

Is the focus on food deserts fruitless? Retail access and food purchases across the socioeconomic spectrum *

Jessie Handbury, Ilya Rahkovsky, and Molly Schnell[†]

July 1, 2017

Abstract

Using comprehensive data describing the healthfulness of household food purchases and the retail landscapes consumers face, we ask whether spatial differences in access are to blame for socioeconomic disparities in nutritional consumption. We find that differences in access, though significant, are small relative to differences in the nutritional content of sales. Household consumption responds minimally to improvements in local retail environments in the short run, and socioeconomic disparities persist among households with equivalent access. Our results indicate that even in the long run, access-improving policies alone can eliminate at most one fifth of existing disparities in nutritional consumption.

*Prottoy Aman Akbar, Yue Cao, and Hae Nim Lee provided us with outstanding research assistance. Marianne Bitler, Anne Case, David Cuberes, Amanda Chuan, Janet Currie, Jan De Loecker, Gilles Duranton, Joe Gyourko, Jakub Kastl, Ephriam Liebttag, Ilyana Kuziemko, Todd Sinai, Diane Whitmore Schanzenbach, Jesse Shapiro, Tom Vogl, and David Weinstein provided helpful comments. We thank participants in seminars at the 2014 Urban Economics Association Meeting, the 2015 American Economic Association Meeting, the 2015 NBER Summer Institute, the Becker Friedman Institute at the University of Chicago, Brown, the Economic Research Service at the USDA, the Federal Reserve Bank of Kansas City, the Federal Trade Commission, NYU, Penn Economics, Penn State, Princeton, the Robert Wood Johnson Foundation, the University of New South Wales, the University of Sydney, the University of Toronto, the University of Warwick, Wharton, and Yale SOM. Jessie Handbury would like to thank the Wharton Social Impact Initiative, the Research Sponsors' Program of the Wharton Zell-Lurie Real Estate Center, and the Economic Research Service at the USDA for generous financial support. This paper was previously circulated under the title "What drives nutritional disparities? Retail access and food purchases across the socioeconomic spectrum," NBER Working Paper No. 21126, April 2015. The views expressed are those of the authors and should not be attributed to the Economic Research Service, the USDA, Nielsen, or IRI.

[†]Handbury: Real Estate Department, The Wharton School, University of Pennsylvania (handbury@wharton.upenn.edu); Rahkovsky: Economic Research Service, USDA (irahkovsky@ers.usda.gov); Schnell: Department of Economics, Princeton University (mollyks@princeton.edu)

1 Introduction

It is well known that there are large nutritional disparities across socioeconomic groups in the United States. Under the assumption that these disparities are caused by differential access to healthy foods, both the federal government and local municipalities spend millions of dollars each year to stimulate supermarket development and to encourage existing retailers to offer healthier foods in underserved communities.^{1,2} If spatial differences in retail access drive socioeconomic disparities in nutritional consumption, such policies will narrow nutritional disparities. If differential diets are instead caused by differences in tastes, price sensitivities, or constraints, policies aimed at improving access to healthy foods alone will do little to improve the diets of disadvantaged populations.

In this paper, we examine the role of access in explaining why more educated households purchase healthier foods.³ Combining panel data on household purchases, store locations, product availability, and prices across the US from 2006 to 2011, we find that improving access to healthy foods alone will do little to close the gap in the nutritional quality of grocery purchases across households with different levels of education. We estimate that even if spatial disparities in access were entirely resolved, over 82% of the existing socioeconomic disparities in nutritional consumption would remain.

There has long been agreement among researchers that both spatial disparities in access and socioeconomic disparities in nutritional consumption exist, but the actual effect of access to healthy foods on food purchases has been heavily contested (Bitler and Haider (2011)). A large public health literature has inferred the impact of food environments on consumption from a cross-sectional correlation between local store density and food purchases with mixed results (see Larson et al. (2009) for a review).⁴ Since socioeconomic disparities in nutritional consumption and access could be driven entirely by differences in demand, the observed correlation between consumption

¹The Agricultural Act of 2014 appropriated \$125 million in federal funds to be spent annually to promote access to healthy foods in disadvantaged communities (Aussenberg (2014)). Many state and local governments have also introduced programs to improve access to nutritious foods by providing loans, grants, and tax credits to qualifying businesses operating in underserved neighborhoods (CDC (2011)).

²First Lady Michelle Obama made improving access to healthy foods a cornerstone of her agenda while in the White House, stating in 2011 that “it’s not that people don’t know or don’t want to do the right thing; they just have to have access to the foods that they know will make their families healthier” (Curtis (2011)).

³We focus primarily on disparities that exist across households with and without a household head with a college degree. Our results are robust to instead proxying for household socioeconomic status using income, a more continuous measure of education, or the interaction between household income and education. Consistent with Ogden et al. (2010), we find that disparities across education groups are larger than those across income groups.

⁴The relationship between disparities in access and nutritional consumption has been largely ignored by economists, who have instead focused on the role of differential price elasticities in generating disparities in consumption (Jones (1997); Bertail and Caillavet (2008); Park et al. (1996)).

and access is not sufficient to uncover the role that access plays in generating nutritional disparities. To highlight this challenge, we present a simple model that nests two mechanisms, one driven by access and one driven by demand, each of which can independently explain socioeconomic disparities in purchases. On the supply side, even if preferences are homothetic and identical, high-socioeconomic status (SES) households will purchase healthier bundles than low-SES households if they are more likely to live in locations where the cost of accessing healthy food is lower. On the other hand, if households sort by SES, demand-side factors that lead high-SES households to purchase healthier bundles—such as non-homothetic preferences or social norms—could themselves generate differences in access via preference externalities.

Our model motivates two complementary analyses that exploit the detailed nature of our data to identify the causal impact of access to healthy foods on nutritious consumption. Our first empirical strategy is a time-series analysis that measures the short-run responses of households to observed changes in access. Recent studies measuring the effects of changes in retail landscapes on food purchases are local in scope, looking at either the entry of a few supermarkets or an intervention to increase the availability of nutritious food products in a single urban food desert, and find modest effects (see, for example, Song et al. (2009), Weatherspoon et al. (2013), and Elbel et al. (2015)).⁵ Generalizing these local estimates, we find that the elasticity of the healthfulness of household food purchases with respect to the density and nutritional quality of retailers in the household's vicinity is positive but close to zero. Our results demonstrate that providing the typical low-SES household with the retail environment of the average high-SES neighborhood would decrease the gap in nutritional consumption across these groups in the short run by less than 5%. Since improvements in access to healthy foods are more likely to occur in close proximity to sample households with growing tastes for these products, we expect the impacts of policy-induced changes in retail environments to be even more limited than the effects we measure using endogenous changes in access.

Despite limited short-run responses to improvements in access, it is possible that nutritional disparities would be reduced over time as low-SES households benefit from continued exposure to expanded retail access. Our second empirical strategy is cross-sectional and measures the long-run effect of equivalent access by comparing the disparities in nutritional consumption that exist across all sample households to the disparities that persist across households living in the same neighbor-

⁵A notable exception is Freedman and Kuhns (2016), who examine whether the federal government's New Markets Tax Credit (NMTC) influenced supermarket entry using a regression discontinuity design based on the program's median family income threshold. Consistent with our results, they find no systematic difference in household purchases across tracts on either side of this threshold. However, without data on household purchases prior to the introduction of the program, they are unable to look at the response of household purchases to supermarket entry.

hood or shopping in the same store. If differential access is entirely to blame for nutritional disparities, then any systematic differences in the nutritional quality of household purchases that we observe when looking across the entire US should disappear when we compare households subject to the same retail environment. On the contrary, even when we control for residential or retail location, we observe socioeconomic disparities that are 82-87% as large as those that exist in the full cross-section. If tastes vary with unobservable household characteristics, and households sort into residential and retail locations according to these tastes, then observed within-location disparities will underestimate the disparities that would persist if retail access were equalized nationwide. Our results therefore indicate that while access-improving policies will have larger effects in the long run than in the short run, even in the long run such policies will eliminate *at most* one fifth of current socioeconomic disparities in nutritional consumption.

Our paper is related to the literature in economics that uses wide-spread changes in built environments to examine the relationship between retail environments and obesity (Currie et al. (2010); Anderson and Matsa (2011); Courtemanche and Carden (2011); Eid et al. (2008)) but departs from these previous studies in three important dimensions. First, we are concerned not just with the relationship between access and nutritional consumption, but rather the interaction between access, nutritional consumption, and household SES.⁶ This is important for evaluating the effectiveness of current policies, as recent efforts to improve access do so with the intent of reducing disparities in consumption across different socioeconomic groups. Second, we look directly at food purchases, the primary mechanism by which we expect changes in retail environments to impact obesity, rather than obesity itself.⁷ Finally, in addition to leveraging time-series variation in retail environments to measure short-run impacts of improvements in retail access, we further use a novel cross-sectional approach to bound the maximal long-run impact of access-improving policies on socioeconomic disparities in nutritional consumption.

Dubois et al. (2014) take a structural approach to study the role that food environments play in generating international gaps in food purchases between the US, UK, and France. In their context, a structural demand system is necessary to measure how the purchases of households of different nationalities would adjust to being placed in the same retail environment. Our setting and data

⁶Currie et al. (2010) examine differences by race and education and find that the impact of fast food entry on weight gain is greatest among African American mothers and mothers with a high school education or less. In our time-series analysis, we find that more educated households respond slightly more to improvements in access to healthful foods. These differential findings are consistent with the evidence presented by Chen et al. (2010) and Volpe et al. (2013) showing that the impact of store entry depends on neighborhood characteristics and the type of store entering.

⁷While there is a large literature in economics on the relationship between SES and various health behaviors that are known to contribute to obesity (e.g., Cutler et al. (2003), Cutler and Lleras-Muney (2010), and Grossman (2015)), grocery purchases are one health behavior which has received little attention.

allow us to use more direct, reduced-form approaches that exploit the fact that we (i) can directly measure how household purchases adjust to measurable time-series variation in retail environments and (ii) observe different household types of interest already subject to the same retail environment. In fact, in recent work, Alcott et al. (2015) estimate the demand system from Dubois et al. (2014) in the context of consumption disparities across income groups in the US and confirm our results.

More broadly, our work contributes to a growing literature that studies the causes and consequences of inequality across different socioeconomic groups in the US. Recent work highlights the role that differences in environmental toxins, school quality, and neighborhoods play in generating socioeconomic disparities in health, education, and labor market outcomes (Aizer et al. (2016); Currie and Walker (2011); Currie et al. (2015); Ludwig et al. (2011); Chetty et al. (2011); Chetty and Hendren (2016)). While these studies show what neighborhoods do, our study instead provides complementary evidence of what neighborhoods do not: contrary to the popular narrative surrounding food deserts, we find that inequality in access to healthy foods is not driving the large socioeconomic disparities in nutritional consumption that we observe.

