the cost effectiveness of healthy eating subsidies, combined with our earlier results on the lack of efficacy of supply-side policies, suggests that policymakers interested in reducing nutritional inequality might redirect efforts from supply-side policies toward means-tested subsidies.

VIII Conclusion

We study the causes of "nutritional inequality": why the wealthy tend to eat more healthfully than the poor in the United States. The public health literature has documented that lower-income neighborhoods suffer from lower availability of healthy groceries and that lower-income households tend to eat less healthfully. In public policy circles and in government, this relationship has been taken as causal, with significant policy attention devoted to improving access to healthy groceries in low-income

bottom income quartile household reduces top minus bottom quartile Health Index difference by $0.005/0.56 \approx 0.9\%$.

neighborhoods.

We test this hypothesis using several complementary empirical strategies. Entry of a new supermarket has economically small effects on healthy grocery purchases, and we can conclude that differential local supermarket density explains no more than about 1.5 percent of the difference in healthy eating between high- and low-income households. The data clearly show why this is the case: Americans travel a long way for shopping, so even people who live in "food deserts" with no supermarkets get most of their groceries from supermarkets. Entry of a new supermarket nearby therefore mostly diverts purchases from other supermarkets. This analysis reframes the discussion of food deserts in two ways. First, the notion of a "food desert" is misleading if it is based on a market definition that understates consumers' willingness-to-travel. Second, any benefits of "combatting" food deserts derive less from healthy eating and more from reducing travel costs.

In a second event study analysis, we find that moving to an area where other people eat more or less healthfully does not affect households' own healthy eating patterns, at least over the several year time horizon that the data allow. In combination with the assumption that any endogeneity would generate upward bias in our estimated "place effects," we can conclude that the partial equilibrium place effects explain no more than three percent of differences in healthy eating between high- and low-income households. In our third reduced-form analysis, we find that even when we look at households living in the same zip codes or shopping at the same retailers, the nutrition-income relationship decreases by no more than one-third. Because of endogenous sorting into neighborhoods and stores, we think of this as an upper bound on the long-run effects of equalizing supply.

Our formal demand model estimates allow us to simulate precise counterfactuals in which low-income households are afforded the product availability and relative prices available to higher-income households. Consistent with the reduced form event study analyses, we find that equalizing supply would close the gap in healthy eating between low- and high-income households by less than 10 percent. By contrast, a means-tested subsidy for healthy groceries could increase low-income households' healthy eating to the level of high-income households at about 15 percent of the cost of the SNAP program. Although it would be crucial to carry out a full analysis of what market failures justify such subsidies, this result suggests that policymakers focused on reducing nutritional inequality might redirect efforts from supply-side policies toward means-tested subsidies.

References

Aizer, Anna and Janet Currie, "The Intergenerational Transmission of Inequality: Maternal Disadvantage and Health at Birth," *Science*, 2014, *344* (6186), 856–861.

Algert, Susan J., Aditya Agrawal, and Douglas S. Lewis, "Disparities in Access to Fresh Produce in Low-Income Neighborhoods in Los Angeles," *American Journal of Preventive Medicine*, 2006, *30* (5), 365–370.

Allcott, Hunt, Benjamin Lockwood, and Dmitry Taubinsky, "Regressive Sin Taxes, with an Application to the Optimal Soda Tax," *Working Paper, NYU*, 2018.

—, **Rebecca Diamond, and Jean-Pierre Dubé**, "The Geography of Poverty and Nutrition: Food Deserts and Food Choices Across the United States," *National Bureau of Economic Research Working Paper 24094*, 2017.

Alwitt, Linda F. and Thomas D. Donley, "Retail Stores in Poor Urban Neighborhoods," *The Journal of Consumer Affairs*, 1997, *31* (1), 139–164.

Anderson, Michael and David Matsa, "Are Restaurants Really Supersizing America?," *American Economic Journal: Applied Economics*, 2011, *3* (1), 152–188.

Atkin, David, "The Caloric Costs of Culture: Evidence from Indian Migrants," *American Economic Review*, 2016, *106* (4), 1144–1181.

—, **Benjamin Faber, and Marco Gonzalez-Navarro**, "Retail Globalization and Household Welfare: Evidence from Mexico," *Journal of Political Economy (forthcoming)*, 2016.

Aussenberg, Randy Alison, "SNAP and Related Nutrition Provisions of the 2014 Farm Bill (P.L. 113-79)," *Congressional Research Service Report*, 2014.

Baker, EA, M. Schootman, E. Barnidge, and C. Kelly, "The Role of Race and Poverty in Access to Foods that Enable Individuals to Adhere to Dietary Guidelines," *Preventing Chronic Disease*, 2006, *3* (3), A76.

Baum, Charles and Shin-Yi Chou, "Why Has the Prevalence of Obesity Doubled?," *Review of Economics of the Household*, 2016, *14* (2), 251–267.

Berry, Steven T., "Estimating Discrete-Choice Models of Product Differentiation," *Rand Journal of Economics*, 1994, *25* (2), 242–262.

Bertail, Patrice and France Caillavet, "Fruit and Vegetable Consumption Patterns: A Segmentation Approach," *American Journal of Agricultural Economics*, 2008, *90* (3), 827–842.

Beydoun, May A., Lisa M. Powell, and Youfa Wang, "The Association of Fast Food, Fruit and Vegetable Prices with Dietary Intakes Among US Adults: Is There a Modification by Family Income?," *Social Science & Medicine*, 2008, *66* (11), 2218–2229.

Bitler, Marianne and Steven Haider, "An Economic View of Food Deserts in the United States," *Journal of Policy Analysis and Management*, 2011, *30* (1), 153–176.

Bronnenberg, Bart J., Jean-Pierre Dubé, and Matthew Gentzkow, "The Evolution of Brand Preferences: Evidence from Consumer Migration," *American Economic Review*, 2012, *102* (6), 2472–2508.

Case, Anne and Angus Deaton, "Rising Morbidity and Mortality in Midlife Among White Non-Hispanic Americans in the 21st Century," *Proceedings of the National Academy of Sciences*, 2015, *112* (49), 15078–15083.

CDC, "State Initiatives Supporting Healthier Food Retail: An Overview of the National Landscape," 2011.

Chetty, Raj, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler, "The Association Between Income and Life Expectancy in the United States, 2001–2014," *Journal of the American Medical Association*, 2016, *315* (16), 1750–1766.

—, **Nathan Hendren, Patrick Kline, and Emmanuel Saez**, "Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States," *Quarterly Journal of Economics*, 2014, *129* (3), 1553–1623.

Chou, Shin-Yi, Michael Grossman, and Henry Saffer, "An Economic Analysis of Adult Obesity: Results from the Behavioral Risk Factor Surveillance System," *Journal of Health Economics*, 2004, *23* (3), 565–587.

Courtemanche, Charles J. and Art Carden, "Competing with Costco and Sam's Club: Warehouse Club Entry and Grocery Prices," *NBER Working Paper 17220*, 2011.

__ , Joshua C. Pinkston, Christopher J. Ruhm, and George L. Wehby, "Can Changing Economic Factors Explain the Rise in Obesity?," *Southern Economic Journal*, 2016, *82* (4), 1266–1310.

Cummins, Steven, Anne Findlay, Mark Petticrew, and Leigh Sparks, "Healthy Cities: The Impact of Food Retail-led Regeneration on Food Access, Choice, and Retail Structure," *Built Environment*, 2005, *31* (4), 288–301.

Currie, Janet, Stefano DellaVigna, Enrico Moretti, and Vikram Pathania, "The Effect of Fast Food Restaurants on Obesity and Weight Gain," *American Economic Journal: Economic Policy*, 2010, *2* (3), 32–63.

Curtis, Colleen, "First Lady Michelle Obama on Making a Difference in Cities with Food Deserts," *The White House*, 2011.

Cutler, David and Adriana Lleras-Muney, "Understanding Differences in Health Behaviors by Education," *Journal of Health Economics*, 2010, *29* (1), 1–28.

Cutler, David M., Edward L. Glaeser, and Jesse M. Shapiro, "Why Have Americans Become More Obese?," *Journal of Economic Perspectives*, 2003, *17* (3), 93–118.

Darmon, Nicole and Adam Drewnoski, "Does Social Class Predict Diet Quality?," *American Journal of Clinical Nutrition*, 2008, *87* (5), 1107–1117.

Davis, Brennan and Christopher Carpenter, "Proximity of Fast-Food Restaurants to Schools and Adolescent Obesity," *American Journal of Public Health*, 2009, *99* (3), 505–510.

DellaVigna, Stefano and Matthew Gentzkow, "Uniform Pricing in US Retail Chains," *Working Paper*, 2017.

Dubois, Pierre, Rachel Griffith, and Aviv Nevo, "Do Prices and Attributes Explain International Differences in Food Purchases?," *American Economic Review*, 2014, *2014* (3), 832–867.

Dunn, Richard A, "The Effect of Fast-Food Availability on Obesity: An Analysis by Gender, Race, and Residential Location," *American Journal of Agricultural Economics*, 2010, *92* (4), 1149–1164.

Eid, Jean, Henry G. Overman, Diego Puga, and Matthew A. Turner, "Fat City: Questioning the Relationship Between Urban Sprawl and Obesity," *Journal of Urban Economics*, 2008, *63* (2), 385–404.

Einav, Liran, Ephraim Leibtag, and Aviv Nevo, "Recording Discrepancies in Nielsen Homescan Data: Are They Present and Do They Matter?," *Quantitative Marketing and Economics*, 2010, *8* (2), 207–239.

Elbel, Brian, Alyssa Moran, L. Beth Dixon, Kamila Kiszko, Jonathan Cantor, Courtney Abrams, and Tod Mijanovich, "Assessment of a Government-Subsidized Supermarket in a High-Need Area on Household Food Availability and Children's Dietary Intakes," *Public Health Nutrition*, February 2015, *26*, 1–10.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams, "Sources of Geographic Variation in Health Care: Evidence From Patient Migration," *Quarterly Journal of Economics*, 2016, *131* (4), 1681–1726.

__ , __ , and __ , "Place-Based Drivers of Mortality: Evidence from Migration," *Working Paper*, 2018.

__ , __ , and __ , "What Drives Prescription Opioid Abuse? Evidence from Migration," *Working Paper*, 2018.

Furnee, CA, W Groot, and HM van den Brink, "The Health Effects of Education: A Meta-Analysis," *European Journal of Public Health*, 2008, *18* (4), 417–421.

Gelbach, Jonah B, "When Do Covariates Matter? And Which Ones, and How Much?," *Journal of Labor Economics*, 2016, *34* (2), 509–543.

Grossman, Gene, "The Relationship Between Health and Schooling: What's New?," *NBER Working Paper 21609*, 2015.

Guenther, P.M., J. Reedy, and S.M. Krebs-Smith, "Development of the Healthy Eating Index," *Journal of the American Dietetic Association*, 2008, *108* (11), 1896–1901.

Handbury, Jessie, Ilya Rahkovsky, and Molly Schnell, "What Drives Nutritional Disparities? Retail Access and Food Purchases across the Socioeconomic Spectrum," *National Bureau of Economic Research Working Paper 21126*, 2015.

Hastings, Justine, Ryan Kessler, and Jesse M. Shapiro, "The Effect of SNAP on the Composition of Purchased Foods: Evidence and Implications," *Working Paper, Brown University*, 2018.

Hausman, Jerry, "Valuation of New Goods Under Perfect and Imperfect Competition," in Timothy Bresnahan and Robert J. Gordon, eds., *The Economics of New Goods*, University of Chicago Press, 1996, pp. 207–248.

Hilmers, Angela, David C. Hilmers, and Jayna Dave, "Neighborhood Disparities in Access to Healthy Foods and Their Effects on Environmental Justice," *American Journal of Public Health*, 2012, *102* (9), 1644–1654.

Horowitz, Carol R., Kathryn A. Colson, Paul L. Hebert, and Kristie Lancaster, "Barriers to Buying Healthy Foods for People With Diabetes: Evidence of Environmental Disparities," *American Journal of Public Health*, 2004, *94* (9), 1549–1554.

Hut, Stefan, "Determinants of Dietary Choice in the US: Evidence from Consumer Migration," *Working Paper, Brown University*, 2018.

Hwang, Minha and Sungho Park, "The Impact of Walmart Supercenter Conversion on Consumer Shopping Behavior," *Management Science*, 2016, *62* (3), 817–828.

Jetter, Karen M. and Diana L. Cassady, "The Availability and Cost of Healthier Food Alternatives," *American Journal of Preventive Medicine*, 2005, *30* (1), 38–44.

Jones, Eugene, "An Analysis of Consumer Food Shopping Behavior Using Supermarket Scanner Data: Differences by Income and Location," *American Journal of Agricultural Economics*, 1997, pp. 1437–1443.

Kliemann, N, J Wardle, F Johnson, and H Croker, "Reliability and Validity of a Revised Version of the General Nutrition Knowledge Questionnaire," *European Journal of Clinical Nutrition*, 2016, *70* (10), 1174–1180.

Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz, "Experimental Analysis of Neighborhood Effects," *Econometrica*, 2007, *75* (1), 83–119.

Lakdawalla, Darius, Tomas Philipson, and Jay Bhattacharya, "Welfare-Enhancing Technological Change and the Growth of Obesity," *American Economic Review, Papers and Proceedings*, 2005, *95* (2), 253–257.

Larson, Nicole, Mary Story, and Melissa Nelson, "Neighborhood Environments: Disparities in Access to Healthy Foods in the U.S.," *American Journal of Preventive Medicine*, 2009, *36* (1), 74–81.

Lipman, Melissa, "FTC May Waste Time Updating Price-Bias Guide, Attys Say," 2012.

Ludwig, Jens, Lisa Sanbonmatsu, Lisa Gennetian, Emma Adam, Greg J. Duncan, Lawrence F. Katz, Ronald C. Kessler, Jeffrey R. Kling, Stacy Tessler Lindau, Robert C. Whitaker, , and Thomas W. McDade, "Neighborhoods, Obesity, and Diabetes—A Randomized Social Experiment," *New England Journal of Medicine*, 2011, *365*, 1509–1519.

Molitor, David, "The Evolution of Physician Practice Styles: Evidence from Cardiologist Migration," *American Economic Journal: Economic Policy*, 2018, *10* (1), 326–356.

Okrent, Abigail and Aylin Kumcu, "U.S. Households' Demand for Convenience Foods," *USDA Economic Research Report*, 2016, (211).

Oster, Emily, "Unobservable Selection and Coefficient Stability: Theory and Evidence," *Journal of Business and Economic Statistics (forthcoming)*, 2016.

Park, John L., Rodney B. Holcomb, Kellie Curry Raper, and Oral Capps, "A Demand Systems Analysis of Food Commodities by U.S. Households Segmented by Income," *American Journal of Agricultural Economics*, 1996, *78* (2), 290–300.

Powell, Lisa M., Sandy Slater, Donka Mirtcheva, Yanjun Bao, and Frank J. Chaloupka, "Food Store Availability and Neighborhood Characteristics in the United States," *Preventive Medicine*, 2007, *44*, 189–195.

Rashad, Inas, "Structural Estimation of Caloric Intake, Exercise, Smoking, and Obesity," *Quarterly Review of Economics and Finance*, 2006, *46* (2), 268–283.

_ , Michael Grossman, and Shin-Yi Chou, "The Super Size of America: An Economic Estimation of Body Mass Index and Obesity in Adults," *Eastern Economic Journal*, 2006, *32* (1), 133–148.

Rehm, Colin D., Jose L. Penalvo, Ashkan Afshin, and Dariush Mozaffarian, "Dietary Intake Among U.S. Adults, 1999–2012," *Journal of the American Medical Association*, 2016, *315* (23), 2542–2553.

CBPP (Center for Budget and Policy Priorities), "A Quick Guide to SNAP Eligibility and Benefits," 2018.

