

NBER WORKING PAPER SERIES

THE GEOGRAPHY OF POVERTY AND NUTRITION:
FOOD DESERTS AND FOOD CHOICES ACROSS THE UNITED STATES

Hunt Allcott
Rebecca Diamond
Jean-Pierre Dubé

Working Paper 24094
<http://www.nber.org/papers/w24094>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2017, Revised January 2018

We thank Catherine Wright for exceptional research assistance, and we thank Charles Courtemanche and Sungho Park for sharing data. We also thank seminar participants at Columbia, Duke, the 2015 and 2016 NBER Summer Institutes, Microsoft Research, New York University, Princeton, Stanford, the Tilburg Christmas Research Camp, the University of Pennsylvania, and USC Marshall for helpful comments. We are grateful for funding from the Chicago Booth Initiative on Global Markets. This paper reflects the authors' own analyses and calculations based in part on data reported by Nielsen through its Homescan, RMS, and PanelViews services for all grocery categories over 2004-2015, for all retail channels in the U.S. market. The conclusions drawn from the Nielsen data are those of the authors 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 views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Hunt Allcott, Rebecca Diamond, and Jean-Pierre Dubé. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Geography of Poverty and Nutrition: Food Deserts and Food Choices Across the United States

Hunt Allcott, Rebecca Diamond, and Jean-Pierre Dubé

NBER Working Paper No. 24094

December 2017, Revised January 2018

JEL No. D12,I12,I14,L81,R20

ABSTRACT

We study the causes of “nutritional inequality”: why the wealthy tend to eat more healthfully than the poor in the U.S. Using event study designs exploiting supermarket entry and households' moves to healthier neighborhoods, we reject that neighborhood environments have meaningful effects on healthy eating. Using a structural demand model, we find that exposing low-income households to the same availability and prices experienced by high-income households reduces nutritional inequality by only 9%, while the remaining 91% is driven by differences in demand. These findings contrast with discussions of nutritional inequality that emphasize supply-side factors such as food deserts.

Hunt Allcott
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
hunt.allcott@nyu.edu

Jean-Pierre Dubé
University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
jdube@chicagobooth.edu

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

The Geography of Poverty and Nutrition: Food Deserts and Food Choices Across the United States

Hunt Allcott, Rebecca Diamond, and Jean-Pierre Dubé*

January 18, 2018

Abstract

We study the causes of “nutritional inequality”: why the wealthy tend to eat more healthfully than the poor in the U.S. Using event study designs exploiting supermarket entry and households’ moves to healthier neighborhoods, we reject that neighborhood environments have meaningful effects on healthy eating. Using a structural demand model, we find that exposing low-income households to the same availability and prices experienced by high-income households reduces nutritional inequality by only 9%, while the remaining 91% is driven by differences in demand. These findings contrast with discussions of nutritional inequality that emphasize supply-side factors such as food deserts.

Keywords: Inequality, food deserts, event study, demand estimation, decomposition.

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

I Introduction

Studies by Aizer and Currie (2014), Case and Deaton (2015), Chetty et al. (2014), Chetty et al. (2016), Saez and Piketty (2003), and many others have drawn increased attention to the causes and consequences of socioeconomic inequality. These consequences play out in many different ways, including educational opportunities, health outcomes, and mass incarceration. In this paper, we study one important correlate of socioeconomic status—what we eat and drink, which in turn affects obesity

*This paper supersedes our original July 2015 working paper, available [here](#). 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. We thank Catherine Wright for exceptional research assistance, and we thank Charles Courtemanche and Sungho Park for sharing data. We also thank seminar participants at Columbia, Duke, the 2015 and 2016 NBER Summer Institutes, Microsoft Research, New York University, Princeton, Stanford, the Tilburg Christmas Research Camp, the University of Pennsylvania, and USC Marshall for helpful comments. We are grateful for funding from the Chicago Booth Initiative on Global Markets. This paper reflects the authors’ own analyses and calculations based in part on data reported by Nielsen through its Homescan, RMS, and PanelViews services for all grocery categories over 2004-2015, for all retail channels in the U.S. market. The conclusions drawn from the Nielsen data are those of the authors 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.

and other health problems—and quantify the economic mechanisms that drive the nutrition-income relationship.¹

A large 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.² Parts of the public health literature, along with many policymakers and advocates, have subsequently argued that the reduced local supply of healthy foods is an important *cause* of unhealthy eating in disadvantaged neighborhoods.³ Through this lens, addressing the problem of food deserts is an important issue of public health and social justice. For example, Hilmers, Hilmers, and Dave (2012, page 1652) write:

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... This principle of fairness and equity needs to be reflected in neighborhood environments that facilitate healthy food choices for all societal strata.

Accordingly, policies such as the U.S. Healthy Food Financing Initiative have been developed to subsidize and assist grocers in underserved areas.⁴ These policies were adopted despite limited causal evidence on the elasticity of healthy eating with respect to local availability of healthy foods. It is possible that higher costs of supplying produce and other healthy foods in low-income neighborhoods drive supply differences and significantly affect consumption. Furthermore, results from the Moving To Opportunity (MTO) experiment suggest that moving to lower-poverty neighborhoods reduces obesity (Kling, Liebman, and Katz, 2007; Ludwig et al., 2011). However, it is also possible that the observed supply differences are simply equilibrium responses to differences in demand, and the MTO obesity results could be driven by increased exercise and reduced stress instead of changes in eating patterns.

This paper combines reduced-form analyses with a structural demand model to quantify the relative importance of local supply versus demand in generating nutritional inequalities. Teasing apart supply- versus demand-side explanations is crucial for understanding whether and how policymakers should intervene. Policies like the Healthy Food Financing Initiative could increase both efficiency and equity if they address market failures, reduce costs, and affect purchases. But if the nutrition-income

¹Obesity is now one of the most important health problems in the U.S. and many other countries. Obesity is estimated to be responsible for 10-27 percent of U.S. medical costs, amounting to several hundred billion dollars annually (Finkelstein et al., 2009; Cawley et al., 2015). Especially for women, obesity rates are higher among those with lower income and education (Drewnowski and Specter, 2004): low-income women are 45 percent more likely than high-income women to be obese, and women who have not completed college are about 70 percent more likely than those who have (Ogden et al., 2010). It is well-established that eating and drinking patterns are a leading contributor to obesity and that these patterns differ across socioeconomic groups (Rehm et al., 2016; U.S. Department of Agriculture, 2014).

²See, 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).

³In 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.”

⁴The Healthy Food Financing Initiative has awarded \$220 million since 2011 in subsidies and technical assistance to grocery stores, farmers markets, and other suppliers of healthy foods in underserved areas (TRF 2017a). Pennsylvania’s Fresh Food Financing Initiative provided another \$85 million in similar grants and loans (TRF 2017b). Projects aimed at “eliminating food deserts” were eligible for the \$100 million in Community Transformation Grants under the Affordable Care Act. In the United Kingdom, the 2001 Food Poverty Eradication Bill required local and national governments to document and take actions to eliminate food deserts.

relationship is primarily driven by differences in demand, it is less clear that improving availability will have any effect on consumption, let alone improve social welfare.

We exploit a rich combination of datasets including Nielsen Homescan, a 60,000-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 that covers about 40 percent of all U.S. grocery purchases. We match the Homescan data to extensive surveys of panelists’ nutrition knowledge gathered by Nielsen for Allcott, Lockwood, and Taubinsky (2017). Finally, we gather data on the entry dates and locations of 1,914 new supermarkets from several different chains, 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.

We begin by illustrating two stylized facts using our Health Index, which scores grocery purchases based on U.S. government recommended daily intakes of healthy and unhealthy macronutrients such as fiber and sugar. First, there is a meaningful nutrition-income relationship: households with income over \$70,000 buy groceries that are 0.29 standard deviations healthier than households with income below \$25,000. Second, the UPCs available in RMS stores are considerably less healthy in low-income neighborhoods than in high-income neighborhoods. This difference is almost entirely explained by the fact that there are different types of stores in low-income neighborhoods—specifically, fewer supermarkets and more drug and convenience stores. This market structure underscores the hypothesis that the entry of new supermarkets might increase healthy eating in low-income neighborhoods.

We then use two event study designs to test whether the local environment has an economically significant effect on healthy eating. The first design looks within households, before versus after the entry of a new supermarket nearby. In both the full sample and in the sample of households in food deserts, we find economically small (and mostly statistically insignificant) effects of store entry on healthy eating. We show that 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 in zip codes with no supermarkets still buy almost 90 percent of their groceries from supermarkets. These additional results explain why local supermarket entry does not materially increase healthy eating, although it does make consumers better off by reducing travel costs. We can bound the effect of differential local access to supermarkets at no more than about five percent of the nutrition-income relationship. Thus, differences in local access to supermarkets offering healthy options do not appear to explain why the wealthy eat more healthfully than the poor.

One might also hypothesize a broader set of place-related factors, including peer effects or supply differences other than supermarket density, that could contribute to nutritional inequality. To test for this broader class of place-related factors, 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 do not converge toward eating patterns in the new location. Any endogeneity in moving decisions likely biases against this null result. 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 about three percent of the nutrition-income relationship.

These reduced-form analyses cannot entirely address the concerns of some advocates and policy-makers, who argue specifically that poor neighborhoods should have the same availability and prices of healthy foods as wealthy neighborhoods (e.g., Hilmers, Hilmers, and Dave, 2012). For example, adding a supermarket to a food desert changes the availability and prices of both healthy and unhealthy foods in ways that may not replicate the supply conditions of wealthy neighborhoods. Similarly, moving to a healthy area affects people’s lives in many ways other than just changing grocery supply. 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 simulate 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 macronutrients (saturated fat, sugar, salt, etc.). We depart from prior work in deriving 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 Shaws. Suppose Safeway is able to supply pizza cheaply, while Shaws can supply apples cheaply. Then, cities dominated by Safeway will have relatively low prices for pizza, while cities dominated by Shaws will have relatively low prices for apples. Our key identifying assumption is that these geographic differences in product group prices due to differences in the presence of specific chains are independent of unobserved differences in local product group tastes. The instrument has a very strong first stage, even after conditioning on category and market fixed effects, and it may be of broader use for research in other multi-product retail settings.

The estimates show a striking and systematic pattern between income and preferences for healthy and unhealthy nutrients. Higher-income households have monotonically stronger preferences for all four nutrients and food groups deemed “healthy” by U.S. government nutritional guidelines: fiber, protein, fruit, and vegetables. Higher-income households also have monotonically weaker preferences for two of the four “unhealthy” nutrients (saturated fat and sugar), while preferences for the two others (sodium and cholesterol) are statistically indistinguishable. The preference differences are economically significant: households with income below \$25,000 are willing to pay an average of \$0.62 per day to consume the U.S. recommended daily intake of healthy nutrients instead of the maximum daily intake of unhealthy nutrients, whereas households with income above \$70,000 are willing to pay almost twice that amount.

We use our demand model to simulate policies in which households with incomes below \$25,000 are exposed to the prices and product availability experienced by households with income above \$70,000.

Consistent with our event study results, we find only small effects of these supply-side factors on consumer 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. This evidence does not support the notion that eliminating food deserts would have material effects on nutritional inequality. Overall, these findings suggest that some of the existing supply-oriented policy initiatives discussed above will have limited effects on healthy eating in disadvantaged neighborhoods.

Having established that demand-side factors explain the bulk of the nutrition-income relationship, we then explore the potential underlying factors driving the differences in demand. In a final section, we use the structural model to isolate variation in healthy eating due to households' *demand* patterns, holding local supply conditions constant across households. We then project these demand differences onto household covariates, asking how much of the income-demand differences are explained by each covariate. We find that educational differences explain about twenty percent of the relationship between income and healthy grocery demand, while about seven percent is explained by the nutrition knowledge scores gathered by Allcott, Lockwood, and Taubinsky (2017).⁵ While these findings are only correlational, they may suggest an opportunity for policies focusing on demand-side factors, such as education and health knowledge, to reduce nutritional inequality.

Our results contribute to a growing literature on nutritional inequality. One particularly relevant paper is the important work by Handbury, Rakhovsky, and Schnell (2015; 2017). Our two papers overlap in two specific ways. First, both papers begin with stylized facts quantifying how purchases of and access to healthy groceries vary by socioeconomic status, which Handbury, Rakhovsky, and Schnell (2017) measure with education and we measure with income. The basic fact that purchases and access vary by income is already well-documented (e.g., Larson et al., 2009; Powell et al., 2007; Sharkey et al., 2010; U.S. Department of Agriculture, 2014), so we do not think of these stylized facts as a primary contribution of our paper.

Second, both papers study how Homescan households' healthy grocery purchases respond to local variation in supply of healthy groceries. However, Handbury, Rakhovsky, and Schnell (2017)'s empirical specifications and data on local supply are quite different than ours, and these differences matter a lot. Handbury, Rakhovsky, and Schnell (2017) find a strong within-household correlation between healthy grocery purchases and the supply of healthy groceries at nearby RMS stores. They do not determine whether this supply variation is generated by store entry and exit, within-store changes in stocking patterns, or changes in the set of stores included in the RMS sample. By contrast, we focus on the event study generated by supermarket entry, which we think of as a large, well-defined, and plausibly exogenous within-household shock. Handbury, Rakhovsky, and Schnell (2017) do also estimate entry event study regressions using a different dataset of new stores that enter the RMS sample. However, their entry estimates have a "lack of precision" (page 27), and they thus relegate these results to an appendix. The rest of our paper—including the migration event study, the structural model, and the decompositions of structurally estimated preferences—does not overlap with their important

⁵This result is consistent with Allcott, Lockwood, and Taubinsky (2017)'s finding that nutrition knowledge declines with income and is an important mediator of why sugary drink consumption declines with income.

work.

Our results also add to the findings from event studies of entry by 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. Wrigley et al., 2003; Cummins et al., 2005, 2015; Elbel et al., 2015). Our household migration event study adds a nutritional aspect to recent work using migration to understand the evolution of brand preferences (Bronnenberg, Dubé, and Gentzkow, 2012), the drivers of geographic variation in health care utilization (Finkelstein, Gentzkow, and Williams, 2016), and the caloric costs of culture (Atkin, 2016). 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 (e.g., Cutler and Lleras-Muney (2010), and see Grossman (2015) and Furnee, Groot, and van den Brink (2008) for reviews), 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. Put simply, our paper’s main contribution is to bring together a unified set of insights on why the wealthy and the poor eat differently in the U.S.

The remainder of the paper is organized as follows. Sections II through VIII, respectively, present data, stylized facts, reduced-form empirical analysis, demand model setup, demand model estimation and results, structural decomposition, and the conclusion.

II Data

II.A Nielsen Homescan and Retail Scanner Data

We use the Nielsen Homescan Panel for 2004-2015 to measure household grocery purchases.⁶ Homescan includes about 39,000 households each year for 2004-2006, and about 61,000 households each year for 2007-2015. 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 grocery purchases, not overall diets, because Homescan does not include data on food purchased in 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.⁸ Appendix Figure A1 shows that about 60 percent of

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

⁷The 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 healthful and unhealthful 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.

⁸The magnet data continue after 2006, but panelists now record only expenditures and not weights purchased. Because

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.⁹ Thus, we emphasize 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), presence of children, race, and the age, educational attainment, employment status, and weekly work hours for male and female household heads. The household income bins become more coarse at higher incomes: the three highest-income bins are \$60-70k, \$70-100k, and more than \$100k. In the usual case where there are two household heads, we use the mean of age, education, and employment variables observed for both heads. 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. In addition to the standard Homescan data, we observe self-reports of the importance of staying healthy and a detailed nutrition knowledge quiz from a Homescan PanelViews add-on survey carried out by Nielsen for Allcott, Lockwood, and Taubinsky (2017) in October 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 stores at 160 retail chains for 2006-2015. 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 gather zip code median income from the American Community Survey 2011 five-year estimates and county mean income from the Regional Economic Information System. We deflate prices 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 open dates to avoid attenuation bias. Some datasets, such as InfoUSA and the National Establishment Time Series, might be reasonable for cross-sectional analyses, but they do not sufficiently precisely record the open dates of new establishments; see Bitler and Haider (2011, page 162) for further discussion. Furthermore, to measure true changes in availability experienced by consumers, we must use actual new establishments, not store locations that continue to operate but change management.

In light of these issues, we measure entry with two datasets. First, for the period between January 2004 and December 2013, we gathered the exact store open dates and addresses for 1,914 large grocery

prices per unit weight can vary substantially across stores and neighborhoods, we do not use these data to construct food purchases.

⁹After excluding canned, frozen, and dried produce, about 39 percent of magnet households' fresh produce calories are from packaged items. Households with incomes less than about \$20,000 buy about five percentage points less of their fresh produce calories from packaged items, but the proportion is constant at moderate and high incomes.

stores and supercenters spanning multiple chains.¹⁰ Second, we 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 Health Index

We purchased UPC-level nutrition facts from a marketing data provider, and we gathered nutrition facts for non-packaged items from the U.S. Department of Agriculture (USDA) National Nutrient Database for Standard Reference (USDA 2014c). Panel C of Table 1 presents nutritional summary statistics across all UPC codes in the Nielsen Homescan and RMS data.

It will be useful to characterize goods and preferences using a one-dimensional index of healthfulness per calorie. Perhaps the most standard index in the literature is the USDA’s Healthy Eating Index (HEI), which was designed to score entire diets on an easily-understandable range from 0-100 (Guenther, Reedy, and Krebs-Smith, 2008; Volpe and Okrent, 2012). However, the HEI has sharp non-linearities because it values a balanced diet, so an item’s contribution to the HEI depends on the consumer’s full diet. This is inappropriate for datasets like Homescan and RMS where consumers’ full diets are not observed.

We therefore construct a modified version of the HEI that is based on the same U.S. government dietary recommendations but is linear and additively separable in macronutrients. The U.S. government Dietary Guidelines are organized around four “healthy” macronutrients and food groups to “increase” (fruits, vegetables, protein, and fiber) versus four “unhealthy” macronutrients to “reduce” (saturated fat, sugar, sodium, and cholesterol). Each healthy nutrient has a recommended daily intake (RDI), and each unhealthy nutrient has a maximum RDI. Our “raw Health Index” for product n is then mechanically generated from these official recommendations: it is the sum of healthy minus unhealthy nutrient contents per 1000 calories, weighting each by its recommended daily intake (RDI): $\tilde{H}_n = \sum_c G_c a_{nc}/r_c$, where a_{nc} is the grams of nutrient c per 1000 calories, r_c is the RDI for a normal adult, and G_c takes value 1 for “healthy” macronutrients and -1 for “unhealthy” nutrients. For example, the maximum RDI of sodium for a normal adult is 2.3 grams per day, and the RDI of fiber is 29.5 grams per day. A 1000-calorie UPC that contained 2.3 grams of sodium and no other “healthy” or “unhealthy” nutrients would have raw Health Index of -1; a 1000-calorie UPC that contained 2.3 grams of sodium, 29.5 grams of fiber, and no other “healthy” or “unhealthy” nutrients would have raw Health Index of 0. Appendix Table A1 details the full list of official increase/reduce recommendations and RDIs used to construct the Health Index.

In the 2004-2015 Homescan data, the mean household-level raw Health Index per 1000 calories purchased is $\mu_{\tilde{H}} = -2.77$, and the within-year standard deviation is $\sigma_{\tilde{H}} = 0.74$. To aid interpretability, we use a normalized Health Index H_n that is adjusted to be mean zero, standard deviation one across households: $H_n = \frac{\tilde{H}_n - \mu_{\tilde{H}}}{\sigma_{\tilde{H}}}$.

¹⁰The Nielsen data use agreement does not permit the release of the chains’ identities.

Collapsing to one single measure of healthy eating risks obscuring important results on other more specific measures. We will also present results using other measures, and for when we do focus on Health Index, Appendix Table A2 shows that the strongest correlates of Health Index in household-by-year Homescan data are purchases of sugar (with correlation coefficient $\rho \approx -0.75$), fiber ($\rho \approx 0.64$), protein ($\rho \approx 0.52$), and fruits and vegetables ($\rho \approx 0.33$ and 0.34).

III Stylized Facts: Purchases and Supply of Healthful Foods

III.A The Nutrition-Income Relationship

Figure 1 presents a map of the estimated Health Index of packaged grocery purchases by county, using the 2006-2015 RMS data.¹¹ The significant variation across geographies raises the question of whether moving to a different county or neighborhood might affect demand, or whether this geographic variation simply reflects sorting of people with similar preferences. The county-level Health Index is highly correlated with county mean income (correlation coefficient $\rho \approx 0.53$) and with Chetty et al. (2016)’s county-level life expectancy measure ($\rho \approx 0.53$), underscoring both the inequities and the potential implications of nutritional decisions.

At the more micro level of the household, the basic result that eating patterns vary by income has been well-documented through the NHANES survey (U.S. Department of Agriculture, 2014), which began in the 1960s. Using the 2004-2015 Homescan panel data, Figure 2 presents four different measures of the nutrition-income relationship: grams of sugar per 1000 calories purchased, share of bread calories from whole grain breads, share of total calories from (packaged) produce, and normalized Health Index across all grocery purchases. These four measures paint a consistent picture: low-income households purchase more sugar, less whole-grain breads, less produce, and lower-Health Index groceries.¹² Interestingly, the relationship between Health Index and income has a steeper slope at incomes below \$20,000. This finding is reminiscent of the Chetty et al. (2016) finding of a sharper drop in life expectancy for people in the lowest income percentiles.

Households with reported (nominal) income above \$70,000 (approximately the top 34 percentiles)

¹¹Since the RMS do not contain the complete census of stores, the distribution of store channel 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 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 RMS versus the true distribution of store channel types observed in Zip Code Business Patterns.

¹²Appendix Figure A2 presents analogues to Figure 2, considering each individual macronutrient. Consistent with NHANES, the most economically and statistically significant relationships in Homescan are that higher-income households tend to get a larger share of calories from protein and fiber, and a smaller share from sugar. Higher-income households also tend to consume more saturated fat, sodium, and cholesterol; although the difference is considerably smaller in percentage terms. Thus, although the Recommended Daily Intakes impose specific weights on macronutrients in our Health Index, higher-income diets would tend to be classified as “more healthy” unless the weights change substantially.

Appendix Figure A3 re-creates the figure using the magnet subsample for 2004-2006, which includes bulk purchases as well as packaged items. The results are qualitatively similar, although the income group differences are attenuated slightly because the nutrition-income relationship is less stark in 2004-2006 compared to the full 2004-2015 sample, as shown below in Figure 3.

buy groceries with a Health Index 0.29 standard deviations higher than households with incomes below \$25,000 (approximately the bottom 24 percentiles). Our main objective is to explain this 0.29 standard deviation difference.

Our 12-year sample window allows us to examine potential trends in the nutrition-income relationship. Figure 3 illustrates that the Health Index increased by 0.20 standard deviations for 2012-2015 relative to 2004-2007 for households with incomes above \$70,000. By contrast, the Health Index increased by only 0.04 standard deviations for households with incomes below \$25,000.¹³ This trend underscores the increasing relevance of the issues we study.¹⁴

To benchmark the potential importance of these differences, note that over the full 2004-2015 sample, households with income above \$70,000 purchase approximately one additional gram of fiber and 3.5 fewer grams of sugar per 1000 calories relative to households with income below \$25,000. Using correlational analysis, Montonen et al. (2003) find that consuming one additional gram of fiber per 1000 calories is conditionally associated with a 9.4 percent decrease in type-2 diabetes, and results from Yang et al. (2014) imply that 3.5 fewer grams of sugar per 1000 calories is conditionally associated with a ten percent decrease in death rates from cardiovascular disease.¹⁵

III.B Availability of Healthy Groceries in Low- versus High-Income Neighborhoods

It is often argued that the nutrition-income relationship is driven by low availability of healthful foods in low-income neighborhoods. We document the premise of this supply-side theory by measuring the healthfulness of the choice sets available in different neighborhoods. The four panels of Figure 4 present the relationship between neighborhood median income and the healthfulness of items offered in RMS stores, again measured by average sugar content, the share of bread UPCs that are whole grain, the share of all UPCs that are produce, and the mean Health Index of UPCs offered. Because this figure weights UPCs only by calories in the package and not by quantity sold, this figure reflects choice sets, not consumption. All four panels show the same qualitative result: stores in higher-income zip codes offer healthier items. Furthermore, the slope of the Health Index-income relationship is as strong or

¹³Appendix Figure A4 shows that wealthy households' differential increases in Health Index are driven by improvements in nearly all macronutrients that comprise the Health Index: relative to low-income households, high-income households have increased purchases of produce, protein, and fiber, and decreased purchases of sugar and sodium, over the 2004-2015 sample period. The one exception is cholesterol, an unhealthy macronutrient that high-income households purchase relatively more of later in the sample.