In the context of products provided in private markets characterized by increasing returns, our results indicate that the direction of causality in the relationship between neighborhood characteristics and outcomes of interest may indeed be the reverse. That is, disparities in retail access are not the result of supply-side market failures that in turn cause socioeconomic disparities in outcomes but are instead due to efficient supply-side responses to spatial differences in demand. Given the role of fixed costs in the US supermarket industry documented by Ellickson (2006) and Hottman (2014), we expect that home-market effects and preference externalities (Helpman and Krugman (1985) and Waldfogel (2003)) explain the existence of the observed disparities in access.⁸

The paper proceeds as follows. In Section 2, we describe the datasets that we use. In Section 3, we document (i) how the nutritional quality of purchases varies with household SES and (ii) how access to nutritious foods varies across markets with different socioeconomic profiles. In Section 4, we present a simple theoretical framework to demonstrate how the detailed nature of our data can be used in two complementary analyses to bound the role that access plays in generating consumption disparities. Section 5.1 implements our time-series approach and examines whether we observe the healthfulness of household purchases responding to changes in local access. Section 5.2 takes an alternative, cross-sectional approach and examines whether consumption disparities persist when we control for residential or retail location. In Section 6, we provide a discussion of

⁸Our evidence on the relevance of demand-side factors in explaining product availability relates to Dingel (2014) who shows that home-market demand explains as much of the positive relationship between local income and the export quality of US cities as other supply-side factors.

our results and conclude.

2 Data

We combine six datasets that together describe the nutritional quality of grocery purchases that households make and the retail environments that they face. We supplement this data with two additional datasets to demonstrate both the relevance and generalizability of our results. We introduce each dataset and highlight the features most relevant for our analysis below. The interested reader may refer to Appendix A for additional details.

Our identification strategies rely on intra-household (time-series) and inter-household (cross-sectional) variation in household food purchases observed in the Nielsen Homescan data. The Homescan data contains transaction-level purchase records for a representative panel of 114,286 households across the US between 2006 and 2011.⁹ Households in the panel use a scanner to record all of their purchases at a wide variety of stores where food is sold. After scanning the Universal Product Code (UPC) of each item purchased, households record the date, store name, quantity purchased, and price. For items that do not have a standard UPC, households record the purchase in the relevant “random weight” category, such as “fish” or “candy.”¹⁰ As these categories are too broad to infer meaningful nutritional information, we only consider products with standard UPCs in our primary analysis. We demonstrate that our results are robust to the inclusion of random weight items in 2006, when the random weight category definitions were more precise.

The Homescan data has three features that are important for our analysis. First, we observe household demographic data reported on an annual basis, allowing us to measure the SES of each sample household. Second, we observe household purchases for up to 72 months, with the typical household appearing in our sample for 26 months. This time-series variation allows us to measure the responsiveness of household consumption to changes in their retail environment. Finally, we observe the census tract in which each household resides. We use this information to measure the degree to which socioeconomic disparities in nutritional consumption persist when we control for each household’s retail environment.

Since the Homescan data only includes the stores in which panelists shop and the products that they purchase at these stores, it provides a limited picture of local retail environments. Two

⁹We clean the sample of households provided by Nielsen (see Appendix A). The final sample used in the majority of our analyses includes 99,524 households; refer to Table A.1 for demographic summary statistics for this sample.

¹⁰Random weight items include both healthful (e.g., fruits) and unhealthful (e.g., baked goods) products. Reassuringly, household expenditure shares on random weight items (all, healthful only, and unhealthful only) do not vary systematically with household SES. We further note that many fresh items, such as cartons of fresh strawberries, have standard UPCs.

additional datasets, both maintained by Nielsen, provide a more comprehensive picture of the retail environments that households face. The Nielsen TDLinx data, a geo-coded census of food stores in the US, enables us to calculate concentration indexes that summarize the total number of stores to which households have access. To calculate indexes that depict both the nutritional quality and the prices of products offered at a subsample of these stores, we use the Nielsen Scantrack data. The Scantrack data contains weekly sales and quantities by UPC collected by point-of-sale systems located in over 30,000 participating retailers across the US. Unfortunately, the Scantrack data over-samples grocery stores and drug stores and does not track sales of random weight products. When our results rely on the Scantrack data, we note how these limitations influence their interpretation.

We merge the Nielsen data with three external datasets to obtain UPC-level nutritional information, tract-level travel times to stores within 40km, and tract-level neighborhood demographics. To measure the healthfulness of products purchased by Homescan panelists and offered in Scantrack stores, we use IRI's nutritional database that contains the quantity of macro-nutrients and vitamins per serving, serving size in weight, and the number of servings per container at the UPC-level. To account for the time that it takes households residing in different census tracts to arrive at stores in the TDLinx data, we collect driving and transit times from Google Maps for all stores within 40km of each census tract centroid. Finally, to examine how both the number of stores and the quality and prices of available products vary with neighborhood SES, we use the five-year pooled (2007-2011) American Community Survey (ACS) to measure demographics across census tracts.

Two additional sets of data help to establish the validity and relevance of our results more generally. To demonstrate that our measures of nutritional quality correlate with health outcomes such as BMI and hypertension, we use three waves of the National Health and Nutrition Examination Survey (NHANES; 2005-2006, 2007-2008, 2009-2010). The NHANES data combines objective measures of health with self-reported recall of two-day food consumption at the individual level. To demonstrate that our results speak to general patterns of nutritional consumption, we use the USDA's National Household Food Acquisition and Purchase Survey (FoodAPS) data. The FoodAPS data contains information on *all* food purchases made during a single week—that is, products with and without standard UPCs for either consumption at or away from home—for a sample of 4,826 households.

3 Socioeconomic Disparities in Nutritional Consumption and Access

In this section, we use data describing the nutritional quality of food purchases made by households across the entire US to provide the most thorough depiction of socioeconomic disparities

in nutritional consumption to date. Combining data on the spatial distribution of stores, availability of nutritious products, and prices of healthy and unhealthy foods, we then provide an equally comprehensive depiction of spatial disparities in access.

3.1 Disparities in Nutritional Consumption

We begin by documenting the extent of socioeconomic disparities in nutritional consumption across households. Throughout, we focus on the *quality* rather than the quantity of food a household purchases since the latter is affected by the extent to which a household eats at restaurants, and a propensity for eating out is likely related both to household characteristics and access to restaurants and food stores.¹¹ We measure the quality of household purchases using two complementary indexes, both of which are calculated at a monthly frequency for each household in our sample. Our first index, the “nutrient score,” measures the extent to which a household’s grocery purchases deviate from the nutrient composition recommended in the federal Dietary Guidelines for Americans (DGA). Our second index, the “expenditure score,” measures the extent to which a household’s grocery purchases deviate from the expenditure shares recommended by the USDA Center for Nutrition Policy and Promotion’s “Thrifty Food Plan” (TFP). The expenditure score follows the measure introduced by Volpe et al. (2013) and recently used by Oster (2017). As results are consistent across indexes, we only present the nutrient score here. The interested reader may refer to Appendix C to view results using the expenditure score.

The nutrient score for the grocery purchases recorded by household h in month t is defined as

$$Nutrient\ Score_{ht} = \left[\sum_{j \in J_{Healthful}} \left(\frac{pc_{jht} - pc_j^{DGA}}{pc_j^{DGA}} \right)^2 \Big|_{pc_{jht} < pc_j^{DGA}} + \sum_{j \in J_{Unhealthful}} \left(\frac{pc_{jht} - pc_j^{DGA}}{pc_j^{DGA}} \right)^2 \Big|_{pc_{jht} > pc_j^{DGA}} \right]^{-1}$$

where j indexes nutrients, pc_{jht} denotes the amount of nutrient j per calorie in household h ’s grocery purchases in month t , and pc_j^{DGA} is the amount of nutrient j in the DGA recommended diet per calorie consumed. These recommendations are summarized in the FDA’s instructions for how to make use of nutritional labels.¹² We assign nutrients for which the recommendation is

¹¹We use the FoodAPS data to demonstrate that socioeconomic disparities in nutritional quality are similar across food purchases for consumption at home and all food purchases.

¹²Refer to <http://www.fda.gov/Food/IngredientsPackagingLabeling/LabelingNutrition/ucm274593.htm> for these

an upper bound to the unhealthful category (total fat, saturated fat, sodium, and cholesterol) and nutrients for which the recommendation is a lower bound to the healthful category (fiber, iron, calcium, Vitamin A, and Vitamin C). Although there is significant debate over which nutrients are important for health, the goal of our measure is to capture nutritional characteristics that are both salient to consumers (i.e., are reported on nutritional labels) and can reasonably impact the consumption behavior of informed consumers (i.e., are included in nutritional recommendations).

The nutrient score penalizes households for purchasing less (more) than the recommended amount of healthful (unhealthful) nutrients per calorie. To account for differences in the units in which nutrients are measured, we normalize the deviations of household nutrient purchases from the DGA’s recommendations. We follow Volpe et al. (2013) and summarize the normalized deviations using an inverse squared loss function with equal weighting across nutrients.

While useful for analysis, one drawback of indexes in general is that they can be difficult to interpret. To demonstrate that our nutrient score accords with intuition, Table 1 shows how this measure of nutritional quality varies across three sample bundles. The first bundle consists of only healthy products (broccoli, low-fat yogurt, boneless chicken breast, etc.); the second bundle contains a mix of healthy and unhealthy products; and the third bundle consists of only unhealthy products (potato chips, bacon, Oreo cookies, etc.). We determine the food products included in each bundle by selecting among the most widely purchased UPCs in each of the TFP’s 13 healthful and 10 unhealthful food categories. Full lists of the products in each bundle and the TFP food categories from which they are drawn are provided in Tables A.4 and A.5, respectively.

Table 1: Healthfulness of Sample Bundles

Sample Bundle:	Nutritional Quality		
	Healthy	Mixed	Unhealthy
Nutrient score	0.85	0.77	0.2
Total calories	12,160	15,343	18,525
Total calories per ounce	25.75	32.84	40.08
Fat (grams per 100 cal.)	3.2	4.61	5.54
Expenditure share: soda	0.00%	4.26%	7.59%
Expenditure share: fruits & vegetables	21.66%	9.49%	0.00%

Notes: The above table shows how measures of nutritional quality vary across the three sample bundles defined in Table A.4. In computing expenditure shares, we use the average national price for each item in the bundle.