HHS (U.S. Department of Health and Human Services), "FoodKeeper App," https://www.foodsafety.gov/keep/foodkeeperapp/index.html 2015.

TRF (The Reinvestment Fund), "The Healthy Food Financing Initiative," http://www.healthyfoodaccess.org/resources/library/healthy-food-financing-initiative-hffi 2017.

USDA (U.S. Department of Agriculture, Agricultural Research Service), "Away from Home: Percentages of Selected Nutrients Contributed by Food and Beverages Consumed Away from Home, by Family Income (in Dollars) and Age." What We Eat in America, NHANES 2011-2012," http://www.ars.usda.gov/Services/docs.htm?docid=18349 2014b.

_ , "USDA Food and Nutrient Database for Dietary Studies 2015-2016," https://www.ars.usda.gov/northeast-area/beltsville-md-bhnrc/beltsville-human-nutrition-research-center/food-surveys-research-group/docs/fndds-download-databases/ 2016.

_ , "USDA National Nutrient Database for Standard Reference, Release 28," http://www.ars.usda.gov/ba/bhnrc/ndl 2018.

Richards, Michael R. and Jody L. Sindelar, "Rewarding Healthy Food Choices in SNAP: Behavioral Economic Applications," *Milbank Quarterly*, 2013, *91* (2), 395–412.

Saez, Emmanuel and Thomas Piketty, "Income Inequality in the United States, 1913–1998," *Quarterly Journal of Economics*, 2003, *118* (1), 1–39.

Schanzenbach, Diane Whitmore, "Pros and Cons of Restricting SNAP Purchases," 2017.

Sharkey, Joseph R., Scott Horel, and Wesley R. Dean, "Neighborhood Deprivation, Vehicle Ownership, and Potential Spatial Access to a Variety of Fruits and Vegetables in a Large Rural Area in Texas," *International Journal of Health Geographics*, 2010, *9* (26), 1–27.

Shenkin, Jonathan D. and Michael F. Jacobson, "Using the Food Stamp Program and Other Methods to Promote Healthy Diets for Low-Income Consumers," *American Journal of Public Health*, 2010, *100* (9), 1562–1564.

Song, Hee-Jung, Joel Gittelsohn, Miyong Kim, Sonali Suratkar, Sangita Sharma, and Jean Anliker, "A Corner Store Intervention in a Low-Income Urban Community is Associated with Increased Availability and Sales of Some Healthy Foods," *Public Health Nutrition*, 2009, *12* (11), 2060–2067.

Volpe, Richard, Abigail Okrent, and Ephraim Leibtag, "The Effect of Supercenter-Format Stores on the Healthfulness of Consumers' Grocery Purchases," *American Journal of Agricultural Economics*, 2013, *95* (3), 568–589.

_ **and** _ , "Assessing the Healthfulness of Consumers' Grocery Purchases," *USDA Economic Information Bulletin*, 2012, (EIB-102).

Wang, Dong D., Cindy W. Leung, Yanping Li, Eric L. Ding, Stephanie E. Chiuve, Frank B. Hu, and Walter C. Willett, "Trends in Dietary Quality Among Adults in the United States, 1999 Through 2010," *JAMA Internal Medicine*, 2014, *174* (10), 1587–1595.

Weatherspoon, Dave, James Oehmke, Assa Dembele, Marcus Coleman, Thasanee Satimanon, and Lorraine Weatherspoon, "Price and Expenditure Elasticities for Fresh Fruits in an Urban Food Desert," *Urban Studies*, 2012, *50* (1), 88–106.

Wrigley, Neil, Daniel Warm, and Barrie Margetts, "Deprivation, Diet, and Food Retail Access: Findings From the Leeds 'Food Deserts' Study," *Environment and Planning A*, 2003, *35*, 151–188.

Yang, Quanhe, Zefeng Zhang, Edward Gregg, Dana Flanders, Robert Merritt, and Frank Hu, "Added Sugar Intake and Cardiovascular Diseases Mortality Among US Adults," *Journal of the American Medical Association Internal Medicine*, 2014, *174* (4), 516–524.

Tables

Table 1: **Descriptive Statistics**

Variable	Mean	Standard deviation
Panel (a): Nielsen Homescan Households		
Household income ($000s)	61.0	43.3
Years education	13.9	2.06
Age	52.3	14.4
Household size	2.38	1.33
White	0.77	0.42
Black	0.12	0.32
Married	0.50	0.50
Employed	0.61	0.44
Weekly work hours	22.9	16.7
Household daily calorie need	5192	2959
Health importance	0	1
Nutrition knowledge	0	1
Body Mass Index (kg/m^2)	29.5	7.0
Diabetic	0.19	0.38
Panel (b): Zip Code Establishment Counts		
Grocery	1.67	4.07
Large grocery (>50 employees)	0.46	1.05
Supercenters/club stores	0.11	0.40
Drug stores	1.10	2.31
Convenience stores	3.14	5.25
Meat/fish/produce stores	0.27	0.96

Notes: Homescan data include 731,994 household-by-year observations for 2004–2016 and are weighted for national representativeness. The U.S. government Dietary Guidelines include calorie needs by age and gender; we combine that with Homescan household composition to get each household's daily calorie need. Household size is the number of household heads plus the total calorie needs of all other household members divided by the nationwide average calorie need of household heads. Health importance, nutrition knowledge, Body Mass Index, and the diabetic indicator are from Homescan add-on surveys carried out by Nielsen for Allcott, Lockwood and Taubinsky (2018); the former two variables are normalized to have a mean of zero and a standard deviation of one. Health importance is the response to the question, "In general, how important is it to you to stay healthy, for example by maintaining a healthy weight, avoiding diabetes and heart disease, etc.?" Nutrition knowledge is from a battery of 28 questions drawn from the General Nutrition Knowledge Questionnaire (Kliemann et al., 2016). If two household members responded to the PanelViews survey, we take the mean of each survey variable across the two respondents. Zip code establishment counts are from Zip Code Business Patterns data for 2004–2016, with 508,951 zip code-by-year observations.

Table 2: **Effects of Supermarket Entry**

(a) **Effects on Expenditure Shares**

	Full sample		Bottom quartile		Food deserts	
	(1)	(2)	(3)	(4)	(5)	(6)
	Entrants	Grocery/ super/club	Entrants	Grocery/ super/club	Entrants	Grocery/ super/club
Post entry: 0-10 minutes	1.496***	0.037	1.966***	-0.034	1.914***	-0.269*
	(0.098)	(0.051)	(0.243)	(0.145)	(0.303)	(0.159)
Post entry: 10-15 minutes	0.543***	-0.057	0.433***	-0.029	0.762***	-0.038
	(0.059)	(0.035)	(0.144)	(0.094)	(0.166)	(0.119)
Observations	2,874,514	2,874,365	538,041	537,998	646,223	646,181
Dependent var. mean	9.9	88.2	7.5	86.2	6.1	87.7

(b) **Effects on Health Index**

	Full sample	Bottom quartile	Food deserts
	(1)	(2)	(3)
Post entry: 0-10 minutes	0.004	0.005	0.007
	(0.003)	(0.007)	(0.008)
Post entry: 10-15 minutes	0.006***	0.001	0.014***
	(0.002)	(0.005)	(0.005)
Observations	2,874,514	538,041	646,223

Notes: This table uses 2004–2016 Nielsen Homescan data at the household-by-quarter level. The "food desert" subsample comprises observations with no grocery stores with 50 or more employees, supercenters, or club stores in the zip code in the first year the household is observed there. Expenditure shares are in units of percentage points. Health Index is our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation one across households. Reported independent variables are the count of supermarkets that have entered within a 0–10 or 10–15 minute drive from the household's census tract centroid. All regressions control for household demographics (natural log of income, natural log of years of education, age indicators, household size, race indicators, a married indicator, employment status, and weekly work hours), census division-by-quarter of sample indicators, and household-by-census tract fixed effects. Observations are not weighted for national representativeness. Robust standard errors, clustered by household and census tract, are in parentheses. *, **, ***: statistically significant with 10, 5, and 1 percent confidence, respectively.

Table 3: **Association of Health Index with Local Area Health Index Using Movers**

	(1)	(2)	(3)	(4)
Zip code average Health Index	0.0511**	0.0487**		
	(0.0247)	(0.0245)		
County average Health Index			0.1067*	0.1100**
			(0.0565)	(0.0560)
Household demographics	No	Yes	No	Yes
Observations	564,944	564,944	570,279	570,279
95% confidence interval upper bound	0.100	0.097	0.217	0.220

Notes: This table uses 2004–2016 Nielsen Homescan data at the household-by-year level. The sample excludes observations where less than 50 percent of trips to RMS stores are not in the household's end-of-year county of residence. The Health Index is our overall measure of the healthfulness of grocery purchases and is normalized to have a mean of zero and a standard deviation of one across households. Household demographics are natural log of income, natural log of years of education, age indicators, household size, race indicators, a married indicator, employment status, and weekly work hours. All regressions also control for year indicators and household fixed effects. Observations are not weighted for national representativeness. Robust standard errors, clustered by household and local area (zip code or county), are in parentheses. *, **, ***: statistically significant with 10, 5, and 1 percent confidence, respectively.

Table 4: **Preferences for Nutrients by Household Income**

Income quartile	(1) Sodium	(2) Whole fruit	(3) Other fruit	(4) Whole grains	(5) Refined grains	(6) Greens, beans	(7) Other veg
Income Q1	-0.242***	-0.252***	0.177***	0.252***	0.022***	0.965***	-0.438***
	(0.009)	(0.017)	(0.004)	(0.004)	(0.001)	(0.009)	(0.014)
Income Q2	-0.351***	-0.092***	0.215***	0.312***	0.036***	1.02***	-0.307***
	(0.011)	(0.012)	(0.002)	(0.006)	(0.001)	(0.009)	(0.010)
Income Q3	-0.432***	-0.057***	0.261***	0.363***	0.045***	1.16***	-0.319***
	(0.012)	(0.011)	(0.001)	(0.006)	(0.001)	(0.011)	(0.010)
Income Q4	-0.634***	-0.006	0.331***	0.477***	0.073***	1.49***	-0.337***
	(0.022)	(0.014)	(0.001)	(0.012)	(0.003)	(0.021)	(0.014)

	(8) Dairy	(9) Sea, plant protein	(10) Meat protein	(11) Added sugar	(12) Solid fats	(13) Unobserved characteristic	(14) WTP for Health Index
Income Q1	0.144***	-0.275***	0.061***	0.0002***	0.0005***	-0.00076***	0.429***
	(0.001)	(0.008)	(0.001)	(0.0001)	(0.00001)	(0.00003)	(0.011)
Income Q2	0.163***	-0.313***	0.055***	-0.0007***	0.0009***	-0.0010***	0.631***
	(0.001)	(0.009)	(0.001)	(0.0001)	(0.00002)	(0.00004)	(0.006)
Income Q3	0.189***	-0.359***	0.057***	-0.0018***	0.001***	-0.0013***	0.820***
	(0.001)	(0.009)	(0.001)	(0.0001)	(0.00002)	(0.00004)	(0.003)
Income Q4	0.245***	-0.450***	0.059***	-0.004***	0.001***	-0.021***	1.141***
	(0.002)	(0.014)	(0.001)	(0.0002)	(0.00003)	(0.00008)	(0.003)

Notes: This table presents GMM estimates of the nutrient preference parameters $\tilde{\beta}_c$ from Equation (12). Magnitudes represent willingness to pay for a unit of the nutrient, where the units are those used in the Healthy Eating Index. Sodium is in grams; whole fruit, other fruit and dairy are in cups; whole grains, refined grains, and both types of protein are in ounces, added sugar is in teaspoons; solid fats are in calories. "WTP for Health Index" in column 14 equals $\sum_c \hat{\tilde{\beta}}_c s_c r_c$, where s_c is the maximum possible score on the Healthy Eating Index for dietary component c, and r_c is the difference in consumption of component c to receive the maximum instead of the minimum score. Standard errors, clustered by household, are in parentheses. *, **, ***: statistically significant with 10, 5, and 1 percent confidence, respectively.

Table 5: **Decomposition of Nutrition-Income Relationship by Household Demographics**

	(1)	(2)	(3)
	Full model	Unconditional relationship	Auxiliary regressions
ln(Income)	0.134	0.267	
	(0.0209)***	(0.0178)***	
ln(Years education)	0.685		1.922
	(0.0939)***		(0.0688)***
1(White)	-0.247		0.0588
	(0.0474)***		(0.0283)**
1(Black)	-0.291		-0.155
	(0.0574)***		(0.0342)***
1(Married)	0.0435		0.431
	(0.0256)*		(0.0209)***
Employed	0.0701		0.635
	(0.0802)		(0.0234)***
Weekly work hours	-0.0000934		0.0193
	(0.00225)		(0.000611)***
Health importance	0.102		0.0691
	(0.0119)***		(0.0109)***
Nutrition knowledge	0.132		0.155
	(0.0128)***		(0.0118)***
Household size	-0.103	-0.119	
	(0.0116)***	(0.0111)***	
Census division indicators	Yes	No	No
Age indicators	Yes	Yes	Yes
Year indicators	Yes	Yes	Yes
Observations	81,839	81,839	81,839

Notes: These regressions use 2004–2016 Nielsen Homescan data at the household-by-year level, using only the subsample that responded to the Homescan add-on survey carried out by Nielsen for Allcott, Lockwood and Taubinsky (2018). Health importance is the response to the question, "In general, how important is it to you to stay healthy, for example by maintaining a healthy weight, avoiding diabetes and heart disease, etc.?" Nutrition knowledge is from a battery of 28 questions drawn from the General Nutrition Knowledge Questionnaire (Kliemann et al., 2016). Health importance and nutrition knowledge are both normalized to have a mean of zero and a standard deviation of one. Columns 1 and 2 present estimates of Equation (28), a regression of the Health Index of demand-only consumption predictions on covariates. Each row of column 3 presents the coefficient from a regression of natural log of household income on the variable listed in each row, controlling for age and year indicators and household size. Observations are weighted using the Homescan sample weights. Robust standard errors, clustered by household, are in parentheses. *, **, ***: statistically significant with 10, 5, and 1 percent confidence, respectively.

Table 6: **Impacts of Means-Tested Healthy Grocery Subsidies**

	(1)	(2)	(3)
	Subsidy to close gap by 0.9% (estimated supermarket entry effect)	Subsidy to close gap by 9% (structural estimate of equal supply conditions)	Subsidy to close gap by 100%
Subsidy parameter s	0.000067	0.000657	0.00601
Mean subsidy for subsidized products	0.60%	5.8%	48.7%
Average subsidy payment per household	\$2.62	\$26.35	\$336.10
Total subsidy, all bottom-quartile households	\$84 million	\$830 million	\$10.57 billion

Notes: This table presents the amounts of healthy grocery subsidies required to reduce the difference in the Health Index between top and bottom income quartile households by given percentages. The subsidies are available to bottom income quartile households in proportion to a product's Health Index, conditional on the Health Index being above the median product's Health Index. "Mean subsidy" is the percent discount among products receiving strictly positive subsidies. "Total subsidy, all bottom-quartile households" is the average subsidy payment per household multiplied by 31.45 million, the number of households in the bottom income quartile.