¹⁴These results on trends in the nutrition-income relationship confirm suggestive results from the NHANES data (Rehm et al., 2016), but since the Homescan sample is an order of magnitude larger, there can now be more confidence in this finding.

¹⁵We 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. If within-household income changes are exogenous and free of measurement error, the results would imply that short-term income changes explain 10-15 percent of the correlation between Health Index and household income. We do not place too much weight on this result, but it at least tentatively 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.

stronger for neighborhood choice sets (in Figure 4) as it is for household purchases (in Figure 2).¹⁶¹⁷

Appendix B presents several additional stylized facts that are relevant for the analyses below. First, Appendix B.B shows that six store channel type indicators (large grocery, small grocery, supercenter/club store, convenience store, drug store, and other mass merchant) explain 73 percent of the variation in Health Index across RMS stores. This is important, because it means that we can characterize a neighborhood retail environment largely by the types of stores that are present, without needing to analyze the specific UPCs available. Furthermore, we show that large grocery stores, supercenters, and club stores offer more produce UPCs and a higher average Health Index than convenience stores, drug stores, and other mass merchants. Indeed, supercenters such as Walmart Supercenter, Target, and Meijer by definition have full lines of grocery and produce items, as do club stores such as Sam’s Club and Costco. In the analysis below, we thus refer to large grocers, supercenters, and club stores as “supermarkets,” with the understanding that supermarkets carry more produce and more healthful items. We define a “food desert” as a zip code with no supermarkets.

Second, Appendix B.C shows that low-income neighborhoods tend to have fewer large grocery stores and more drug and convenience stores per capita. The ZBP data show that 24 percent of zip codes (weighted by population) are “food deserts.” By contrast, 55 percent of zip codes with median income below \$25,000 are food deserts. This lower concentration of healthy stores has generated support among policymakers for incentivizing entry of new supermarkets into low-income neighborhoods. We will estimate the effects of this type of entry in Section IV.A.

Third, Appendix B.D shows 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.

IV Event Studies: Retailer Entry and Household Moves

In the previous section, we documented that low-income neighborhoods have less local supply of healthy foods, and low-income households consume less healthy groceries. Do the differences in supply cause the differences in consumption? The key identification challenge is the possibility that household preferences differ systematically by income. If that is the case, then 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 unmeasured additional factors that systematically affect both supply

¹⁶This pattern of less-healthy choice sets and store types in low-income neighborhood stores is broadly consistent with a large public health literature (e.g., Larson et al., 2009; Sharkey et al., 2010; Powell et al., 2007). Interestingly, Handbury, Rahkovsky, and Schnell (2017) find qualitatively different results: they emphasize that the differences in local availability of healthy groceries in high- versus low-education neighborhoods are small relative to the differences in purchases between such neighborhoods.

¹⁷Unlike the nutrition-income relationship in consumption from Figure 2, the weights in the Health Index could matter for the conclusion that choice sets are more healthful in higher-income neighborhoods. Appendix Figure A5 shows that while stores in higher-income neighborhoods offer UPCs with more protein and fiber and less sugar per 1000 calories, they also offer more of three “bad” macronutrients, saturated fat, sodium, and cholesterol.

and demand in low- versus high-income neighborhoods.

We therefore use two event studies that look within-household across changes in local environments. First, we compare grocery purchases before and after the entry of a nearby supermarket. Second, we compare purchases before and after households move across zip codes and counties.

IV.A Effects of Supermarket Entry

We use household-by-quarter Homescan data in an event study framework to measure the effects of supermarket entry on grocery purchases. Using the google maps application program interface (API), we downloaded the driving time (assuming no congestion delay) between each Census tract centroid and the address of each of our 1,914 entering supermarkets. S_{dct} is the count of supermarket entries that have previously occurred within driving distance band d of Census tract c as of quarter t or earlier. We use two distance bands, $b \in [0, 10)$ minutes and $b \in [10, 15)$ minutes. 10 and 15 minutes are the median and 75th percentile of shopping travel times in the 2009 NHTS. Almost all households experience either zero or one entry in our data: for $b \in [0, 10)$ minutes, S_{bct} takes value 0 for 87 percent of observations, 1 for 11 percent of observations, 2 for 1.4 percent of observations, and 3, 4, or 5 for the remaining 0.3 percent. Since the set of households exposed to local entry are not nationally representative, we do not use the Homescan sample weights for this analysis.

Let \mathbf{X}_{it} denote the vector of potentially time-varying household covariates presented in Table 1: natural log of income, natural log of years of education, indicators for each integer age from 23-90, an indicator for the presence of children, race indicators, a married indicator, employment status, weekly work hours, and household daily calorie need. Let Y_{ict} denote an outcome for household i in tract c in quarter t .

We then run the following regression:

$$Y_{ict} = \sum_b \tau_b S_{bct} + \gamma \mathbf{X}_{it} + \mu_{dt} + \phi_{ic} + \varepsilon_{ict}, \quad (1)$$

where μ_{dt} is a vector of Census division-by-quarter of sample 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. The τ_b coefficients measure the effect of entry under the identifying assumption that store entry is exogenous to within-household preference changes over time. While retailers carefully plan entry and exit in response to local population growth and changes in local demographics, it is reasonable to assume that entry is uncorrelated with within-household demand changes conditional on division-by-quarter fixed effects and household demographics. In this and all other regressions in this section, standard errors are clustered by household.

Before estimating Equation (1), we first show graphical results of the event study. We define E_{cbqt} as an indicator variable for whether one supermarket entered in distance band b of Census tract c , q quarters before quarter t . B_{ibt} is an indicator variable for whether observation it is part of a balanced panel around one supermarket entry in distance band b : the household is observed in the same Census

tract continuously for all four quarters before and all eight quarters after one supermarket entry, with no other entries in that distance band during that window. We run the following regression:

$$Y_{ict} = \sum_b B_{ibt} \left(\sum_q \tau_{bq} E_{cbqt} + \alpha_b \right) + \gamma \mathbf{X}_{it} + \mu_{dt} + \phi_{ic} + \varepsilon_{ict}. \quad (2)$$

Figure 5 presents the $\tau_{[0,10)q}$ coefficients and 95 percent confidence intervals, illustrating how one supermarket entry within ten minutes of household i affects purchases in event time, using the balanced panel around the entry. The omitted category is $q = -1$, so all coefficients are relative to the outcome in the last quarter before entry.

The top two panels present expenditure shares (in units of percentage points), showing that entry clearly affects purchasing patterns. The dependent variable in the top left panel is the combined expenditure share across the several retail chains we observe in our entry dataset. Expenditures at entering retailers eventually increase by about six percentage points, and most of the adjustment occurs within the first four quarters after entry.¹⁸ Note that these effects on the combined expenditures at the retailers in our entry dataset understate the expenditure share increase at the specific entering store, because that entrant store also diverts expenditures from other stores owned by the retailers in our entry dataset.

The top right figure shows combined expenditures at grocery stores, supercenters, and club stores, which as we have shown above tend to offer a wider variety of produce and healthier items overall. To the extent that supermarket entry simply diverts sales from these store types, the actual changes in healthy product availability that shoppers experience—and thus the possible effects on healthful purchases—are more limited. Indeed, the combined expenditure share at supermarkets increases by only one percentage point. Thus, the primary effect of supermarket entry is to divert sales from other supermarkets.¹⁹

The bottom two panels present results with the Health Index of purchased groceries as the dependent variable. The bottom left panel uses the full sample, while the bottom right panel restricts the sample to the 25 percent of the sample that lives in “food deserts”—that is zip codes with no supercenters, club stores, or large (>50 employee) grocery stores as of 2003. Both figures show no increase in healthy purchases after supermarket entry. In any given quarter, we can reject Health Index increases of more than about 0.04 standard deviations with 95 percent confidence. Appendix Figure A8 presents the analogous figures for the $\tau_{[10,15)q}$ coefficients. The expenditure share changes are attenuated, as would be expected given that the entering stores are 10-15 minutes away instead of 0-10 minutes away, and there is again no apparent change in Health Index.²⁰

¹⁸This gradual adjustment of purchases is consistent with the results of Atkin, Faber, and Gonzalez-Navarro (2016), who study retail expansion in Mexico.

¹⁹These results are consistent with those of Hwang and Park (2016), who look at a subset of Walmart Supercenters that opened between 2003 and 2006.

²⁰There are two slight anomalies in Figure 5. First, in the top left panel, entrant expenditure share is slightly but statistically significantly different between the fourth quarter before entry and the last quarter before entry. Second, in the top right panel, expenditures are slightly but statistically significantly lower in the quarter of entry than they are in the quarter before entry. There are various possible explanations. These anomalies are small relative to the eventual expenditure share changes, and they do not appear in the estimates for distance band $b \in [10, 15)$ minutes in Appendix

Table 2 presents estimates of Equation (1), using the same dependent variables as Figure 5. Panel A considers the effects on expenditure shares, first at the entrant retailers and then at all grocery stores, supercenters, and club stores. Unsurprisingly, all effects are significantly larger for stores entering within ten 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 low-income households, and columns 5 and 6 limit to households in “food deserts.” The expenditure share changes are generally larger for low-income households and households in food deserts. However, consistent with Figure 5, 75 to 85 percent of entrant chains’ expenditure share increase consists of diverted sales from other grocery stores, supercenters, and club stores, while less than one-quarter is diverted from the other store channel types that typically offer less healthy groceries. Appendix Table A5 shows that most of this diversion from less healthy channel types is from other mass merchants. Even in food deserts, one supermarket entry causes expenditure shares at drug and convenience stores to drop by only a fraction of a percentage point.

Panel B of Table 2 presents effects on Health Index, which again is normalized to standard deviation 1 across households. Consistent with Figure 5, entry has no statistically significant effect, with one exception: there is a statistically significant but economically small effect of supermarket entry 10-15 minutes away from households in food deserts. Appendix Table A5 repeats these estimates using three alternative definitions of food deserts; again, there is a tightly estimated zero effect of entry within 0-10 minutes and a statistically significant but economically small effect of entry 10-15 minutes away.

We can use these estimates to determine the share of the difference in Health Index between low-income and high-income households that can be explained by having more local supermarkets. From Section III.A, we know that households with income above \$70,000 buy groceries with Health Index 0.29 standard deviations higher than those with income below \$25,000. The upper bound of the 95 percent confidence interval from column 2 of Panel B implies that one supermarket entry increases Health Index by no more than 0.036 standard deviations for low-income households. Using Zip Code Business Patterns, we calculate that the same high-income Homescan households have an average of 2.4 supermarkets in their zip code, while the low-income households have an average of 2.0, for an average difference of 0.4 supermarkets. Thus, we can conclude that local access to supermarkets explains no more than $0.036 \times 0.4 / 0.29 \approx 4.8\%$ of the high- versus low-income difference in Health Index.

These regressions consider entry by only a limited set of retailers, which could reduce both power and external validity. Appendix C.B presents parallel estimates using zip code-by-year counts of large grocery stores, supercenters, and club stores from Zip Code Business Patterns. These estimates are closely consistent with the results in Table 2: even in food deserts, the main effect of supermarket entry is to divert expenditures from other supermarkets, so consumers’ choice sets are largely unchanged. Entry has tightly-estimated zero effects on Health Index for the full sample as well as the low-income and “food desert” subsamples. In sum, differences in local access to supermarkets offering healthy options do not appear to be driving the nutrition-income relationship in household purchases.

Figure A8.

IV.A.1 Why Doesn't Entry Matter?

Especially given the academic and policy attention to local access to healthy grocery options, it is remarkable that supermarket entry seems to matter so little. We now document two key facts to help explain why this is the case.

First, Americans travel long distances to shop, typically in cars. This fact can be seen with the 2009 National Household Travel Survey (NHTS), a nationally-representative survey that gathers demographics, vehicle ownership, and “trip diaries” from 150,000 households. Figure 6 shows average one-way distances for trips beginning or ending in “buying goods: groceries/clothing/hardware store.” The mean trip is 5.2 miles, the median trip is 3.0 miles, and 90 percent of shopping trips are by car. People with household income less than \$25,000 (labeled as “poor” on the figure) travel a mean of 4.8 miles. People who live in “food deserts,” again defined as zip codes with no large grocery stores, supercenters, or club stores at the time of the NHTS survey, travel almost seven miles on average.

The Homescan sample is limited to urban areas, where travel distances may be shorter. The fourth through sixth bars in Figure 6 thus exclude rural areas (Census places with population less than 2500), showing that mean travel distances are still 4.5 miles, and 5.0 miles in urban food deserts. Finally, low-income households who live in food deserts and do not own a car, a disadvantaged group that would be heavily affected by lack of local access to healthy groceries, represent only 0.6 percent of the population and travel a mean of 2.0 miles. These long travel distances suggest that people may still shop in supermarkets even if they don't have a supermarket nearby. Appendix Figure A10 presents median travel distances and the share of trips by auto for the same subgroups.

Second, Figure 7 shows that low-income households and households in food deserts spend only slightly less in supermarkets.²¹ Households with income below \$25,000 spend about 87 percent of their grocery dollars at supermarkets, while households with incomes above \$70,000 spend 91 percent. For households in our “food deserts,” the supermarket expenditure share is only a fraction of a percentage point lower.²²

Of course, households in food deserts do benefit from more variety and reduced travel costs when a new supermarket enters nearby. However, these results show that because most consumers already travel to shop in supermarkets, supermarket entry does not significantly change choice sets, and thus doesn't affect healthy eating.

IV.B “Place Effects” Identified by Movers

While we have shown that supermarket entry has no 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.

²¹This is consistent with Broda, Leibtag, and Weinstein (2009), who report relatively small expenditure share differences across income groups.

²²Appendix Figure A9 presents expenditure shares by income category for all channel types.

To test for place effects, we measure the within-household changes in grocery purchases of Homescan households that move during our sample period. Homescan panelists that move typically report their moves to Nielsen within a few days or weeks of moving. We observe each of the household’s grocery shopping trips, along with the county of the visited store for those stores in the RMS data. Nielsen also reports the household’s county of residence as of the end of every year. In the end-of-year data, Homescan households change zip codes 17,956 times between 2004 and 2015, and they change counties 10,498 times.

We test whether household Health Index is associated with the (time-invariant) sample average Health Index of groceries purchased in the household’s current geographic location, conditional on household fixed effects. By conditioning on household fixed effects, we isolate the grocery purchase changes 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.²³

Define H_m as the estimated local Health Index of packaged groceries purchased in geographic area m , where m will be either a zip code or a county.²⁴ For county-level H_m , we use the same estimates mapped in Figure 1, and we use the same approach to estimate H_m by zip code.²⁵ To ensure that H_m reflects the actual local environment that each household faces, we drop the 15 percent of observations where less than 50 percent of trips to RMS stores are not in the household’s end-of-year county of residence. For these regressions, μ_t represents year indicators, ϕ_i is a household fixed effect, and \mathbf{X}_{it} is again the vector of potentially time-varying household covariates described in Table 1. We estimate the following regression in household-by-year Homescan data:

$$Y_{imt} = \tau H_m + \gamma \mathbf{X}_{it} + \mu_t + \phi_i + \varepsilon_{imt}. \quad (3)$$

Since the set of movers are not nationally representative, we again do not use the Homescan sample weights for this analysis.

Before estimating Equation (3), we first show graphical results of the event study. B_{it} is an indicator for whether observation it is part of a balanced sample around a move: we observe a household in four consecutive years indexed by y , beginning with $\underline{y} = -1$ (the year before the move) and ending in $\bar{y} = 2$ (two years after the move), with at least 50 percent of trips to RMS stores being in the household’s end-of-year county of residence in all three years other than the year of the move. The balanced panels for $\underline{y} = -1$ and $\bar{y} = 2$ include 2,548 cross-zip moves and 2,019 cross-county moves. We let H_f and H_o , respectively, denote the average Health Index of grocery purchases in the final and original locations, respectively, and we define $\Delta_{fo} = H_f - H_o$.

²³One way to study longer-term trends is to exploit retrospective data on where Homescan panelists previously lived, as in Bronnenberg, Dubé, and Gentzkow (2012). We have collected Homescan panelists’ lifetime histories of Census place of residence and could integrate these data into our analysis, although we have not yet done so because this would require careful consideration of endogeneity concerns that are less relevant for Bronnenberg, Dubé, and Gentzkow (2012)’s study of brand choice.

²⁴We have also run analyses that isolate supply differences using the Health Index of groceries *available* in area m , and the results are similar.

²⁵The average Homescan panelist lives in a zip code with 4.0 RMS stores and a county with 95 RMS stores.

For the graphical results, we estimate the following regression in household-by-year data:

$$Y_{it} = B_{it} \cdot \left(\alpha \Delta_{fo} + \sum_y (\tau_y \Delta_{fo} + \omega_y) \right) + \gamma \mathbf{X}_{it} + \mu_t + \phi_i + \varepsilon_{it}. \quad (4)$$

We estimate τ_y coefficients from year y to \bar{y} , with $y = -1$ as the omitted category. The ω_y coefficients are intercepts for each year, α measures the association between Health Index and the change in local environment in the year before the move ($y = -1$), and τ_y measures the differences in that association between $y = -1$ and each other year in the event study window. The interaction with B_{it} means that we identify these coefficients using only the households in the balanced sample, 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 8 presents the event study of cross-county moves. The top left panel shows the share of shopping trips to RMS stores that are in the new versus old county for households with $B_{it} = 1$. This illustrates that after excluding households that do less than half of their shopping in their end-of-year county of residence, almost 100 percent of trips are in the old county before the move, and almost 100 percent of trips are in the new county after the move. The top right panel presents the distribution across households of the local environment change Δ_{fo} , in units of normalized Health Index. The median cross-county mover experiences a local Health Index change of 0.15 standard deviations; this variation in Δ_{fo} is what identifies τ_y .

The bottom two panels show the estimated τ_y coefficients and 95 percent confidence intervals. The bottom left panel shows results excluding household demographics \mathbf{X}_{it} , while the bottom right panel includes \mathbf{X}_{it} . In both cases, there is no statistically significant post-move Health Index change associated with Δ_{fo} . In other words, there is no statistical evidence that households purchase more (less) healthy groceries when they move to counties where other households purchase more (less) healthy groceries. The 95 percent confidence intervals rule out post-move τ_y coefficients of larger than about 0.15 to 0.20. This means that when a household moves between counties where average healthy eating patterns differ by amount x , we can reject that the household’s own eating patterns change by $0.15x$ to $0.20x$ within the next two years.

Appendix C.C presents various alternative specifications, including a repeat of Figure 8 for cross-zip code moves and results for balanced panel windows that include more years before and after moves. There is no 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 ideal experiment to estimate the effects of place would consist of randomly assigning households to different neighborhoods. By contrast, households in our data move for reasons that may create endogeneity concerns. The ω_y parameters control for temporary changes in eating patterns due to the move itself or systematic differences between movers and non-movers. On the other hand, if some time-varying unobserved factor such as a better job or desire for a lifestyle change both causes healthier (or less healthy) eating and causes moves to a healthier place, the τ coefficients may not capture the causal effect of place. In Appendix Figure A13 and Appendix Table A9, for example, we show that

moving to healthier counties is statistically significantly associated with income increases, although moving to healthier zip codes is not.

We address this concern in two ways. First, we include controls for observable household demographics \mathbf{X}_{it} , which includes natural log of current-year income, natural log of years of education, indicators for each integer age from 23-90, an indicator for the presence of children, race indicators, a married indicator, employment status, weekly work hours, and household daily calorie need. This helps to control for observed changes in income, job responsibilities, household composition, and marriage status that could generate endogeneity. For both the graphical estimates of Equation (4) and the regression tables from Equation (3), the τ coefficient estimates are very similar regardless of whether \mathbf{X}_{it} is included or excluded. However, adding \mathbf{X}_{it} only slightly increases the regression R^2 , so unobserved within-household changes could be relevant (Oster, 2016).

Second, we assume that any remaining endogeneity would bias τ_y upward. This is the natural direction of bias for unobserved lifestyle changes that cause people to move to healthier places and also cause healthier eating, or unreported salary increases that cause moves to higher-income (and healthier) places and also allow healthier eating. Under this direction-of-bias assumption, our $\hat{\tau}$ is an upper bound on the causal place effect, which biases against our finding of no place effects.

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 1-2 because H_n is missing for zip codes with no RMS stores. In all four columns, $\hat{\tau}$ is both statistically and economically insignificant. At the 5 percent significance level, we can reject values of τ greater than about 0.02 in columns 1-2, and about 0.08 in columns 3-4. Thus, moving to a place with x units higher Health Index is associated with less than a $0.02x$ to $0.08x$ increase in a household’s Health Index within the next few years. Including controls for household demographics \mathbf{X}_{it} has very little impact on the results.

Our tight bound on the medium-term effects of place on healthy eating contrasts with the strong immediate impact of migration on health care utilization (Finkelstein, Gentzkow, and Williams, 2016) and on brand choice (Bronnenberg, Dubé, and Gentzkow, 2012). As an example of Bronnenberg, Dubé, and Gentzkow (2012)’s immediate brand choice effect, we estimate Equation (3) using county-level Coke market share for H_m , where “Coke market share” is Coke calories purchased/(Coke+Pepsi calories purchased). As shown in Appendix Table A10 we estimate a highly statistically significant $\hat{\tau} \approx 0.14$: moving to a county with (say) a 10 percentage point higher Coke market share is associated with about a 1.4 percentage point increase in the share of household Coke+Pepsi purchases that are Coke. Thus, using the same regressions with cross-county moves, we can rule out a place effect on healthy eating of anything more than about 56 percent the size of the place effect on Coke/Pepsi brand choice.

Using the above results, we can bound the extent to which location explains the nutrition-income relationship. We consider a partial equilibrium thought experiment where an individual household moves from a low-income health environment to a high-income health environment, leaving aside general equilibrium effects that would occur if this happened at large scale. The average household

with income above \$70,000 lives in a zip code (county) with a Health Index 0.13 (0.10) higher than households with incomes below \$20,000. The upper bounds of the confidence intervals on τ for zip codes (counties) from Table 3 are about 0.02 (0.08), and the difference between the high and low-income Health Index is 0.29 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 explain no more than $0.13 \times 0.02 / 0.29 \approx 0.75\%$ of the high- versus low-income difference in Health Index using cross-zip code moves, and no more than $0.10 \times 0.08 / 0.29 \approx 3\%$ using cross-county moves.

V A Model of Nutrient Demand

From the model-free analyses in the previous sections, we know that supermarket entry and moving to areas with healthier eating patterns both have limited effects on households' healthy eating patterns. While compelling, these event studies are conceptually imprecise because store entry and moving change many dimensions of the consumption environment. For example, supermarket entry changes both pricing and access for healthy and unhealthy foods and need 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. To isolate more precisely the effects of changing *only* the supply of groceries from conditions in poor neighborhoods to conditions in wealthy neighborhoods, we need to estimate a structural model and simulate specific counterfactuals. This type of counterfactual exercise is the only way to evaluate the effects of equalizing the store and product choices between the rich and poor, while holding all other aspects of the local environment constant. In addition, the structural analysis allows us to isolate the role of preferences and explore potential correlates with preference heterogeneity across income groups, an exercise we carry about below in section VII.B.

We build our structural approach on the modeling 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 nutrients. Let $\mathbf{y} = (y_1, \dots, y_{\mathcal{N}})$ denote the quantities purchased (measured in calories) of each of the \mathcal{N} food products (UPCs) available across all the stores, and let $\mathbf{p} = (p_1, \dots, p_{\mathcal{N}})$ denote the corresponding prices paid per calorie. Let $\mathbf{\Psi} = (\Psi_1, \dots, \Psi_{\mathcal{N}})$ 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 $n = 1, \dots, \mathcal{N}$ products is characterized by C nutrient characteristics $\{a_{n1}, \dots, a_{nC}\}$.