Table 1 yields three takeaways. First, as expected, the healthy bundle has a higher nutrient score than the mixed bundle, which in turn has a higher nutrient score than the unhealthy bundle.

recommendations; last accessed on December 4, 2014.

Second, the nutrient score correlates well with measures of nutrition commonly used in the literature, including fat per calorie and expenditure share on fruits and vegetables.¹³ Finally, lower nutrient scores are associated with higher calorie bundles. Though we prefer our nutrient score as a comprehensive and objective index of the nutritional quality of calories households purchase, it is reassuring that our index is correlated with these recognizable measures of healthfulness.

We are interested in the extent to which the nutritional quality of household purchases varies systematically with household SES. In most of our analyses, we use an indicator denoting whether at least one household head has a college degree to proxy for household SES. We make this empirical choice for three reasons. First, as will be shown below, the college divide reveals significant disparities in our sample that highlight the relationship between household SES, nutritional consumption, and access. This is consistent with a large literature in health economics documenting the relationship between education, health behaviors, and numerous health outcomes (see, for example, Cutler and Lleras-Muney (2010) and Grossman (2015)). Second, without exogenous shocks to income or education, we cannot speak to the relative causal contributions of these household characteristics to nutritional consumption. By focusing on a single dimension of household SES, we highlight disparities along socioeconomic lines while respecting the limitations of our data. Finally, measurement issues are of greater concern when working with household income. Not only is education more precisely measured in the Homescan data (education is reported in years whereas income is reported in bins), but there is more intra-year variability in income than education and Homescan panelists only report socioeconomic variables at an annual frequency.

The raw averages reveal significant disparities: the average nutrient score of college-educated households is 1.29, 18% of a standard deviation higher than the average nutrient score of households without a college education at 1.04.¹⁴ To absorb the effects of seasonality and nationwide trends, in Table 2 we regress log household-month nutrient scores on measures of household SES and year-month fixed effects. Column (1) shows that there is a statistically significant association between college attainment and nutritional consumption: college-educated households purchase healthier bundles than households in which neither household head has a college degree. Column (2) shows that this relationship persists—in magnitude and significance—conditional on household demographics.

As previously noted, we use the college divide as our primary proxy for household SES. To demonstrate that this empirical choice captures meaningful differences that are comparable to those

¹³Refer to Table A.3 for correlations between the household nutrient scores used in our analysis and common measures of nutritional quality.

¹⁴Table A.2 provides summary statistics of the nutritional quality of household purchases, both in aggregate and by SES. Refer to Figure A.1 for average household nutrient scores by deciles of household education.

Table 2: Household Characteristics and Nutritional Quality of Purchases

	Ln(Nutrient Score)						
	Homescan						FoodAPS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
College-Educated	0.168*** (0.0046)	0.161*** (0.0046)			0.138*** (0.0047)	0.0942*** (0.0077)	0.241** (0.050)
Education			0.0928*** (0.0023)				
Ln(Income)				0.0713*** (0.0023)	0.0471*** (0.0024)		
Observations	2,553,494	2,553,494	2,553,494	2,553,494	2,553,494	292,283	3,800
R^2	0.012	0.017	0.013	0.009	0.013	0.004	0.013
Random Weight	No	No	No	No	No	Yes	N/A
Demo. Controls	No	Yes	No	No	No	No	No

Notes: Standard errors are in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. All variables are standardized by the variable's standard deviation. Columns (1)-(6) use the Homescan data; column (7) uses the FoodAPS data. Column (6) uses nutrient scores for 2006 only that include random weight purchases. When the Homescan data is used, observations are at the household-month level, standard errors are clustered by household, and year-month fixed effects are included. Observations in the FoodAPS data are at the household level. Column (2) includes controls for household demographics, including household size dummies, average head of household age, a dummy for marital status of household heads, dummies for households with either a female or male household head only, a dummy for the presence of children, and dummies for whether the household reports being white, black, Asian, or Hispanic. All specifications include expenditure weights. Refer to Table A.6 for regression results showing disparities across individual nutrients.

measured using alternative proxies for household SES, columns (3)-(5) replicate the analysis from column (1) using continuous education, log income, or both log income and the college-educated indicator as independent variables, respectively.¹⁵ Regardless of the proxy for household SES that we use, we see that the nutritional quality of purchases is increasing in household SES. Furthermore, comparing the R-squareds, we see that the college-educated indicator explains a similar amount of the variation in household purchases as does log income or continuous education. Interestingly, the standardized coefficients reported in column (5) indicate that college attainment explains more of the variation in household nutrition than income: the nutritional quality of household purchases varies more across college groups conditional on income than across income groups conditional on college attainment.

Columns (1)-(5) of Table 2 only use nutrient scores computed for non-random weight purchases (that is, for products with standard UPCs). In column (6), we replicate the analysis from column (1) using nutrient scores computed using both random weight and non-random weight purchases

¹⁵Since income is log-normally distributed whereas education is normally distributed, we include income in logs and education in levels.

in 2006.¹⁶ Despite a much smaller sample size, we still see that households in which at least one household head has a college degree purchase food products that accord more closely with recommendations for nutritional intake than less educated households.

The Homescan data can only speak to nutritional disparities in food for consumption at home. If nutritional disparities in food away from home (at restaurants, schools, etc.) follow a different pattern, then the disparities documented in columns (1)-(6) of Table 2 may not speak to food purchases more generally. In column (7) of Table 2, we replicate the analysis from column (1) using information on the quality of *all* food purchases documented in the FoodAPS data. We see that nutritional disparities are comparable whether we consider food for consumption at home or all food purchases: more educated households purchase food products with higher nutrient scores. If anything, socioeconomic disparities in nutritional consumption are even more pronounced when we consider food for consumption both at and away from home.

The consumption disparities documented in Table 2 translate into meaningful differences in health outcomes. Using the NHANES data, we calculate nutrient scores for respondents based on their two-day food recall. We then regress indicators for various health conditions on individual-level nutrient scores and demographic controls. Dividing the standardized coefficient estimates shown in Table 3 by the mean levels of each dependent variable, we see that a one standard deviation increase in nutrient scores is associated with an 11% reduction in the probability of being obese, a 25% reduction in the probability of having diabetes, and a 5% reduction in the probability of having hypertension. Furthermore, we see that our nutrient score explains an order of magnitude more of the variation in these health outcomes than the Healthy Eating Index (HEI)—a measure of diet quality commonly used to measure conformance to federal dietary guidelines.¹⁷ While a two-day food recall need not be indicative of an individual's regular diet, we take these results as evidence that the socioeconomic disparities in nutritional consumption that we observe are important for understanding differences in health outcomes across these groups.

¹⁶For the 43 random weight categories in 2006, we use the average nutritional characteristics of products with UPCs in each category to infer the nutritional content of purchases. Recall that the random weight categories in 2007-2011 are too broad to infer meaningful nutritional information for random weight items.

¹⁷In the NHANES data, our nutrient score and the HEI are highly correlated (correlation coefficient of 0.55). While the HEI is commonly used in the literature, we prefer our nutrient score for two reasons: (i) the HEI conflates differences in quantity and quality and (ii) our nutrient score explains relatively more of the variation in outcomes of public health concern.

Table 3: Nutritional Quality of Consumption and Health Outcomes in NHANES

	Obesity		Diabetes		Hypertension	
	(1)	(2)	(3)	(4)	(5)	(6)
Nutrient Score	-0.039*** (0.005)	-0.036*** (0.006)	-0.027*** (0.003)	-0.040*** (0.004)	-0.017*** (0.005)	-0.010* (0.005)
Healthy Eating Index		-0.001 (0.000)		0.002*** (0.000)		-0.001** (0.000)
Observations	9,527	9,527	9,527	9,527	9,527	9,527
R^2	0.045	0.045	0.091	0.095	0.238	0.239
Mean Dep. Var.	0.364	0.364	0.108	0.108	0.331	0.331

Notes: Standard errors are in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The above table displays the output from a linear probability model of indicators for various health outcomes on individual-level nutrient scores and the Healthy Eating Index. These measures of nutritional quality are calculated using two-day food recalls as reported in NHANES. Explanatory variables are standardized by the variable’s standard deviation. All regressions include controls for log income, education, a cubic in age, and indicators for whether the respondent is male, white, black, or Hispanic. We exclude children aged 15 and under. To proxy for expenditure weights, we weight by total calories in all specifications; results are robust to the use of alternative weights.

3.2 Spatial Disparities in Access

We now turn to documenting disparities in access to healthy foods across neighborhoods with different socioeconomic profiles. We characterize retail environments using indexes that reflect the number of stores consumers have access to, the healthfulness of the products available in these stores, and the prices of both healthy and unhealthy products offered by these stores. Analogous to our household-level analysis, we primarily use the share of college-educated residents to proxy for neighborhood SES.¹⁸ We document that our results are robust to instead using either (i) median household income or (ii) an indicator denoting neighborhoods that have both an above median share of college-educated residents and above median income to measure SES.

3.2.1 Store Concentration

We begin with simple concentration indexes that reflect the spatial distribution of retail food stores in and around each census tract in the US. The concentration indexes are kernel densities based on store locations from the TDLinx data and driving times computed using Google Maps. Let d_{sl} denote the driving time between store s and the centroid of census tract l and S_t the universe of

¹⁸In many tables and figures, we divide tracts into “high” and “low” education groups. Tracts are considered high education if their share of college-educated residents falls above the median share across all tracts (21.4%) and low education otherwise. 52% of tracts are high education and 48% are low education (Table A.7).

stores in our sample in year t . We define the concentration index for census tract l in year t as a Gaussian kernel with a bandwidth of 10 minutes of driving time ignoring all stores further than 40km from the tract centroid:¹⁹

$$Concentration\ Index_{lt} = \sum_{s \in S_t} w_{sl} \quad \text{where} \quad w_{sl} = \begin{cases} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2} \left(\frac{d_{sl}}{10}\right)^2} & \text{if } distance_{sl} \leq 40\text{km} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

To examine how store concentration varies systematically with neighborhood characteristics, we combine these indexes with tract demographics from the ACS. Raw averages already suggest significant spatial correlation between education and store concentration: households in tracts with an above versus below median share of college-educated residents face concentration indexes that are on average over 18% of a standard deviation higher.²⁰

In Table 4, we show how these concentration indexes vary with the share of college-educated residents in a tract. As shown in column (1), local education is positively associated with store concentration.²¹ Columns (2)-(7) display the relationship between the share of college-educated residents and concentration indexes that each reflect the density of stores of a certain type. The results in column (1) do not mask significant heterogeneity across most store types. In fact, high-SES neighborhoods have a greater concentration of all store types other than dollar stores, whose sales make up less than 1% of sales for in-home food consumption over our sample period, and club stores. Since club stores draw from a geographically disperse market, their location decisions are more likely driven by land costs and road accessibility than smaller and more local grocery,

¹⁹This bandwidth was selected to match the expenditure-weighted distribution of household trips observed in the Homescan data. It implies, as is the case in the data, that a household is approximately 65% as likely to visit a store that is a 10 minute drive from the centroid of their residential census tract as they are to visit a store at the census tract centroid. Refer to Figure A.2 for a comparison of the implied bandwidth weights and the observed distribution of household shopping trips.