Figures

Figure 1: **Healthfulness of Grocery Purchases by Household Income**

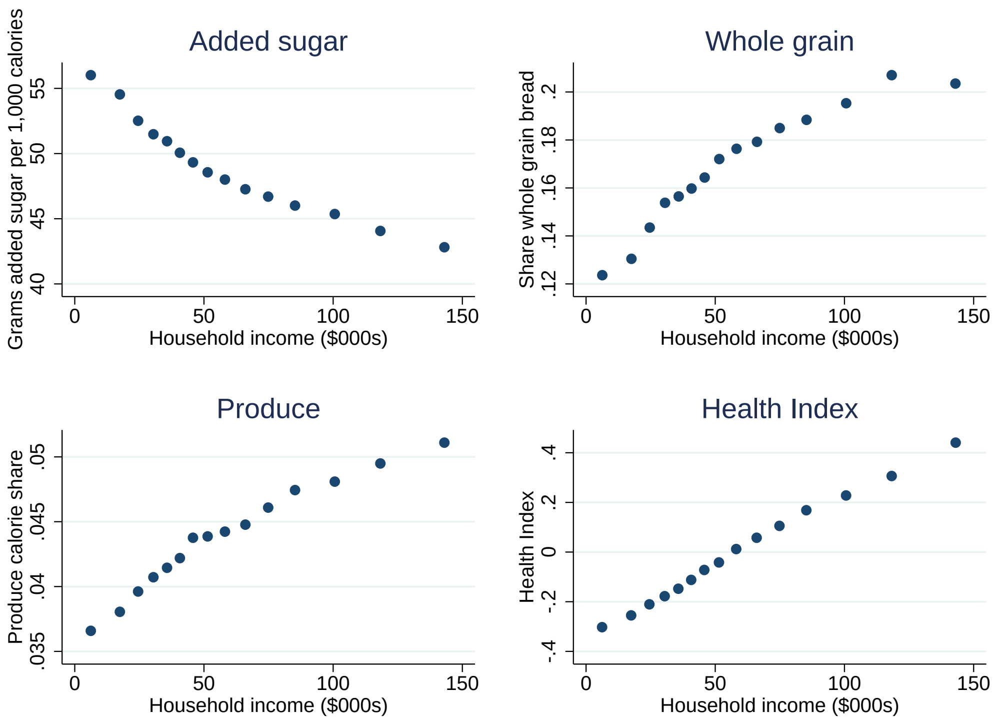

Notes: This figure presents Nielsen Homescan data for 2004–2016. Each panel presents a binned scatterplot of a grocery healthfulness measure against average household income across all years the household is observed in Homescan, residual of age and year indicators and household size. Added sugar is the grams of added sugar per 1,000 calories purchased; whole grain is the calorie-weighted average share of bread, buns, and rolls purchases that are whole grain; produce is the share of calories from fresh, canned, dried, and frozen fruits and vegetables; and the Health Index is our overall measure of the healthfulness of grocery purchases and is normalized to have a mean of zero and a standard deviation of one across households. Observations are weighted for national representativeness.

Figure 2: **Store Average Healthfulness and Size by Zip Code Median Income**

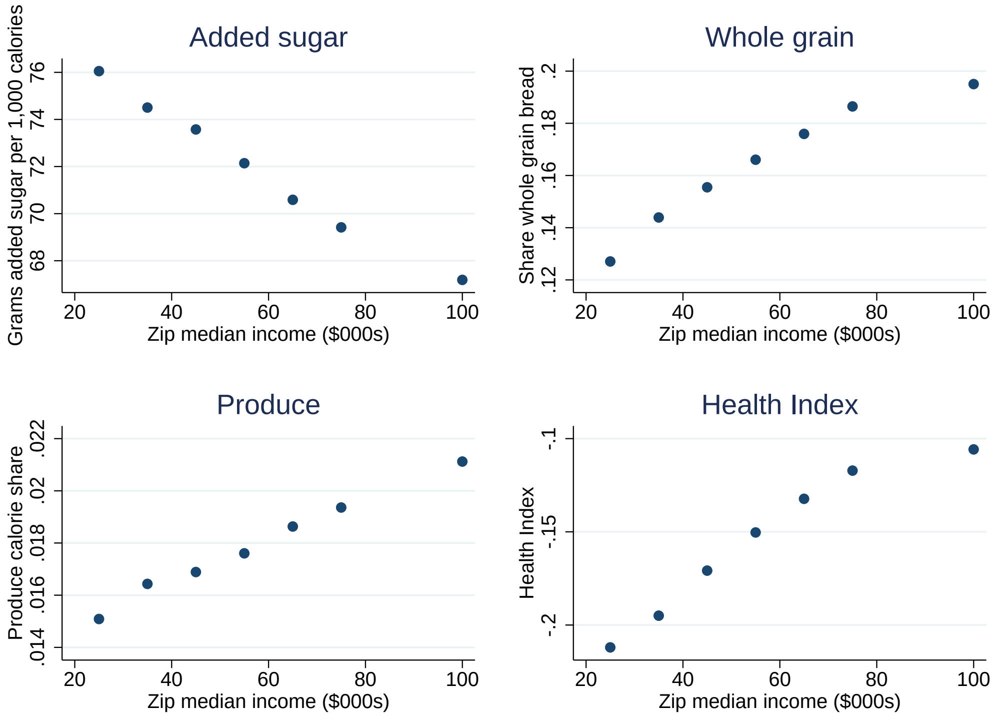

Notes: This figure uses Nielsen RMS data for 2006–2016. To measure store healthfulness, we construct the following across all UPCs offered in each store: average grams of added sugar per 1,000 calories; the calorie-weighted share of bread, buns, and rolls UPCs that are whole grain; the calorie-weighted share of UPCs that are produce; and the calorie-weighted mean Health Index. This figure presents the means of these variables across stores within categories of zip code median income.

Figure 3: **Store Counts by Zip Code Median Income**

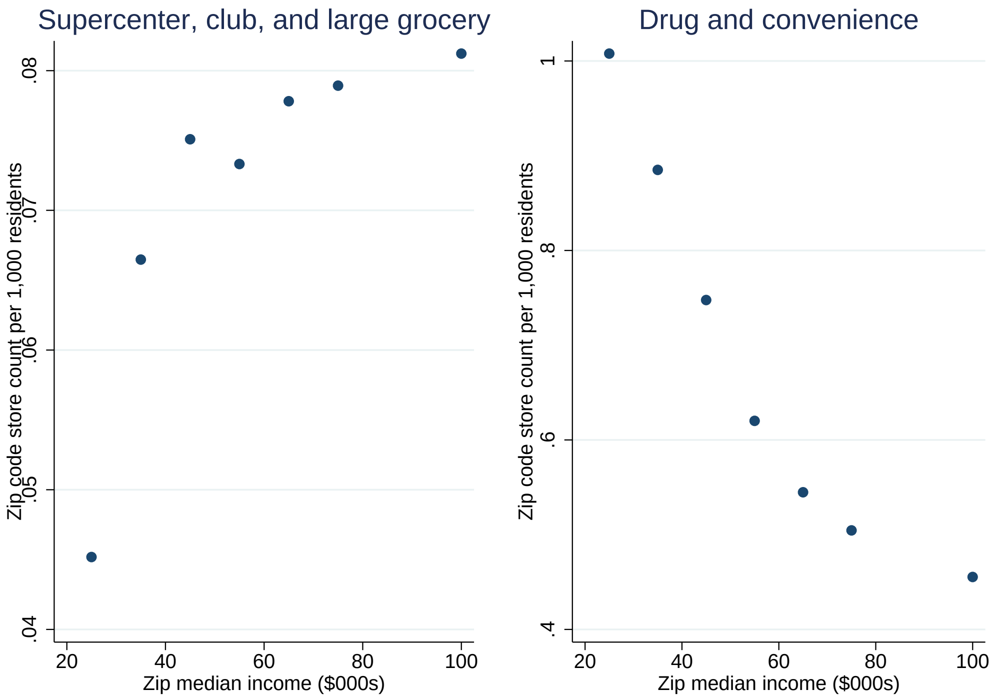

Notes: This figure presents counts of stores per 1,000 residents for the average zip code in each income category, using 2004–2016 Zip Code Business Patterns data. Large (small) grocers are defined as those with 50 or more (fewer than 50) employees.

Figure 4: **Event Study of Supermarket Entry**

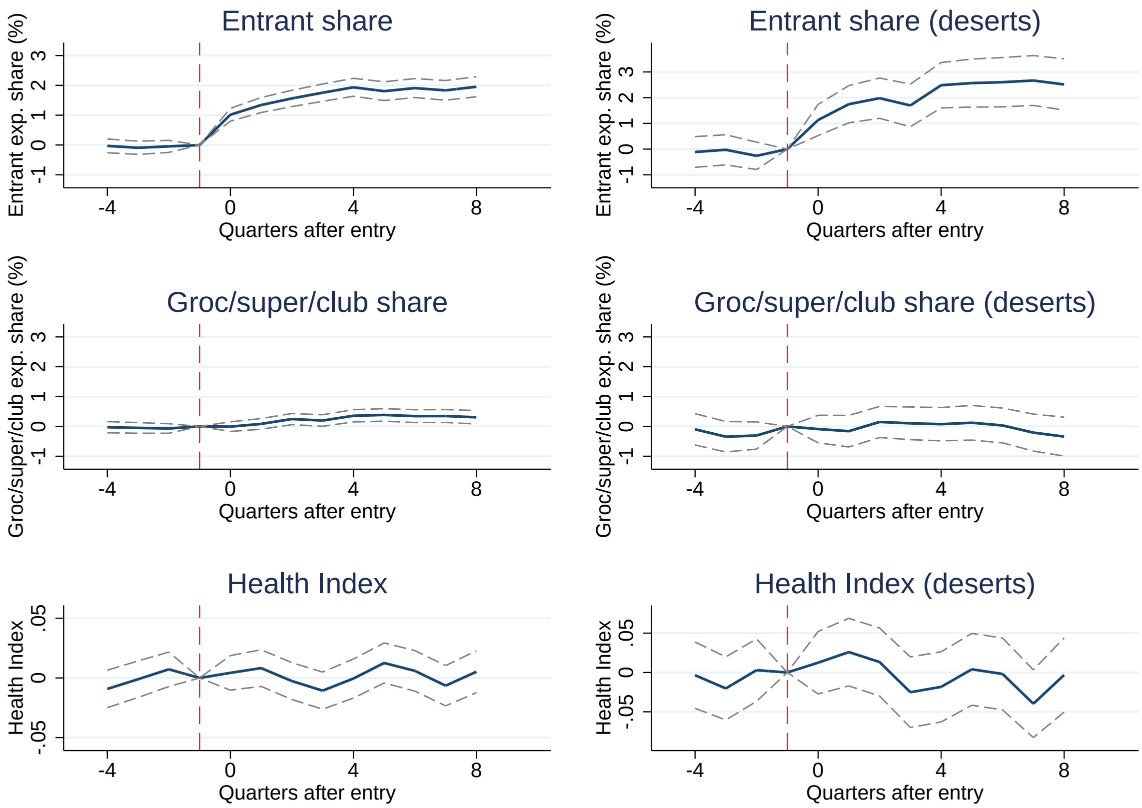

Notes: This figure presents the $\tau_{[0,10)q}$ parameters and 95 percent confidence intervals from estimates of Equation (2)—the effects of supermarket entry—using 2004–2016 household-by-quarter Homescan data. All regressions control for household demographics (natural log of income, natural log of years of education, age indicators, household size, race indicators, a married indicator, employment status, and weekly work hours), census division-by-quarter of sample indicators, and household-by-census tract fixed effects. The top two panels present effects on expenditure shares (in units of percentage points) across all retailers with stores that have entered within a 15-minute drive of the household. The middle two panels present effects on the combined expenditure share of grocery stores, supercenters, and club stores. We keep the y-axis on the same scale between the top and middle panels so that the magnitudes can be compared easily. The bottom two panels present effects on the Health Index, our overall measure of the healthfulness of grocery purchases which is normalized to have a mean of zero and a standard deviation of one across households. The left panels include the full sample, while the right panels include only the "food desert" subsample: observations with no grocery stores with 50 or more employees, supercenters, or club stores in the zip code in the first year the household is observed there. Observations are not weighted for national representativeness.

Figure 5: **Average Health Index of Store Purchases by County**

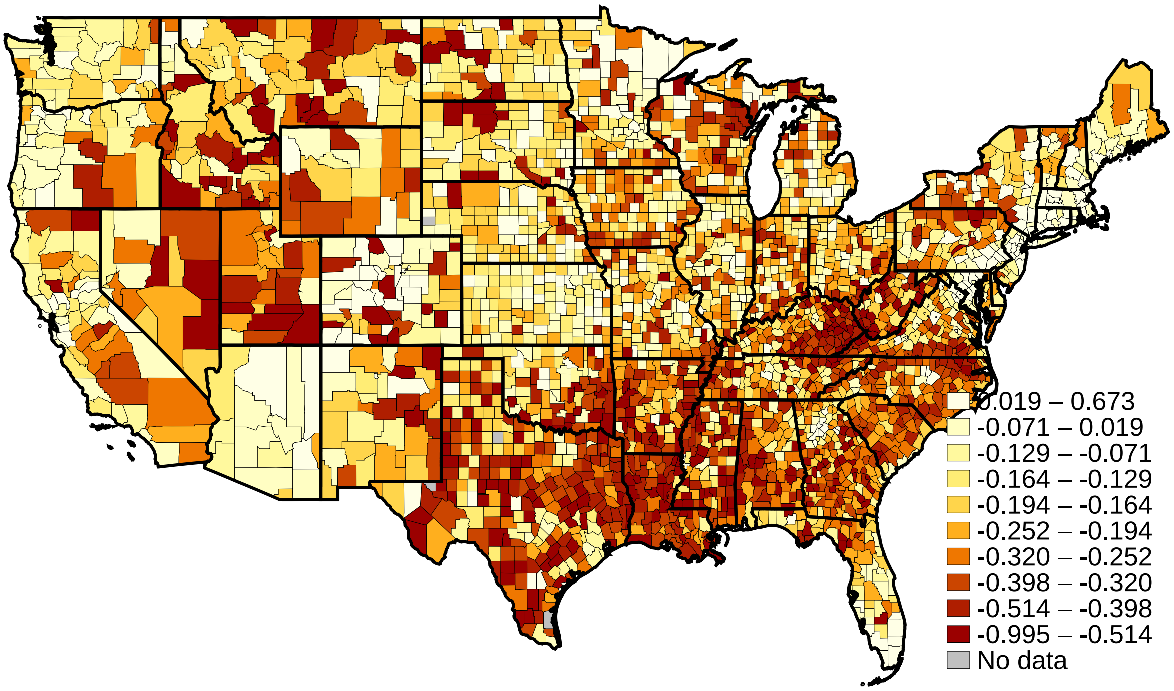

Notes: This figure presents the calorie-weighted average normalized Health Index of packaged grocery purchases by county, using 2006–2016 Nielsen RMS data. The Health Index is our overall measure of the healthfulness of grocery purchases and is normalized to have a mean of zero and a standard deviation of one across households. Note that purchases in RMS are less healthful than in Homescan, so the average normalized Health Index on this map is less than zero.