Define the $\mathcal{N} \times C$ matrix $\mathbf{A} = \begin{Bmatrix} a_{11}, \dots, a_{1C} \\ \dots \\ a_{n1}, \dots, a_{nC} \end{Bmatrix}$, which measures the nutrient content (in kilograms)

per calorie for each of the C nutrients in each of the \mathcal{N} different goods. The $C \times 1$ vector $\mathbf{z} = \mathbf{A}'\mathbf{y}$ denotes the total nutrient consumption associated with the household's bundle of purchases.

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

$$\begin{aligned} & \max_{s \in \mathbb{P}(S), x, \mathbf{y}} U(x, \mathbf{z}, \mathbf{y}; \Theta, \Psi) - \alpha d(s) \\ & s.t. \\ & \sum_{n=1}^{\mathcal{N}} y_n p_n + x \leq w. \end{aligned} \tag{5}$$

We assume that the utility function $U(x, \mathbf{z}, \mathbf{y}; \Theta, \Psi)$ 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.A Indirect Utility and Store Choice

The household faces a different set of prices and available products depending on which set of stores it visits. Define \mathbf{p}_s as the price vector at a given set of stores s . Without loss of generality, suppose we partition all the products into $n = 1, \dots, N$ goods that the household purchases, and $n = N+1, \dots, \mathcal{N}$ goods that the household does not purchase. Thus, $y_n > 0$ ($n = 1, \dots, N$) and $y_n = 0$ ($n = N+1, \dots, \mathcal{N}$). Conditional on a set of stores visited, the optimal consumption for each of the $n = 1, \dots, N$ purchased goods, $y_{sn}^*(\mathbf{p}_s; \Theta, \Psi)$, satisfies the following system of first-order necessary conditions:

$$\sum_{c=1}^C a_{nc} \frac{\partial U(x^*, \mathbf{z}^*, \mathbf{y}^*; \Theta, \Psi)}{\partial z_c} - \frac{\partial U(x^*, \mathbf{z}^*, \mathbf{y}^*; \Theta, \Psi)}{\partial x} p_n + \frac{\partial U(x^*, \mathbf{z}^*, \mathbf{y}^*; \Theta, \Psi)}{\partial y_n} = 0 \quad , \quad n = 1, \dots, N. \tag{6}$$

The household's outside good expenditure is $x^* = w - \sum_{n=1}^N y_{sn}^*(\mathbf{p}_s; \Theta, \Psi) p_n$.

Let $\mathbf{y}_s^*(\mathbf{p}_s; \Theta, \Psi)$ denote the vector of demands conditional on visiting the set of stores s . The conditional indirect utility associated with visiting the set s is:

$$v_s(\mathbf{p}_s; \Theta, \Psi) = U(I - \mathbf{p}'_s \mathbf{y}_s^*(\mathbf{p}_s; \Theta, \Psi), \mathbf{A}' \mathbf{y}_s^*(\mathbf{p}_s; \Theta, \Psi), \mathbf{y}_s^*(\mathbf{p}_s; \Theta, \Psi); \Theta, \Psi) - \alpha d(s). \tag{7}$$

The household's optimal store choice problem can then be written as the following discrete choice problem:

$$s^*(\mathbf{p}_s; \Theta, \Psi) = \max_{s \in \mathbb{P}(S)} \{v_s(\mathbf{p}_s; \Theta, \Psi)\}_{\forall s \in \mathbb{P}(S)}. \tag{8}$$

Although we will not estimate store choice, we discuss its role in our ability to estimate demand below.

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 nutrients:

$$U(x, \mathbf{z}, \mathbf{y}; \Theta, \Psi) = \sum_{j=1}^J \mu_j \ln \left(\sum_{k=1}^{\mathcal{K}_j} \Psi_{kj} y_{kj}^{\theta_j} \right) + \sum_{c=1}^C \beta_c z_c + \lambda x. \quad (9)$$

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 $k = 1, \dots, \mathcal{K}_j$ 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 nutrient c , capturing both how the nutrient tastes immediately when eaten and the eventual health implications.

Let K_j denote the number of products that the household actually purchases in category j , where $\sum_j K_j = N$. The first-order conditions from Equation (6) can be summed over all products purchased in category j and re-written as follows:

$$\sum_{k=1}^{K_j} p_{kj} y_{kj}^* = \sum_{c=1}^C \frac{\beta_c}{\lambda} \sum_{k=1}^{K_j} a_{kjc} y_{kj}^* + \frac{\mu_j \theta_j}{\lambda}, \quad j = 1, \dots, J. \quad (10)$$

The term $\frac{\mu_j \theta_j}{\lambda}$ represents what the household would spend on product group j if products in that group had no nutritional value. The term $\sum_{c=1}^C \frac{\beta_c}{\lambda} \sum_{k=1}^{K_j} a_{kjc} y_{kj}^*$ captures the household's additional expenditures in group j due to the products' nutrient contents. 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 nutrients that it values from that category (a_{kjc} is larger for nutrients where β_c is more positive). 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 $i = 1, \dots, I$ index households, let $t = 1, \dots, T$ index years, and let τ index weeks in year t . The first-order conditions from Equation (10) can be aggregated across weeks to the household-by-year level:

$$\sum_{\tau \in t} \sum_{k=1}^{K_j} p_{kjt} y_{ikjt} = \sum_{c=1}^C \frac{\beta_c}{\lambda_i} \sum_{k=1}^{K_j} \sum_{\tau \in t} a_{kjc} y_{ikjt} + \sum_{\tau \in t} \frac{\mu_{ij\tau} \theta_{ij\tau}}{\lambda_i}. \quad (11)$$

Equation (11) illustrates the appeal of the model, in that it allows for estimation of nutrient 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 $Y_{ijt} = \sum_{\tau \in t} \sum_{k \in J} y_{ikj\tau}$. Define \tilde{p}_{ijt} and \tilde{a}_{ijct} , respectively, as the calorie-weighted average price paid and average amount of nutrient c for household i 's purchases in category j in year t . Define $\tilde{\delta}_{ijt} = \sum_{\tau \in t} \frac{\mu_{ij\tau} \theta_{ij\tau}}{\lambda_i}$ as the strength of household i 's preferences for category j in year t . Finally, define $\tilde{\beta}_c = \frac{\beta_c}{\lambda_i}$, the money-metric marginal utility of each nutrient. Equation (11) can now be written more compactly as:

$$\tilde{p}_{ijt} Y_{ijt} = \sum_{c=1}^C \tilde{\beta}_c \tilde{a}_{ijct} Y_{ijt} + \tilde{\delta}_{ijt}. \quad (12)$$

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

Ideally, we could directly estimate the average nutrient preference parameters β_c using a version of Equation (12):

$$\tilde{p}_{ijt} Y_{ijt} = \sum_{c=2}^C \tilde{\beta}_c \tilde{a}_{ijct} Y_{ijt} + \epsilon_{ijt}. \quad (13)$$

However, the consumption of nutrient c , $\tilde{a}_{ijct} Y_{ijt}$, could be econometrically endogenous due to unobserved characteristics or preferences. One might therefore want to instrument for $\tilde{a}_{ijct} Y_{ijt}$ using an instrument such as variation in local product availability. The premise behind this regression would be that households change their group-level calorie purchases Y_{ijt} in response to local variation in the nutrient content of available goods within the product group. The magnitude of these responses would identify the strength of preferences for different nutrients. However, 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.

To resolve this mechanical endogeneity problem, we depart from the empirical strategy in Dubois, Griffith, and Nevo (2014). Instead of estimating equation (12) directly, we solve for total calories Y_{ijt} and then take logs of both sides, giving

$$\ln Y_{ijt} = -\ln \left(\tilde{p}_{ijt} - \sum_{c=2}^C \tilde{\beta}_{ic} \tilde{a}_{ijct} + \xi_i \right) + \ln \tilde{\delta}_{ijt}. \quad (14)$$

We separate the household’s category preferences $\ln \tilde{\delta}_{ijt}$ into a product group fixed effect δ_j , a local geographic market fixed effect ϕ_m , and the household-specific deviation ε_{ijt} , so $\ln \tilde{\delta}_{ijt} = \delta_j + \phi_m + \varepsilon_{ijt}$.²⁶ Our final estimating equation is thus

$$\ln Y_{ijt} = -\ln \left(\tilde{p}_{ijt} - \sum_{c=2}^C \tilde{\beta}_c \tilde{a}_{ijct} - \xi \right) + \delta_j + \phi_m + \varepsilon_{ijt}. \quad (15)$$

Equation (15) uses intuitive variation to estimate the preference parameters. The independent variable inside the parentheses is an “implicit price”: the actual price adjusted for the utility value of the nutrients in category j . As this implicit price increases, quantity purchased decreases. This equation also highlights a disadvantage of the Cobb-Douglas functional form: the group-level implicit price elasticity is restricted to unity, and there is no parameter that directly translates to a price elasticity of demand. Instead, 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} .

VI.B Identifying Nutrient Preferences

In Section VI.D, we detail a GMM estimator of Equation (15). A key moment condition for identifying nutrient preferences $\tilde{\beta}_c$ will be

$$E((\delta_j + \varepsilon_{ijt}) \tilde{a}_{ijct}) = 0. \quad (16)$$

This moment condition generates two types of variation with which to identify nutrient preferences $\tilde{\beta}_c$. First, there is variation *between* the $J = 45$ product groups: if consumers tend to buy more of product groups that have lots of sodium, we infer that $\tilde{\beta}_c$ for sodium is high. Second, there is variation across households *within* product groups: if, in years when consumers purchase especially salty products within a group, they also tend to buy more calories in that group, we also infer that $\tilde{\beta}_c$ for sodium is high.

The identifying assumption in Equation (16) has two types of economic implications, which parallel the two types of identifying variation described above. First, we must assume that product group preferences δ_j are uncorrelated with nutrient contents \tilde{a}_{ijct} . This assumption could be violated if, for example, people like foods that are more convenient to prepare, and convenience and salt content are correlated. In Appendix Table A11, we show that our $\tilde{\beta}_c$ for nutrients are largely unchanged when we explicitly include perishability and convenience as additional product characteristics.²⁷ Furthermore,

²⁶In practice, a local geographic market is a three-digit zip code.

²⁷It is unlikely that supply-side instrumental variables could be useful in specifically isolating nutrient preferences, due to the challenge in isolating variation in nutrient contents separately from other factors such as taste or shelf life.

our empirical results in Section VI.E below indicate a very systematic pattern of preferences across income groups, with higher-income households having stronger preferences for healthy nutrients and weaker preferences for unhealthy nutrients. While we certainly cannot rule out other unobserved confounders, it is difficult to imagine a set of confounders that could generate such a consistent pattern of results across many different nutrients.

The second economic implication is that 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 nutrients \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 the potential correlation between a household's idiosyncratic product group preferences ε_{ijt} and the household's average price paid \tilde{p}_{ijt} :

$$E(\varepsilon_{ijt}, \tilde{p}_{ijt}) \neq 0. \tag{17}$$

Price endogeneity could arise from both the demand and supply sides. On the demand side, recall that households choose stores endogenously using Equation (8). Households potentially shop at stores that have lower prices for product categories they prefer. Price endogeneity can also arise on the supply side due to standard simultaneity bias: retailers should set higher markups in response to higher demand. Price endogeneity can also arise if category preferences ε_{ijt} are correlated with demand for higher-quality (and thus higher-priced) products.

To address the possibility of price endogeneity, we instrument the prices using a new instrument. 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. Since 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 Shaws. Suppose Safeway is able to supply pizza cheaply, while Shaws can supply apples cheaply. Then, markets dominated by Safeway will face relatively low prices for pizza, while markets dominated by Shaws will face relatively low prices for apples.

We construct our instrument as follows. For retail chain r in market (i.e. three-digit zip code) 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 the total number of establishments of retailer r in market m , let N_{jrt} denote retailer r 's total nationwide sales in product group j , and let N_{kt} denote the total nationwide units sold 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 :

$$P_{jmt} = \frac{\sum_{r \in m} N_{rmt} N_{jrt} \cdot \sum_{k=1}^{\mathcal{K}_j} N_{kt} \Delta \ln(p_{krt,-m})}{\sum_{r \in m} N_{rmt} N_{jrt} \cdot \sum_{k=1}^{\mathcal{K}_j} N_{kt}}. \quad (18)$$

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

$$E(\varepsilon_{ijt} P_{jmt}) = 0. \quad (19)$$

The variation in our instrument reflects geographic variation in the presence of each chain across markets. Since our model includes product group and market fixed effects, our instrument relies on variation in the relative prices across product groups within a market.²⁸ The key assumption is that a chain with a comparative pricing advantage in supplying product group j does not base its entry decisions on a given market's tastes for j . We cannot explicitly rule out such an entry pattern. However, chains have diverse geographic networks, and it is unlikely that the match between a market's product group preferences and the chains' relative pricing advantages plays a large role in entry decisions.

Our conversations with people in 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.

To examine the variation in our instrument more closely, Appendix Figure A14 plots the standard deviation in our instrument for the six product "departments" (broadly aggregated product categories), after residualizing against year, market, and product group fixed effects. We see the largest variation in fresh produce, followed in order by dairy, frozen foods, packaged meats and deli items, and dry grocery. This is exactly what one might expect: since produce is grown more in certain areas of the country, transportation costs will 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 they are thus easy to transport.

The instrument is very powerful. Appendix Figure A15 shows that there is a robust linear rela-

²⁸DellaVigna 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.

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 of 83 to 177 in the four income groups.

This instrument is novel in the literature, and it can be used in situations where 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 is useful in identifying a store’s residual demand elasticity but 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 GMM Estimation

For estimation, we construct separate datasets for each of four household income groups, with nominal income cutoffs at \$25,000, \$50,000, and \$70,000.²⁹ 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 product 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.³⁰ For each income group, we estimate four parameter vectors: the $(C - 1) \times 1$ vector $\tilde{\beta}$ of preferences for observed nutrients, the scalar ξ representing the unobserved nutrient, the $J \times 1$ vector δ of product group fixed effects, and the $M \times 1$ vector ϕ of market fixed effects.

To specify the moment conditions, let \mathbf{D}_j be a $J \times 1$ vector of dummy variables for whether the observation is in category j , and let \mathbf{D}_m be an $M \times 1$ vector of dummy variables for whether household i is in market m . The model estimation relies on the following set of $(C + J + M)$ identifying moments:

$$\begin{aligned}
 E((\delta_j + \varepsilon_{ijt})\tilde{a}_{ijct}) &= 0 \quad , \quad c = 1, \dots, C \\
 E(\varepsilon_{ijt}P_{jmt}) &= 0 \\
 E(\varepsilon_{ijt}D_{ijt}) &= 0 \quad , \quad j = 1, \dots, J \\
 E(\varepsilon_{ijt}D_{im}) &= 0 \quad , \quad m = 1, \dots, M
 \end{aligned} \tag{20}$$

Loosely, the first set of moments identifies the $\tilde{\beta}$ parameters, the second identifies ξ , the third set identifies δ , and the fourth set identifies ϕ . We construct the sample analogue of the moment

²⁹We choose these cutoffs because these cutoffs give four groups of approximately equal sizes, within the constraints of the income ranges that Homescan panelists report.

³⁰10.6% of observations at the household-by-product group-by-year level have zero purchases and are thus dropped. “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 nutrient contents, suggesting that these dropped observations do not affect our estimated preference heterogeneity across income groups.

conditions from Equation (20) in the following $(C + J + M) \times 1$ vector: $\mathbf{g}_{ijt} = \mathbf{g}_{ijt}(\delta, \phi, \tilde{\beta}, \xi)$. We can then write the GMM estimator as:

$$\left(\hat{\delta}, \hat{\phi}, \hat{\beta}, \hat{\xi}\right) = \arg \min_{(\delta, \phi, \tilde{\beta}, \xi)} \left(\frac{1}{IJT} \sum_i \sum_j \sum_t \mathbf{g}_{ijt} \right)' \mathbf{W} \left(\frac{1}{IJT} \sum_i \sum_j \sum_t \mathbf{g}_{ijt} \right) \quad (21)$$

where \mathbf{W} is a $(C + J + M) \times (C + J + M)$ weight matrix. Appendix D.A presents details of the GMM estimator and its standard errors.

VI.E Estimation Results

Table 4 reports the estimated nutrient preference parameters $\hat{\beta}_c$ for each of the four income groups, in units of dollars per kilogram. We have nine observed characteristics: the eight macronutrients that enter the Health Index, plus unsaturated fat. Because fat, carbohydrates, and protein mechanically comprise all of calorie intake, we cannot include all three of those macronutrients as characteristics affecting total calorie demand. We exclude the nutrient carbohydrates and, accordingly, $\hat{\beta}_c$ measures the average willingness-to-pay (WTP) to have a kilogram of nutrient c instead of a kilogram of carbohydrate. This normalization also removes differences in λ , the marginal utility of a dollar, across income groups.

For each individual nutrient, the $\hat{\beta}_c$ parameters all have the same sign for all income groups. In other words, all income levels share common directional preferences for each nutrient relative to the omitted nutrient (carbohydrates). All four income groups strongly prefer saturated fat to carbohydrates, all dislike sodium and cholesterol relative to carbohydrates, and all have positive $\hat{\beta}_c$ coefficients for all four healthy nutrients.

While the signs of $\hat{\beta}_c$ are common across income groups, there is a striking pattern in the magnitudes. Higher-income households have monotonically stronger preferences for all four “healthy” nutrients and food groups (fiber, protein, fruit, and vegetables), and the differences between the lowest- and highest-income groups are statistically different. Point estimates indicate that higher-income households also have monotonically weaker preferences for three of the four unhealthy nutrients (saturated fat, sodium, and sugar), although preferences for cholesterol and sodium are statistically indistinguishable across income groups, and the point estimates for cholesterol are almost identical. The magnitudes of some preference differences are large: the highest-income group values fruit (relative to carbohydrate) nearly three times as much as the lowest-income group, and vegetables about twice as much. Higher-income households also purchase more of the unobserved nutrient, suggesting that higher-income households purchase higher-quality products.

Combining higher WTP for healthy nutrients with lower WTP for unhealthy nutrients, it is clear that high-income households have consistently stronger preferences for health. To summarize this result, column 12 in Table 4 reports the WTP to consume the recommended daily amount of “healthy” nutrients and food groups (protein, fiber, fruit, and vegetables) minus the maximum daily recommended amount of “unhealthy” nutrients (saturated fat, cholesterol, sodium, sugar). Specifically, if

$G_c = 1$ for “healthy” nutrients, $G_c = -1$ for “unhealthy” nutrients, and r_c is the official recommended daily intake of nutrient c , column 12 reports $\sum_c \hat{\beta}_c G_c r_c$. All income groups value healthy groceries, but the highest-income group is willing to pay the most, making health a normal good. The lowest-income group is willing to pay \$0.62 per day to consume the healthy bundle instead of the unhealthy bundle, while the highest-income group is willing to pay \$1.18 per day.

A potential concern is that omitted variables could generate the systematic monotonicity of nutrient preferences in income. While we cannot randomly assign nutrients, we can assess how much the point estimates change when adding additional covariates. As one key example, foods differ in perishability and preparation time, and these characteristics might be correlated with nutrients and might be valued differently by low- versus high-income households. To consider this concern, 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 A11 reports the model estimates including these two controls. Higher-income households more strongly value convenience, while all income groups prefer fresher foods with shorter shelf lives. The nutrient preferences across the income distribution are quite similar to the main model estimates. The lowest income group is willing to pay \$0.52 per day to consume the healthy bundle instead of the unhealthy bundle, while the highest-income group is willing to pay \$1.12 per day.

VII Explaining Nutritional Inequality

VII.A Decomposing Consumption Differences into Supply versus Demand Factors

Using the model estimates from section VI.E, we decompose the observed nutrition purchase differences observed across the income groups into underlying supply-side factors (prices and nutrient availability) and demand-side factors (preferences for product groups and nutrients). Since 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 by $g = 1, 2, 3, 4$. For a given set of prices \tilde{p}_{gj} , observed and unobserved nutrients \tilde{a}_{gjc} and ξ_g , product group preferences δ_{gj} , and nutrient preferences $\tilde{\beta}_{gc}$, we can predict \hat{Y}_{gj} , the total calories that income group g consumes within product group j :

$$\hat{Y}_{gj} = \frac{\exp(\delta_{gj})}{\tilde{p}_{gj} - \sum_{c=2}^C \tilde{\beta}_{gc} \tilde{a}_{gjc} - \xi_g}. \quad (22)$$

For simplicity, this equation excludes the market fixed effect ϕ_m . To first order, this 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 nutrients \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:

$$H_{gj} = \frac{\sum_k Q_{kj} \sum_s N_{ge} \mathbf{1}(kj \in e) H_{kj}}{\sum_k Q_{kj} \sum_s N_{ge} \mathbf{1}(kj \in e)}. \quad (23)$$

The representative price and observed nutrients 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:

$$\hat{H}_g = \sum_j \hat{Y}_{gj} H_{gj}. \quad (24)$$

For this subsection, we re-normalize the Health Index so that the initial difference between the highest and lowest income groups equals unity.

The left-most points on Figure 9 display each income group's initial Health Index level, calculated by substituting the predictions of Equation (22) into Equation (24). Interpretation of the rest of this figure is aided by the fact that this initial difference between the highest- and lowest-income groups is normalized to one.

We can now simulate the changes in healthy eating that would occur under different counterfactual scenarios for supply and demand parameters. 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, which is $g = 4$. Thus, for each income group g , we calculate the product group calorie demand as:

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

We then compute the corresponding Health Index by substituting the quantities in Equation (25) into Equation (24). Figure 9 shows that prices do not appear to explain much of the difference in Health Index across income groups. Recall from Section B.D that healthful foods are in fact slightly more expensive in the high-income groups. Accordingly, this equalization of prices actually increases the Health Index difference between the highest and lowest income groups by about two percent.

Our second counterfactual explores the part of the food desert discussion related to availability, by setting the nutrient 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:³¹

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

Figure 9 shows that equalizing the availability of nutrients decreases the Health Index difference between the highest and lowest income groups, but by only 11 percent relative to baseline. In combination, equalizing prices and availability decreases the difference by nine percent. These first two counterfactuals confirm our findings from Section IV that 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 nutrient availability, we now also set the nutrient preferences $\tilde{\beta}_{gc}$ in each income group to those of the highest-income group, generating calorie demand

$$\hat{Y}_{gj} = \frac{\exp(\delta_{gj})}{\tilde{p}_{4j} - \sum_{c=2}^C \tilde{\beta}_{4c} \tilde{a}_{4jc} - \xi_4}. \quad (27)$$

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 84 percent relative to baseline.

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

$$\hat{Y}_{gj} = \frac{\exp(\delta_{4j})}{\tilde{p}_{4j} - \sum_{c=2}^C \tilde{\beta}_{4c} \tilde{a}_{4jc} - \xi_4}. \quad (28)$$

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 ten percent is explained by the supply side. 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. We now explore some of the potential underlying factors

³¹The unobserved nutrient ξ represents a combination of amount of the unobserved nutrient and the income group’s preference for it. While this is a mix of supply and demand forces, we will attribute it to supply. This biases against our argument that supply doesn’t matter, and in any event, results are similar if we attribute ξ to demand by resimulating it with the $\tilde{\beta}_c$ parameters below.

driving the heterogeneous preferences for nutrients. We first analyze the observable characteristics that predict demand for healthy groceries. Then, we show which observables explain the correlation between income and healthy grocery demand.