²⁰Refer to Table A.7 for summary statistics of the variables we use to measure store concentration, both in aggregate and by tract SES. Figure A.3 displays concentration indexes by deciles of local college-educated shares.

²¹Socioeconomic disparities in store concentration grow as we consider the number of stores within rings that are further from each census tract centroid. For example, while neighborhoods with an above median share of college-educated residents have 5% of a standard deviation more stores within 1-2km of the census tract centroid than tracts with a below median share of college-educated residents, this socioeconomic disparity rises to 30% of a standard deviation when we instead consider the number of stores within 16-32km of tracts (Table A.7). In fact, when considering stores within just 0.5km, low-SES neighborhoods have slightly more stores than high-SES neighborhoods on average (1.38 versus 1.35, respectively). Therefore, if one uses a very small bandwidth when computing the concentration indexes defined in Equation 1, low-SES neighborhoods can appear to have a higher concentration of stores than high-SES neighborhoods. This is consistent with the results of Powell et al. (2007) who find that low-income zip codes have more non-chain supermarkets and grocery stores than higher income zip codes. Despite these patterns of store concentration within very small geographic bands, the evidence suggests that households primarily travel beyond their nearest store when making food purchases (Ver Ploeg et al. (2015); Rahkovsky and Snyder (2015)). We therefore believe that very small bandwidths overstate access in low-SES relative to high-SES neighborhoods.

Table 4: Neighborhood Characteristics and Store Concentration

	(1) All	(2) Grocery	(3) Conven.	(4) Drug	(5) Club	(6) Dollar	(7) Mass
Ln(College Share)	0.13*** (0.005)	0.16*** (0.005)	0.11*** (0.005)	0.15*** (0.005)	-0.0002 (0.005)	-0.07*** (0.007)	0.15*** (0.006)
Observations	36,951	36,470	36,489	36,398	36,640	19,860	31,175
R^2	0.018	0.024	0.013	0.024	0.000	0.005	0.023

Notes: Standard errors are in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Observations are at the tract-year level. All variables are standardized by the variable's standard deviation. These results are for 2010; they are representative of other years in the TDLinx sample. Refer to Table A.10 for analogous results using median income or the interaction of indicators denoting tracts with above median income and above median college-educated shares as alternative proxies for neighborhood SES.

convenience, and drug stores. Note that the spatial distribution of stores across neighborhoods looks similar whether we use the share of college-educated residents (Table 4), median income (Table A.10, Panel A), or the interaction of indicators denoting tracts with above median income and above median college-educated shares (Table A.10, Panel B) to proxy for neighborhood SES.

3.2.2 Store Inventory

Our concentration indexes allow us to examine disparities in the spatial distribution of retailers but fail to account for the fact that not all stores are equal. Importantly, stores may differ in the products they offer, even within store types. Before proceeding to a formal analysis of store inventory, in Table 5 we explore how availability differs across neighborhoods for the three sample bundles introduced in Section 3.1. The left panel of Table 5 shows the percentage of census tracts in which the entirety of a given bundle can be found across stores, conditional on there being at least one Scantrack store in the tract that we are able to match to location information in the TDLinx data.²² We see that the unhealthy bundle is more likely than the healthy bundle to be available across all tracts. Comparing availability across tracts with different socioeconomic compositions, we further see that each bundle is most likely to be available in tracts with a higher share of college-educated residents. In fact, both the healthy and unhealthy bundles are over six percentage points more likely to be available in tracts with above versus below median share of college-educated residents.

²²Appendix A provides details on this match.

Table 5: Availability and Cost of Sample Bundles

Bundle:	Availability (% of Tracts)			Avg. Cost Per 100 Cals. (Std. Dev.)		
	Healthy	Mixed	Unhealthy	Healthy	Mixed	Unhealthy
All Tracts	48.3%	48.3%	79.8%	0.43 (0.05)	0.35 (0.03)	0.3 (0.03)
High Educ.	51.0%	51.0%	82.4%	0.45 (0.05)	0.36 (0.03)	0.3 (0.03)
Low Educ.	44.5%	44.4%	76.1%	0.41 (0.05)	0.34 (0.03)	0.29 (0.02)

Notes: The above table presents the availability and cost at the tract level of the bundles defined in Table A.4. Bundle availability is calculated as the share of tracts that offer all the items (or similar products) listed in the corresponding bundles, while bundle cost is calculated as the expenditure required to buy the bundle where the price of each component is equal to the average price charged for similar products for each bundle item in a given tract-month. Similar products are defined as products in the same product module whose description contains the same key words as the description of the exact item in the bundle. For example, similar products for “Tuna-Chunk Light” are products in the “SEAFOOD-TUNA-SHELF STABLE” module with a description containing the key words “TUNA WTR CHK LT”. See Footnote 18 for a description of how tracts are separated according to education levels.

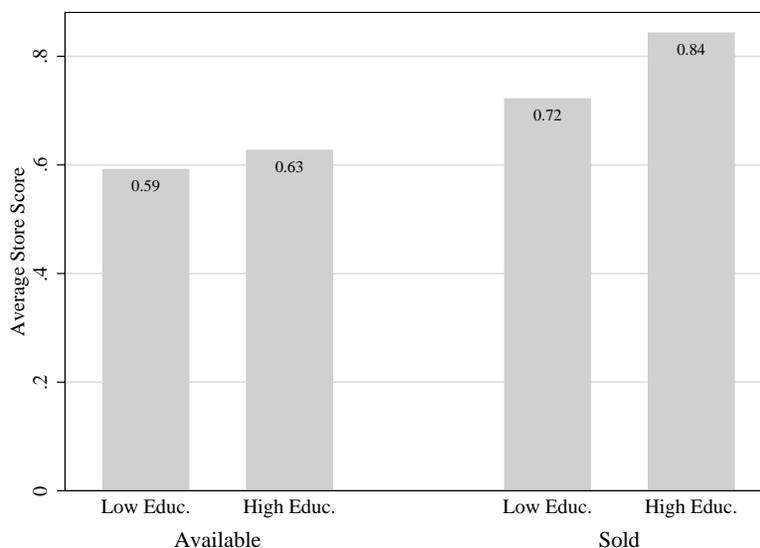
To measure spatial differences in inventory, we compute healthfulness indexes for each of the stores in the Scantrack-TDLinx matched sample. These indexes summarize the nutritional content of the products offered in a given store in a given month using a store-level variant of the nutrient score defined for households in Section 3.1 (see Appendix D for the formula). This store-level nutrient score reflects the per calorie nutrients that a representative household would purchase in store s in month t . The representative household purchases *all* of the products available in a store such that their relative UPC-level expenditure shares for that store reflect the national average. Variation across store-level nutrient scores therefore comes only from differences in the mix of UPCs available across stores, which we infer from whether a product was sold by a store in a given month, not from continuous differences in the quantities sold.

Differences in the store-level nutrient indexes are small. The average nutrient score of stores located in census tracts with an above median share of college-educated residents is 0.63 versus 0.59 across stores in census tracts with a below median share of college-educated residents.²³ To assess the magnitude of this difference, we benchmark the gap in nutritional availability to the gap in sales measured using nutrient scores calculated with store-sales rather than national-sales weights (that is, store-level nutrient scores that reflect differences in local demand as well as product availability). In Figure 1, we see that differences in the healthfulness of products sold across neighborhoods are much more pronounced than differences in the healthfulness of the products available. The disparity in the nutritional quality of products sold across neighborhoods with an above versus below median share of college-educated residents is *three times* as large as the dis-

²³Refer to Table A.8 for summary statistics of the variables we use to assess product availability and store sales, both in aggregate and by tract SES. Figure A.3 displays kernel densities of store-level nutrient scores by deciles of local college-educated shares.

parity in the nutritional quality of products offered in stores across these neighborhoods.

Figure 1: Nutrient Scores Across Census Tracts: Available vs. Sold



Notes: The above figure presents raw store-level nutrient score averages, computed using either national-sales weights (left) or store-sales weights (right), across census tracts with different socioeconomic compositions. See Footnote 18 for a description of how tracts are separated according to education levels. These results are for January 2010; they are representative of other months in the Scantrack sample. A meticulous reader may wonder whether it is possible for the nutrient scores of a nationally representative consumer to be lower than the nutrient scores of bundles actually sold across all neighborhoods. This is not an error but rather an artifact of a skewed distribution of store-level nutrient purchases combined with an index that does not reward healthy deviations.

Recall that the Scantrack sample under-represents convenience stores and other small retailers. Table 4 shows that high-SES neighborhoods have a greater number of stores of nearly all types (including convenience stores). The fact that the estimated disparity is consistent across store types indicates that the composition of store types is very similar across neighborhoods. Therefore, including more convenience stores—small retailers that have lower nutrient scores on average than other store types (Figure A.4)—should impact all neighborhoods similarly, thereby affecting the level of average nutrient scores but not the disparity across neighborhoods with different socioeconomic compositions. If anything, since there is less variability in nutrient scores across convenience stores than other store types (Figure A.4), better coverage of these retailers would condense the distribution of average store-level nutrient scores across tracts. In fact, if we use the average store-level nutrient scores of each store type in the Scantrack data to impute store-level nutrient scores for all stores in the TDLinx data, the disparity between high-SES and low-SES neighborhoods is even more limited (see Figure A.5).

Further recall that the Scantrack data does not track random-weight UPC sales. The direction of the bias introduced by this data limitation will depend on whether healthy products without standard UPCs are more likely to be offered in low-SES or high-SES neighborhoods. Note that the

majority of our main results below do not rely on disparities estimated using the Scantrack data and thus are not subject to this issue. Where appropriate, we note where our results do rely on these estimated disparities and provide a variety of robustness checks to help mitigate any concerns.

3.2.3 Store Pricing

Given the disconnect between the nutritional quality of products available and the nutritional quality of products sold across neighborhoods, it is unlikely that differences in product availability alone drive the observed differences in sales. That being said, there are other store policies, such as pricing, that may also influence household purchases. For example, even though a product is on the shelf in a low-SES neighborhood, the product may be prohibitively expensive such that the item is not truly “accessible” to households in that neighborhood. The Scantrack data includes both revenue and quantity information across products, allowing us to use inferred prices to examine whether differential store pricing patterns might explain purchase disparities.