Figure 6: **Event Study of Moves Across Counties**

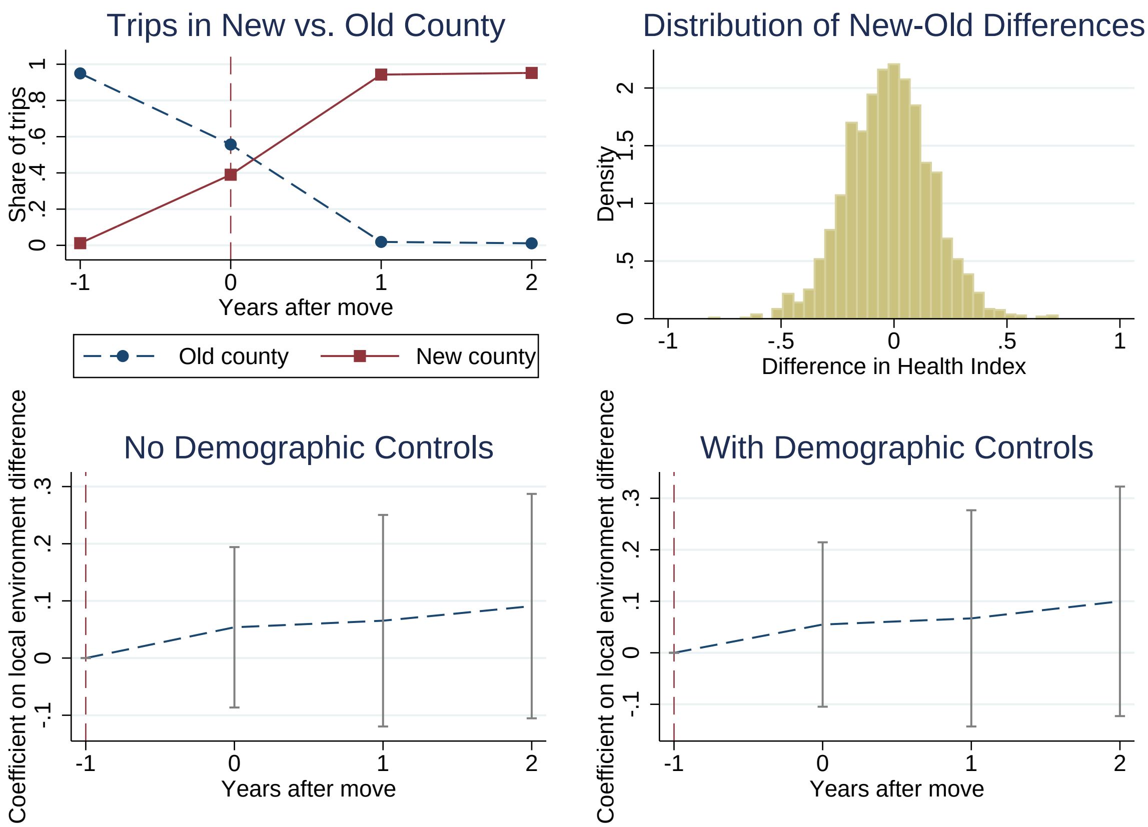

Notes: Using 2004–2016 Homescan data, these figures present results for the event study of moves across counties. The top left panel presents the share of shopping trips that are in the new versus old county. The top right panel presents the distribution across balanced panel households of the difference in the Health Index between the new and old county. The bottom panels present the τ_y parameters and 95 percent confidence intervals from estimates of Equation (4): associations between the household-level Health Index and the difference in the average local Health Index between post-move and pre-move locations. The bottom right panel includes controls for household demographics (natural log of income, natural log of years of education, age indicators, household size, race indicators, a married indicator, employment status, and weekly work hours). The Health Index is our overall measure of the healthfulness of grocery purchases and is normalized to have a mean of zero and a standard deviation of one across households. Observations are not weighted for national representativeness.

Figure 7: **Healthfulness of Grocery Purchases by Household Income Within Markets**

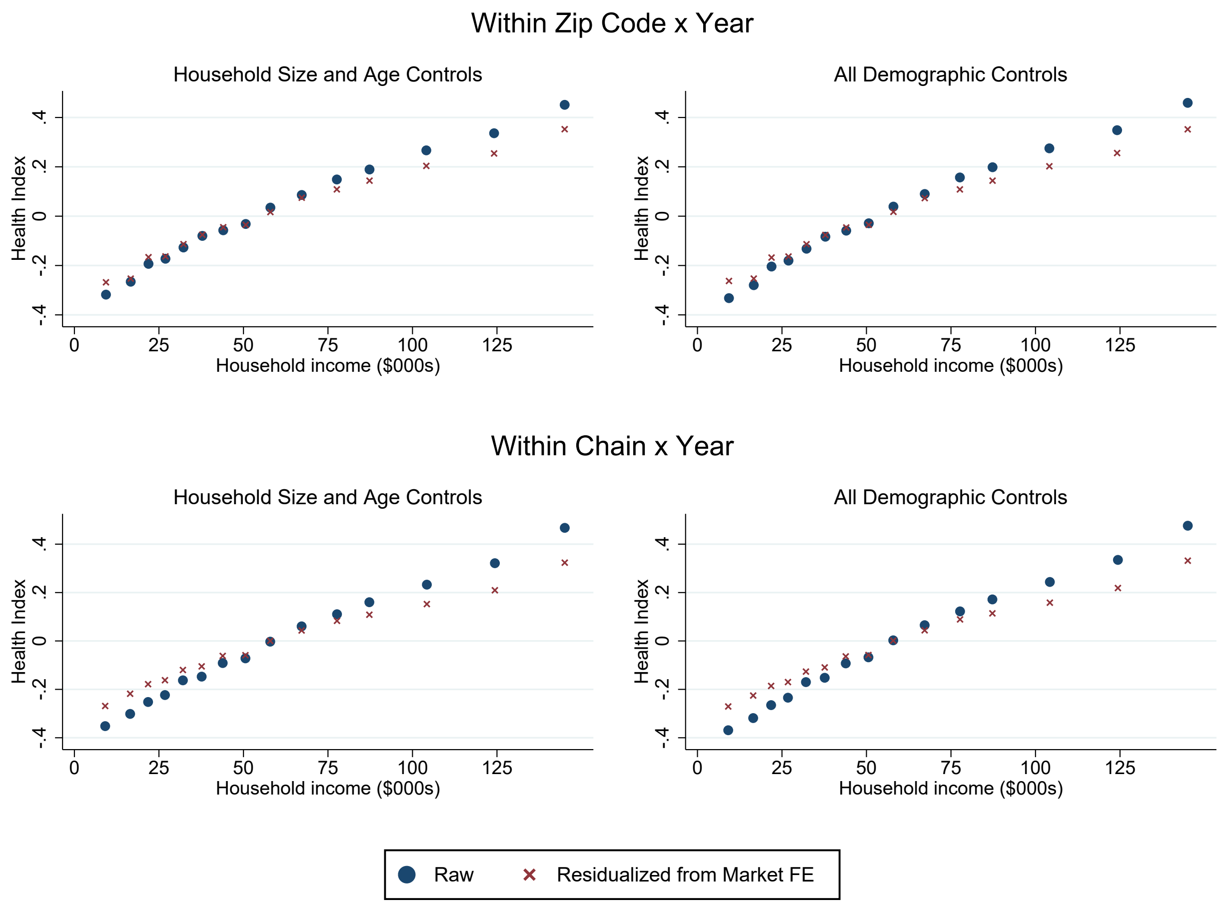

Notes: This figure presents Nielsen Homescan data for 2004–2016. Each panel presents a binned scatterplot of a grocery healthfulness measure against average household income across all years the household is observed in Homescan. The data are aggregated to the household-year level for the top two panels and to the household-retail chain-year level for the bottom two panels, where the sample is restricted to observations that reflect more than half of a household's reported calories purchased in a given year (such that there is only one observation per household-year reflecting the household's purchases in their primary store for that year). Each point depicts the mean Health Index across observations in each income bin controlling for a set of demographic and/or location controls. The circles depict the mean estimated in a regression of the household-year-level Health Index against income-bin fixed effects as well as household size, age, and panel year controls in the left-hand panel and with a full set of demographic and panel year controls in the right-hand panels. The crosses depict the mean estimated in a regression of the household-year-level health index against income-bin fixed effects, zip code-year fixed effects for the top panels and retail chain-year fixed effects for the bottom panels. The sample is limited to those sets of households that identify the within-location means (that is, households that reside in zip codes with observations from panelists in more than one income bin or households shopping in retail chains with observations from panelists in more than one income bin).

Figure 8: **Correlation of Price Instrument with Product Health Index: Differential between High- and Low-Income Counties**

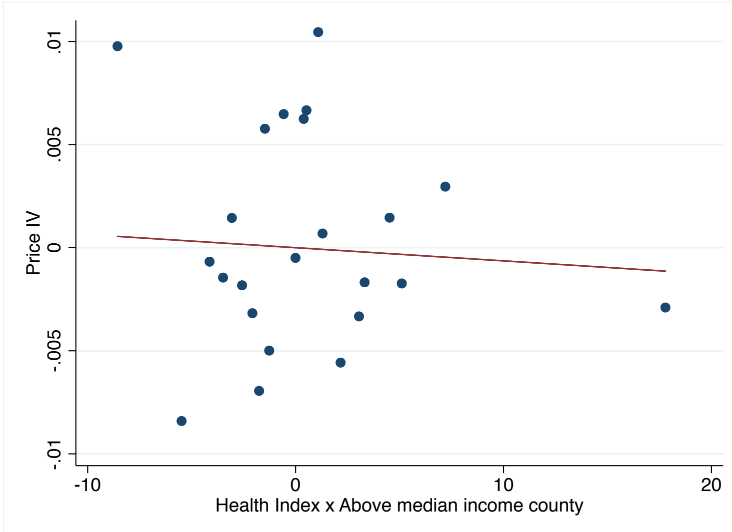

Notes: This figure presents a binned scatterplot of Equation (17) showing the relationship between our price instrument and the interaction between product group Health Index and an indicator for whether the county average income in 2010 was above the national median, residual of product group and county fixed effects. Data are at the product group-by-county-by-year level. The estimated slope of the regression plotted is -0.0000639, with a standard error of 0.0000553.

Figure 9: **Predicted Health Index for Each Income Group**

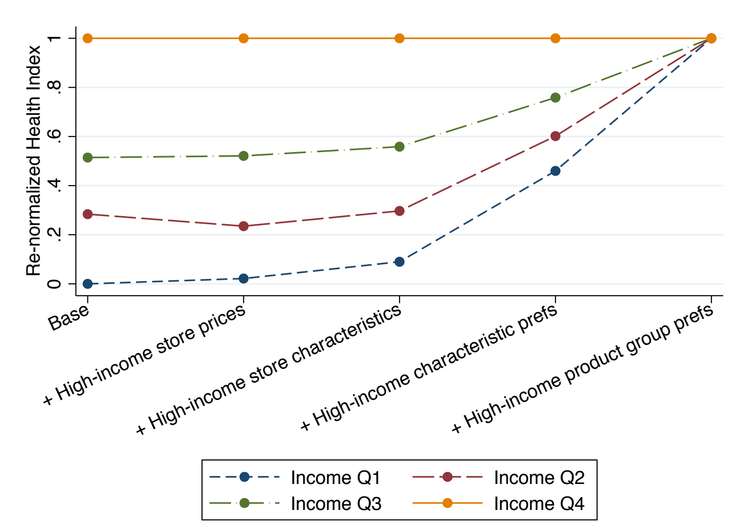

Notes: Each category on the x-axis represents a separate counterfactual calculation. Income groups are quartiles of income, residual of household size and age and year indicators. The base category measures the Health Index for each income group when each group retains their own preferences and faces their own local supply conditions. The second category sets all prices to those observed for the high-income group. The third category sets all prices and product nutrient characteristics to those in the high-income group. The fourth and fifth categories, respectively, set nutrient preferences and product group preferences equal to those for the high-income group. The Health Index presented on the y-axis is re-normalized so that the base difference between the highest- and lowest-income groups equals one.

Figure 10: **Explaining the Relationship Between Income and Healthy Grocery Demand**

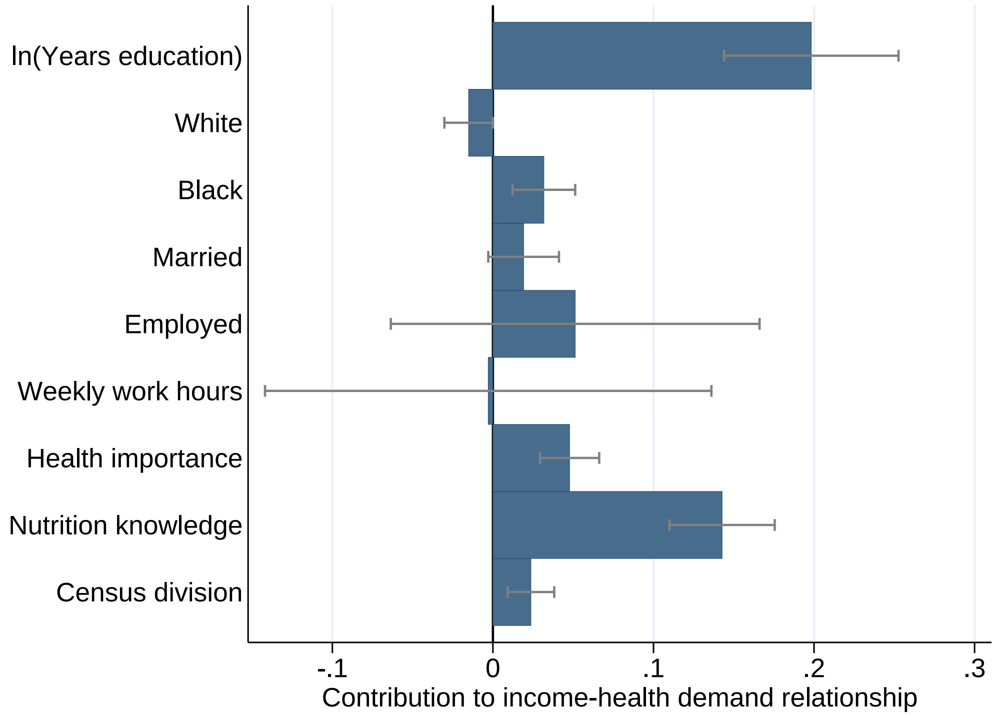

Notes: This figure presents the $\tilde{\pi}_v$ parameters and 95 percent confidence intervals from Equation (29), representing the share of the correlation between income and demand for healthy groceries that is explained by each variable.

Online Appendix: Not for Publication

Food Deserts and the Causes of Nutritional Inequality

Hunt Allcott, Rebecca Diamond, Jean-Pierre Dubé, Jessie Handbury, Ilya Rahkovsky, and Molly Schnell

A Appendix to Data Section

A.A Magnet Calorie Shares

Figure A1: **Magnet Data: Share of Produce from Packaged Items**

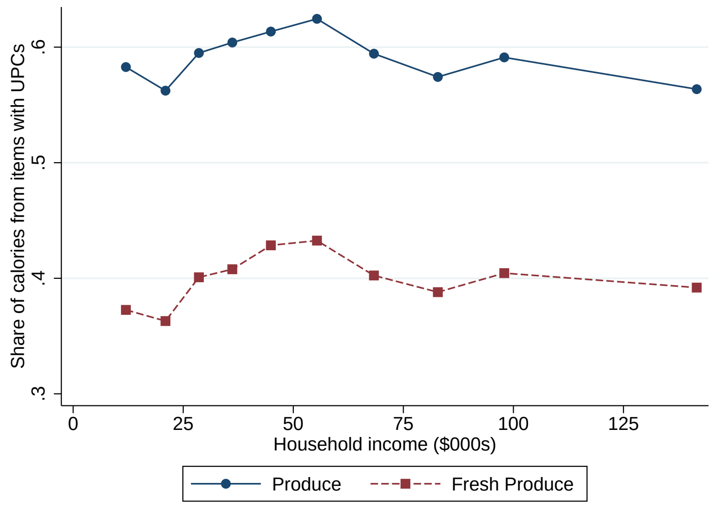

Notes: This figure uses the Nielsen Homescan "magnet" subsample for 2004–2006 to show the share of produce and fresh produce calories coming from packaged items with UPCs, which are the items that we observe outside the magnet subsample. The x-axis presents bins of average household income across all years the household is observed in Homescan. "Produce" includes fresh, dried, canned, and frozen produce. Observations are weighted for national representativeness.

A.B Health Index

Table A1: **Correlations Between Health Index and Its Components in Homescan**

Component	Correlation with Health Index
Adequacy ("healthy") components	
Total fruits	0.56
Whole fruits	0.57
Vegetables	0.41
Greens and beans	0.47
Whole grains	0.50
Dairy	0.25
Total protein	0.42
Sea and plant protein	0.64
Monounsaturated fat	0.11
Polyunsaturated fat	0.07
Moderation ("unhealthy") components	
Refined grains	-0.33
Sodium	-0.09
Added sugar	-0.41
Saturated fat	-0.21
Solid fats	-0.44

Notes: Using Homescan household-by-year data for 2004–2016, this table presents the correlation coefficients between the Health Index and its components, using components in units per 1,000 calories consumed. Observations are weighted for national representativeness.