Instead of simply analyzing equilibrium *purchases* of healthy groceries, we would like to isolate each household’s *demand* for healthy groceries, while holding supply conditions constant. To do this, we first compute the sample average prices and observed nutrients for each product group, denoted \bar{p}_j and \bar{a}_{jc} , and the sample average unobserved nutrient $\bar{\xi}$. As in the decomposition in Section VII.A, we interpret these three parameters as reflecting “supply.” Using the parameter estimates, we also back out each observation’s fitted error term, $\hat{\varepsilon}_{ijt}$, 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:

$$\hat{Y}_{ijt} = \frac{\exp(\hat{\delta}_{gj} + \hat{\phi}_{gm} + \hat{\varepsilon}_{ijt})}{\bar{p}_j - \sum_{c=2}^C \hat{\beta}_{gc} \bar{a}_{jc} - \bar{\xi}}. \quad (29)$$

Using these “demand-only” consumption predictions, we then compute the “demand-only” raw Health Index of resulting overall grocery consumption, using a formula analogous to the original Health Index defined in Section II.C:

$$\tilde{H}_{it}^D = \sum_j \sum_c \hat{Y}_{ijt} \frac{G_c \bar{a}_{jc}}{r_c}, \quad (30)$$

where r_c is the recommended daily intake for a normal adult and G_c takes value 1 for “healthy” macronutrients and -1 for “unhealthy” nutrients. We again normalize to mean zero, standard deviation one within the sample to get normalized demand-only Health Index H_{it}^D . This approach to specifically isolate and analyze correlates of “health demand” (separately from other supply-related factors that influence equilibrium outcomes) is a novel feature of this analysis.

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

$$\hat{H}_{it}^D = \alpha \ln w_{it} + \gamma \mathbf{X}_{it} + \mu_t + \varepsilon_i, \quad (31)$$

where w_{it} denotes household income, μ_t are year indicators, and \mathbf{X}_{it} now denotes other non-income household covariates. Specifically, \mathbf{X}_{it} contains demographic variables from the Homescan data that are conjectured to be associated with healthy eating, as well as two additional variables from the new Homescan add-on survey carried out by Nielsen for Allcott, Lockwood, and Taubinsky (2017): the self-reported importance of staying healthy, and the share of correct answers on 28 questions from the General Nutrition Knowledge Questionnaire (Kliemann et al., 2016).³² Both survey variables are

³²The 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. The General Nutrition Knowledge Questionnaire is standard in the public health literature. 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,” and the correct answer is “Plain yogurt.” The average score on the full questionnaire was 71 percent correct.

normalized into standard deviation units. 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 natural log years education is 0.558, meaning that at the sample mean education level, a one year increase in education is associated with a 0.04 standard deviation increase in demand for healthy groceries. A long literature has shown that education is highly predictive of health behaviors and outcomes; see Grossman (2015), Furnee et al. (2008), and ?’s analysis of the Homescan data. Our results additionally suggest that in the context of healthy eating, this correlation acts through demand, not through differences in local supply that might be correlated with education.

The results in column 1 also show that a one standard deviation increase in nutrition knowledge is associated with a 0.0440 standard deviation increase in demand for healthy groceries. In separate regressions where we exclude nutrition knowledge, the education coefficient rises by about 8.5 percent, suggesting that health knowledge is one mechanism through which education may influence health behavior. Similarly, Cutler and Lleras-Muney (2010) find that knowledge explains 10-20 percent of the relationship between education and drinking and smoking.

The second goal of this section is to determine which observables explain the relationship between income and demand for healthy groceries. To do this, we use the approach of Gelbach (2016), which is to conceptualize covariates \mathbf{X}_{it} as “omitted variables” in the relationship between health demand \hat{H}_{it}^D and household income $\ln w_{it}$, and then calculate the “omitted variables bias” from excluding each specific covariate. Column 2 presents the unconditional relationship between \hat{H}_{it}^D and $\ln w_{it}$, giving a coefficient we denote as $\tilde{\alpha}$. The income coefficient in column 1 is $0.0847/0.135 \approx 0.63$ of the coefficient in column 2. Therefore, these observables explain slightly more than one-third of the relationship between income and healthy grocery demand.

Column 3 of Table 5 presents estimates of seven separate auxiliary regressions, $\ln w_{it} = \Gamma X_{vit} + \epsilon_i$, which give $\hat{\Gamma}$, the covariance between natural log 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 variables 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:

$$\tilde{\pi}_v = \frac{\hat{\Gamma}_v \hat{\gamma}_v}{\tilde{\alpha}}. \quad (32)$$

As shown in Gelbach (2016), these contributions can be aggregated to jointly consider vectors of indicator variables, and we do this for the vector of age indicators and the vector of Census division indicators.³³ 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

³³Age has a clear non-monotonic relationship with Health Index, so we let age enter non-parametrically instead of linearly or log-linearly.

income, at about 21 percent. Nutrition knowledge explains the second-largest share, at about seven percent. These results are correlations, so they do not reflect the causal effect of additional education or nutrition knowledge interventions. At a minimum, these results are consistent with arguments that offering improved general education and health education to low-income families could play an important role in reducing nutritional inequality.

Census divisions are the fourth most important mediator. The explanatory power of the Census divisions was foreshadowed by Figure 1: the “south central” states (from Texas to Kentucky) have lower income and purchase less healthy groceries, while the coastal states have higher incomes and eat more healthfully. This suggests a separate role for culture and tastes that vary across regions of the United States—a counterpart to Atkin (2016)’s analysis of geographic taste variation in India, although the health costs in Atkin’s study come from consuming too few calories, while the possible health costs in our analysis come from consuming unhealthy kinds of calories.

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 some circles, this relationship has been taken as causal, with significant policy attention devoted to improving access to healthy groceries in low-income neighborhoods.

We test this hypothesis using several complementary empirical strategies. Entry of a new supermarket has a tightly estimated zero effect on healthy grocery purchases, and we can conclude that differential local access to supermarkets explains no more than about five 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 households 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 entire 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 “combating food deserts” derive less from healthy eating and more from reducing travel costs.

In our 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.

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 ten percent. After separating out supply variation, the descriptive correlations in our final section show that education and nutrition knowledge predict healthy grocery demand and explain non-negligible shares of the relationship between income and healthy grocery demand. For a policymaker who wants to help low-income families to eat more healthfully, the analyses in this paper suggest an opportunity for future research to explore the demand-side benefits of improving health education—if possible through effective interventions—rather than changing local supply.

References

- ACHESON, S. D. (1998): “Independent Inquiry into Inequalities in Health Report,” Tech. rep., London: The Stationery Office.
- AIZER, A. AND J. CURRIE (2014): “The intergenerational transmission of inequality: maternal disadvantage and health at birth,” *Science*, 344, 856–861.
- ALGERT, S. J., A. AGRAWAL, AND D. S. LEWIS (2006): “Disparities in Access to Fresh Produce in Low-Income Neighborhoods in Los Angeles,” *American Journal of Preventive Medicine*, 30, 365–370.
- ALLCOTT, H., B. LOCKWOOD, AND D. TAUBINSKY (2017): “A Theory of Regressive Sin Taxes, with an Application to the Optimal Soda Tax,” *NBER Working Paper 23085*.
- ALWITT, L. F. AND T. D. DONLEY (1997): “Retail Stores in Poor Urban Neighborhoods,” *The Journal of Consumer Affairs*, 31, 139–164.
- ANDERSON, M. AND D. MATSA (2011): “Are Restaurants Really Supersizing America?” *American Economic Journal: Applied Economics*, 3, 152–188.
- ATKIN, D. (2016): “The Caloric Costs of Culture: Evidence from Indian Migrants,” *American Economic Review*, 106, 1144–1181.
- ATKIN, D., B. FABER, AND M. GONZALEZ-NAVARRO (2016): “Retail Globalization and Household Welfare: Evidence from Mexico,” *Journal of Political Economy* (forthcoming).
- BAKER, E., M. SCHOOTMAN, E. BARNIDGE, AND C. KELLY (2006): “The role of race and poverty in access to foods that enable individuals to adhere to dietary guidelines,” *Preventing Chronic Disease*, 3, A76.
- BERRY, S. T. (1994): “Estimating Discrete-Choice Models of Product Differentiation,” *Rand Journal of Economics*, 25, 242–262.
- BITLER, M. AND S. HAIDER (2011): “An Economic View of Food Deserts in the United States,” *Journal of Policy Analysis and Management*, 30, 153–176.
- BRODA, C., E. LEIBTAG, AND D. WEINSTEIN (2009): “The Role of Prices in Measuring the Poor’s Living Standards,” *Journal of Economic Perspectives*, 23, 77–97.
- BRONNENBERG, B. J., J.-P. DUBÉ, AND M. GENTZKOW (2012): “The Evolution of Brand Preferences: Evidence from Consumer Migration,” *American Economic Review*, 102, 2472–2508.
- CASE, A. AND A. DEATON (2015): “Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century,” *Proceedings of the National Academy of Sciences*, 112, 15078–15083.

- CAWLEY, J., C. MEYERHOEFER, A. BIENER, M. HAMMER, AND N. WINTFELD (2015): “Savings in Medical Expenditures Associated with Reductions in Body Mass Index Among US Adults with Obesity, by Diabetes Status,” *PharmacoEconomics*.
- CHETTY, R., N. HENDREN, P. KLINE, AND E. SAEZ (2014): “Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States,” *Quarterly Journal of Economics*, 129, 1553–1623.
- CHETTY, R., M. STEPNER, S. ABRAHAM, S. LIN, B. SCUDERI, N. TURNER, A. BERGERON, AND D. CUTLER (2016): “The Association Between Income and Life Expectancy in the United States, 2001-2014,” *Journal of the American Medical Association*, 315, 1750–1766.
- COURTEMANCHE, C. J. AND A. CARDEN (2011): “Competing with Costco and Sam’s Club: Warehouse Club Entry and Grocery Prices,” *NBER Working Paper 17220*.
- CUMMINS, S., A. FINDLAY, M. PETTICREW, AND L. SPARKS (2005): “Healthy Cities: The Impact of Food Retail-led Regeneration on Food Access, Choice, and Retail Structure,” *Built Environment*, 31, 288–301.
- CUMMINS, S., E. FLINT, AND S. MATTHEWS (2015): “New Neighborhood Grocery Store Increased Awareness of Food Access But Did Not Alter Dietary Habits or Obesity,” *Health Affairs*, 33, 283–291.
- CURRIE, J., S. DELLAVIGNA, E. MORETTI, AND V. PATHANIA (2010): “The Effect of Fast Food Restaurants on Obesity and Weight Gain,” *American Economic Journal: Economic Policy*, 2, 32–63.
- CUTLER, D. AND A. LLERAS-MUNEY (2010): “Understanding Differences in Health Behaviors by Education,” *Journal of Health Economics*, 29, 1–28.
- DAVIS, B. AND C. CARPENTER (2009): “Proximity of Fast-Food Restaurants to Schools and Adolescent Obesity,” *American Journal of Public Health*, 99, 505–510.
- DELLAVIGNA, S. AND M. GENTZKOW (2017): “Uniform Pricing in US Retail Chains,” *Working Paper*.
- DREWNOWSKI, A. AND S. SPECTER (2004): “Poverty and obesity: the role of energy density and energy costs,” *American Journal of Clinical Nutrition*, 79, 6–16.
- DUBOIS, P., R. GRIFFITH, AND A. NEVO (2014): “Do Prices and Attributes Explain International Differences in Food Purchases?” *American Economic Review*, 2014, 832–867.
- DUNN, R. A. (2010): “The Effect of Fast-Food Availability on Obesity: An Analysis by Gender, Race, and Residential Location,” *American Journal of Agricultural Economics*, 92, 1149–1164.
- EINAV, L., E. LEIBTAG, AND A. NEVO (2010): “Recording Discrepancies in Nielsen Homescan Data: Are They Present and Do They Matter?” *Quantitative Marketing and Economics*, 8, 207–239.
- ELBEL, B., A. MORAN, L. B. DIXON, K. KISZKO, J. CANTOR, C. ABRAMS, AND T. MIJANOVICH (2015): “Assessment of a Government-Subsidized Supermarket in a High-Need Area on Household Food Availability and Children’s Dietary Intakes,” *Public Health Nutrition*, 26, 1–10.
- FINKELSTEIN, A., M. GENTZKOW, AND H. WILLIAMS (2016): “Sources of Geographic Variation in Health Care: Evidence From Patient Migration,” *Quarterly Journal of Economics*, 131, 1681–1726.
- FINKELSTEIN, E. A., J. G. TROGDON, J. W. COHEN, AND W. DIETZ (2009): “Annual Medical Spending Attributable to Obesity: Payer-And-Service-Specific Estimates,” *Health Affairs*, 28, w822–w831.
- FURNEE, C., W. GROOT, AND H. VAN DEN BRINK (2008): “The health effects of education: a meta-analysis,” *European Journal of Public Health*, 18, 417–421.
- GELBACH, J. B. (2016): “When Do Covariates Matter? And Which Ones, and How Much?” *Journal of Labor Economics*, 34, 509–543.

- GROSSMAN, G. (2015): “The Relationship Between Health and Schooling: What’s New?” *NBER Working Paper 21609*.
- GUENTHER, P., J. REEDY, AND S. KREBS-SMITH (2008): “Development of the Healthy Eating Index,” *Journal of the American Dietetic Association*, 108, 1896–1901.
- HANDBURY, J., I. RAHKOVSKY, AND M. SCHNELL (2015): “What Drives Nutritional Disparities? Retail Access and Food Purchases across the Socioeconomic Spectrum,” *National Bureau of Economic Research Working Paper 21126*.
- (2017): “Is the focus on food deserts fruitless? Retail access and food purchases across the socioeconomic spectrum,” *Working Paper, Wharton*.
- HAUSMAN, J. (1996): “Valuation of New Goods Under Perfect and Imperfect Competition,” in *The Economics of New Goods*, ed. by T. Bresnahan and R. J. Gordon, University of Chicago Press, 207–248.
- HILMERS, A., D. C. HILMERS, AND J. DAVE (2012): “Neighborhood Disparities in Access to Healthy Foods and Their Effects on Environmental Justice,” *American Journal of Public Health*, 102, 1644–1654.
- HOCH, S. J., B.-D. KIM, A. L. MONTGOMERY, AND P. E. ROSSI (1995): “Determinants of Store-Level Price Elasticity,” *Journal of Marketing Research*, 32, 17–29.
- HOROWITZ, C. R., K. A. COLSON, P. L. HEBERT, , AND K. LANCASTER (2004): “Barriers to Buying Healthy Foods for People With Diabetes: Evidence of Environmental Disparities,” *American Journal of Public Health*, 94, 1549–1554.
- HWANG, M. AND S. PARK (2016): “The Impact of Walmart Supercenter Conversion on Consumer Shopping Behavior,” *Management Science*, 62, 817–828.
- JETTER, K. M. AND D. L. CASSADY (2005): “The Availability and Cost of Healthier Food Alternatives,” *American Journal of Preventive Medicine*, 30, 38–44.
- KAUFMAN, P. R., J. M. MACDONALD, S. M. LUTZ, AND D. M. SMALLWOOD (1997): “Do the Poor Pay More for Food? Item Selection and Price Differences Affect Low-Income Household Food Costs,” Agricultural economic report no. 759, Food and Rural Economics Division, Economic Research Service, U.S. Department of Agriculture.
- KLIEMANN, N., J. WARDLE, F. JOHNSON, AND H. CROKER (2016): “Reliability and validity of a revised version of the General Nutrition Knowledge Questionnaire,” *European Journal of Clinical Nutrition*, 70, 1174–1180.
- KLING, J. R., J. B. LIEBMAN, AND L. F. KATZ (2007): “Experimental Analysis of Neighborhood Effects,” *Econometrica*, 75, 83–119.
- LARSON, N., M. STORY, AND M. NELSON (2009): “Neighborhood Environments: Disparities in Access to Healthy Foods in the U.S,” *American Journal of Preventive Medicine*, 36, 74–81.
- LIPMAN, M. (2012): “FTC May Waste Time Updating Price-Bias Guide, Attys Say,” .
- LUDWIG, J., L. SANBONMATSU, L. GENNETIAN, E. ADAM, G. J. DUNCAN, L. F. KATZ, R. C. KESSLER, J. R. KLING, S. T. LINDAU, R. C. WHITAKER, , AND T. W. MCDADE (2011): “Neighborhoods, Obesity, and Diabetes - A Randomized Social Experiment,” *New England Journal of Medicine*, 365, 1509–1519.
- MONTONEN, J., P. KNEKT, R. JARVINEN, A. AROMAA, AND A. REUNANEN (2003): “Whole-grain and fiber intake and the incidence of type 2 diabetes,” *American Journal of Clinical Nutrition*, 77, 622–629.
- OGDEN, C. L., M. M. LAMB, M. D. CARROLL, AND K. M. FLEGAL (2010): “Obesity and Socioeconomic Status in Adults: United States, 2005-2008,” <http://www.cdc.gov/nchs/data/databriefs/db50.pdf>.

- OKRENT, A. AND A. KUMCU (2016): “U.S. Households’ Demand for Convenience Foods,” *USDA Economic Research Report*.
- OSTER, E. (2016): “Unobservable Selection and Coefficient Stability: Theory and Evidence,” *Journal of Business and Economic Statistics (forthcoming)*.
- POWELL, L. M., S. SLATER, D. MIRTICHEVA, Y. BAO, AND F. J. CHALOUKKA (2007): “Food store availability and neighborhood characteristics in the United States,” *Preventive Medicine*, 44, 189–195.
- RAO, M., A. AFSHIN, G. SINGH, AND D. MOZAFFARIAN (2013): “Do Healthier Foods and Diet Patterns Cost More than Less Healthy Options? A Systematic Review and Meta-Analysis,” *BMJ Open*, 3:e004277.
- REHM, C. D., J. L. PENALVO, A. AFSHIN, AND D. MOZAFFARIAN (2016): “Dietary Intake Among US Adults, 1999-2012,” *Journal of the American Medical Association*, 315, 2542–2553.
- HHS (U.S. DEPARTMENT OF HEALTH AND HUMAN SERVICES) (2015): “FoodKeeper App,” <https://www.foodsafety.gov/keep/foodkeeperapp/index.html>.
- TRF (THE REINVESTMENT FUND) (2017a): “The Healthy Food Financing Initiative,” <http://www.healthyfoodaccess.org/resources/library/healthy-food-financing-initiative-hffi>.
- (2017b): “Pennsylvania Fresh Food Financing Initiative,” <https://www.reinvestment.com/success-story/pennsylvania-fresh-food-financing-initiative/>.
- USDA (U.S. DEPARTMENT OF AGRICULTURE, AGRICULTURAL RESEARCH SERVICE) (2014b): “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>.
- (2014c): “USDA National Nutrient Database for Standard Reference, Release 26,” <http://www.ars.usda.gov/ba/bhnrc/ndl>.
- SAEZ, E. AND T. PIKETTY (2003): “Income Inequality in the United States, 1913-1998,” *Quarterly Journal of Economics*, 118, 1–39.
- SHARKEY, J. R., S. HOREL, AND W. R. DEAN (2010): “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*, 9, 1–27.
- U.S. DEPARTMENT OF AGRICULTURE, A. R. S. (2014): *What We Eat in America, NHANES 2011-2012*, chap. Away from Home: Percentages of Selected Nutrients Contributed by Food and Beverages Consumed Away from Home, by Family Income (in Dollars) and Age.
- VOLPE, R. AND A. OKRENT (2012): “Assessing the Healthfulness of Consumers’ Grocery Purchases,” *USDA Economic Information Bulletin*.
- VOLPE, R., A. OKRENT, AND E. LEIBTAG (2013): “The Effect of Supercenter-Format Stores on the Healthfulness of Consumers’ Grocery Purchases,” *American Journal of Agricultural Economics*, 95, 568–589.
- WRIGLEY, N., D. WARM, AND B. MARGETTS (2003): “Deprivation, Diet, and Food Retail Access: Findings From the Leeds ‘Food Deserts’ Study,” *Environment and Planning A*, 35, 151–188.
- YANG, Q., Z. ZHANG, E. GREGG, D. FLANDERS, R. MERRITT, AND F. HU (2014): “Added Sugar Intake and Cardiovascular Diseases Mortality Among US Adults,” *Journal of the American Medical Association Internal Medicine*, 174, 516–524.

Tables

Table 1: **Descriptive Statistics**

Variable	Mean	Standard deviation
Panel A: Nielsen Homescan Households		
Household income (\$000s)	62.2	45.4
Years education	13.9	2.06
Age	52.2	14.4
Have children	0.33	0.47
White	0.77	0.42
Black	0.12	0.32
Married	0.49	0.50
Employed	0.61	0.44
Weekly work hours	21.0	15.0
Household daily calorie need	5282	2987
Health importance	0	1
Nutrition knowledge	0	1
Panel B: Zip Code Establishment Counts		
Grocery	1.67	4.03
Large grocery (>50 employees)	0.46	1.04
Supercenters/club stores	0.11	0.39
Drug stores	1.09	2.28
Convenience stores	3.13	5.23
Meat/fish/produce stores	0.28	0.96
Panel C: UPC Characteristics		
Calories	1204	3226
Grams per 1000 calories		
Fat	34.2	29.2
Saturated fat	11.2	13.2
Cholesterol	0.085	0.28
Sodium	9.85	225
Carbohydrates	153	84.4
Fiber	12.2	31.3
Sugar	73.6	76.0
Protein	31.8	38.5

Notes: Homescan data include 668,844 household-by-year observations for 2004-2015 and are weighted for national representativeness. Health importance and nutrition knowledge are from Homescan add-on surveys carried out by Nielsen for Allcott, Lockwood, and Taubinsky (2017); each is normalized to mean zero, standard deviation 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). Zip code establishment counts are from Zip Code Business Patterns data for 2004-2015, with 470,229 zip code-by-year observations. UPC characteristics are for all 1.57 million UPCs that ever appear in the Nielsen Homescan or RMS data.

Table 2: **Effects of Supermarket Entry**

	(1)	(2)	(3)	(4)	(5)	(6)
Sample:	Full Sample		Income < \$25,000		"Food Desert" Zip Codes	
Panel A: Effects on Expenditure Shares						
Expenditure shares at store type:	Entrants	Grocery/ Super/Club	Entrants	Grocery/ Super/Club	Entrants	Grocery/ Super/Club
Post entry: 0-10 minutes	3.782 (0.149)***	0.634 (0.0905)***	4.738 (0.460)***	1.191 (0.294)***	4.034 (0.445)***	0.842 (0.293)***
Post entry: 10-15 minutes	1.076 (0.0924)***	0.184 (0.0603)***	1.020 (0.275)***	0.172 (0.193)	1.392 (0.248)***	0.272 (0.157)*
Observations	2,627,947	2,627,947	404,868	404,868	640,498	640,498
Dependent var. mean	19	88	22	85	23	88
Panel B: Effects on Health Index						
	(1)	(2)	(3)			
Sample:	Full Sample		Income < \$25,000		"Food Desert" Zip Codes	
Post entry: 0-10 minutes	0.00538 (0.00506)	0.0103 (0.0131)	-0.00675 (0.0146)			
Post entry: 10-15 minutes	0.00285 (0.00354)	0.0115 (0.00950)	0.0265 (0.00921)***			
Observations	2,627,947	404,868	640,498			

Notes: This table uses 2004-2015 Nielsen Homescan data at the household-by-quarter level. "Food desert" zip codes are those with no grocery stores with 50 or more employees, supercenters, or club stores in 2003. 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 indicators for whether a specific retailer has 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, an indicator for whether the household includes children, race indicators, employment status, weekly work hours, and total calorie need), 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, 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	-0.0110 (0.0139)	-0.00885 (0.0138)		
County average Health Index			-0.0144 (0.0464)	-0.00990 (0.0461)
Household demographics	No	Yes	No	Yes
Observations	560,492	560,492	565,914	565,914
95% confidence interval upper bound	0.016	0.018	0.077	0.081

Notes: This table uses 2004-2015 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. Health Index is our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation one across households. Household demographics are natural log of income, natural log of years of education, age indicators, an indicator for whether the household includes children, race indicators, employment status, weekly work hours, and total calorie need. All regressions also control for year indicators and household fixed effects. 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.