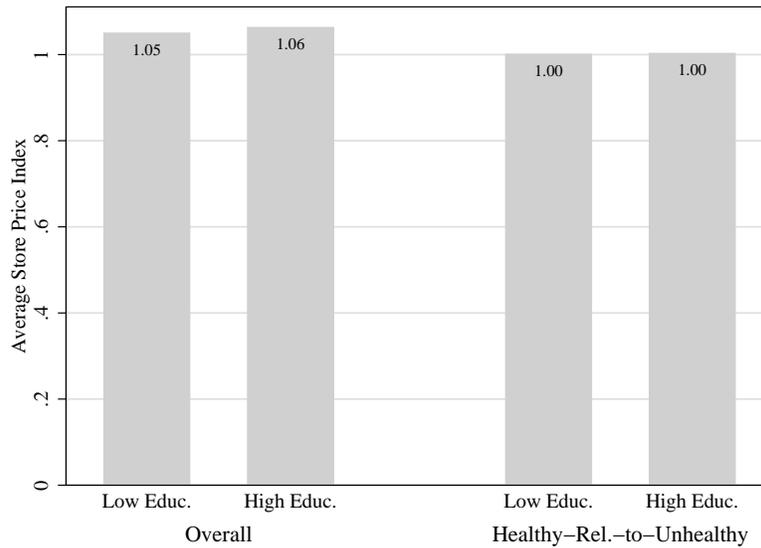
One commonly cited hypothesis for why low-income consumers eat less healthy foods is that unhealthy calories are less expensive than healthy calories.²⁴ Since low-income consumers face tighter budget constraints and food is a necessity good, they will allocate more of their expenditure towards cheaper, less healthful foods than high-income consumers. While relative prices of healthy versus unhealthy food products may be a key driver of nutritional disparities in general, they are only relevant for this paper insofar as the pricing practices of the stores in low-SES neighborhoods lead low-SES households to purchase more unhealthy foods than they would if they had access to the prices offered by stores in high-SES neighborhoods.

If store pricing is to blame for the relative unhealthfulness of sales in low-SES neighborhoods, it must be the case that either (i) these stores charge higher prices for all food products, limiting their customers’ consumption possibilities and forcing them to allocate more of their expenditure than they would otherwise towards cheaper products, or (ii) these stores charge relatively more than stores in high-SES neighborhoods for healthful relative to unhealthful food products.²⁵ We explore these hypotheses by looking at both the spatial distribution of prices for all food products and the distribution of prices for healthy relative to unhealthy foods (see Appendix D for a detailed

²⁴In the majority of product groups, we observe that the national average price per calorie of products in healthful TFP food categories is, on average, higher than the national average price per calorie of products in unhealthful TFP food categories. For the sample baskets introduced in Section 3.1, Table 5 displays the average cost per 100 calories of each basket across all tracts and across tracts with different education profiles. Comparing the cost per calorie across bundles, we see that the healthy bundle is on average more expensive than the mixed bundle which in turn is more expensive than the unhealthy bundle.

²⁵Neither of these hypotheses are upheld in our samples bundles: there are no significant differences in the average prices of any of our sample bundles across tracts with different socioeconomic compositions.

Figure 2: Prices Across Census Tracts: All Products and Healthy Relative to Unhealthy



Notes: The above figure presents average store-level price indexes, computed using either all products (left) or the ratio of price indexes for healthy and unhealthy products (right), across census tracts with different socioeconomic compositions. See Footnote 18 for a description of how tracts are separated according to education levels. These results are for January 2010; they are representative of other months in the Scantrack sample.

overview of how these prices indexes are computed). Figure 2 shows that differences in pricing alone cannot be driving consumption disparities: stores in neighborhoods with an above median share of college-educated residents charge slightly more than stores in neighborhoods with a below median share of college-educated residents for all products on average, and healthful foods are no more expensive relative to unhealthful foods in these neighborhoods. If anything, pricing patterns should cause store sales in low-SES neighborhoods to be more, as opposed to less, healthful than store sales in neighborhoods with more educated residents.

We finally note that the manner in which healthful products are presented, including their shelf space and department cleanliness, may also make these products relatively less attractive in certain stores (see, for example, Zenk et al. (2011)). Unfortunately, we do not have the qualitative data required to assess whether these differences can help explain socioeconomic differences in store sales. In Section 5.2, we use fixed effects to control for *all* differences in access across neighborhoods and even across stores to obtain an upper bound on the role that these factors jointly play in explaining socioeconomic differences in household purchases.

4 Theoretical Framework

We have demonstrated that there are large socioeconomic disparities in the nutritional content of household grocery purchases as well as significant, yet more limited, spatial disparities in access

to healthy foods. The direction of causality here is undetermined. It is possible that the disparities in nutritional consumption are due entirely to the fact that more educated households have better access to healthy food than less educated households. It is also possible that households sort into locations where they have access to the food products that they prefer to purchase or, more likely, that households sort into locations based on factors correlated with their demand for grocery products (e.g., willingness to pay for housing, proximity to employment opportunities, schools, etc.), and spatial disparities in product availability arise entirely because stores cater to local demand.

In this section, we introduce a simple and quite general theoretical framework in which household SES and local supply both influence food purchases. This framework formally demonstrates the challenge that the previous literature has faced in identifying the causal link between access and the nutritional quality of household purchases. It also suggests two ways in which we can use the detailed nature of our data to overcome this challenge.

4.1 Set-up

Consider a model with M locations indexed by l . Each location l has a population of equal size N composed of heterogeneous households whose SES, indexed by h , can take one of two values: low (L) or high (H). We rank locations by their share of high-SES households, with higher l locations having larger shares of high-SES households. We assume that the share of high-SES households in a neighborhood is exogenously determined.

Each household decides how much to consume of each product in a set of grocery products indexed by nutritional quality $q = 1, \dots, Q$, where a higher q is associated with a more healthful product. The household can also choose to consume an outside good, z . The household selects a consumption bundle to maximize utility subject to their budget constraint, which is determined by the household's income, y_h , and the cost of accessing food products of each quality in their location l , $p_l(q)$. The cost of access reflects the retail price of food products, travel costs, and storage.

The household's problem is therefore given by

$$\max_{\mathbf{x}, z} U_h(\mathbf{x}, z) \text{ subject to } \mathbf{p}'_l \mathbf{x} + p_l(z)z \leq y_h$$

where \mathbf{x} is a $Q \times 1$ vector of quantities of the differentiated grocery varieties and \mathbf{p}_l is a $Q \times 1$ vector of the prices of these products in location l .

The solution to the household's problem yields a Marshallian demand curve for products of each quality q , $x_h(q, \mathbf{P}_l)$, where $\mathbf{P}_l = (\mathbf{p}_l, p_l(z))$ is the vector of access costs in location l . The possibility that both utility and demand are a function of household SES is accounted for by the

fact that these functions are indexed by h . This accommodates the possibility that there are either SES-specific tastes and/or non-homothetic preferences.

Denote by $s_h(q)$ the share of total grocery expenditures that households with SES h allocate to products of quality q . We can express this across-location expenditure share as

$$s_h(q) = \sum_{l=1}^M \theta_h(l) s_h(q, \mathbf{P}_l) \quad (2)$$

where $\theta_h(l)$ is the share of SES- h households that reside in location l and $s_h(q, \mathbf{P}_l) = \frac{x_h(q, \mathbf{P}_l)}{y_h - p_l(z) z_h(\mathbf{P}_l)}$ is the within-grocery expenditure share for products of quality q for SES- h households residing in location l . The sales-weighted average quality of food products consumed by SES- h households across all locations is therefore given by

$$Q_h = \sum_{q=1}^Q s_h(q) q \quad (3)$$

Fact 1. *If $s_h(q)$ is supermodular in SES h and product quality q (i.e., $\partial s_H(q)/\partial q > \partial s_L(q)/\partial q$), then the average quality of food consumption, Q_h , is increasing in household SES.*

4.2 Mechanisms

Equation 2 highlights the separate roles that access and demand can play in generating the socioeconomic disparities in nutritional consumption that we observed in Section 3.1. We present these distinct mechanisms in the propositions below.

Proposition 1 (Supply-side mechanism). *If (i) demand does not vary with SES (i.e. $s_h(q, \mathbf{P}_l) = s(q, \mathbf{P}_l) \forall h$ in any given market l) and (ii) the spatial distribution of high-SES households is correlated with access to healthful food products (i.e., $\text{Corr}(\theta_H(l), \partial s(q, \mathbf{P}_l)/\partial q) > 0$), then the average quality of food consumption, Q_h , will be increasing in household SES.*

Proof. If $s_h(q, \mathbf{P}_l) = s(q, \mathbf{P}_l) \forall h$ in any given market l , Equation 2 reduces to $s_h(q) = \sum_{l=1}^M \theta_h(l) s(q, \mathbf{P}_l)$. That is, across-location expenditure shares only vary with SES through differences in the spatial distribution of households by SES, $\theta_h(l)$. Since $\theta_H(l)$ and $\partial s(q, \mathbf{P}_l)/\partial q$ are positively correlated across locations (by assumption), $s_h(q)$ is supermodular in SES h and product quality. Therefore, by Fact 1, Q_h is increasing in household SES. \square

Therefore, if high-SES households tend to live in locations where the costs of accessing food products incentivize all households, regardless of SES, to purchase healthier foods, then high-

SES households will buy healthier foods than low-SES households on average, even if the two sets of households have the same tastes. This will hold whenever $\partial^2 p_l(q)/\partial q \partial l > 0$, i.e., whenever healthier food products are sold at lower prices or are more available in neighborhoods with a larger share of high-SES residents. In practice, if tastes do not vary with SES, such cost differences could arise as the result of differences in wholesale and retailing costs. For example, if healthy foods cost more and rents are higher in high-SES neighborhoods, then firms in high-SES neighborhoods will have a comparative advantage in the distribution of nutritious products. In reality, however, it is likely that tastes do vary with SES to some extent, and this pricing and availability pattern arises at least in part because local firms in high-SES neighborhoods cater to local high-SES tastes for these products.