Table A2: **Associations Between Health Outcomes and Dietary Quality Measures**

(a) **Predicting Body Mass Index**

	(1)	(2)	(3)	(4)
Healthy Eating Index (standard)	-1.272***	-1.251***		
	(0.066)	(0.400)		
Health Index (linearized)			-1.230***	-1.149***
			(0.065)	(0.365)
Controls	No	Yes	No	Yes
R^2	0.029	0.882	0.032	0.882
N	10,993	10,993	10,993	10,993
Dependent var. mean	29.47	29.47	29.47	29.47

(b) **Predicting Diabetes Status**

	(1)	(2)	(3)	(4)
Healthy Eating Index (standard)	-0.0265***	-0.0298		
	(0.0032)	(0.0196)		
Health Index (linearized)			-0.0297***	-0.0323*
			(0.0029)	(0.0173)
Controls	No	Yes	No	Yes
R^2	0.005	0.891	0.007	0.891
N	11,067	11,067	11,067	11,067
Dependent var. mean	0.17	0.17	0.17	0.17

Notes: This table presents regressions of health outcomes on dietary quality measures, using household-level Nielsen Homescan data. Body Mass Index is weight (in kilograms) divided by the square of height (in meters). Diabetic takes value 1 if the panelist reported that she had been diagnosed with diabetes, and 0 otherwise. If two household members responded to the PanelViews survey, we take the mean of each survey variable across the two respondents. Dietary quality measures are normalized to have a mean of zero and a standard deviation of one across households; we then take the calorie-weighted average across all years the household was observed in the Homescan sample. "Controls" are household-by-census tract fixed effects and household demographics (natural log of income, natural log of years of education, age indicators, household size, race indicators, a married indicator, employment status, and weekly work hours). Observations are not weighted for national representativeness. Robust standard errors, clustered by household, are in parentheses. *, **, ***: statistically significant with 10, 5, and 1 percent confidence, respectively.

B Appendix to Stylized Facts Section

B.A Additional Figures and Tables

Figure A2: **Healthy Eating Index Components by Household Income**

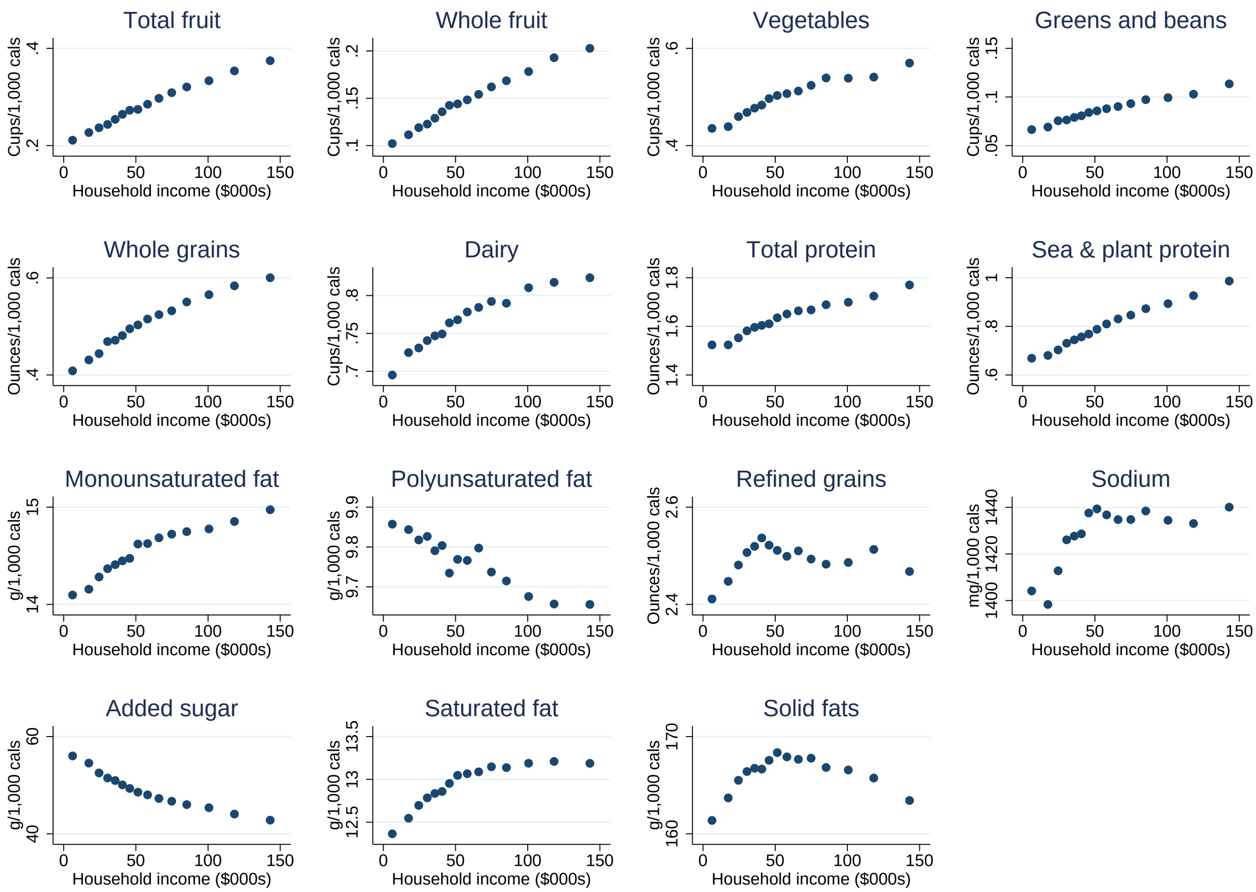

Notes: This figure presents Nielsen Homescan data for 2004–2016. Each panel presents a binned scatterplot of a dietary quality measure against average household income across all years the household is observed in Homescan, residual of age and year indicators and household size. The first 10 panels present the "healthy" dietary components of the Healthy Eating Index, while the final five panels present the "unhealthy" components. Observations are weighted for national representativeness.

Figure A3: **Magnet Subsample: Healthful Purchases by Household Income**

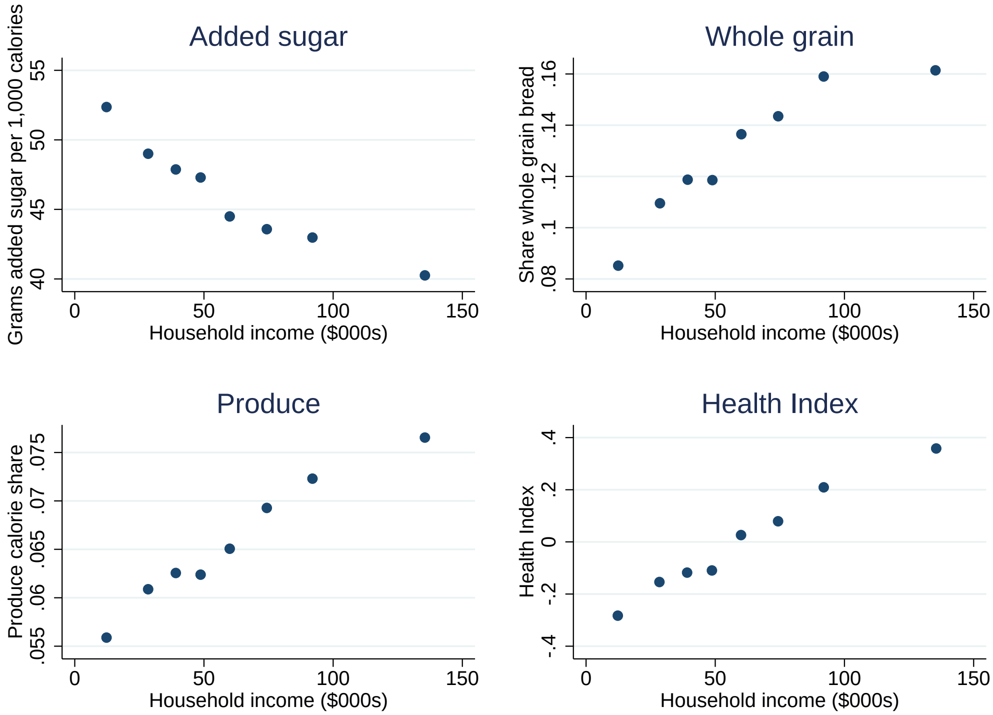

Notes: This parallels Figure 1, except it uses the 2004–2006 magnet subsample, which also records purchases of non-UPC items such as bulk produce. Each panel presents a binned scatterplot of a grocery healthfulness measure against average household income across all years the household is observed in Homescan, residual of age and year indicators and household size. Added sugar is the grams of added sugar per 1,000 calories purchased; whole grain is the calorie-weighted average share of bread, buns, and rolls purchases that are whole grain; produce is the share of calories from fresh, canned, dried, and frozen fruits and vegetables; and the Health Index is our overall measure of the healthfulness of grocery purchases and is normalized to have a mean of zero and a standard deviation of one across households. Observations are weighted for national representativeness.

Table A3: **Pooled OLS versus Within-Household Income Variation**

	(1)	(2)	(3)
ln(Household income)	0.1556***	0.0131**	0.0134**
	(0.0052)	(0.0051)	(0.0055)
Household-by-Census tract fixed effects	No	Yes	Yes
Household demographics	No	No	Yes
Observations	659,456	659,456	659,456
Income coefficient/column 1 coefficient	1.00	0.08	0.09

Notes: This table presents regressions of the Health Index on the natural log of household income and year indicators using Nielsen Homescan data for 2004–2016. Columns 2 and 3 also include household-by-census tract fixed effects, and column 3 also includes household demographics (natural log of years of education, age indicators, household size, race indicators, a married indicator, employment status, and weekly work hours). The sample is restricted to households observed in two or more years. The Health Index is our overall measure of the healthfulness of grocery purchases and is normalized to have a mean of zero and a standard deviation of one across households. Observations are weighted for national representativeness. Robust standard errors, clustered by household, are in parentheses. *, **, ***: statistically significant with 10, 5, and 1 percent confidence, respectively.

Table A4: **Correlates of the Count of Produce UPCs Available in RMS Stores**

	(1)	(2)	(3)
ln(Zip median income)	0.405***	0.000	0.094***
	(0.009)	(0.005)	(0.003)
ln(Annual revenue)		0.366***	
		(0.001)	
1(Large grocery)			1.488***
			(0.004)
1(Small grocery)			1.036***
			(0.008)
1(Supercenter/club)			0.883***
			(0.012)
1(Convenience store)			-0.143***
			(0.002)
1(Drug store)			-0.083***
			(0.002)
Observations	369,903	369,903	369,903
R^2	0.04	0.80	0.92

Notes: This table presents regressions of the count of produce UPCs available in RMS stores on store characteristics, using 2006–2016 Nielsen RMS data at the store-by-year level. ln(Annual revenue) is revenue from packaged grocery items with UPCs. “Large” (“small”) grocery stores are those with at least (less than) $5 million in annual revenue. Mass merchants other than supercenters and club stores are the omitted store type in column 3. Robust standard errors, clustered by zip code, are in parentheses. *, **, ***: statistically significant with 10, 5, and 1 percent confidence, respectively.

C Appendix to Reduced-Form Event Studies

C.A Additional Figures and Tables for Entry Event Study

Figure A4: **Event Study of Supermarket Entry Between 10 and 15 Minutes from Home**

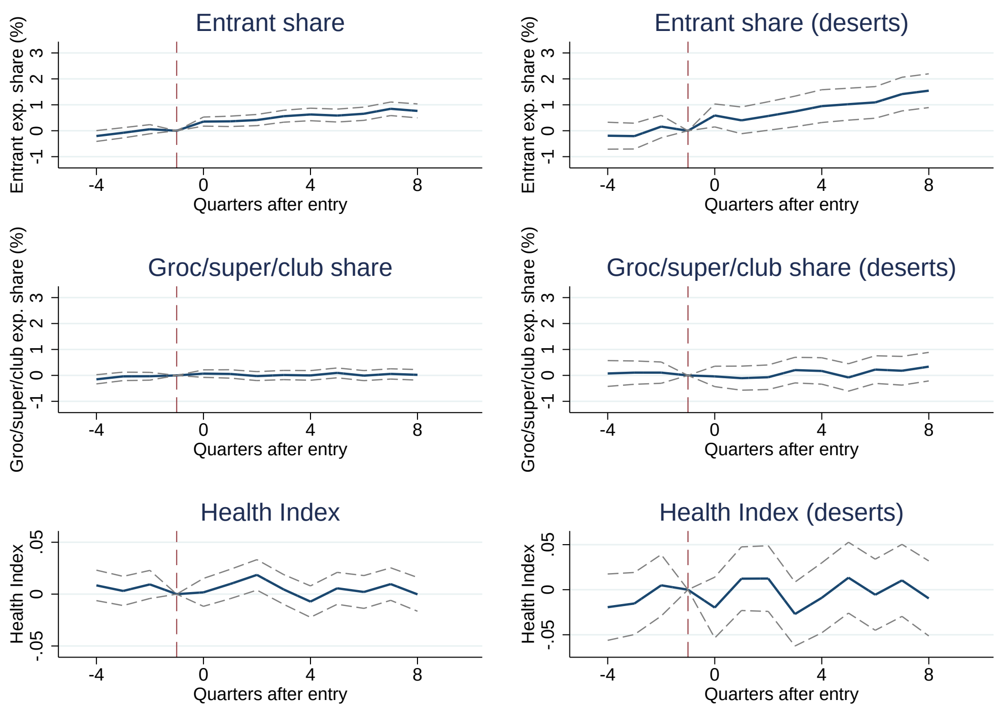

Notes: This figure presents the $\tau_{[10,15)q}$ parameters and 95 percent confidence intervals from estimates of Equation (2): the effects of supermarket entry, using 2004–2016 household-by-quarter Homescan data. All regressions control for household demographics (natural log of income, natural log of years of education, age indicators, household size, race indicators, a married indicator, employment status, and weekly work hours), census division-by-quarter of sample indicators, and household-by-census tract fixed effects. The top two panels present effects on expenditure shares (in units of percentage points) across all retailers with stores that have entered within a 15-minute drive of the household. The middle two panels present effects on the combined expenditure share of grocery stores, supercenters, and club stores. We keep the y-axis on the same scale between the top and middle panels so that the magnitudes can be easily compared. The bottom two panels present effects on the Health Index, our overall measure of the healthfulness of grocery purchases which is normalized to have a mean of zero and a standard deviation of one across households. The left panels include the full sample, while the right panels include only the "food desert" subsample: observations with no grocery stores with 50 or more employees, supercenters, or club stores in the zip code in the first year the household is observed there. Observations are not weighted for national representativeness.