Table 4: Preferences for Nutrients by Household Income

Income group	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Carbs	Unsat. Fat	Sat. Fat	Fiber	Protein	Sugar	Sodium	Cholest.	Fruit	Veg	Unobs. Nutrient	WTP for Health Index
$Income \leq \$25k$	-	-0.11 (0.11)	15.90 (0.19)	9.90 (0.19)	5.95 (0.20)	1.87 (0.06)	-29.78 (2020)	-13.45 (4.24)	0.43 (0.03)	0.38 (0.01)	-0.32 (0.02)	0.62 (0.02)
$\$25k < Income \leq \$50k$	-	-1.51 (0.03)	13.70 (0.04)	10.19 (0.12)	6.63 (0.13)	1.30 (0.02)	-29.93 (0.97)	-13.45 (2.38)	0.67 (0.01)	0.53 (0.004)	-0.11 (0.005)	0.86 (0.01)
$\$50k < Income \leq \$70k$	-	-1.95 (0.04)	13.03 (0.05)	10.29 (0.08)	6.83 (0.08)	1.16 (0.02)	-29.97 (0.69)	-13.45 (1.55)	0.84 (0.01)	0.63 (0.003)	-0.05 (0.005)	0.99 (0.01)
$\$70k < Income$	-	-2.30 (0.32)	12.50 (0.30)	10.38 (0.38)	7.32 (1.16)	1.06 (0.15)	-30.03 (7.62)	-13.45 (21.58)	1.11 (0.06)	0.74 (0.01)	0.00005 (0.02)	1.18 (0.06)

Notes: This table presents GMM estimates of the nutrient preference parameters $\hat{\beta}_c$ from Equation (15). Magnitudes represent willingness to pay for a kilogram of the nutrient instead of a kilogram of carbohydrates. Value of fruit and vegetables accounts for value over and beyond macronutrient characteristics of the fruit and vegetables. “WTP for Health Index” in column 12 equals $\sum_c \hat{\beta}_c G_c r_c$, where $G_c = 1$ for “healthy” nutrients, $G_c = -1$ for “unhealthy” nutrients, and r_c is the recommended daily intake of nutrient c detailed in Appendix Table A1. Standard errors, clustered by household, are in parentheses. *, **, ***: Statistically significant with 10, 5, and 1 percent confidence, respectively.

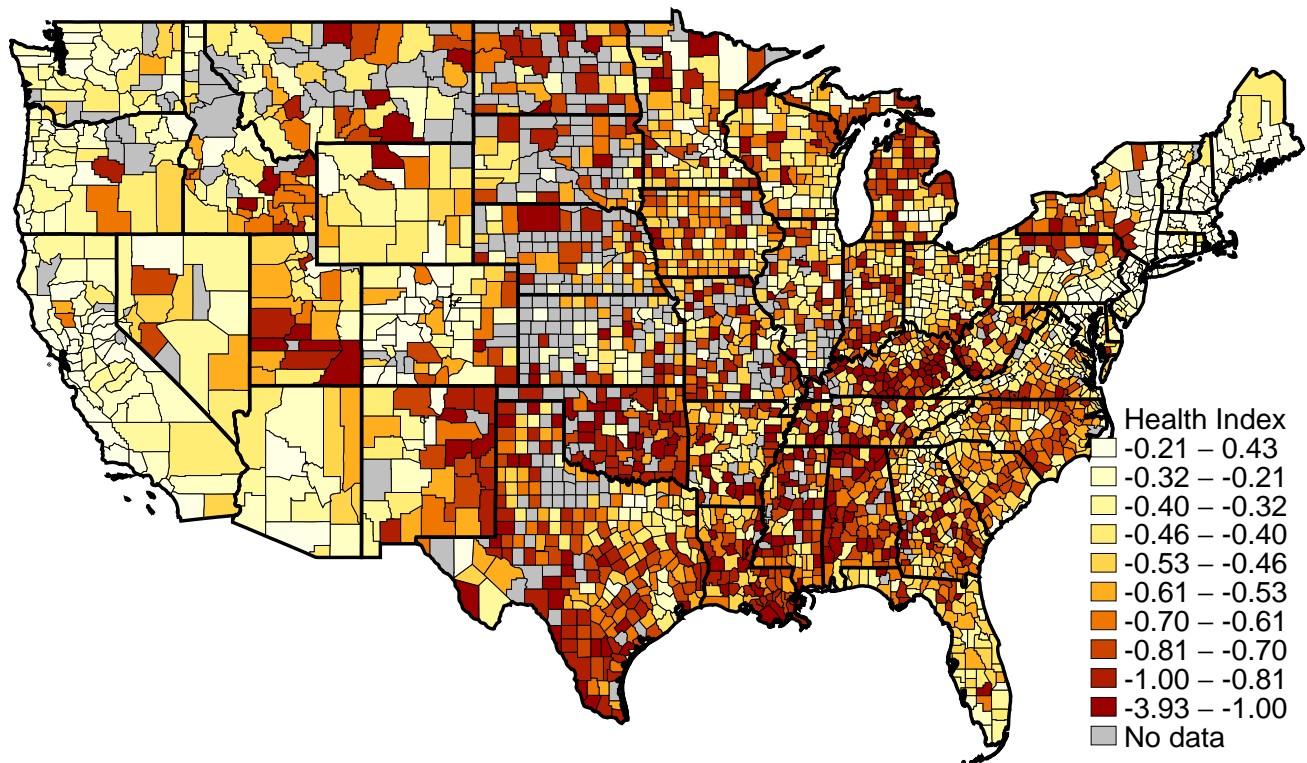
Table 5: **Decomposition of Nutrition-Income Relationship by Household Demographics**

	(1)	(2)	(3)
	Full Model	Unconditional Relationship	Auxiliary Regressions
ln(Household income)	0.0847 (0.0130)***	0.135 (0.0116)***	
ln(Years education)	0.558 (0.0745)***		1.527 (0.0378)***
1(Have children)	-0.0261 (0.0227)		0.216 (0.0130)***
1(White)	0.0245 (0.0423)		-0.0247 (0.0160)
1(Black)	-0.330 (0.0518)***		-0.0606 (0.0200)***
1(Married)	0.0178 (0.0199)		0.510 (0.0105)***
Health importance	0.0128 (0.0101)		0.0565 (0.00595)***
Nutrition knowledge	0.0447 (0.0111)***		0.125 (0.00616)***
Age indicators	Yes	No	
Census division indicators	Yes	No	
Year indicators	Yes	Yes	
Observations	117,941	117,941	117,941

Notes: These regressions use 2004-2015 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 (2017). 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 mean zero, standard deviation one. Columns 1 and 2 present estimates of Equation (31), a regression of the Health Index of demand-only consumption predictions on covariates. Each row of column 3 presents the coefficient from a univariate regression of natural log of household income on the variable listed in each row. 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.

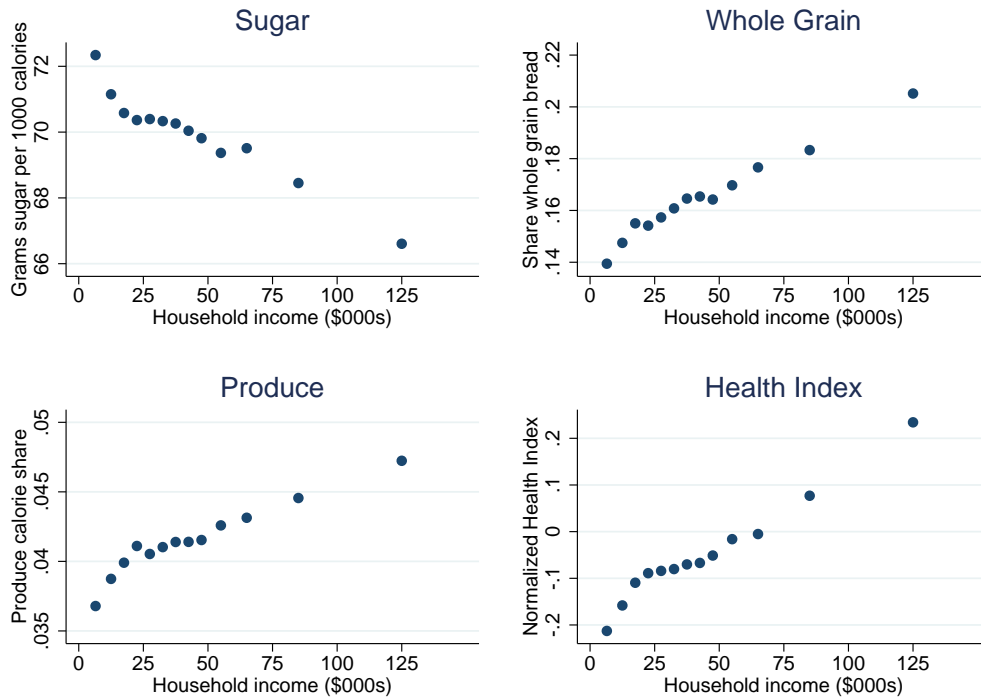
Figures

Figure 1: Average Health Index of Store Purchases by County



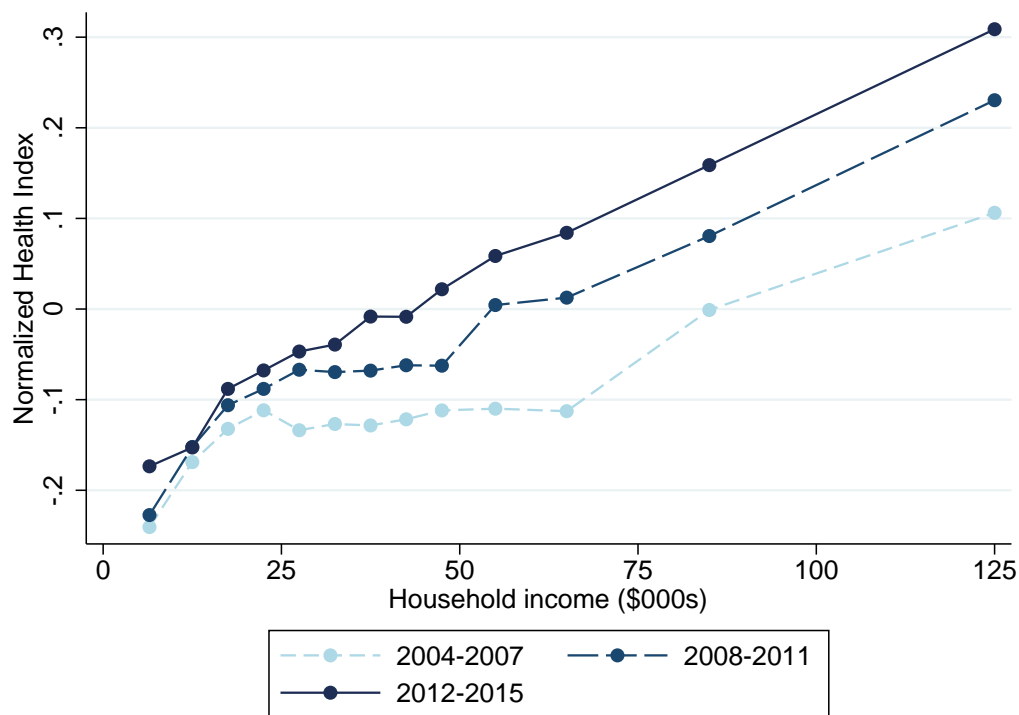
Notes: This figure presents the calorie-weighted average normalized Health Index of packaged grocery purchases by county, using 2006-2015 Nielsen RMS data. Health Index is our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation one across households. 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 2: **Healthfulness of Grocery Purchases by Household Income**



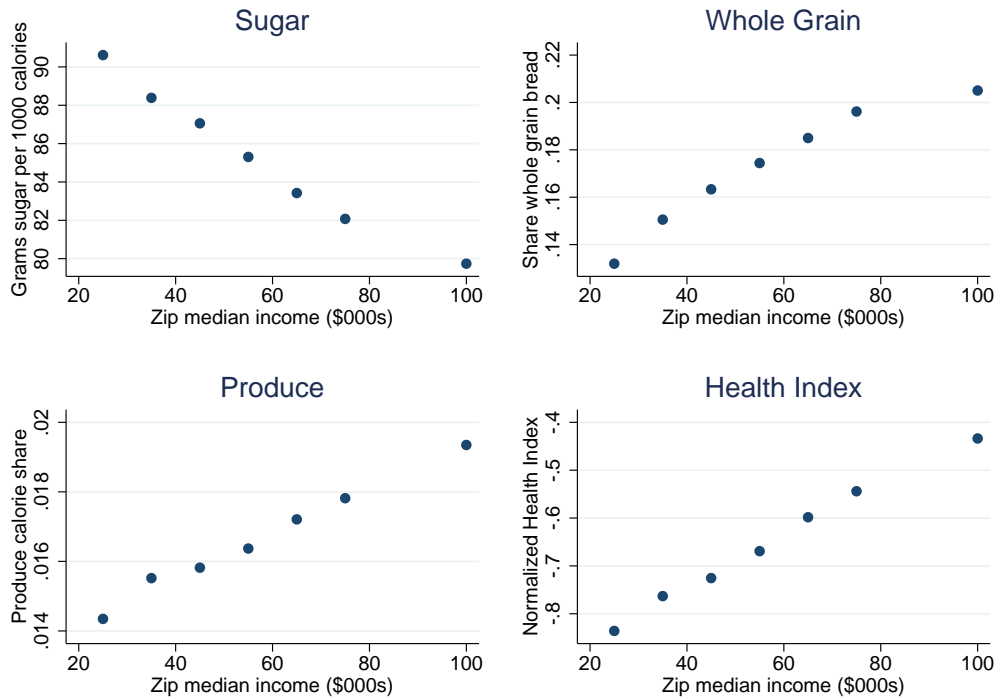
Notes: This figure presents Nielsen Homescan data for 2004-2015. The x-axis presents nominal income bins; household incomes larger than \$100,000 are coded as \$125,000. Sugar is the grams of sugar per 1000 calories purchased, whole grain is the calorie-weighted average share of bread, buns, and rolls purchases that are whole grain, produce is the share of calories from fresh, canned, dried, and frozen fruits and vegetables, and Health Index is our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation one across households. Observations are weighted for national representativeness.

Figure 3: Trends in Healthfulness of Grocery Purchases by Household Income



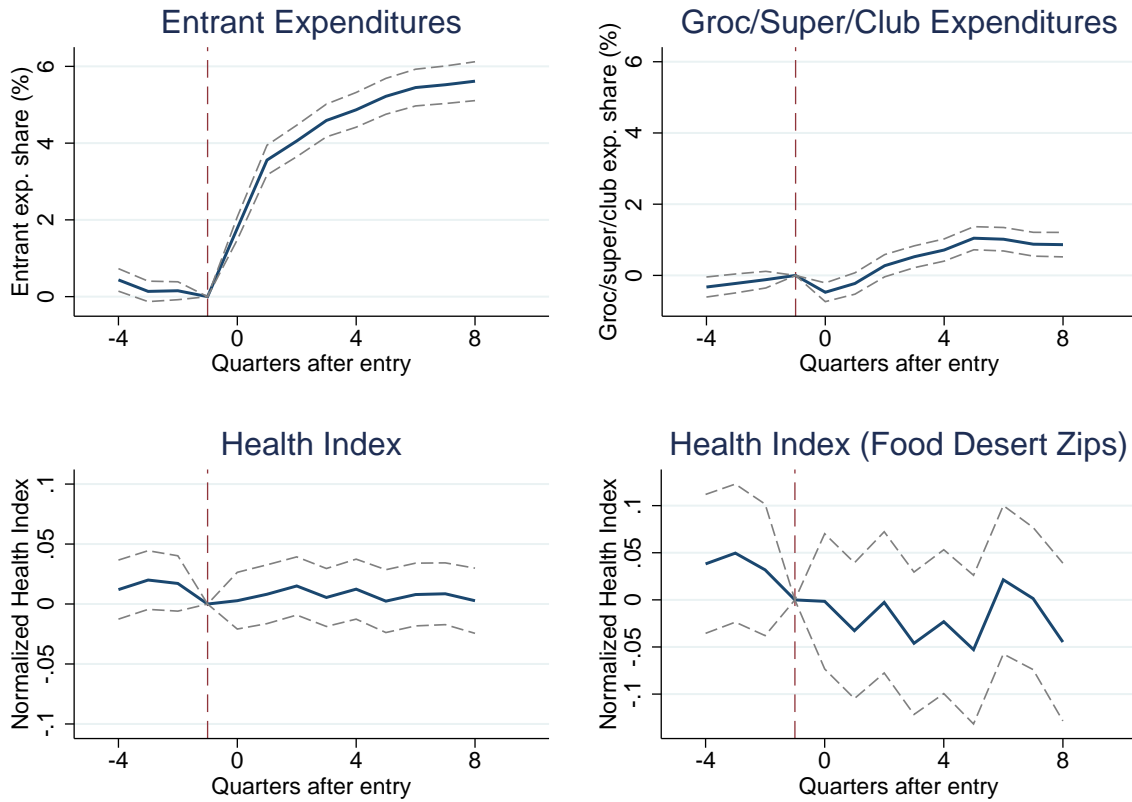
Notes: This figure presents Nielsen Homescan data for 2004-2015. The x-axis presents nominal income bins; household incomes larger than \$100,000 are coded as \$125,000. Health Index is our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation one across households. Observations are weighted for national representativeness.

Figure 4: **Store Average Healthfulness and Size by Zip Code Median Income**



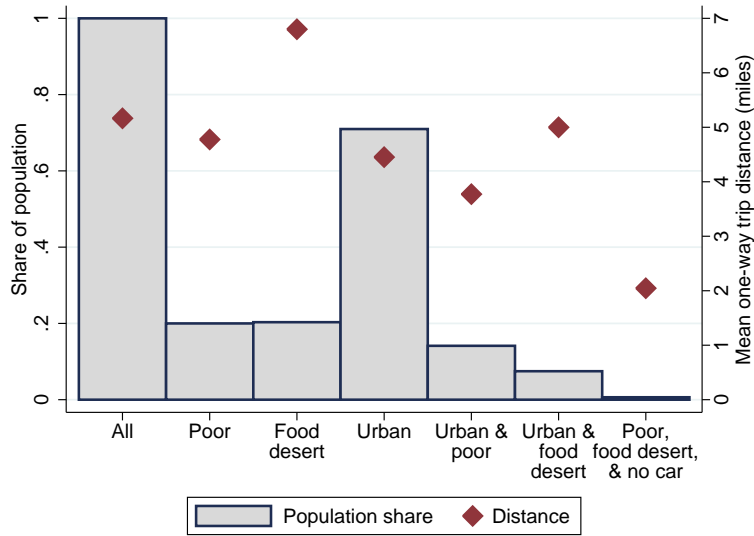
Notes: This figure uses Nielsen RMS data for 2006-2015. We constructed average grams of sugar per 1000 calories, the calorie-weighted share of bread, buns, and rolls UPCs that are whole grain, the calorie-weighted share of UPCs that are produce, and the calorie-weighted mean Health Index, across all UPCs offered in each store. This figure presents the means of these variables across stores within categories of zip code median income.

Figure 5: Event Study of Supermarket Entry



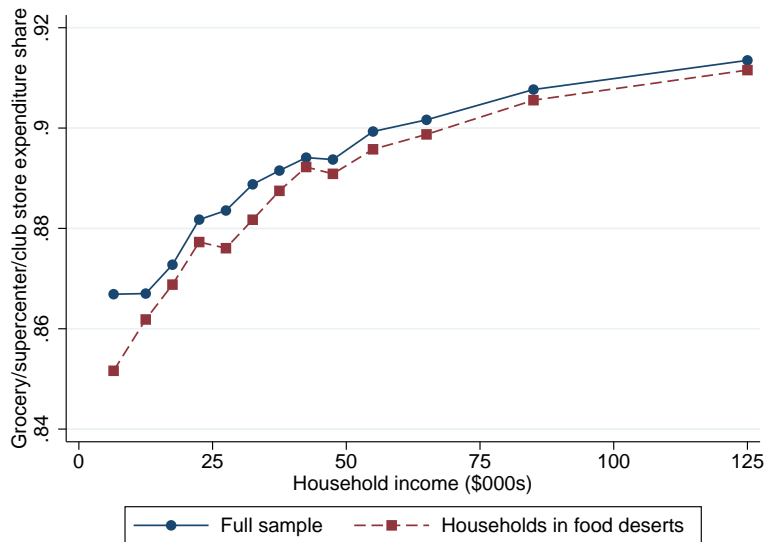
Notes: This figure presents the $\tau_{[0,10]q}$ parameters and 95 percent confidence intervals from estimates of Equation (2): the effects of entry by several large supermarket chains, using 2004-2015 household-by-quarter Homescan data. All regressions control for household demographics (natural log of income, natural log of years of education, age indicators, an indicator for whether the household includes children, race indicators, employment status, weekly work hours, and total calorie need), 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. The bottom two panels present effects on Health Index, our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation one across households. Observations are not weighted for national representativeness.

Figure 6: Shopping Trip Distances by Household Income



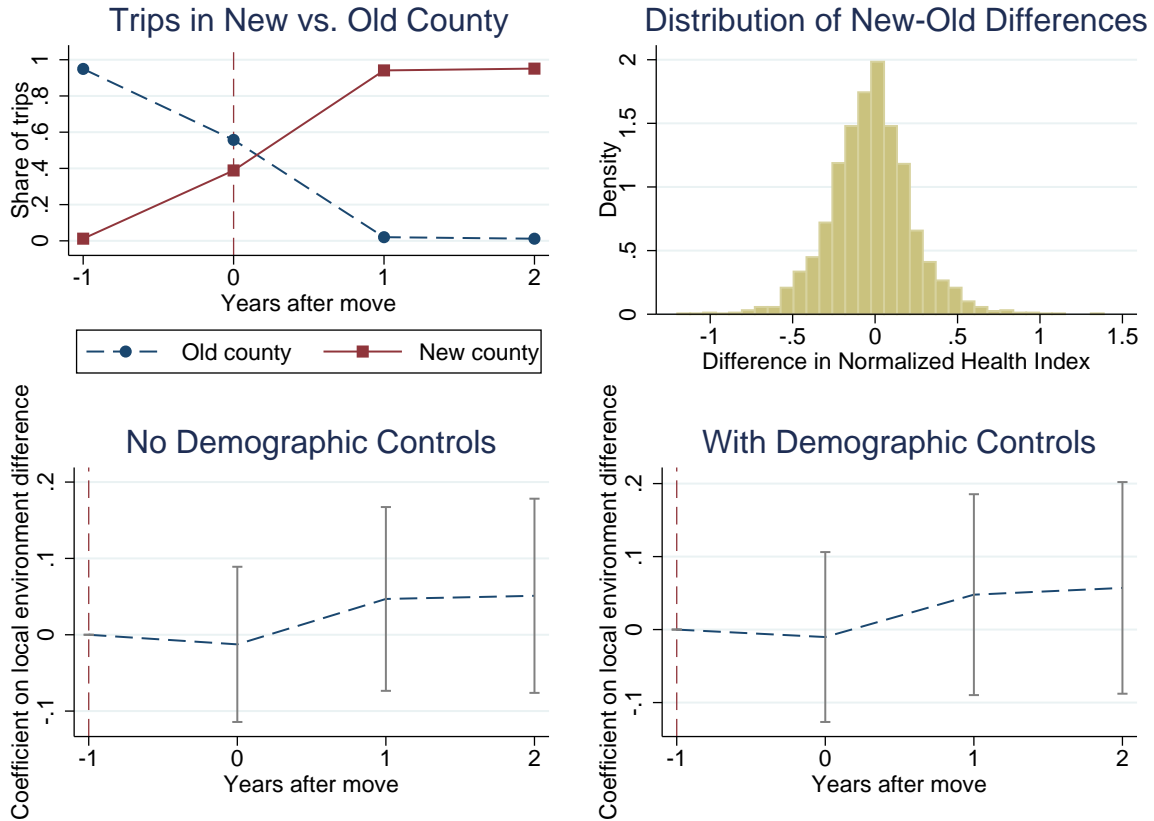
Notes: Data are from the 2009 National Household Travel Survey. Diamonds represent the mean one-way trip distance for trips beginning or ending in “buying goods: groceries/clothing/hardware store.” “Poor” means household income less than \$25,000. “Food desert” means that the household is in a zip code 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 7: Supermarket Expenditure Shares by Household Income



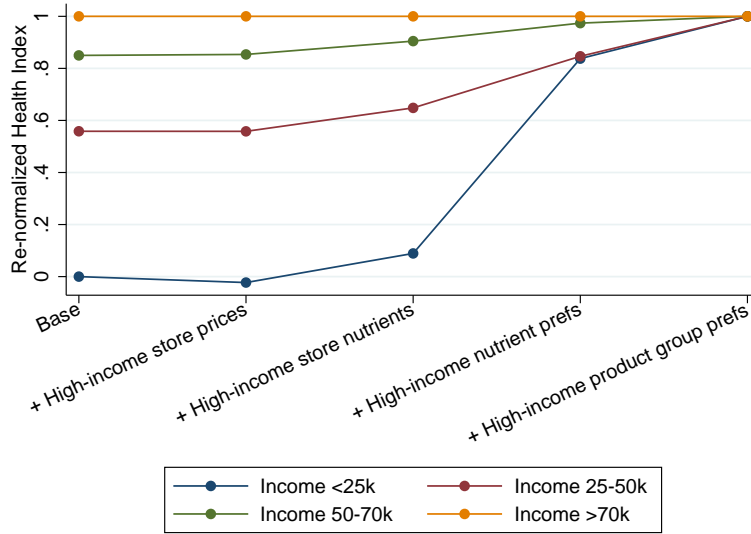
Notes: This figure uses Nielsen Homescan data for 2004-2015. The x-axis presents nominal income bins; household incomes larger than \$100,000 are coded as \$125,000. A household-by-year observation is in a “food desert” if its zip code does not have any grocery stores with 50 or more employees, supercenters, or club stores in that year. Observations are weighted for national representativeness.