Proposition 2 (Demand-side mechanism). *If (i) supply does not vary across locations (i.e., $\mathbf{P}_l = \mathbf{P} \forall l$) and (ii) high-SES households purchase relatively more healthy products than low-SES households in all locations regardless of access (i.e., $\partial s_H(q, \mathbf{P})/\partial q > \partial s_L(q, \mathbf{P})/\partial q \forall \mathbf{P}$), then the average quality of food consumption, Q_h , will be increasing in household SES.*

Proof. If $\mathbf{P}_l = \mathbf{P} \forall l$, Equation 2 reduces to $s_h(q) = \sum_{l=1}^M \theta_h(l) s_h(q, \mathbf{P}) = s_h(q, \mathbf{P})$. That is, across-location expenditure shares equal within-location expenditure shares for each SES group. Since high-SES households purchase relatively more healthy products than low-SES households in all locations regardless of access (by assumption), $\partial s_H(q, \mathbf{P})/\partial q = \partial s_H(q)/\partial q > \partial s_L(q, \mathbf{P})/\partial q = \partial s_L(q)/\partial q \forall \mathbf{P}$. Therefore, by Fact 1, Q_h is increasing in household SES. \square

High-SES households may purchase relatively more healthy products than low-SES households for a variety of reasons. For $y_H > y_L$, this could be the result of income effects. That is, households with lower incomes may spend more on low-quality products either because they cost less or because there are complementarities between consumption of the outside good z and the quality of grocery products. High-SES households might also spend more on high-quality products because they attain more utility from these products, regardless of their expenditure on the outside good (due to complementarities between education and nutrition, for example). For the purposes of this paper, we remain agnostic as to why high-SES households spend more on healthy foods within locations. We simply seek to measure the role that these demand-side factors, relative to supply-side factors, play in generating the differences in purchases that we observe across households.

4.3 Empirical Approaches

In Section 5, we disentangle these supply-side and demand-side forces with two empirical approaches motivated by the model above. The first, time-series approach turns off any time-invariant

demand-side sources of heterogeneity by estimating how the purchases of households with constant education change over time in response to varying retail environments. The second, cross-sectional approach instead turns off the supply-side source of heterogeneity, i.e., differences in access, by measuring the socioeconomic disparities in nutritional quality that persist across households either living or shopping in the same retail environment.

Time-series approach Consider the framework above with locations recast as markets that are separated by time instead of space. The change in the quality of products purchased by a household with SES h between time t and $t + 1$ is given by

$$Q_h(t + 1) - Q_h(t) = \sum_{q=1}^Q (s_h(q, \mathbf{P}_{t+1}) - s_h(q, \mathbf{P}_t)) q$$

Assuming that a household’s income and tastes are constant over time—or at least over the time horizon that we consider empirically—we can estimate the elasticity of healthfulness with respect to access by regressing changes in the healthfulness of household purchases against variables that summarize store concentration and product availability.

It is possible that tastes vary over time, however, and we expect that changes in availability across markets will be correlated with unobserved changes in the prevalent tastes of local residents. While the tastes of any one household in our panel might not reflect the prevalent local tastes, i.e., a household’s tastes may not change or may change in the opposite direction, we expect that the tastes of our sample households will be correlated and covary with local tastes on average. As a result, our estimate of the elasticity of household purchases with respect to changes in their retail environment is subject to an upward omitted variable bias. We therefore interpret these elasticities as an upper bound for the true elasticity that we expect to govern the response of purchases to improved access that is driven by policy as opposed to endogenous firm responses to changes in market fundamentals.

Cross-sectional approach Within a location, the average quality of products purchased by SES- h households is given by

$$Q_h(l) = \sum_{q=1}^Q s_h(q, \mathbf{P}_l) q$$

Comparing the average quality across high-SES and low-SES households, we have that

$$Q_H(l) - Q_L(l) = \sum_{q=1}^Q (s_H(q, \mathbf{P}_l) - s_L(q, \mathbf{P}_l)) q$$

If high-SES households have relatively higher expenditure shares on high-quality products, then $Q_H(l) > Q_L(l)$ on average across locations. That is, we will observe socioeconomic disparities in nutritional consumption even though the households we are comparing have equivalent access. To the extent that these differences in demand yield preference externalities or home-market effects, differences in aggregate local demand will play a role in generating the correlation between $\theta_H(l)$ and $\partial s_h(q, \mathbf{P}_l)/\partial q$. Looking within locations we ignore these effects, whereby potentially underestimating the role of demand-side factors and, in turn, providing an upper-bound for the role of access. As the correlation between time-variant components of a given household’s demand and their retail access is likely more limited than the correlation between unobservable components of the household’s tastes and where they choose to live or shop, we expect our cross-sectional results to yield a more conservative bound than our time-series approach.

5 Role of Access in Explaining Consumption Disparities

We now implement the empirical strategies suggested by the theory above to identify the causal role of access in explaining consumption disparities across households with different levels of education.²⁶ We first measure how nutritional consumption responds to a changing retail environment by leveraging observed changes in households’ retail environments over our panel. This analysis provides an upper bound on the potential impact of access-improving policies on socioeconomic disparities in consumption in the short run. To examine the maximal long-run potential of policies that equalize access, we then take a cross-sectional approach and compare the disparities that persist across households living in the same residential location or shopping in the same store.

5.1 Changing Retail Environments

Over the six years in our sample, we observe changes in the retail environments of households for three reasons: (i) a household moves to a different census tract with different access, (ii) stores enter and/or exit a household’s neighborhood, and (iii) the stores in a household’s neighborhood change the products they offer. Since household moves are endogenous, and sample households

²⁶Analogous results using alternative proxies for household SES are available upon request.

only report their residential location on an annual basis, we limit our attention to households that did not move throughout the sample.

We measure how the healthfulness of incumbent household purchases responds to continuous changes in retail environments driven by either store entry and exit or changing product mixes. To do so, we regress household nutrient scores against two separate time-varying kernel densities that capture changes in local retail environments.²⁷ The first is the concentration index that we introduced in Section 3.2. Recall that these indexes are kernel densities that measure the driving time-weighted number of stores in a household’s vicinity. As the TDLinx data contains a snapshot of retail environments as of June each year, these concentration indexes are at the tract-year level.²⁸ To measure the nutritional quality of the products offered in local stores, we construct kernel densities of store nutrient scores that measure the driving time-weighted average of healthful product availability in a household’s vicinity.²⁹ We begin by considering nutrient score densities at a monthly frequency.

The results of this time-series analysis are provided in Table 6. In column (1), we see that even after controlling for the concentration and healthfulness of surrounding stores, household purchase quality is still higher among households in which at least one household head has a college degree. Perhaps unsurprisingly, household nutrient scores are positively associated with both store concentration and the healthfulness of product offerings. We cannot interpret these coefficients causally, however, as households may sort spatially by unobservable characteristics that are correlated with tastes for healthy foods. If households with stronger tastes for healthy foods sort into locations where these foods are more accessible, then our coefficients on store concentration and store nutrient scores will be biased upwards. If we assume that household preferences are fixed over the time period that we study (up to six years), we can control for this sorting behavior by including household fixed effects.^{30,31} In this specification, the coefficients are identified from time-series variation in purchases and retail environments (column (2)). Here, we observe that the nutritional quality of the average household’s purchases responds neither to changes in the concentration of retail outlets

²⁷As there is limited variation in price indexes across neighborhoods (Figure 2), we exclude prices from this analysis.

²⁸In this section, years are defined July-June to accord with the TDLinx snapshots.

²⁹The nutrient score densities weight store scores using the same Gaussian kernel used to construct the concentration indexes (that is, a bandwidth of 10 minutes of driving time ignoring stores further than 40km from the tract centroid). Letting S_t denote the universe of stores in time t and N_{slt} the average nutrient score of store s in census tract l in time t , the weighted average nutrient score for census tract l in time t is given by $\frac{1}{\sum_{s=1}^{S_t} w_{sl}} \sum_{s=1}^{S_t} w_{sl} N_{slt}$, where w_{sl} is defined as in Equation (1).

³⁰Since education is nearly constant across our sample period for a given household, we do not control for whether either household head has a college degree when we include household fixed effects.

³¹When we include household fixed effects, we use Conley Spatial HAC standard errors that allow for both temporal and spatial correlation. We use the `reg2hdfespatial` package written by Thiemo Fetzter; available at <http://www.trfetzter.com/conley-spatial-hac-errors-with-fixed-effects>.

Table 6: Response of Nutritional Quality of Household Purchases to Changes in Retail Access

	Ln(Nutrient Score)		
	(1)	(2)	(3)
College-Educated	0.158*** (0.0016)		
Ln(Store Concentration)	0.0120*** (0.00043)	-0.00227 (0.0078)	
* College			-0.00174 (0.0079)
* Non-College			-0.00534 (0.0080)
Ln(Avg. Store Score)	0.136*** (0.0058)	0.0149 (0.0079)	
* College			0.0314*** (0.0092)
* Non-College			0.00230 (0.0087)
Observations	1,601,848	1,580,395	1,580,395
R^2	0.011	0.305	0.305
Household FEs	No	Yes	Yes

Notes: Standard errors are in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Observations are at the household-month level. Column (1) reports two-way clustered standard errors by tract and month; columns (2) and (3) report Conley Spatial HAC standard errors. All regressions include month-year fixed effects. The sample only includes households that resided in the same census tract throughout their entire participation in the Homescan panel.

in the household's vicinity nor to improvements in the nutrient composition of products sold by stores in their neighborhood.

To explore whether the responsiveness of household purchases to changes in retail environments varies by household SES, we interact the access kernel densities with indicators denoting whether the household does or does not have a college-educated household head in column (3). We see that the statistically insignificant average response of household-level nutrient scores to the concentration of stores in their vicinity does not mask any differences across households with different levels of education. We do however observe a socioeconomic disparity in the responsiveness of household nutrient scores to local product offerings, with more educated households improving their consumption more than less educated households when they are offered a more nutritionally-balanced mix of food products in their neighborhood stores.³²

³²This is in contrast to the results of Oster (2017), who finds that individuals with higher levels of education do

Table 7: Response of Household Purchases to Changes in Retail Access: Varying Data Frequency

	Ln(Nutrient Score)				
	(1)	(2)	(3)	(4)	(5)
Ln(Store Concentration)					
* College	-0.00174 (0.0079)	-0.0193 (0.012)	-0.0461* (0.023)	-0.00657 (0.0092)	0.00478 (0.013)
* Non-College	-0.00534 (0.0080)	-0.0237 (0.012)	-0.0434 (0.023)	-0.0102 (0.0094)	0.00430 (0.013)
Ln(Avg. Store Score)					
* College	0.0314*** (0.0092)	0.120*** (0.012)	0.0638** (0.022)	0.0909*** (0.011)	0.171*** (0.018)
* Non-College	0.00230 (0.0087)	0.0940*** (0.012)	0.0541* (0.021)	0.0533*** (0.011)	0.140*** (0.017)
Observations	1,580,395	707,942	255,346	707,100	254,750
R^2	0.305	0.343	0.446	0.455	0.659
Data Frequency	Month	Quarter	Year	Quarter	Year
Collapse Method		Last Month	Last Month	Average	Average

Notes: Conley Spatial HAC standard errors are in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Observations are at the household-data frequency level. All regressions include household and data frequency fixed effects. The sample only includes households that resided in the same census tract throughout their entire participation in the Homescan panel.