Figure A5: **Event Study of Supermarket Entry with Additional Leads**

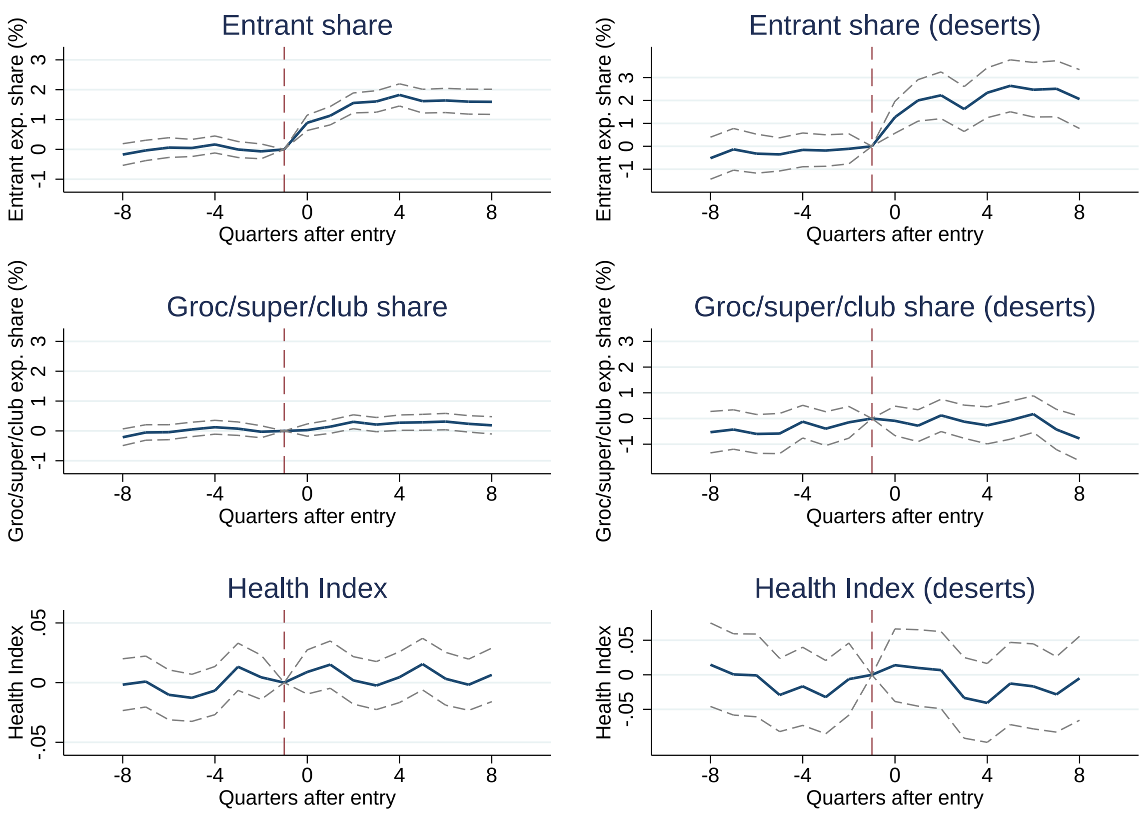

Notes: This figure presents the $\tau_{[0,10)q}$ parameters and 95 percent confidence intervals from estimates of Equation (2): the effects of supermarket entry, using 2004–2016 household-by-quarter Homescan data. The figure parallels Figure 4, except with additional leads of the entry. All regressions control for household demographics (natural log of income, natural log of years of education, age indicators, household size, race indicators, a married indicator, employment status, and weekly work hours), census division-by-quarter of sample indicators, and household-by-census tract fixed effects. The top two panels present effects on expenditure shares (in units of percentage points) across all retailers with stores that have entered within a 15-minute drive of the household. The middle two panels present effects on the combined expenditure share of grocery stores, supercenters, and club stores. We keep the y-axis on the same scale between the top and middle panels so that the magnitudes can be easily compared. The bottom two panels present effects on the Health Index, our overall measure of the healthfulness of grocery purchases which is normalized to have a mean zero and a standard deviation of one across households. The left panels include the full sample, while the right panels include only the "food desert" subsample: observations with no grocery stores with 50 or more employees, supercenters, or club stores in the zip code in the first year the household is observed there. Observations are not weighted for national representativeness.

Table A5: **Effects of Supermarket Entry**

(a) **Effects on Expenditure Shares at Drug and Convenience Stores**

	Full sample	Bottom quartile	Food deserts
	(1)	(2)	(3)
Post entry: 0-10 minutes	-0.015	-0.048	0.058
	(0.022)	(0.065)	(0.067)
Post entry: 10-15 minutes	-0.033*	-0.053	-0.112*
	(0.017)	(0.047)	(0.064)
Observations	2,874,365	537,998	646,181
Dependent var. mean	2.6	3.1	2.4

(b) **Effects on Health Index Using Alternative Food Desert Definitions**

	<1000 produce UPCs	No medium groceries	Three-mile radius
	(1)	(2)	(3)
Post entry: 0-10 minutes	0.004	0.013	0.017
	(0.011)	(0.012)	(0.014)
Post entry: 10-15 minutes	0.006	0.004	0.019*
	(0.007)	(0.007)	(0.010)
Observations	411,654	378,682	490,551

Notes: This table uses 2004–2016 Nielsen Homescan data at the household-by-quarter level. The table parallels Table 2, except Panel (a) presents effects on expenditure shares at drug and convenience stores, and Panel (b) uses alternative definitions of a "food desert." In Panel (b), column 1 defines food deserts as zip codes with fewer than 1,000 produce UPCs, as predicted by projecting produce UPC counts in RMS stores from column 3 of Table A4 onto Zip Code Business Patterns store count data; column 2 uses the primary food desert definition but also excludes any zip codes with grocery stores employing between 10 and 49 employees; and column 3 defines a zip code as a food desert only if all zip codes with centroids within three miles have no grocery stores with 50 or more employees, supercenters, or club stores. Expenditure shares are in units of percentage points. The Health Index is our overall measure of the healthfulness of grocery purchases and is normalized to have a mean of zero and a standard deviation of one across households. Reported independent variables are the count of supermarkets that have entered within a 0–10 or 10–15 minute drive from the household's census tract centroid. All regressions control for household demographics (natural log of income, natural log of years of education, age indicators, household size, race indicators, a married indicator, employment status, and weekly work hours), census division-by-quarter of sample indicators, and household-by-census tract fixed effects. Observations are not weighted for national representativeness. Robust standard errors, clustered by household and census tract, are in parentheses. *, **, ***: statistically significant with 10, 5, and 1 percent confidence, respectively.

Table A6: **Effects of Supercenter Entry**

(a) **Effects on Expenditure Shares**

	Full sample		Bottom quartile		Food deserts	
	(1)	(2)	(3)	(4)	(5)	(6)
	Entrants	Grocery/ super/club	Entrants	Grocery/ super/club	Entrants	Grocery/ super/club
Post entry: 0-10 minutes	3.786***	0.439***	5.224***	0.535*	3.629***	0.227
	(0.226)	(0.107)	(0.629)	(0.314)	(0.624)	(0.302)
Post entry: 10-15 minutes	1.601***	0.033	1.902***	0.238	1.976***	0.205
	(0.114)	(0.071)	(0.345)	(0.209)	(0.292)	(0.183)
Observations	2,874,514	2,874,365	538,041	537,998	646,223	646,181
Dependent var. mean	3.7	88.2	3.4	86.2	2.6	87.7

(b) **Effects on Health Index**

	Full sample	Bottom quartile	Food deserts
	(1)	(2)	(3)
Post entry: 0-10 minutes	0.008	0.019	0.018
	(0.006)	(0.015)	(0.016)
Post entry: 10-15 minutes	0.006	-0.002	0.020*
	(0.004)	(0.011)	(0.011)
Observations	2,874,514	538,041	646,223

Notes: This table uses 2004-2016 Nielsen Homescan data at the household-by-quarter level. It parallels Table 2, except it considers entry by supercenters only, excluding other types of grocery stores. The "food desert" subsample comprises observations with no grocery stores with 50 or more employees, supercenters, or club stores in the zip code in the first year the household is observed there. Expenditure shares are in units of percentage points. The Health Index is our overall measure of the healthfulness of grocery purchases and is normalized to have a mean of zero and a standard deviation of one across households. Reported independent variables are the count of supermarkets that have entered within a 0–10 or 10–15 minute drive from the household's census tract centroid. All regressions control for household demographics (natural log of income, natural log of years of education, age indicators, household size, race indicators, a married indicator, employment status, and weekly work hours), census division-by-quarter of sample indicators, and household-by-census tract fixed effects. Observations are not weighted for national representativeness. Robust standard errors, clustered by household and census tract, are in parentheses. *, **, ***: statistically significant with 10, 5, and 1 percent confidence, respectively.

Figure A6: **Shopping Trip Distances by Household Income**

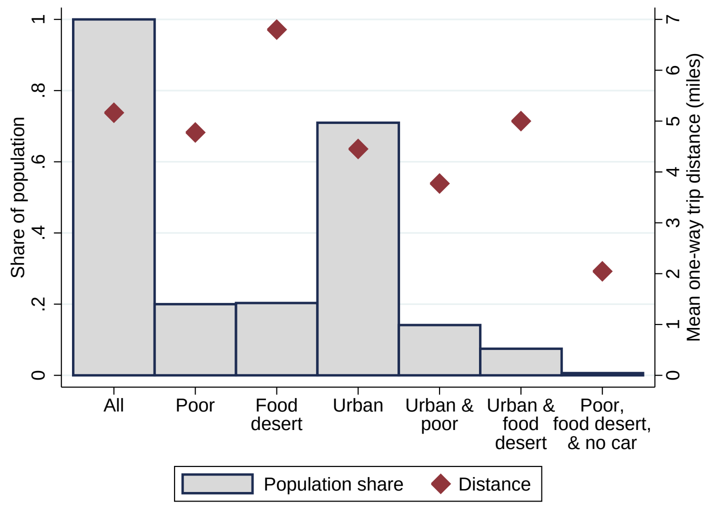

Notes: Data are from the 2009 National Household Travel Survey. Diamonds represent the mean one-way trip distance for trips beginning or ending in "buying goods: groceries/clothing/hardware store." "Poor" means households in the bottom income quartile. "Food desert" means that the household is in a zip code with no grocery stores with 50 or more employees, supercenters, or club stores. "Urban" includes urbanized areas or urban clusters of at least 2500 people, using the U.S. Census Bureau definition. "No car" means that the household does not own a car.

Figure A7: **Median Shopping Trip Distances by Household Income**

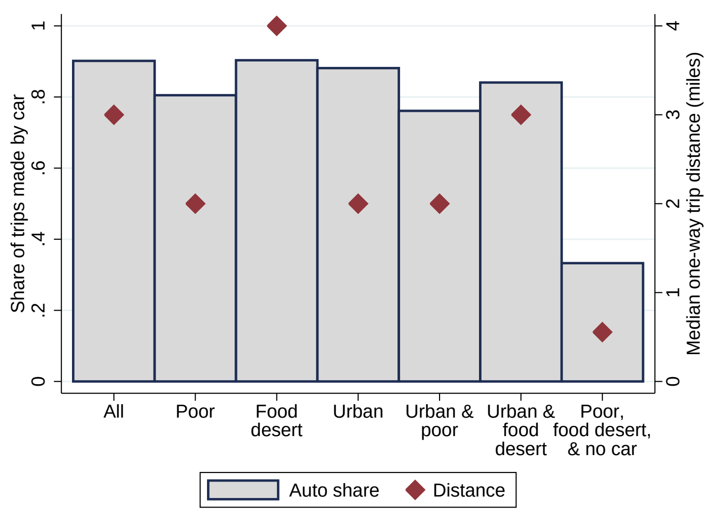

Notes: Data are from the 2009 National Household Travel Survey. Diamonds represent the median one-way trip distance for trips beginning or ending in "buying goods: groceries/clothing/hardware store." "Poor" means households in the bottom income quartile. "Food desert" means that the household is in a zip code with no grocery stores with 50 or more employees, supercenters, or club stores. "Urban" includes urbanized areas or urban clusters of at least 2500 people, using the U.S. Census Bureau definition. "No car" means that the household does not own a car.

Figure A8: **Supermarket Expenditure Shares by Household Income**

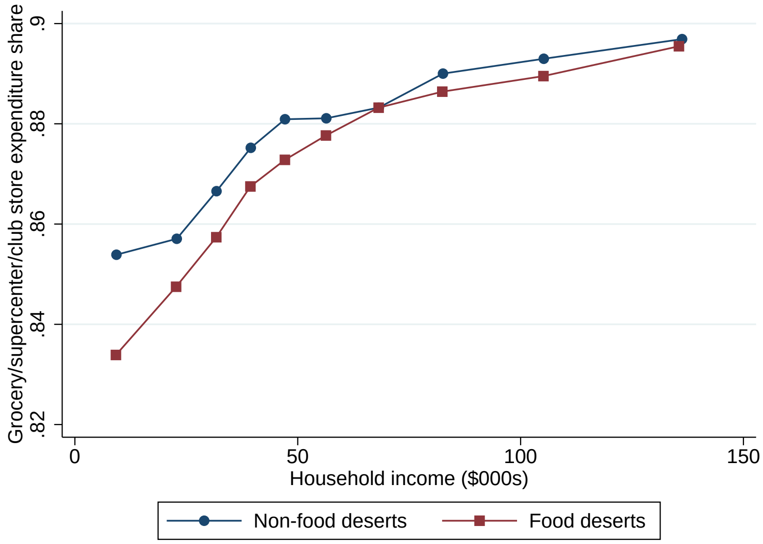

Notes: This figure presents the share of grocery expenditures that are made at grocery stores, supercenters, and club stores against average household income across all years the household is observed in Homescan, residual of age and year indicators and household size, using Nielsen Homescan data for 2004–2016. A household-by-year observation is in a "food desert" if its zip code does not have any grocery stores with 50 or more employees, supercenters, or club stores in that year. Observations are weighted for national representativeness.

C.B Appendix to Movers Event Study

Figure A9: **Event Study of Moves Across Zip Codes**

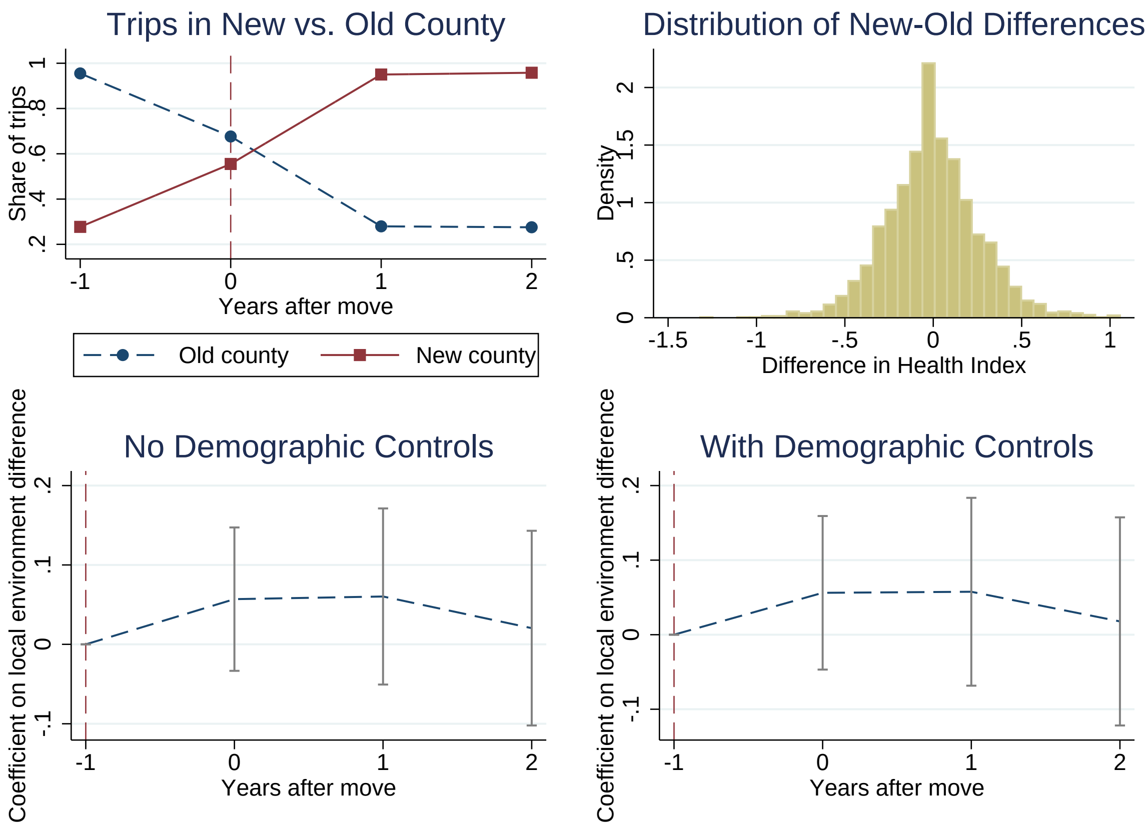

Notes: Using 2004–2016 Homescan data, these figures present results for the event study of moves across zip codes. The top left panel presents the share of shopping trips that are in the new versus old county. The top right panel presents the distribution across balanced panel households of the difference in the Health Index between the new and old zip code. The bottom panels present the τ_y parameters and 95 percent confidence intervals from estimates of Equation (4): associations between the household-level Health Index and the difference in the average local Health Index between post-move and pre-move locations. The bottom right panel includes controls for household demographics (natural log of income, natural log of years of education, age indicators, household size, race indicators, a married indicator, employment status, and weekly work hours). The Health Index is our overall measure of the healthfulness of grocery purchases and is normalized to have a mean of zero and a standard deviation of one across households. Observations are not weighted for national representativeness.