Figure 8: **Event Study of Moves Across Counties**



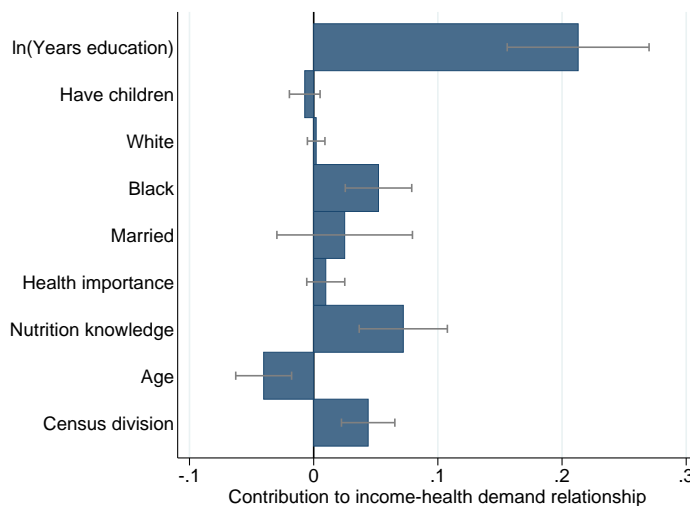
Notes: Using 2004-2015 Homescan data, these figures present results for the event study of moves across counties. The top left panel presents the share of shopping trips that are in the new versus old county. The top right panel presents the distribution across balanced panel households of the difference in 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 household-level Health Index and the difference in 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, an indicator for whether the household includes children, race indicators, employment status, weekly work hours, and total calorie need). Health Index is our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation one across households. Observations are not weighted for national representativeness.

Figure 9: Predicted Health Index for Each Income Group



Notes: Each category on the x-axis represents a separate counterfactual calculation. The base category measures the Health Index for each income group when each group retains their own preferences and face their own local supply conditions. The second category sets all prices to those observed for the high-income group—that is, households with income greater than \$70,000. 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 (32), representing the share of the correlation between income and demand for healthy groceries that is explained by each variable.

Online Appendix: Not for Publication

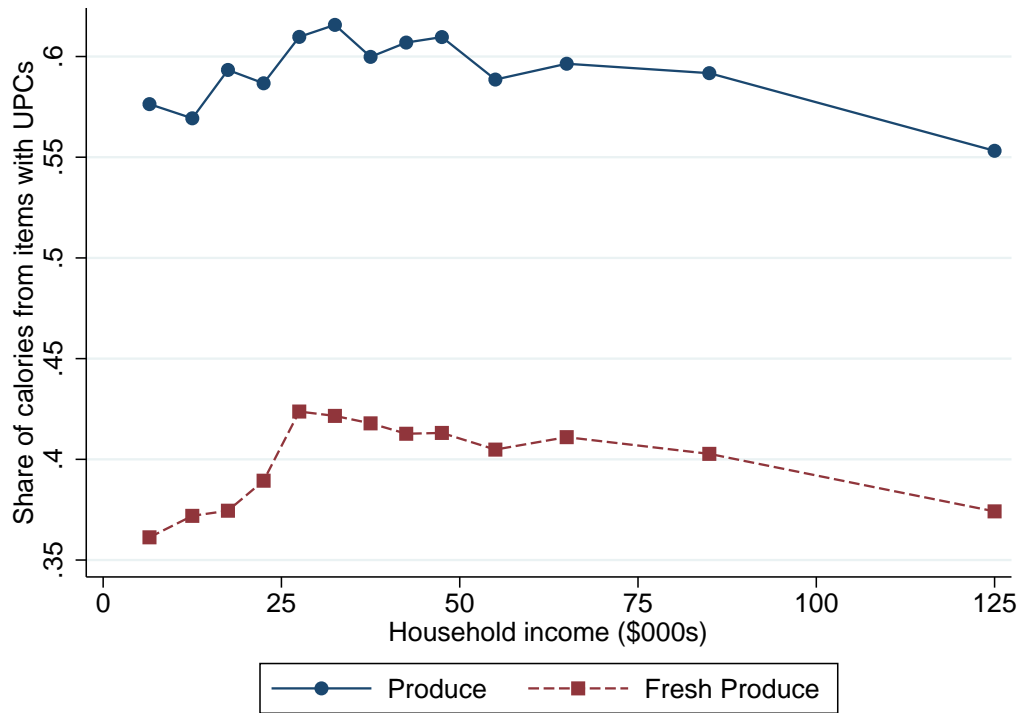
The Geography of Poverty and Nutrition: Food Deserts and Food Choices Across the United States

Hunt Allcott, Rebecca Diamond, and Jean-Pierre Dubé

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 items with UPCs, which are the items that we observe outside the Magnet subsample. “Produce” includes fresh, dried, canned, and frozen produce. Observations are weighted for national representativeness.

A.B Health Index

Table A1: **Health Index Function**

Nutrient	Recommendation	Recommended Daily	
		Intake (grams)	Explanation
<i>Fruits and vegetables</i>			
Fruits	Increase	320	Two cups/day (Food Patterns); 160 g/cup
Vegetables	Increase	390	Three cups/day (Food Patterns); 130 g/cup
<i>All other items</i>			
Protein	Increase	51	51 grams/day (DRI)
Fiber	Increase	29.5	29.5 grams/day (DRI)
Sugar	Reduce	32.8	45% of 282 calories/day from sugar+sat. fat (Food Patterns)
Saturated fat	Reduce	17.2	55% of 282 calories/day from sugar+sat. fat (Food Patterns)
Sodium	Reduce	2.3	2300 mg/day (Dietary Guidelines)
Cholesterol	Reduce	0.3	300 mg/day (Dietary Guidelines)

Notes: Our “raw Health Index” for product n is the sum of healthy minus unhealthy nutrient contents per 1000 calories, weighting each by its recommended daily intake (RDI): $\hat{H}_n = \sum_c G_c a_{nc}/r_c$, where a_{nc} is the grams of nutrient c per 1000 calories, r_c is the RDI for a normal adult, and G_c takes value 1 for “healthy” macronutrients to “increase” and -1 for “unhealthy” nutrients to “reduce.” This table presents the increase/reduce recommendation G_c and RDI r_c used to construct the raw Health Index.

Table A2: **Correlations Between Health Index and Its Components in Homescan**

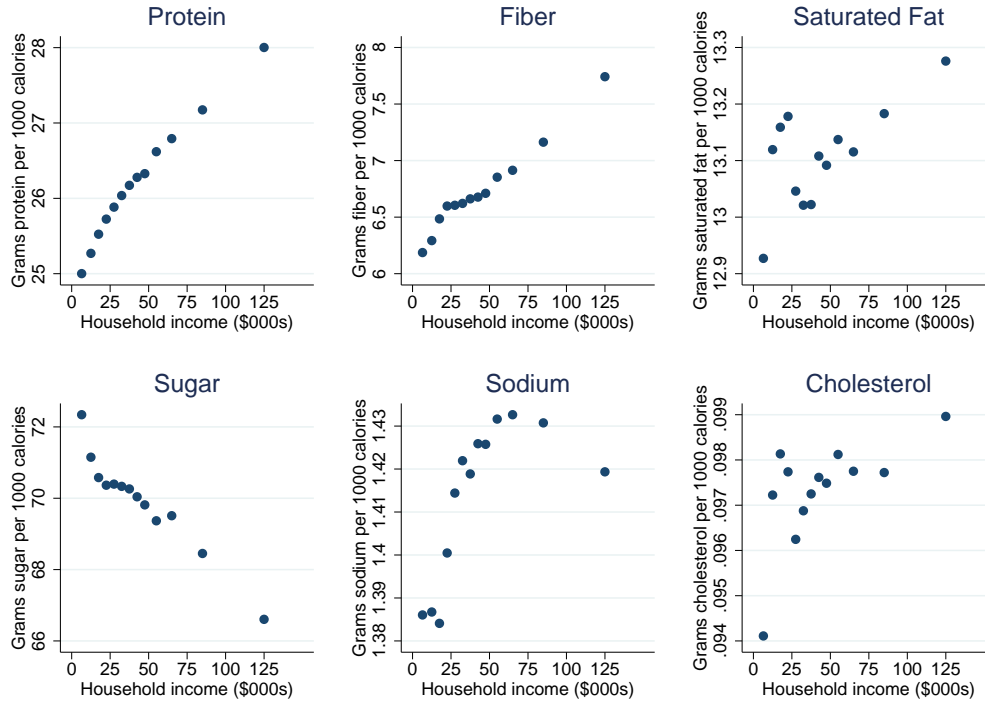
Attribute	Correlation with Health Index
<i>Fruits and vegetables</i>	
Fruits	0.33
Vegetables	0.34
<i>All other items</i>	
Protein	0.52
Fiber	0.64
Sugar	-0.75
Saturated fat	-0.05
Sodium	-0.19
Cholesterol	-0.008

Notes: Using Homescan household-by-year data for 2004-2015, this table presents the correlation coefficients between Health Index and its components, using data in units of grams per 1000 calories consumed. Observations are weighted for national representativeness.

B Appendix to Stylized Facts Section

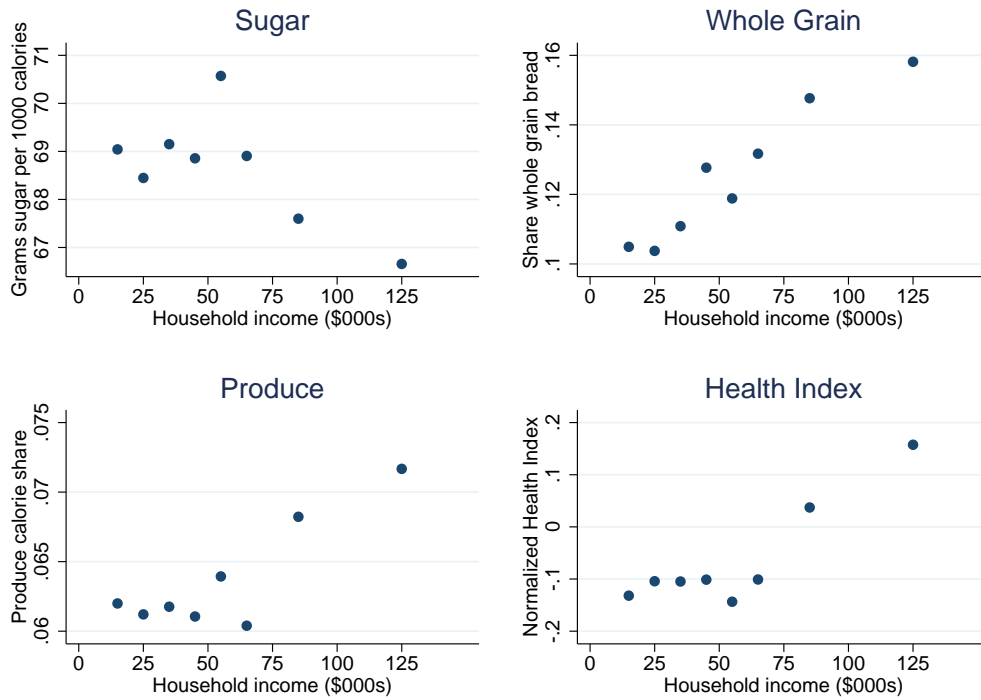
B.A Additional Figures and Tables

Figure A2: Macronutrient Purchases by Household Income



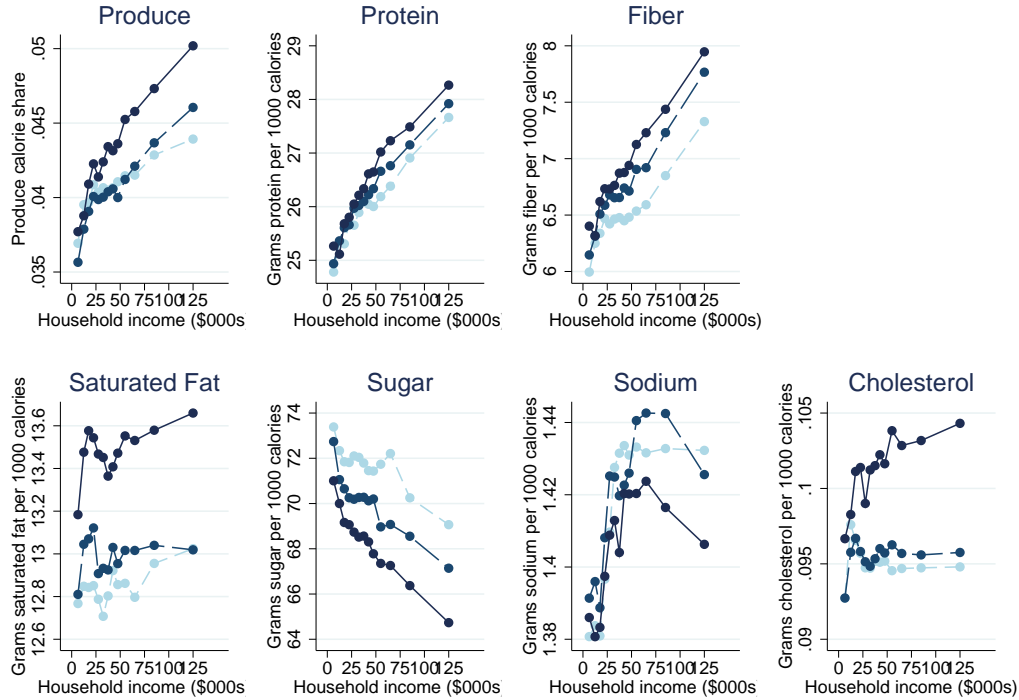
Notes: This figure presents calorie-weighted average macronutrient contents of purchases using Nielsen HomeScan data for 2004-2015. The x-axis presents nominal income bins; household incomes larger than \$100,000 are coded as \$125,000. Observations are weighted for national representativeness.

Figure A3: Magnet Subsample: Healthful Purchases by Household Income



Notes: Nielsen Homescan data, magnet subsample, for 2004-2006. The x-axis presents nominal income bins; household incomes larger than \$100,000 are coded as \$125,000. This parallels Figure 2, except using the magnet subsample which also records purchases of non-UPC items such as bulk produce. Sugar is the grams of sugar per 1000 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 Health Index is our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation one across households. Observations are weighted for national representativeness.

Figure A4: Trends in Macronutrient Purchases by Household Income



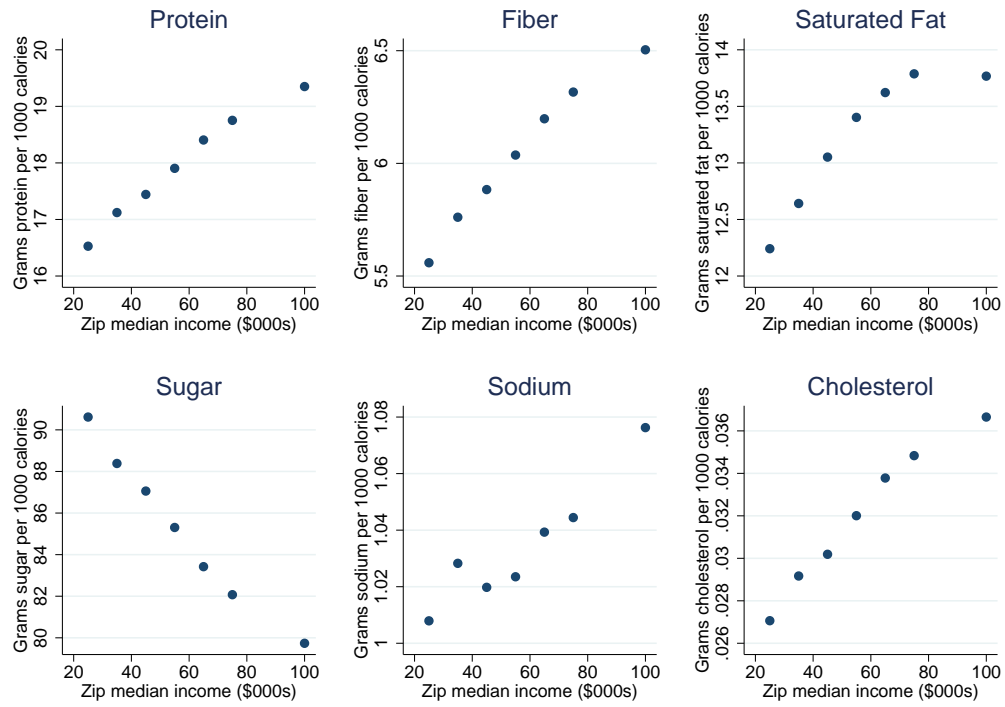
Notes: This figure presents calorie-weighted average macronutrient contents of purchases using Nielsen Homescan data for 2004-2015. The x-axis presents nominal income bins; household incomes larger than \$100,000 are coded as \$125,000. On each plot, the three lines plot 2004-2007, 2008-2011, and 2012-2015 averages, respectively, in light, medium, and dark lines. Observations are weighted for national representativeness.

Table A3: Pooled OLS versus Within-Household Income Variation

	(1)	(2)	(3)	(4)
Sample:	Households Observed in Two or More Years		Exclude Income > \$100k	
ln(Household income)	0.130 (0.00506)***	0.0198 (0.00484)***	0.0160 (0.00492)***	0.0134 (0.00580)**
Household-by-Census tract fixed effects	No	Yes	Yes	Yes
Household demographics	No	No	Yes	Yes
Observations	603,230	603,230	603,230	516,170
Income coefficient/column 1 coefficient	1	0.15	0.12	0.10

Notes: This table presents regressions of Health Index on natural log of household income and year indicators using Nielsen Homescan data for 2004-2015. Columns 2-4 also include household-by-Census tract fixed effects, and columns 3 and 4 also include household demographics (natural log of years of education, age indicators, and an indicator for whether the household includes children, race indicators, employment status, weekly work hours, and total calorie need). The sample is restricted to households observed in two or more years; column 4 additionally excludes observations with household income above \$100,000. Health Index is our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation 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.

Figure A5: Store Average Healthfulness by Zip Code Median Income



Notes: Using Nielsen RMS data for year 2006-2015, we constructed calorie-weighted mean macronutrient content across all UPCs offered in each store. This figure presents the means of these variables across stores within categories of zip code median income.

B.B Healthfulness by Store Type

Appendix Table A4 presents estimates of the relationship between healthy grocery availability, zip code median income, and store characteristics using the 2006-2015 RMS data. We consider two measures of healthy grocery availability H_{jt} at store j in year t : the count of (packaged) produce UPCs offered and the calorie-weighted mean Health Index of UPCs offered. The table presents regressions of H_{jt} on the natural log of zip code median income and additional store covariates \mathbf{X}_{jt} :

$$H_{jt} = \alpha \text{Zip Median Income}_j + \gamma \mathbf{X}_{jt} + \varepsilon_{jt}. \quad (33)$$

Columns 1 and 4 confirm that stores in higher-income zip codes offer substantially more produce UPCs and overall healthier items. There are two potential explanations of this result that healthy food is less available in low-income zip codes. First, stores in higher-income neighborhoods are also considerably larger than stores in low-income neighborhoods, and larger stores naturally offer more variety and could offer healthier UPCs. Second, stores in low-income neighborhoods could stock less healthy options, even after conditioning on store size.

The data are consistent with the first explanation. Columns 2 and 5 show that revenues explain almost all of the relationship between neighborhood income and healthy grocery availability: large stores systematically offer healthier groceries and are much more prevalent in higher-income zip codes.³⁴ Columns 3 and 6 show that even when excluding the revenue variable, retail “channel type” indicators explain 75 to 80 percent of the neighborhood income-store healthfulness relationship. In separate regressions, we find that channel type indicators alone explain 92 and 73 percent of the variance in produce UPC counts and mean Health Index, respectively. Thus, the reason why low-income neighborhoods have less supply of healthy UPCs is that there are fewer large stores, not because the same types of stores stock different types of products. This also means that we can characterize a neighborhood retail environment largely by the types of stores that are present, without needing to analyze the specific UPCs available.

Appendix Table A4 shows that grocery stores, supercenters, and club stores offer more produce UPCs and a higher average Health Index than convenience stores, drug stores, and other mass merchants. Indeed, supercenters such as Walmart Supercenter, Target, and Meijer by definition have full lines of grocery and produce items, as do club stores such as Sam’s Club and Costco. In the paper, we thus refer to large grocers, supercenters, and club stores as “supermarkets,” with the understanding that supermarkets carry more produce and more healthful items.

³⁴The retailing literature typically uses ACV to measure store size (Hoch et al., 1995).

Table A4: Correlates of Healthfulness of UPCs Available at RMS Stores

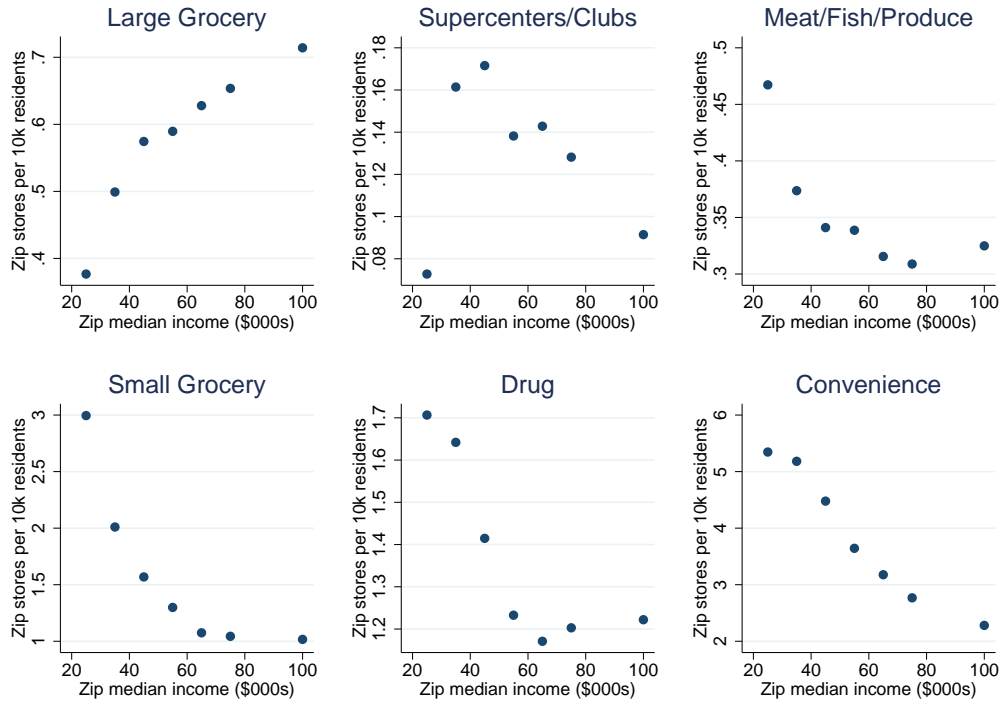
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	Count of Produce UPCs Offered			Mean Health Index of UPCs Offered		
ln(Zip median income)	389.7 (8.85)***	-1.16 (5.30)	85.1 (3.13)***	0.30 (0.0075)***	-0.016 (0.0050)***	0.077 (0.0054)***
ln(Annual revenue)		353.4 (1.10)***			0.29 (0.0011)***	
1(Large grocery)			667.0 (34.7)***			-0.69 (0.059)***
1(Small grocery)			249.1 (34.1)***			-0.87 (0.055)***
1(Supercenter/club)			75.7 (36.5)**			-0.91 (0.060)***
1(Convenience store)			-916.2 (34.1)***			-1.79 (0.058)***
1(Drug store)			-856.8 (34.0)***			-1.89 (0.058)***
1(Other mass merchant)			-781.7 (33.1)***			-1.78 (0.057)***
Observations	339,113	339,113	339,113	339,113	339,113	339,113
R^2	0.043	0.80	0.95	0.035	0.70	0.89

Notes: This table uses 2006-2015 Nielsen RMS data at the store-by-year level. Health Index is our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation one across households. 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. There is no omitted store type in columns 3 and 6. Robust standard errors, clustered by zip code, are in parentheses. *, **, ***: Statistically significant with 10, 5, and 1 percent confidence, respectively.