It is possible that the results in Table 6 are attenuated by noise in monthly measures of the nutritional quality of household purchases and store offerings. In Table 7, we replicate the analysis from column (3) of Table 6 aggregating the data to either a quarterly (columns (2) and (4)) or annual (columns (3) and (5)) frequency. For ease of comparison, the monthly results from column (3) of Table 6 and reproduced in the first column of Table 7. While the exact estimated elasticities are sensitive both to the frequency of the data and the method of aggregation, we observe households with and without a college education improving their nutritional consumption in response to improvements in the nutritional offerings of local stores when data is either quarterly or yearly. Consistent with Table 6, the response among college-educated households is greater than the response of less-educated households.

Despite being statistically different than zero, the improvements in nutritional consumption not adjust their nutritional consumption more in response to a diabetes diagnosis than individuals with lower levels of education.

documented in Table 7 are small when compared to the existing socioeconomic disparities in nutritional consumption. To demonstrate this, we consider how a household without a college-educated household head would respond to a change in their retail environment equivalent to moving from the average low-SES to the average high-SES neighborhood (as proxied by neighborhoods with an above versus below median share of college-educated residents). For a conservative estimate, we use the largest estimated responses (column (5) of Table 7): 0.004 and 0.14 for the elasticities of household nutrient scores with respect to store concentration and average store score, respectively. Moving from the average low-SES to the average high-SES neighborhood translates to increases of 0.72 and 0.032 in the log store concentration index and the log average store score, respectively. Combined with the estimated elasticities, these improvements in access imply that the nutrient score of a typical low-SES household currently residing in the average low-SES neighborhood would improve by 0.008 log units if they were to instead face the same store concentration and product availability as households living in the average high-SES neighborhood. Comparing this change to the socioeconomic disparity in log household scores (0.16), we see that only 4.6% of the gap in nutrient scores would be removed by closing the gap in access to healthy foods. That is, even if low-SES households faced the products and concentration of stores found in high-SES neighborhoods, over 95% of existing socioeconomic disparities in nutritional consumption would remain.

Improvements in access to healthy foods are more likely to occur in close proximity to sample households with growing tastes for these products. The changes in the purchases of households that we measure therefore reflect not only changes in access but also changes in tastes. This correlation between the time-variant component of demand and changes in access yield an upward-biased estimate of the effect of access-improving policies that are implemented independent of changes in local demand conditions. Therefore, while our estimates indicate that the nutritional quality of household purchases responds minimally to changes in retail environments, it is likely that the impact of policy-induced changes on nutritional consumption would be even smaller.

Since some policies subsidize store entry specifically, it is interesting to look at how households respond to changes in access driven by store entry alone. To do so, we use event study specifications to examine how household store and product choices respond to stores entering within various distances of their census tract centroid.³³ To take into account that household behavior may

³³We define a store as entering in a given month if (i) the store is first observed in the Scantrack data in that month, (ii) the store's parent company already appeared at least once in the Scantrack data prior to that month, and (iii) there was no store in the same sub-channel in the same census block in the TDLinx data in the previous month. We require the parent company to already be in the Scantrack data to avoid confusing sample growth with actual store entry. We require that no store in the same sub-channel be in the same census block in the TDLinx data in the month prior to

respond more when (i) a store enters in a neighborhood where there were previously fewer stores and when (ii) the entering store is relatively healthier than the surrounding stores, we interact event time indicators denoting store entry with the entering store's marginal impact on either the tract-level concentration index or the tract-level kernel density-weighted average of store-level nutrient scores. Though we see households shopping at entering stores, we do not see any statistically significant change in the healthfulness of household purchases in response to store entry. While households with a college education are more likely to shop in entering stores that have a larger marginal impact on the local availability of nutritious products, there is no evidence that even these households respond to store entry by improving their nutritional consumption.³⁴ We attribute this indiscernable response to lack of precision, and therefore focus on the results above that exploit the full variation in supply observed in our data. See Figures A.6 and A.7 for these results.

Robustness The results from a variety of robustness checks are summarized in Table 8. For each robustness check, we rerun the specification in column (5) of Table 7 and use the estimated coefficients to recalculate the implied change in household consumption that would be observed by moving a household without a college-educated head from the average low-SES to the average high-SES neighborhood. The first column of Table 8 reports the implied log change in the household score resulting from such a move, the second column reports the disparity between households with and without a college-educated head, and the third column reports the implied change as a percent of the existing disparity. The first row of Table 8 presents our base case, where we again see that our base specification is associated with a 4.6% reduction in the disparity between high-SES and low-SES households.

It is possible that a household without a college education originally residing in a neighborhood with below average access would respond more from moving to the average high-SES neighborhood than would a low-SES household originally living in the average low-SES neighborhood. To address the possibility of non-linear responses, we estimate the responsiveness of households in un-

avoid categorizing store re-branding due to merger and acquisition activity as store entry. We observe 2,106 entries between 2006 and 2011 in the Scantrack data that satisfy these requirements (66% of potential entries with location information).

³⁴In an attempt to increase the precision of our estimates and to verify the robustness of our results, we experimented with a variety of specifications. In particular, we used alternative data frequencies (monthly, quarterly, yearly) and alternative bandwidths (entry within 0.5km, 2km, 4km, 6km) to define events. We also experimented with only considering households that actually shop in the entering store, weighting observations by the share of expenditures the household spends in the entering store, only considering entries of stores with a nutrient score above the existing median, only considering the responses of households with below median access before the store entry, and considering the broader set of entries represented in the full TDLinx data. Regardless of the specification, the implications are the same: there is no measurable response of the nutritional quality of household purchases to a single store entry.

derserved neighborhoods to changes in their retail environments, where a neighborhood is defined as underserved if either its concentration index or nutrient score kernel density falls in the lowest quartile across all census tracts. In the first row of Panel A, we see that low-SES households in underserved neighborhoods respond nearly the same to improvements in their retail environments as comparable households living in the average low-SES neighborhood.

Table 8: Response of Household Purchases to Changes in Retail Access: Robustness

	Total Implied Log Change (t-stat) (1)	Log Disparity (College vs. Non-College) (2)	% Change in Gap (3) = -(1)/(2)
Base Case	0.008 (0.779)	0.16	-4.64%
A. Different Household Samples			
Underserved households	0.008 (0.749)	0.16	-4.84%
Alt. outliers (dropping >P99 & <P1)	0.001 (0.103)	0.16	-0.75%
Excl. WIC	0.007 (0.772)	0.16	-4.61%
Excl. WIC & below SNAP inc.	0.008 (0.785)	0.16	-4.82%
Excl. WIC & below SNAP/CNP inc.	0.007 (0.665)	0.16	-4.10%
Excl. WIC, SNAP/CNP, and with kids	0.007 (0.659)	0.16	-4.12%
B. Alternative Nutrition Measures			
Expenditure score	0.011 (2.928)	0.08	-13.90%
Exp. share on fruits & vegetables	0.002 (2.419)	0.01	-13.30%
Exp. share on soda	0.000 (-0.482)	-0.01	-4.77%
Total calories	0.011 (1.672)	-0.06	16.75%
C. Different Kernel Densities			
Driving time weight, 5 min bw	0.006 (1.120)	0.16	-3.88%
Driving time weight, 15 min bw	0.007 (0.389)	0.16	-4.08%
Distance weight, 3km bw	0.001 (0.135)	0.16	-0.35%
Distance weight, 5km bw	0.005 (0.557)	0.16	-2.86%
Transit time weight, 30 min bw	0.004 (3.026)	0.16	-2.24%

Notes: The above table reports estimated changes in the nutritional quality of household purchases by replicating the regression in column (5) of Table (7) with various household samples, measures of nutritional quality, and kernel density bandwidths. Column (1) reports the estimated log change in the nutritional quality of a household without a college-educated household head were they to move from the average low-SES to the average high-SES neighborhood. Column (2) reports the observed log difference in nutritional quality between college and non-college households. Column (3) reports the estimated log change as a percent of the observed log difference between college and non-college households. Our “base case” excludes outliers (more than twice the distance between the 50th and 90th percentiles from the median), uses a bandwidth of 10 minutes of driving time, and includes WIC households, SNAP/CNP households, and households with kids.

The remainder of Panel A demonstrates that our results are robust to the use of other alternative samples. In particular, we verify that our results are not sensitive to our definition of outliers by

replicating our main analysis on a sample that excludes household-months with nutrient scores that are above (below) the 99th (1st) percentile across all household-month observations rather than excluding household-months that are more than twice the distance between the 50th and 90th percentiles from the median. Next, to verify that our results are not driven by households that either report or are eligible for various types of food assistance that may alter shopping behaviors, we replicate our analysis on four additional sub-samples. These samples are designed to exclude reported WIC recipients, households that are eligible for either SNAP or free and reduced lunch, and households with kids who may receive some of their food at school.³⁵

To make our results more comparable to previous studies, we then replicate our analysis using alternative measures of nutritional quality in Panel B of Table 8. While the precise predicted impact of moving a household without a college-educated head from the average low-SES to the average high-SES neighborhood unsurprisingly depends on the measure of nutritional quality used, the overall story is the same: equating access alone will not be sufficient to erase socioeconomic disparities in nutritional consumption. While we observe the greatest response when nutrition is measured using our expenditure score defined in Appendix C or expenditure shares on fruits and vegetables, over 85% of the disparities in these measures across low-SES and high-SES households remains when low-SES households are subject to the access of the average high-SES neighborhood.³⁶

Interestingly, the final row of Panel B in Table 8 shows that non-college households actually increase their calorie intake when their food environment improves. This adjustment might reflect substitution from consumption away from home to consumption at home. Since the nutrient scores of households without a college-educated household head are typically over 40% higher for food at home versus food away from home in the FoodAPS data (0.854 vs. 0.600; Table A.12), this substitution could yield improvements in the healthfulness of the household's overall nutrient consumption that our in-home nutrient score does not capture. While the FoodAPS data is not rich enough to explore how households substitute between consumption at home and consumption away from home when local access changes, we can use the elasticity of calories with respect to access estimated in the Homescan data to get an idea of how this substitution could impact the interpretation of our results. The bias appears to be small. We estimate that the average low-SES household would increase their calories consumed at home by 1% if moved from the average

³⁵Refer to Appendix A for details on the income cut-offs used to identify eligible households.