Figure A10: **Event Study of Movers with Different Balanced Panel Windows**

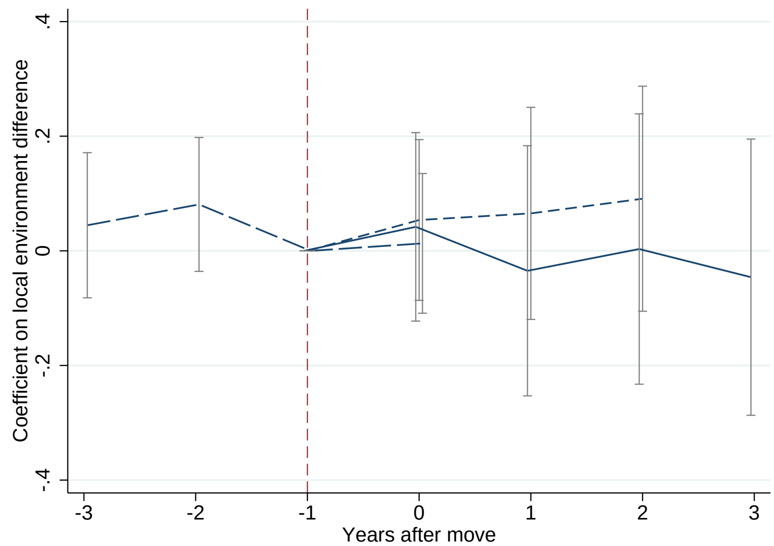

(a) **Without Controls**

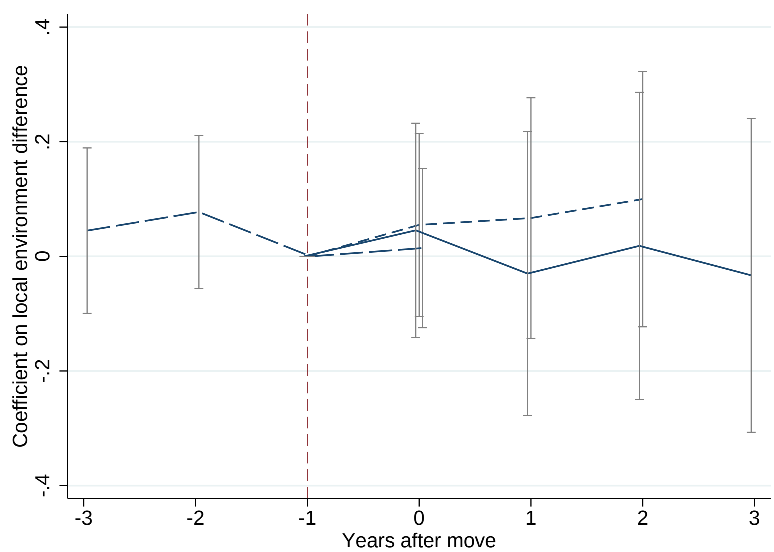

(b) **With Controls**

Notes: Using 2004–2016 Homescan data, these figures present the τ_y parameters and 95 percent confidence intervals from estimates of Equation (4): associations between the household-level Health Index and the difference in the average local Health Index between post-move and pre-move locations. Each figure superimposes three different estimates identified off of balanced panels for different windows around the move. Panel (b) includes controls for household demographics (natural log of income, natural log of years of education, age indicators, household size, race indicators, a married indicator, employment status, and weekly work hours). The Health Index is our overall measure of the healthfulness of grocery purchases and is normalized to have a mean of zero and a standard deviation of one across households. Observations are not weighted for national representativeness.

Figure A11: **Event Study: Income Changes in Mover Households**

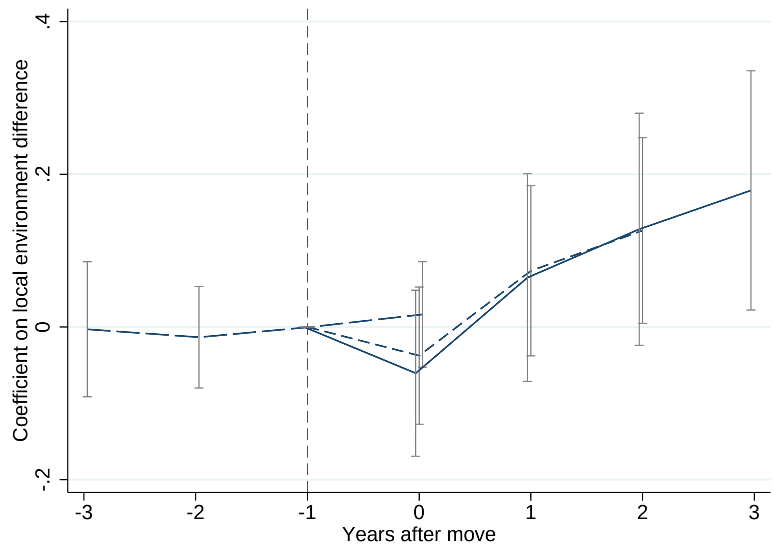

(a) **Moves Across Counties**

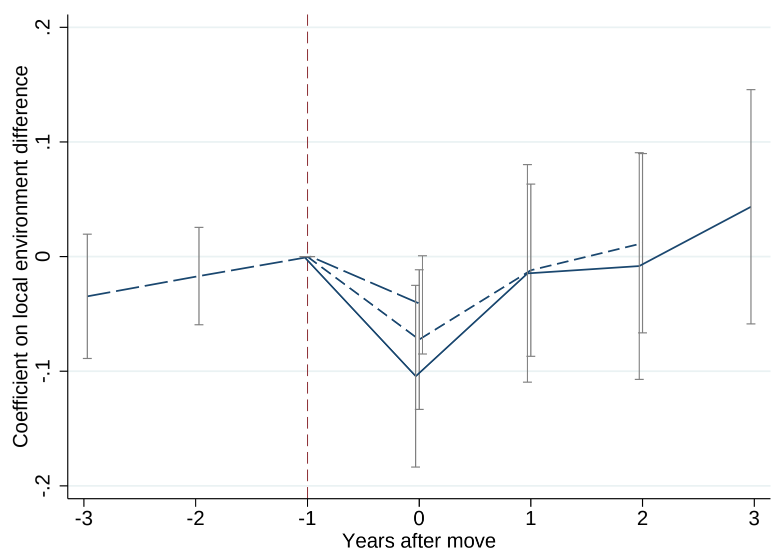

(b) **Moves Across Zip Codes**

Notes: Using 2004–2016 Homescan data, these figures present the τ_y parameters and 95 percent confidence intervals from estimates of Equation (4): associations between natural log of household income and the difference in the average local Health Index between post-move and pre-move locations. All regressions control for year indicators and household fixed effects. Each figure superimposes three different estimates identified off of balanced panels for different windows around the move. The Health Index is our overall measure of the healthfulness of grocery purchases and is normalized to have a mean of zero and a standard deviation of one across households. Observations are not weighted for national representativeness. The regressions are the same as in Figure 6, except with natural log of household income as the dependent variable and no controls for household demographics.

Table A7: **Association of Income with Local Area Health Index Using Movers**

	(1)	(2)
Zip code average Health Index	0.00590	
	(0.36)	
County average Health Index		0.125***
		(3.25)
Household demographics		
Observations	564,944	570,279
95% confidence interval upper bound	0.038	0.200

Notes: This table uses 2004–2016 Nielsen Homescan data at the household-by-year level. The sample excludes observations where less than 50 percent of trips to RMS stores are not in the household's end-of-year county of residence. The dependent variable is the natural log of household income. The Health Index is our overall measure of the healthfulness of grocery purchases and is normalized to have a mean of zero and a standard deviation of one across households. All regressions control for year indicators and household fixed effects. Observations are not weighted for national representativeness. Robust standard errors, clustered by household and local area (zip code or county), are in parentheses. *, **, ***: statistically significant with 10, 5, and 1 percent confidence, respectively.

Table A8: **Association of Coke Market Share with Local Area Coke Market Share Using Movers**

	(1)	(2)
County average Coke market share	0.1620***	0.1613***
	(0.0418)	(0.0417)
Household demographics	No	Yes
Observations	306,714	306,714

Notes: This table uses 2004–2016 Nielsen Homescan data at the household-by-year level. The sample excludes observations where less than 50 percent of trips to RMS stores are not in the household's end-of-year county of residence. Coke market share equals Coke calories purchased / (Coke + Pepsi calories purchased). Household demographics are natural log of income, natural log of years of education, age indicators, household size, race indicators, a married indicator, employment status, and weekly work hours. All regressions also control for year indicators and household fixed effects. Observations are not weighted for national representativeness. Robust standard errors, clustered by household and county, are in parentheses. *, **, ***: statistically significant with 10, 5, and 1 percent confidence, respectively.

D Appendix to Demand Model Estimation

D.A Method of Moments Estimation

Our method of moments estimator is defined as follows:

$$\left(\hat{\boldsymbol{\delta}}, \hat{\boldsymbol{\phi}}, \hat{\tilde{\boldsymbol{\beta}}}, \hat{\xi}\right) = \underset{(\boldsymbol{\delta},\boldsymbol{\phi},\tilde{\boldsymbol{\beta}},\xi)}{\arg\min} \left(\frac{1}{IJT}\sum_i\sum_j\sum_t \boldsymbol{g}_{ijt}\right)' \left(\frac{1}{IJT}\sum_i\sum_j\sum_t \boldsymbol{g}_{ijt}\right).$$

Define $\boldsymbol{Y}$ as the vector of product group calorie consumption Y_{ijt}, $\boldsymbol{F}\left(\tilde{\boldsymbol{\beta}}, \xi\right)$ as the vector of implicit prices $F_{ijt} = \left(\tilde{p}_{ijt} - \sum_{c=2}^{C} \tilde{\beta}_c \tilde{a}_{ijct} - \xi\right)$, $\boldsymbol{D}$ as a stacked matrix of the two dummy variable matrices ($\boldsymbol{D}_j$ and $\boldsymbol{D}_m$), $\boldsymbol{Z}$ as a matrix with all of our vectors of instruments ($\boldsymbol{D}$, the nutrient content $\tilde{\boldsymbol{a}}$, and the price instruments $\boldsymbol{P}$), and $\boldsymbol{Pr_D} = (\boldsymbol{D'ZZ'D})^{-1}\boldsymbol{D'ZZ'}$ as a projection matrix. We can simplify the estimation problem by solving for our vectors of linear coefficients, $\boldsymbol{\delta}$ and $\boldsymbol{\phi}$, as analytic functions of $\tilde{\boldsymbol{\beta}}$ and $\boldsymbol{\xi}$:

$$(\boldsymbol{\delta}, \boldsymbol{\phi}) = \boldsymbol{Pr}_D\left(\ln(\boldsymbol{Y}) - \boldsymbol{F}\left(\tilde{\boldsymbol{\beta}}, \xi\right)\right). \tag{32}$$

Substituting Equation (32) back into Equation (19), we can re-write the MOM estimator in terms of $\tilde{\boldsymbol{\beta}}$ and $\hat{\xi}$:

$$\left(\hat{\tilde{\boldsymbol{\beta}}}, \hat{\xi}\right) = \underset{(\tilde{\boldsymbol{\beta}},\xi)}{\arg\min} \left(\frac{1}{IJT}\sum_i\sum_j\sum_t g_{ijt}\left(\tilde{\boldsymbol{\beta}}, \xi\right)\right)' \left(\frac{1}{IJT}\sum_i\sum_j\sum_t g_{ijt}\left(\tilde{\boldsymbol{\beta}}, \xi\right)\right).$$

At the true value, the gradient for this problem is:

$$-2\boldsymbol{G}\left(\tilde{\boldsymbol{\beta}}, \xi\right)'\boldsymbol{G}\left(\tilde{\boldsymbol{\beta}}, \xi\right) = 0$$

where the Jacobian of the moments, $\boldsymbol{G}\left(\tilde{\boldsymbol{\beta}}, \xi\right)$, is

$$\boldsymbol{G}\left(\tilde{\boldsymbol{\beta}}, \xi\right) = \frac{1}{IJT}\begin{bmatrix} \tilde{\boldsymbol{a}}'\left(\boldsymbol{I} - \boldsymbol{DPr_D}\right) \\ \\ \boldsymbol{P}'\left(\boldsymbol{I} - \boldsymbol{D}_m\boldsymbol{Pr}_{\boldsymbol{D}_m}\right) \\ \\ \boldsymbol{D}'\left(\boldsymbol{I} - \boldsymbol{DPr_D}\right) \end{bmatrix} \nabla_\beta \boldsymbol{F}\left(\tilde{\boldsymbol{p}}, \tilde{\boldsymbol{a}}; \tilde{\boldsymbol{\beta}}\right). \tag{33}$$

In the above equation, $\boldsymbol{I}$ is the identity matrix, and $\boldsymbol{Pr}_{\boldsymbol{D}_m}$ is a projection matrix using $\boldsymbol{D}_m$.