B.C Low-Income Neighborhoods Have More Unhealthy Store Types

Using the Zip Code Business Patterns data for 2004-2015, Appendix Figure A6 plots the average count of stores by channel type for zip codes by median income category. Zip codes vary substantially in area and population, so this figure normalizes store counts per 10,000 residents; the mean zip code has 12,000 residents. Lower-income zip codes have more stores per capita of all channel types, with two exceptions. First, the concentration of large grocery stores per capita is sharply monotonically increasing in median income, consistent with Powell et al. (2007). Second, the concentration of supercenters and club stores takes an inverted-U shape, with many fewer per capita in the very lowest-income zip codes.

Figure A6: Store Counts by Zip Code Median Income



Notes: This figure presents mean store counts per 10,000 residents by zip code income category using data from Zip Code Business Patterns, averaged over 2004-2015. Large (small) grocers are defined as those with 50 or more (fewer than 50) employees.

B.D The Relative Price of Healthy Foods

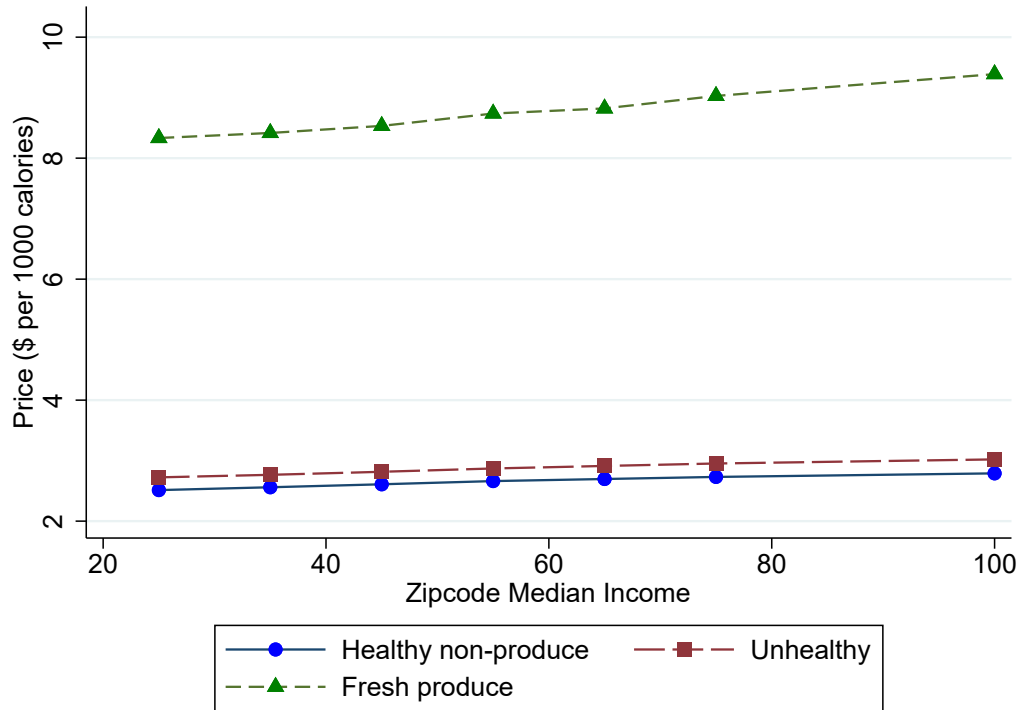
The relative price of healthy foods could matter in two ways. First, healthy foods could be more expensive per calorie than unhealthy foods, leading lower-income households to eat unhealthy foods in order to economize on calories. Second, the price of healthy foods relative to unhealthy foods could be higher in low-income neighborhoods, which could cause consumers with the same preferences to buy more unhealthy foods in low-income neighborhoods.

Appendix Figure A7 presents both of these relative prices in a simplified way. To construct the figure, we take all the RMS grocery sales for 2012 and divide them into three groups: (packaged) fresh produce, all other categories with above-median Health Index, and all categories with below-median Health Index. We calculate the mean price per calorie for available UPCs in each of these three groups for RMS stores in zip codes with different median incomes.

Public health studies typically find that healthy diets cost more per calorie; see Rao et al. (2013) for a meta-analysis of 27 studies. Unsurprisingly, fresh produce is significantly more expensive per calorie than all other categories, regardless of neighborhood income. However, the figure shows that after excluding fresh produce, healthy foods are actually about eight percent less expensive than unhealthy foods.

Broda, Leibtag, and Weinstein (2009) show that higher-income Homescan consumers pay slightly more for the same UPCs. However, the existing evidence on the relative price of healthy versus unhealthy foods in low- versus high-income areas is mixed (Acheson, 1998; Kaufman et al., 1997). Handbury, Rahkovsky, and Schnell (2015) find that the relative price of healthy groceries is the same in low-education versus high-education neighborhoods. In Appendix Figure A7, the upward slope of price per calorie in income reflects Broda, Leibtag, and Weinstein (2009)’s finding. Relative to stores in the lowest-income zip codes, stores in the highest-income zips charge 13 percent more for fresh produce, 11 percent more for healthy non-produce, and 11 percent more for unhealthy foods. Thus, the relative cost between healthy and unhealthy food is quite similar across high and low income zip codes, with fresh produce actually costing a bit more in high income areas.

Figure A7: **Relative Prices of Healthy and Unhealthy UPCs by Zip Code Median Income**

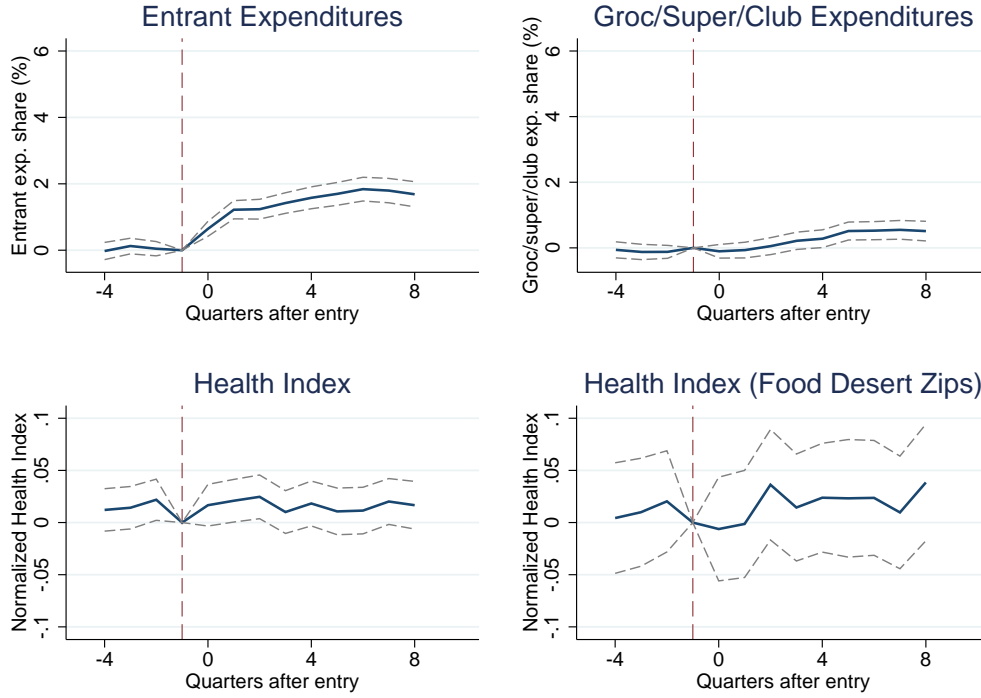


Notes: This figure shows the average price per calorie charged in stores within categories of zip code median income, using 2012 RMS data. “Produce” refers to packaged fresh produce, “Healthy non-produce” refers to all product groups (other than fresh produce) with above-median Health Index, and “Unhealthy” refers to all product groups with below-median Health Index.

C Appendix to Reduced-Form Event Studies

C.A Additional Figures and Tables for Entry Event Study

Figure A8: Event Study of Supermarket Entry Between 10 and 15 Minutes from Home



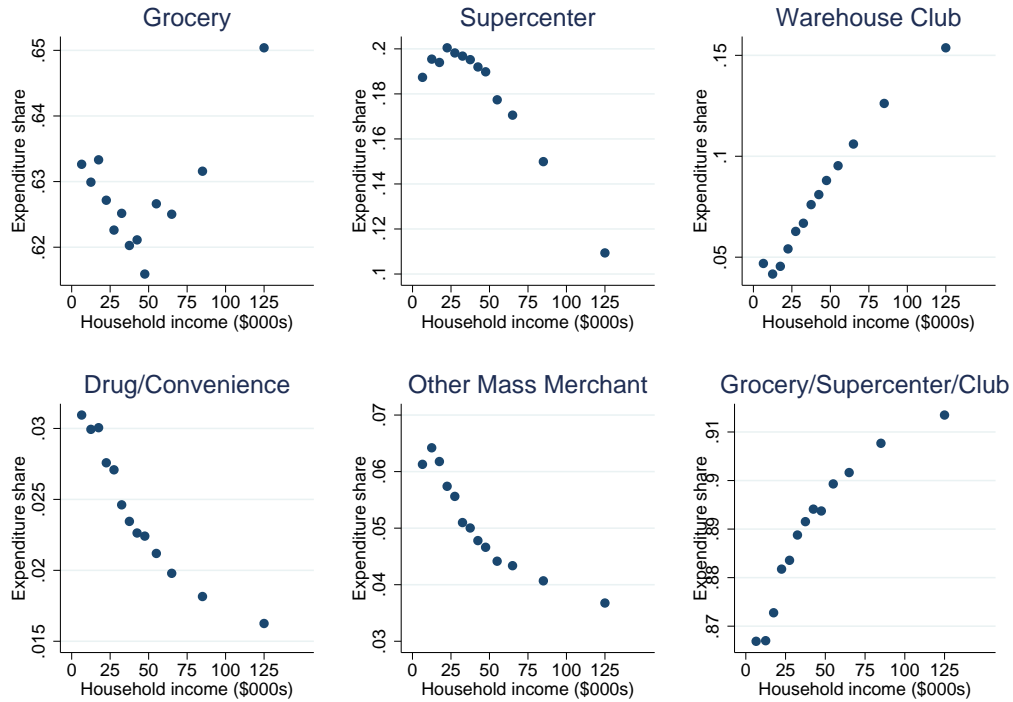
Notes: This figure presents the $\tau_{[10,15]_q}$ parameters and 95 percent confidence intervals from estimates of Equation (2): the effects of entry by several large supermarket chains, using 2004-2015 household-by-quarter Home-scan data. All regressions control for household demographics (natural log of income, natural log of years of education, age indicators, an indicator for whether the household includes children, race indicators, employment status, weekly work hours, and total calorie need), 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. The bottom two panels present effects on Health Index, our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation one across households. The dashed vertical line is the last quarter before entry. Observations are not weighted for national representativeness.

Table A5: **Effects of Supermarket Entry**

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Effects on Expenditure Shares at Other Store Types						
Sample:	Full Sample		Income < \$25,000		"Food Desert" Zip Codes	
Expenditure shares at store type:	Conv/ Drug Stores	Other Mass Merchants	Conv/ Drug Stores	Other Mass Merchants	Conv/ Drug Stores	Other Mass Merchants
Post entry: 0-10 minutes	-0.0888 (0.0352)**	-0.515 (0.0642)***	-0.265 (0.125)**	-0.856 (0.223)***	-0.106 (0.129)	-0.598 (0.164)***
Post entry: 10-15 minutes	-0.0541 (0.0254)**	-0.113 (0.0434)***	-0.120 (0.0904)	-0.0331 (0.143)	-0.116 (0.0782)	-0.233 (0.107)**
Observations	2,627,947	2,627,947	404,868	404,868	640,498	640,498
Dependent var. mean	2.6	5.3	3.5	7.1	2.4	5.6
Panel B: Effects on Health Index Using Alternative Food Desert Definitions						
	(1)	(2)	(3)			
Sample:	< 1000 Produce UPCs	No Medium Groceries	Three-Mile Radius			
Post entry: 0-10 minutes	-0.00968 (0.0182)	-0.0129 (0.0192)	0.00564 (0.0201)			
Post entry: 10-15 minutes	0.0338 (0.0109)***	0.0304 (0.0115)***	0.0463 (0.0133)***			
Observations	408,160	380,868	487,646			

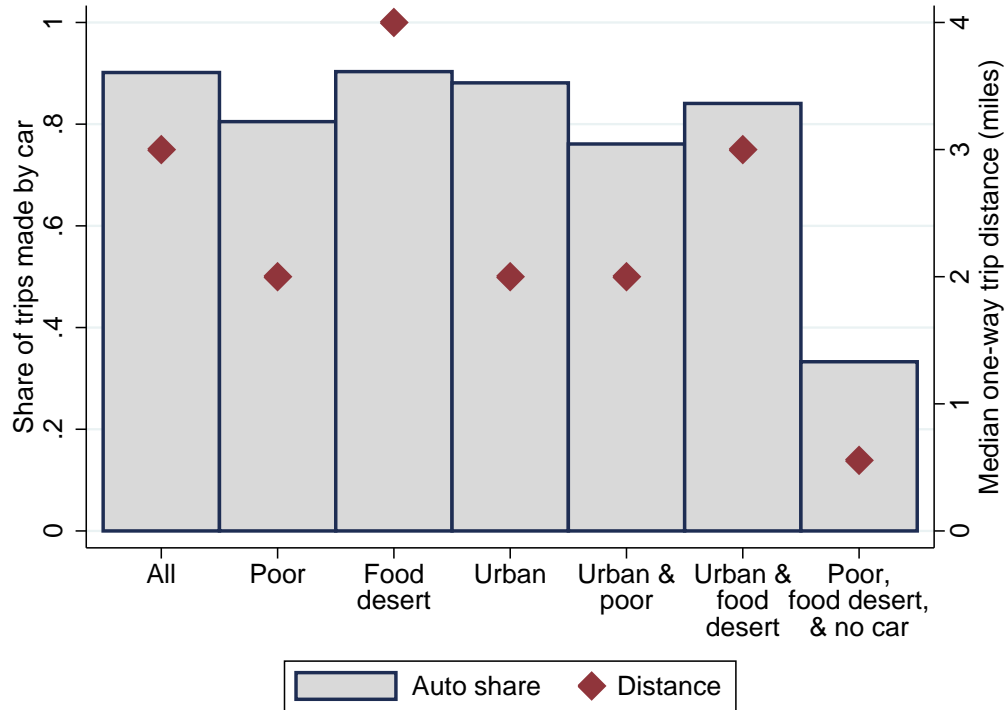
Notes: This table uses 2004-2015 Nielsen Homescan data at the household-by-quarter level. The table parallels Table 2, except Panel A presents effects on expenditure shares at alternative channel types, and Panel B uses alternative definitions of a "food desert." In Panel B, columns 1 and 2 limit the sample to zip codes with fewer than 1000 produce UPCs available in 2003, as predicted by applying RMS data from Table A4 to Zip Code Business Patterns data; columns 3 and 4 also exclude any zip codes with grocery stores employing between 10 and 49 employees in 2003; columns 5 and 6 define 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 in 2003. 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 indicators for whether a specific retailer has 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, an indicator for whether the household includes children, race indicators, employment status, weekly work hours, and total calorie need), 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, are in parentheses. *, **, ***: Statistically significant with 10, 5, and 1 percent confidence, respectively.

Figure A9: Channel Type Expenditure Shares by Household Income



Notes: This figure uses Nielsen Homescan data for 2004-2015. The x-axis presents nominal income bins; household incomes larger than \$100,000 are coded as \$125,000. Another 5-6 percent of expenditures are at channels not plotted, including bakeries, butchers, candy stores, liquor stores, fruit stands, and fish markets; this proportion is fairly constant by income. Observations are weighted for national representativeness.

Figure A10: 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 household income less than \$25,000. “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.

C.B Entry by All Retailers Using Zip Code Business Patterns

To complement the event study estimates in the body of the paper, we present alternative specifications that measure entry using the Zip Code Business Patterns (ZBP) data.

Panel A of Appendix Table A6 shows that the Zip Code Business Patterns data date openings of specific supercenters in the correct year 50 to 80 percent of the time, although they are sometimes recorded a year later and sometimes in a broader “general merchandise” NAICS code (452) instead of the specific “supercenter and club store” NAICS code (452910).

The entry event study regression is analogous to Equation (1). Define S_{zt} and G_{zt} , respectively, as the count of supercenters/club stores and large (at least 50 employee) grocery stores in zip code z in year t . Using household-by-year data and now denoting μ_{dt} as Census division-by-year indicators, the regression is:

$$Y_{izt} = \tau_S S_{zt} + \tau_G G_{zt} + \beta \mathbf{X}_{it} + \mu_{dt} + \phi_{ic} + \varepsilon_{izt} \quad (34)$$

Standard errors are again clustered by household, and observations are again all weighted equally.

Appendix Table A7 presents results. The structure is similar to that of Table 2: Panel A presents effects on expenditure shares, while Panel B presents effects on healthful eating.

Columns 1-3 present estimates for the full sample. Columns 1 and 2 confirm that the ZBP data contain meaningful information. Column 1 shows that conditional on household fixed effects, a larger count of large grocery stores and/or a smaller count of supercenters and club stores in the zip code are both strongly positively associated with higher expenditure share at chain groceries. Column 2 shows the opposite: fewer grocery stores and more supercenters are strongly positively associated with higher expenditures at supercenters and club stores. Column 3 presents effects on combined expenditure shares for all grocery stores, supercenters, and club stores. Columns 4-6 and 7-9 present estimates for the low-income and food desert subsamples. As in Table 2, effects of entry on expenditures generally larger in food deserts. Also as in Table 2, we see that entry by a large grocery retailer substantially diverts sales from other supermarkets, so the effects on combined expenditures at grocery stores, supercenters, and club stores are limited. Appendix Table A8 shows that most of this diversion is from other mass merchants; there is no statistically significant diversion from drug and convenience stores.

The bottom panel shows no statistically significant effect of the number of large grocers and supercenters/clubs on Health Index. With 95 percent confidence, we can bound the effects on low-income households’ Health Index at less than 0.011 standard deviations per large grocery store and 0.052 standard deviations per supercenter or club store. Appendix Table A8 shows that under all alternative definitions of “food deserts,” the number of local large grocers, supermarkets, and club stores has no statistically or economically significant effect on Health Index.

One reason to prefer the earlier regressions with specific known retailers is that we have high confidence that entry dates are correctly measured. We can also imagine using the true supercenter entry dates as an instrument for ZBP data, which are measured with error. Panel B of Appendix

Table A6 shows that the “first stages” of such a regression have coefficients around 0.9 and 0.66 for two different supercenter chains. If the average retailer in ZBP is measured with equal or perhaps somewhat more error than the less well-measured supercenter chain, this suggests that our bounds in the paragraph above should be increased by 50 to 100 percent due to measurement error. Even after this adjustment, however, our results in Tables 2 and A7 suggest that having a supermarket nearby explains at most only a small share of the differences in nutritional decisions between low- and high-income households.

Table A6: **Zip Code Business Patterns Accuracy Check with Known Entry Dates**

	(1)	(2)
Dependent variable is	General	
count of channel type:	Supercenter	Merchandise
Panel A: Difference Estimator		
<u>Supercenter Chain 1:</u>		
2-year lead	0.00478 (0.00622)	0.00629 (0.0166)
1-year lead	0.0375 (0.00978)***	0.0577 (0.0164)***
Entry year	0.562 (0.0199)***	0.821 (0.0210)***
1-year lag	0.208 (0.0169)***	0.0821 (0.0191)***
2-year lag	0.0777 (0.0127)***	0.0133 (0.0145)
<u>Supercenter Chain 2:</u>		
2-year lead	0.0158 (0.0298)	-0.0172 (0.0421)
1-year lead	0.0133 (0.0222)	-0.0451 (0.0462)
Entry year	0.0621 (0.0327)*	0.480 (0.0623)***
1-year lag	0.172 (0.0413)***	0.0701 (0.0594)
2-year lag	0.133 (0.0514)***	0.0918 (0.0599)
Observations	264,734	264,734
Panel B: Fixed Effects Estimator		
Post entry: chain 1	0.902 (0.0138)***	0.932 (0.0227)***
Post entry: chain 2	0.667 (0.0365)***	0.665 (0.0659)***
Observations	297,966	297,966

Notes: Data are at the zip code-by-year level. All regressions include year indicators; fixed effects regressions have zip code fixed effects. Robust standard errors, clustered by zip code, are in parentheses. *, **, ***: Statistically significant with 10, 5, and 1 percent confidence, respectively.

Table A7: Effects of Supermarket Entry Using Zip Code Business Patterns

Panel A: Effects on Expenditure Shares									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sample:	Full Sample			Income < \$25,000			"Food Desert" Zip Codes		
Expend. shares at store type:	Chain Grocers	Super/Club	Grocery/Super/Club	Chain Grocers	Super/Club	Grocery/Super/Club	Chain Grocers	Super/Club	Grocery/Super/Club
Large grocers	0.329 (0.0718)***	-0.382 (0.0615)***	-0.0259 (0.0414)	0.256 (0.207)	-0.430 (0.173)**	-0.0772 (0.136)	0.522 (0.299)*	-0.705 (0.242)***	-0.129 (0.167)
Supers/clubs	-1.896 (0.172)***	2.903 (0.160)***	0.692 (0.0939)***	-2.023 (0.480)***	3.553 (0.451)***	1.057 (0.311)***	-2.472 (0.591)***	3.452 (0.570)***	0.629 (0.291)**
Observations	664,302	664,302	664,302	102,462	102,462	102,462	163,747	163,747	163,747
Dep. var. mean	58	26	88	56	23	86	53	28	88

Panel B: Effects on Health Index			
	(1)	(2)	(3)
Sample:	Full Sample	Income < \$25,000	"Food Desert" Zip Codes
Large grocers	0.00102 (0.00254)	-0.00303 (0.00738)	-0.00518 (0.0104)
Supers/clubs	0.00667 (0.00551)	0.0217 (0.0155)	0.0184 (0.0168)
Observations	664,302	102,462	163,747

Notes: This table uses 2004-2015 Nielsen Homescan data at the household-by-year level. "Food Desert" zip codes are those with no grocery stores with 50 or more employees, supercenters, or club stores in 2003. 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 stores by channel type in the household's zip code. All regressions control for household demographics (natural log of income, natural log of years of education, age indicators, an indicator for whether the household includes children, race indicators, employment status, weekly work hours, and total calorie need), Census division-by-year indicators, and household-by-Census tract fixed effects. 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.

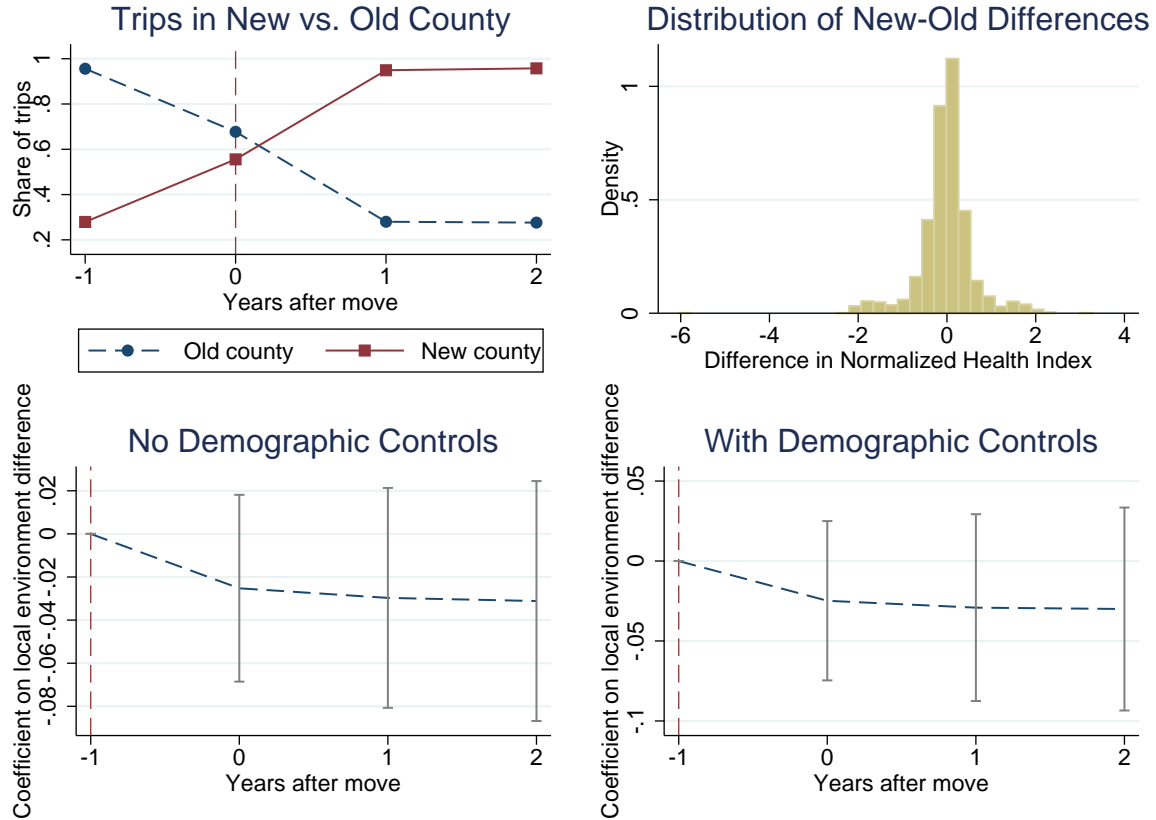
Table A8: Effects of Supermarket Entry Using Zip Code Business Patterns

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Effects on Expenditure Shares at Other Store Types						
Sample:	Full Sample		Income < \$25,000		"Food Desert" Zip Codes	
Expenditure shares at store type:	Conv./ Drug Stores	Other Mass Merchants	Conv./ Drug Stores	Other Mass Merchants	Conv./ Drug Stores	Other Mass Merchants
Large grocers	-0.0145 (0.0173)	0.0737 (0.0278)***	-0.0352 (0.0613)	0.168 (0.0953)*	-0.00259 (0.0643)	0.116 (0.106)
Supercenters/clubs	-0.0216 (0.0348)	-0.717 (0.0641)***	-0.172 (0.130)	-1.079 (0.213)***	0.142 (0.116)	-0.507 (0.187)***
Observations	664,302	664,302	102,462	102,462	163,747	163,747
Dependent var. mean	2.6	5.2	3.5	7.0	2.4	5.6
Panel B: Effects on Health Index Using Alternative Food Desert Definitions						
	(1)	(2)	(3)			
Sample:	< 1000 Produce UPCs	No Medium Groceries	Three-Mile Radius			
Large grocers	-0.00617 (0.0148)	-0.00660 (0.0161)	-0.0144 (0.0129)			
Supers/clubs	0.0158 (0.0200)	0.0133 (0.0204)	0.0192 (0.0211)			
Observations	104,451	98,256	125,399			

Notes: This table uses 2004-2015 Nielsen Homescan data at the household-by-year level. The table parallels Table A7, except Panel A presents effects on expenditure shares at alternative channel types, and Panel B uses alternative definitions of a "food desert." In Panel B, columns 1 and 2 limit the sample to zip codes with fewer than 1000 produce UPCs available in 2003, as predicted by applying RMS data from Table A4 to Zip Code Business Patterns data; columns 3 and 4 also exclude any zip codes with grocery stores employing between 10 and 49 employees in 2003; columns 5 and 6 define 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 in 2003. 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 stores by channel type in the household's zip code. All regressions control for household demographics (natural log of income, natural log of years of education, age indicators, an indicator for whether the household includes children, race indicators, employment status, weekly work hours, and total calorie need), Census division-by-year indicators, and household-by-Census tract fixed effects. Observations are not weighted for national representativeness. 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.

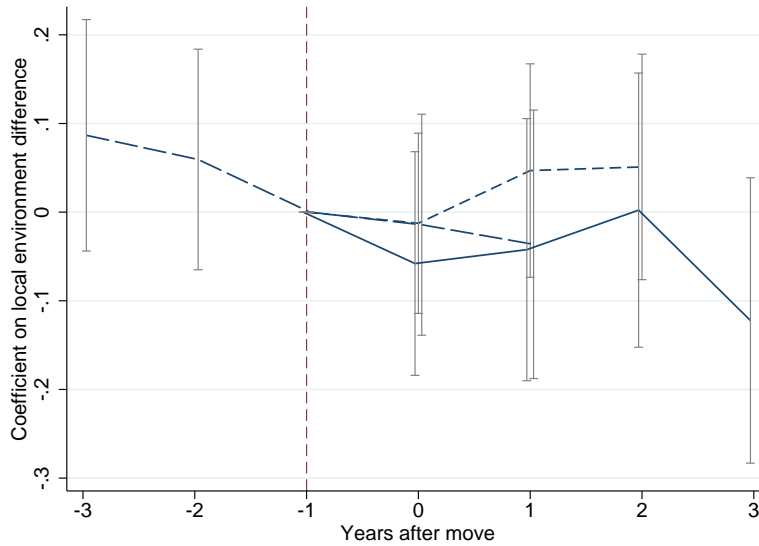
C.C Appendix to Movers Event Study

Figure A11: Event Study of Moves Across Zip Codes

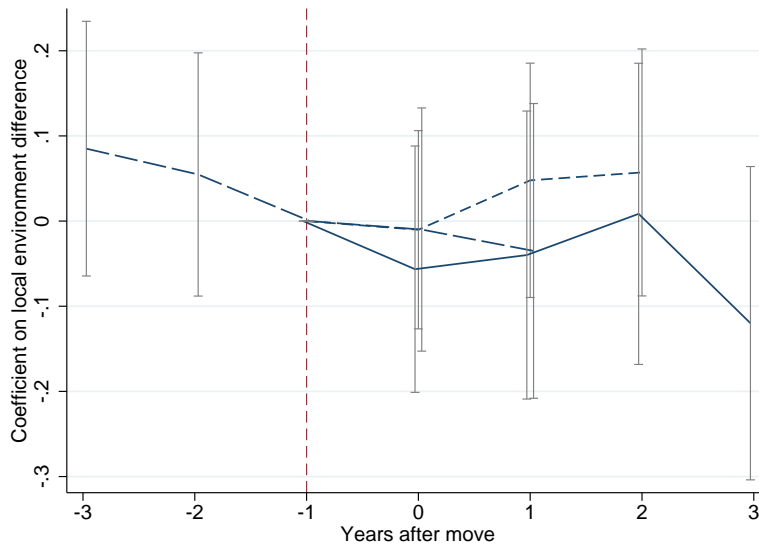


Notes: Using 2004-2015 Homescan data, these figures present results for the event study of moves across zip codes. The top left panel presents the share of shopping trips that are in the new versus old county. The top right panel presents the distribution across balanced panel households of the difference in 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 household-level Health Index and the difference in 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, an indicator for whether the household includes children, race indicators, employment status, weekly work hours, and total calorie need). Health Index is our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation one across households. Observations are not weighted for national representativeness.

Figure A12: **Event Study of Movers with Different Balanced Sample Windows**



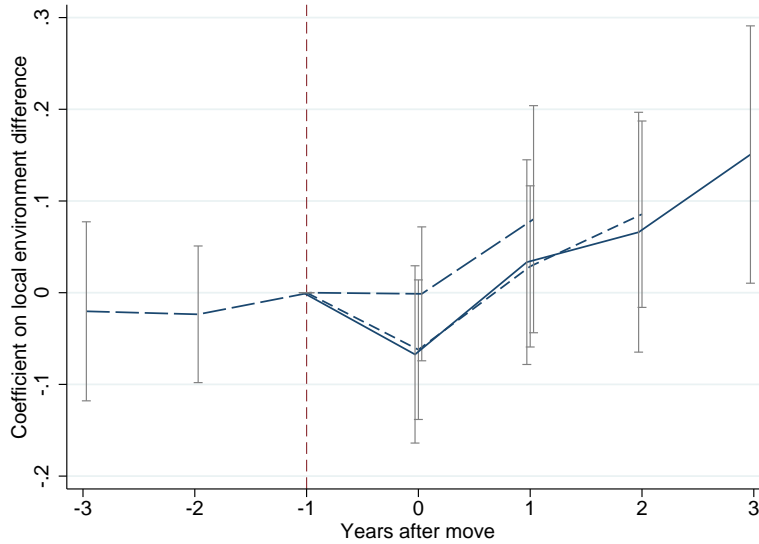
(a) **Without Controls**



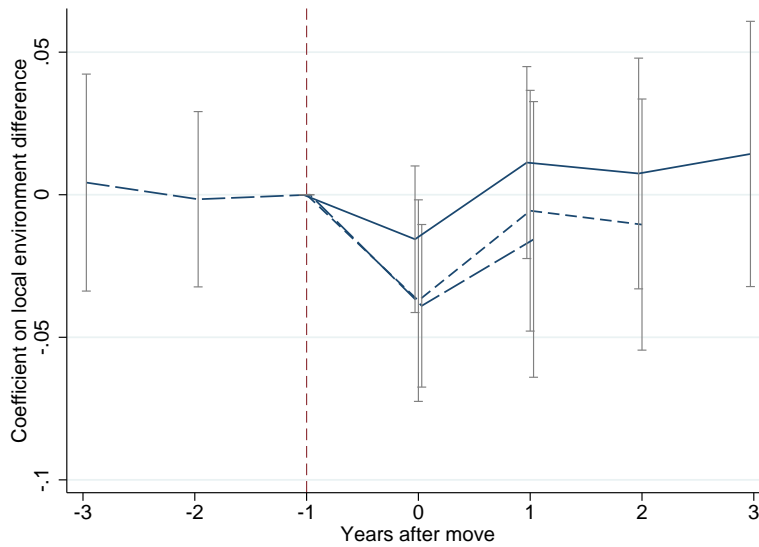
(b) **With Controls**

Notes: Using 2004-2015 Homescan data, these figures present the τ_y parameters and 95 percent confidence intervals from estimates of Equation (4): associations between household-level Health Index and the difference in 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, an indicator for whether the household includes children, race indicators, employment status, weekly work hours, and total calorie need). Health Index is our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation one across households. Observations are not weighted for national representativeness.

Figure A13: **Event Study: Income Changes in Mover Households**



(a) **Moves Across Counties**



(b) **Moves Across Zip Codes**

Notes: Using 2004-2015 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 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. Health Index is our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation one across households. Observations are not weighted for national representativeness. The regressions are the same as in Figure 8, except with natural log of household income as the dependent variable and no controls for household demographics.

Table A9: Association of Income with Local Area Health Index Using Movers

	(1)	(3)
Zip code average Health Index	-0.00658 (0.0100)	
County average Health Index		0.0971 (0.0335)***
Observations	560,492	565,914

Notes: This table uses 2004-2015 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. Health Index is our overall measure of the healthfulness of grocery purchases, normalized to mean zero, standard deviation 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, are in parentheses. *, **, ***: Statistically significant with 10, 5, and 1 percent confidence, respectively.

Table A10: Association of Coke Market Share with Local Area Coke Market Share Using Movers

	(1)	(2)
County average Coke market share	0.138 (0.0547)**	0.136 (0.0547)**
Household demographics	No	Yes
Observations	323,710	323,710

Notes: This table uses 2004-2015 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, an indicator for whether the household includes children, race indicators, employment status, weekly work hours, and total calorie need. All regressions also control for year indicators and household fixed effects. 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.

D Appendix to Demand Model Estimation

D.A GMM Estimation

Our GMM estimator is defined as follows:

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

Define \mathbf{Y} as the vector of product group calorie consumption Y_{ijt} , $\mathbf{F}(\tilde{\boldsymbol{\beta}}, \boldsymbol{\xi})$ as the vector of implicit prices $F_{ijt} = \left(\tilde{p}_{ijt} - \sum_{c=2}^C \tilde{\beta}_c \tilde{a}_{ijct} - \boldsymbol{\xi}\right)$, \mathbf{D} as a stacked matrix of the two dummy variable matrices (\mathbf{D}_j and \mathbf{D}_m), \mathbf{Z} as a matrix with all of our vectors of instruments (\mathbf{D} , the nutrient content $\tilde{\mathbf{a}}$, and the price instruments \mathbf{P}), and $\mathbf{Pr}_D = (\mathbf{D}'\mathbf{Z}\mathbf{W}\mathbf{Z}'\mathbf{D})^{-1}\mathbf{D}'\mathbf{Z}\mathbf{W}\mathbf{Z}'$ 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 $\boldsymbol{\xi}$:

$$(\boldsymbol{\delta}, \boldsymbol{\phi}) = \mathbf{Pr}_D \left(\ln(\mathbf{Y}) - \mathbf{F}(\tilde{\boldsymbol{\beta}}, \boldsymbol{\xi}) \right). \quad (35)$$

Substituting Equation (35) back into Equation (21), we can re-write the GMM estimator in terms of $\tilde{\boldsymbol{\beta}}$ and $\boldsymbol{\xi}$:

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

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

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

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

$$\mathbf{G}(\tilde{\boldsymbol{\beta}}, \boldsymbol{\xi}) = \frac{1}{IJT} \begin{bmatrix} \tilde{\mathbf{a}}'(\mathbf{I} - \mathbf{D}\mathbf{Pr}_D) \\ \mathbf{P}'(\mathbf{I} - \mathbf{D}_m\mathbf{Pr}_{D_m}) \\ \mathbf{D}'(\mathbf{I} - \mathbf{D}\mathbf{Pr}_D) \end{bmatrix} \nabla_{\boldsymbol{\beta}} \mathbf{F}(\tilde{\boldsymbol{p}}, \tilde{\mathbf{a}}; \tilde{\boldsymbol{\beta}}). \quad (36)$$

In the above equation, \mathbf{I} is the identity matrix, and \mathbf{Pr}_{D_m} is a projection matrix using \mathbf{D}_m . The covariance matrix of our full GMM estimator, $\boldsymbol{\Theta}^{GMM} \equiv \left(\hat{\boldsymbol{\delta}}, \hat{\boldsymbol{\phi}}, \hat{\tilde{\boldsymbol{\beta}}}, \hat{\boldsymbol{\xi}}\right)$, is $\text{cov}(\boldsymbol{\Theta}^{GMM}) =$

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

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

and covariance matrix

$$\boldsymbol{\Omega} = E\left(\mathbf{g}_{ijt}(\boldsymbol{\Theta}^{GMM})\mathbf{g}_{ijt}(\boldsymbol{\Theta}^{GMM})'\right).$$

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

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

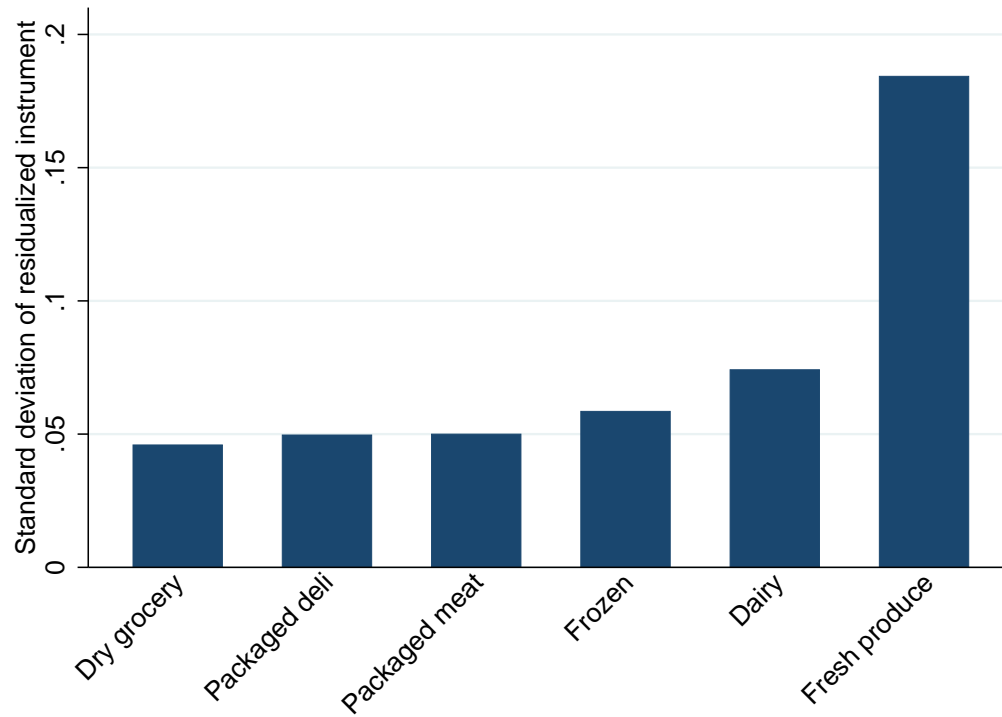
D.B Additional Tables and Figures

Table A11: Preferences for Nutrients by Household Income

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Income group	Unsat. Fat	Sat. Fat	Fiber	Protein	Sugar	Sodium	Cholest.	Fruit	Veg	Unobs. Nutrient	Conve- nience	Shelf Life	WTP for Health Index
$Inc \leq 25k$	-0.83*** (0.11)	9.82*** (0.15)	9.98*** (0.17)	4.40*** (0.29)	1.94*** (0.04)	-29.76*** (1.86)	-13.47*** (4.04)	0.08*** (0.02)	0.42*** (0.02)	0.09*** (0.02)	0.17*** (0.002)	-0.90*** (0.02)	0.52*** (0.012)
$25k < Inc \leq 50k$	-1.77*** (0.07)	13.35*** (0.08)	10.17*** (0.06)	6.61*** (0.14)	1.91*** (0.02)	-29.93*** (0.81)	-13.45*** (1.84)	0.24*** (0.01)	0.65*** (0.004)	-0.05*** (0.01)	0.22*** (0.001)	-0.80*** (0.01)	0.72*** (0.01)
$50k < Inc \leq 70k$	-3.15*** (0.07)	11.12*** (0.09)	10.478*** (0.10)	7.350*** (0.13)	1.39*** (0.02)	-30.06*** (0.85)	-13.45*** (1.91)	0.38*** (0.01)	0.75*** (0.004)	0.09*** (0.01)	0.24*** (0.001)	-0.82*** (0.01)	0.91*** (0.01)
$70k < Inc$	-3.57*** (0.58)	10.50*** (0.55)	10.59*** (0.41)	8.24*** (0.96)	1.26*** (0.15)	-30.14*** (5.41)	-13.45 (11.57)	0.60*** (0.04)	0.93*** (0.03)	0.11 (0.07)	0.28*** (0.01)	-0.88*** (0.05)	1.12*** (0.05)

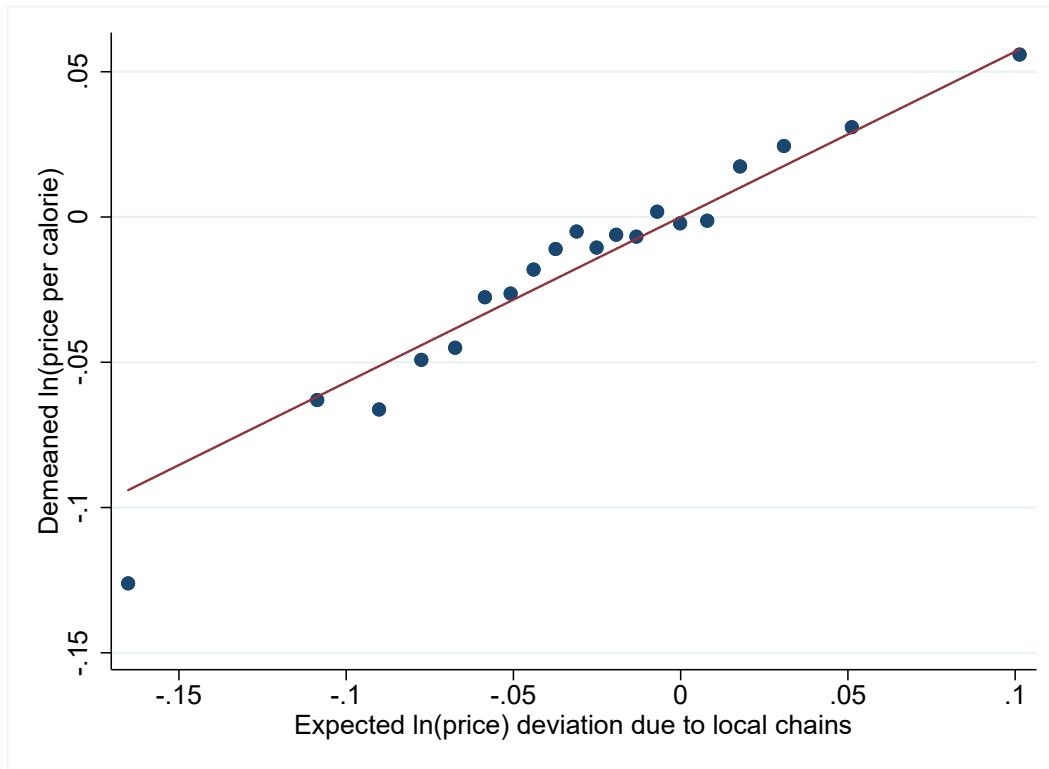
Notes: This table presents GMM estimates of the preference parameters $\tilde{\beta}_c$ from Equation (15), adding convenience and shelf life as additional product characteristics. Shelf life is measured in years and top-coded at one year. Convenience is a score ranging from 0 to 3, defined as follows. 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 of nutrient estimates represent willingness to pay for a kilogram of the nutrient instead of a kilogram of carbohydrates. Value of fruit and vegetables accounts for value over and beyond macronutrient characteristics of the fruit and vegetables. “WTP for Health Index” in column 13 equals $\sum_c \hat{\beta}_c G_c r_c$, where $G_c = 1$ for “healthy” nutrients, $G_c = -1$ for “unhealthy” nutrients, and r_c is the recommended daily intake of nutrient c detailed in Appendix Table A1. Standard errors, clustered by household, are in parentheses. *, **, ***: Statistically significant with 10, 5, and 1 percent confidence, respectively.

Figure A14: Standard Deviation of Price Instrument by Product Department



Notes: This figure presents the standard deviation of our price instrument, after residualizing against year, product group, and zip-3 fixed effects. This residual standard deviation is the variation used to instrument for prices in our model estimates. The standard deviation are 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. Zip-3 and product group fixed effects are residualized out before plotting. There are 20 equally sized bins, and all income groups are included.