The covariance matrix of our full MOM estimator, $\boldsymbol{\Theta}^{MOM} \equiv \left(\hat{\boldsymbol{\delta}}, \hat{\boldsymbol{\phi}}, \hat{\tilde{\beta}}\hat{\xi}\right)$, is $cov\left(\boldsymbol{\Theta}^{MOM}\right) =$

$(\boldsymbol{G}'\boldsymbol{G})^{-1}\boldsymbol{G}'\boldsymbol{\Omega}\boldsymbol{G}(\boldsymbol{G}'\boldsymbol{G})^{-1}$, with Jacobian matrix

$$G = \frac{1}{IJT}\sum_i\sum_j\sum_t \begin{bmatrix} \overrightarrow{0}'_J & -\tilde{a}_{ijt}D'_m & -\tilde{a}_{ijt}\nabla_\beta F'_{ijt} \\ -P_{jmt}D'_j & -P_{jmt}D'_m & -P_{jmt}\nabla_\beta F'_{ijt} \\ -D_jD'_j & -D_jD'_m & -D_j\nabla_\beta F'_{ijt} \\ -D_mD'_{ijt} & -D_mD'_m & -D_m\nabla_\beta F'_{ijt} \end{bmatrix}$$

and covariance matrix

$$\boldsymbol{\Omega} \;=\; \mathbb{E}\left(\boldsymbol{g}_{ijt}\left(\boldsymbol{\Theta}^{MOM}\right)\boldsymbol{g}_{ijt}\left(\boldsymbol{\Theta}^{MOM}\right)'\right).$$

When computing our standard errors, we cluster by household as follows:

$$\hat{\Omega} = \frac{1}{IJT}\sum_i\sum_{j,j'}\sum_{t,t'}\boldsymbol{g}_{ijt}\left(\hat{\tilde{\boldsymbol{\beta}}},\hat{\xi}\right)\boldsymbol{g}_{ij't'}\left(\hat{\tilde{\boldsymbol{\beta}}},\hat{\xi}\right)'.$$

D.B Additional Tables and Figures

Table A9: **Preferences for Nutrients by Household Income**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Income quartile	Sodium	Whole fruit	Other fruit	Whole grains	Refined grains	Greens, beans	Other veg	Dairy
Income Q1	-0.178***	-0.324***	-0.064***	0.177***	-0.014***	0.840***	-0.290***	0.084***
	(0.009)	(0.020)	(0.008)	(0.003)	(0.0004)	(0.007)	(0.016)	(0.003)
Income Q2	-0.299***	-0.225***	-0.086***	0.228***	-0.010***	0.861***	-0.382***	0.063***
	(0.012)	(0.015)	(0.007)	(0.005)	(0.0005)	(0.007)	(0.017)	(0.003)
Income Q3	-0.384***	-0.233***	-0.093***	0.269***	-0.0096***	0.970***	-0.436***	0.068***
	(0.014)	(0.015)	(0.007)	(0.006)	(0.001)	(0.009)	(0.019)	(0.003)
Income Q4	-0.585***	-0.246***	-0.115***	0.365***	0.0031*	1.216***	-0.605***	0.084***
	(0.025)	(0.020)	(0.010)	(0.011)	(0.0017)	(0.018)	(0.031)	(0.004)

	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	Sea, plant protein	Meat protein	Added sugar	Solid fats	Unobserved characteristic	Shelf life	Con-venience	WTP for Health Index
Income Q1	-0.240***	0.027***	0.0003***	-0.0005***	0.0002***	-0.109***	0.270***	0.202***
	(0.009)	(0.002)	(0.00005)	(0.00002)	(0.00001)	(0.003)	(0.001)	(0.021)
Income Q2	-0.294***	0.0021	-0.0009***	-0.0004***	0.00001	-0.150***	0.316***	0.267***
	(0.009)	(0.002)	(0.0001)	(0.00001)	(0.00001)	(0.003)	(0.001)	(0.017)
Income Q3	-0.342***	-0.0056**	-0.0024***	-0.0005***	-0.0002***	-0.175***	0.400***	0.402***
	(0.011)	(0.003)	(0.0001)	(0.00002)	(0.00002)	(0.004)	(0.002)	(0.015)
Income Q4	-0.433***	-0.021***	-0.005***	-0.0005***	-0.0007***	-0.223***	0.542***	0.630***
	(0.017)	(0.004)	(0.0002)	(0.00002)	(0.00004)	(0.006)	(0.005)	(0.013)

Notes: This table presents GMM estimates of the preference parameters $\tilde{\beta}_c$ from Equation (12), adding convenience and shelf life as additional product characteristics. Shelf life is measured in months per 1,000 calories and top-coded at one year. Convenience is a score per 1,000 calories ranging from 0 to 3, defined as follows. 0: basic ingredients. These are raw or minimally processed foods used in producing a meal or snack that are generally composed of a single ingredient, such as milk, dried beans, rice, grains, butter, cream, fresh meat, poultry, and seafood. 1: complex ingredients, such as bread, pasta, sour cream, sauce, canned vegetables, canned beans, pickles, cereal, frozen meat/poultry/seafood, canned meat/poultry/seafood, and lunch meat. 2: ready-to-cook meals and stacks. These are foods that require minimal preparation involving heating, cooking, or adding hot water, such as frozen entrees, frozen pizzas, dry meal mixes, pudding mixes, soup, chili, and powdered drinks. 3: ready-to-eat meals and snacks. These are foods that are intended to be consumed as is and require no preparation beyond opening a container, including refrigerated entrees and sides, canned and fresh fruit, yogurt, candy, snacks, liquid drinks, and flavored milk. Shelf life data are from Okrent and Kumcu (2016), while convenience data are from the U.S. government's FoodKeeper app (HHS 2015). Magnitudes represent willingness to pay for a unit of the nutrient, where the units are those used in the Health Index. Sodium is in grams; whole fruit, other fruit and dairy are in cups; whole grains, refined grains, and both types of protein are in ounces, added sugar is in teaspoons; solid fats are in calories. "WTP for Health Index" in column 16 equals $\sum_c \hat{\tilde{\beta}}_c s_c r_c$, where s_c is the maximum possible score on the Healthy Eating Index for dietary component c, and r_c is the difference in consumption of component c to receive the maximum instead of the minimum score. Standard errors, clustered by household, are in parentheses. *, **, ***: statistically significant with 10, 5, and 1 percent confidence, respectively.

Figure A12: **Geographic Variation in Presence of Large Retail Chains**

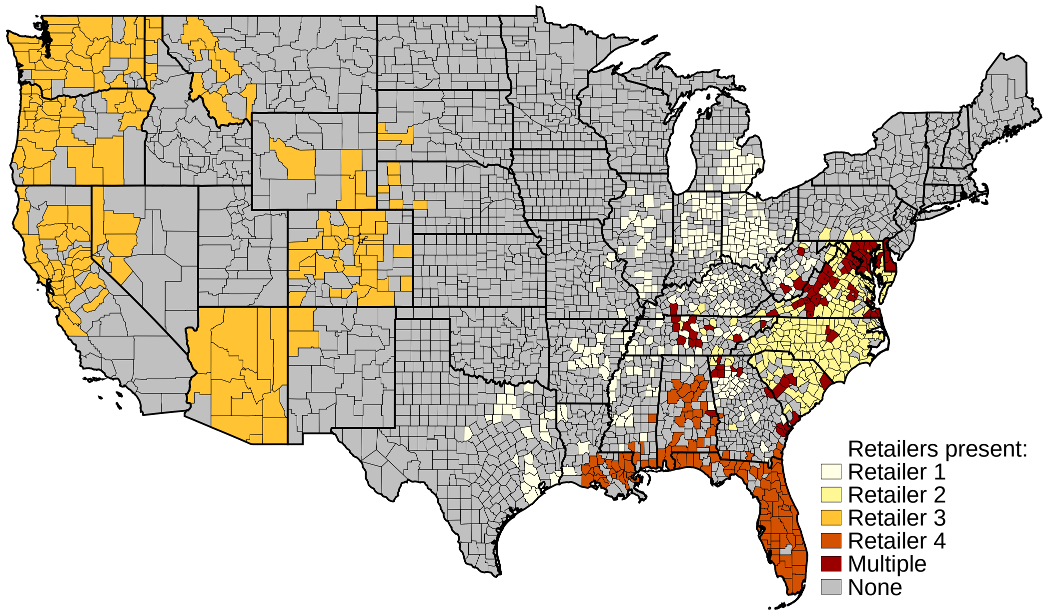

Notes: This figure shows the counties where four large retail chains in RMS had stores in 2015.

Figure A13: **Geographic Variation in the Price Instrument**

(a) **Yogurt**

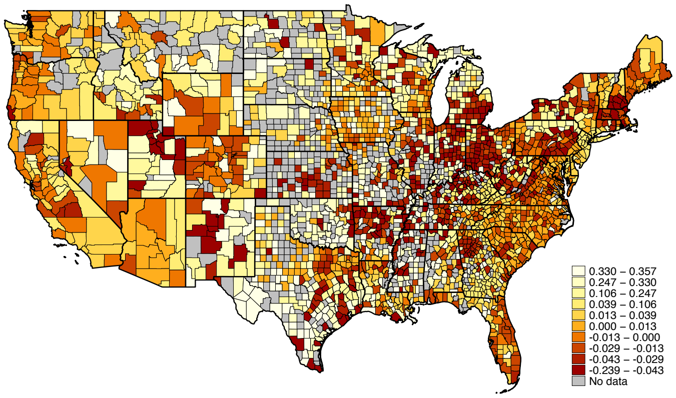

(b) **Fresh Produce**

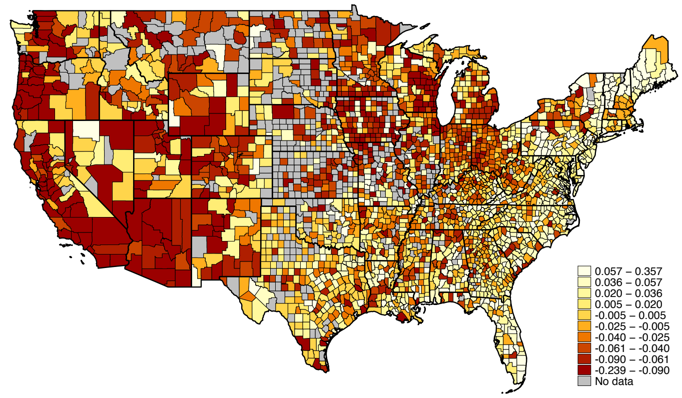

(c) **Frozen Pizza**

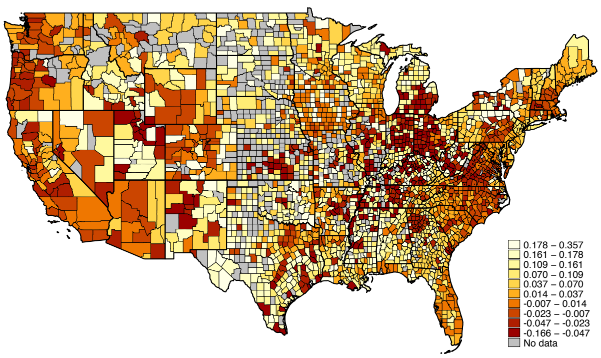

(d) **Cookies**

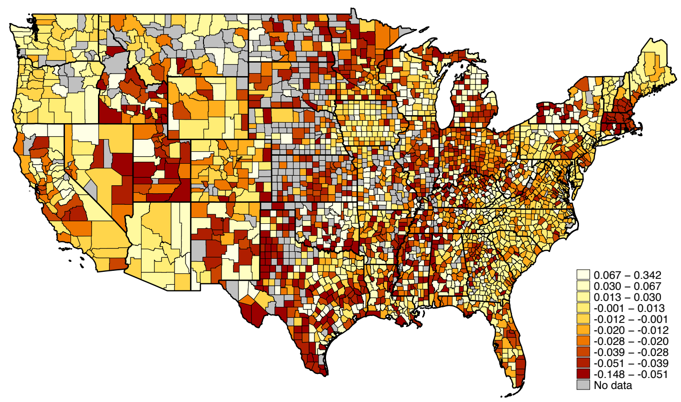

Notes: These figures present the county averages (over the years of our sample) of the price instrument P_{jmt} for example product groups.

Figure A14: **Standard Deviation of Price Instrument by Product Department**

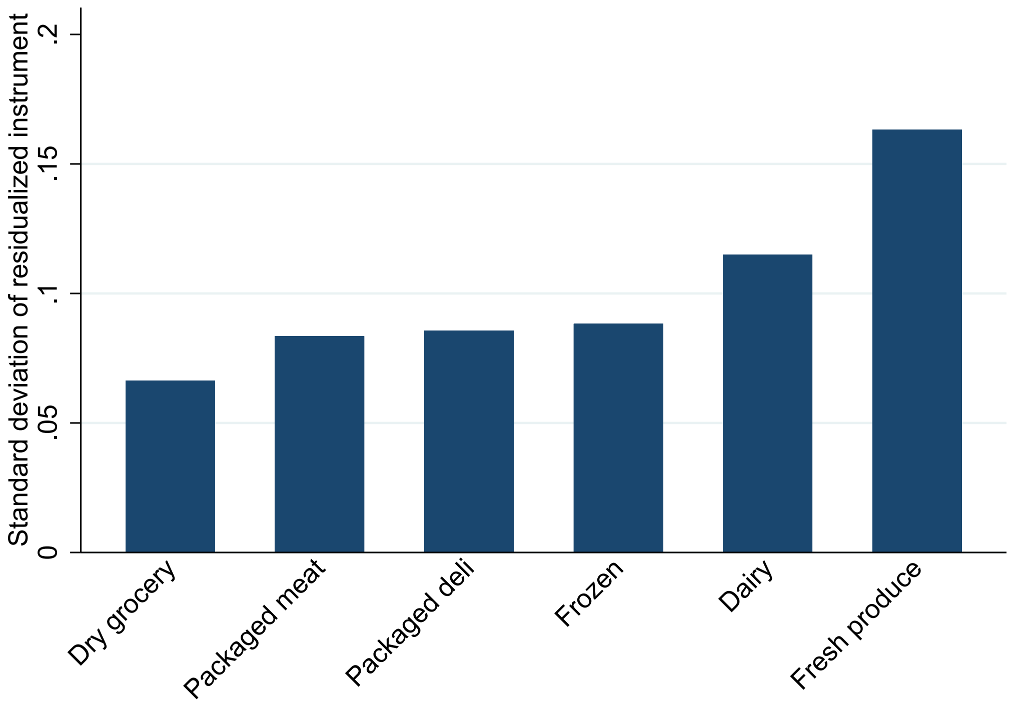

Notes: This figure presents the standard deviation of our price instrument P_{jmt}, after residualizing against year, product group, and county fixed effects. The instrument is in units of log price per calorie.

Figure A15: **Binned Scatterplot of First Stage Price Regression**

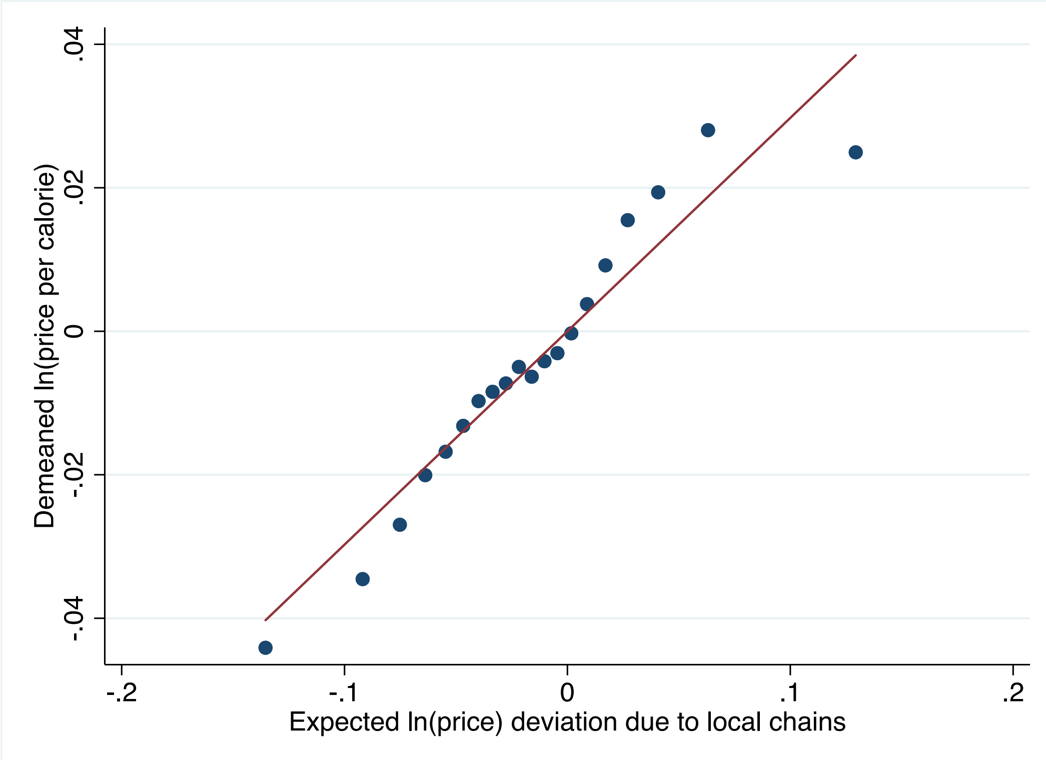

Notes: This figure presents a binned scatterplot of a regression of natural log price per calorie on our price instrument P_{jmt}, using Homescan data at the household-by-product group-by-year level. County-income quartile and product group-income quartile fixed effects are residualized out before plotting. There are 20 equally sized bins, and all income groups are included.