³⁶One interpretation of these results is that it takes time for households to learn how to eat according to the nutritional recommendations of the DGA, so household nutrient scores may take more time to respond. The healthfulness of fruits and vegetables is perhaps more salient to consumers from the outset, so the response of household expenditures on these products is closer to the long-run response that we document in the cross-sectional results in Section 5.2.

low-SES to the average high-SES neighborhood. Since all households in the FoodAPS typically divide their calories 80-20 across food at home and food away from home, moving 1% of calories to in-home consumption would improve the nutrient scores of low-SES households by only 0.4% ($0.8 \cdot 0.854 + 0.2 \cdot 0.6 = 0.803$ versus $0.81 \cdot 0.854 + 0.19 \cdot 0.6 = 0.806$). We therefore do not believe that our results would change if we had richer information on food for consumption both at and away from home.

While our base specification uses driving times to weight stores in constructing our access kernel densities, it is possible that distance or transit times are more relevant for low-SES households. In Panel C of Table 8, we replicate our analysis using access kernel densities constructed using a bandwidth of either 3 or 5km of distance or 30 minutes of transit time in place of a driving time bandwidth of 10 minutes. Using these alternative bandwidths, we predict a reduction in nutritional disparities in consumption of less than 3% by granting households without a college education access to the store concentration and product offerings of the average high-SES neighborhood. Using indexes that are again based on driving times but with bandwidths of either 5 or 15 minutes, we still estimate an improvement in the nutritional consumption of a low-SES household that reduces the gap in nutritional consumption between low-SES and high-SES households by less than 5%.

One final concern with these results is that they are based changes in nutrient score kernel densities measured using the selected Scantrack sample of stores. Suppose that bodegas and other small stores that we do not observe in the Scantrack data stop offering healthy foods when a healthy grocery store enters. If this is the case, we will underestimate the response of households to changes in their retail environments because we will overestimate the extent of changes in access. If, however, small retailers increase the healthy products they offer in order to compete with new entrants, we will overestimate household responses. There is little that we can do to test the robustness of our results to this omission, but we expect the resulting bias on our counterfactual measure to be small. When we use the estimated elasticity of household purchases to environmental changes to measure how the nutritional consumption of the typical low-SES household would change from moving from a low-SES to a high-SES neighborhood, any bias in the estimated elasticity will be offset by a related bias in the estimated environmental change resulting from such a move. If, for example, bodegas in low-SES neighborhoods offer relatively more healthy foods before a large retailer enters, then this limitation of the Scantrack data leads us to (i) underestimate the household elasticity but (ii) overestimate differences in access across neighborhoods.

5.2 Looking Within Locations and Stores

Despite limited nutritional responses to improvements in access in the short run, it is possible that nutritional disparities would be reduced over time as low-SES households benefit from continued exposure to expanded retail access. In the analysis that follows, we measure the component of existing nutritional disparities that persists in the long run across households with equivalent access. Since households sort into neighborhoods and stores based on their demand for grocery products (or factors correlated with their demand), the difference between the disparity that exists in the full cross-section and the disparity that persists conditional on access provides an upper bound for the share of the raw disparity that can potentially be explained by access.

This analysis is presented in Table 9. Column (1) of Panel A reproduces the disparity in household nutrient scores first documented in Table 2. To estimate the component of this disparity that persists among households living in the same location, column (2) controls for residential location by including census tract fixed effects.³⁷ Comparing the estimated disparities in columns (1) and (2), we see that much of the nutritional disparity that exists across households with and without a college-educated household head persists when we control for residential location.

One concern with this within-location analysis is that households living in the same neighborhood may still have differential access. Even within a census tract, distance to retail outlets varies depending on the location of the household and factors such as car ownership or proximity to public transportation may yield differences in the ability of households to travel to stores. To control for these factors, we take a step further and examine how the nutritional quality of purchases varies across households shopping in the *exact* same store. To do this, we first calculate household-store-month nutrient scores that reflect the nutritional quality of the purchases that a given household makes in a specific store in a given month. In column (3) of Panel A of Table 9, we first reproduce column (1) using household-store-month nutrient scores in place of household-month scores. Consistent with our previous results, the healthfulness of household-store purchases is higher among households in which at least one household head has a college degree. To control for exact shopping location, in column (4) we include store fixed effects. As in the within-location analysis, much of the disparity that exists across households with and without a college degree in the full cross-section persists when we compare the purchases of households shopping in the exact same store.

While we do observe households with and without a college education living in the same census tracts and shopping in the same stores, the samples used to estimate the disparities in the full cross-

³⁷Census tracts are relatively small areas with approximately 4,000 residents.

Table 9: Nutritional Quality of Household Purchases: Controlling for Access

A. Full Sample	Geographic Controls		Store Controls	
	(1)	(2)	(3)	(4)
College-Educated	0.169*** (0.0046)	0.122*** (0.0047)	0.120*** (0.0043)	0.0987*** (0.0037)
Observations	2,553,494	2,553,080	5,820,238	5,818,525
R^2	0.012	0.171	0.007	0.122
Access FEs	None	Tract	None	Store
B. FE Identification Sample	Geographic Controls		Store Controls	
	(1)	(2)	(3)	(4)
College-Educated	0.139*** (0.0056)	0.122*** (0.0047)	0.120*** (0.0043)	0.0987*** (0.0037)
Observations	1,638,121	1,638,121	5,683,802	5,683,802
R^2	0.009	0.133	0.007	0.118
Access FEs	None	Tract	None	Store

Notes: Standard errors are in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. In columns (1) and (2) (columns (3) and (4)), observations are at the household-month (household-store-month) level. All variables are standardized by the variable's standard deviation, all regressions include year-month fixed effects, and standard errors are clustered by household. The top panel includes all observations; the bottom panel only includes observations that correspond to census tracts or store without variation in household SES. The top (bottom) panel includes expenditure (expenditure-share) weights.

section are not identical to the samples used to identify the disparities that persist within locations or within stores.³⁸ To make the within-location and within-store disparities directly comparable to the disparities that exist in the full cross-section, we replicate Panel A of Table 9 in Panel B of the same table dropping observations that correspond to census tracts or stores without variation in household SES. Comparing the disparities in the full cross-section in columns (1) and (3) to those in columns (2) and (4), we see that the disparities are only reduced by 13% and 18% when we look within residential locations and within stores, respectively. This indicates that differential access can explain at most 18% of the disparity in nutritional consumption across households with and without a college degree. As households shopping in the exact same store likely have preferences that are more correlated than households living in the same census tract, it is not surprising that we estimate a tighter bound when looking within locations (13%) than within stores (18%).

³⁸Over 50% of household-month observations are of households residing in tracts with both high-SES and low-SES sample residents; over 80% of household-store-month observations are for store-months with both high-SES and low-SES sample customers (Figure A.8).

We stress that even though socioeconomic disparities diminish when we control for residential or retail location, it is unlikely that resolving spatial disparities in access will reduce disparities across the entire US to the same extent. If households sort into retail environments on unobservables that are correlated with their tastes for healthy foods, then the socioeconomic disparities that we observe among households living in the same location or shopping in the same store will be smaller, on average, than the socioeconomic disparities that would persist across the full cross-section of households if access were equated. Therefore, while our estimates indicate that equating access across the entire US could not reduce existing socioeconomic disparities by more than one fifth, it is likely that the true impact would be even smaller.

Robustness In Table 10, we summarize the results of a range of robustness checks. As results are similar for our within-location and within-store analyses, we only report robustness results for our more conservative, within-store analysis here; refer to Table A.11 for robustness results for our within-location analysis. For each alternative specification, we replicate the analysis presented in Table 9. Recall that the difference between the raw disparity (column (2) of Table A.11) and the residualized disparity (column (3) of Table A.11) provides an upper bound for the share of the raw disparity that can potentially be explained by access. This upper bound and the relevant t-statistic are provided in columns (4) and (5) of Table A.11, respectively. The first and second rows of Table 10 present the results for our base case, where we saw that socioeconomic disparities in nutritional consumption are only reduced by 18% when we compare the purchases of households shopping in the same store.

There are a variety of reasons why disparities in nutritional consumption may persist among households with equivalent access. One potential reason is socioeconomic differences in budget constraints. Even though households shopping in the same store face the same product offerings and prices, low-SES households may purchase less healthful bundles because they can only afford a subset of the available products. In Panel A of Table 10, we explore the relative importance of these economic frictions by controlling for monthly food expenditures per 100 calories at the household level. We find that more educated households purchase healthier bundles than less educated households within the same store even conditional on food expenditures. This indicates that the strong relationship between education and the nutritional quality of household purchases that we find can neither be explained by differences in access nor differences in expenditures.³⁹

³⁹Of course, food expenditure is an endogenous outcome variable. In order to identify the role of budget constraints and product prices on within-store purchase disparities, one needs exogenous price variation. In ongoing work, we are exploring the relative importance of budget constraints, price sensitivities, and tastes in explaining the socioeconomic disparities in nutritional consumption that persist within a given store.

Table 10: Socioeconomic Disparities in Nutritional Consumption within Stores: Robustness

	College Coefficient			T-stat of Raw Diff (5)
	N (1)	No FEs (2)	Store FEs (3)	
Full Sample	5,820,240	0.120***	0.099***	-3.73
Fixed Effect Identification Sample	5,683,802	0.120***	0.099***	-3.69
A. Different Controls				
Non-socioeconomic demographics	5,683,802	0.102***	0.081***	-3.70
Ln(Exp. per 100 cal)	5,683,799	0.098***	0.086***	-2.18
Ln(Trips per month)	5,683,802	0.119***	0.099***	-3.58
B. Different Household Samples				
Alternative outliers (dropping >P99 & <P1)	5,902,613	0.136***	0.113***	-3.37
Excluding WIC	5,642,570	0.120***	0.099***	-3.72
Excluding WIC & below SNAP inc.	5,111,492	0.120***	0.099***	-3.53
Excluding WIC & below SNAP/CNP inc.	4,937,419	0.121***	0.100***	-3.45
Excluding WIC, SNAP/CNP, & with kids	4,021,536	0.088***	0.070***	-2.74
C. Alternative Nutrition Measures				
Ln(Absolute loss nutrient score)	5,683,802	0.101***	0.082***	-5.55
Ln(Expenditure score)	5,653,458	0.066***	0.062***	-0.90
Expenditure share on fruits & vegetables	5,654,161	0.011***	0.008***	-5.30
Expenditure share on soda	5,653,458	-0.010***	-0.008***	3.74

Notes: The above table presents average raw and residualized household-store-level scores across households with different socioeconomic profiles. Residualized scores are obtained by subtracting store fixed effects estimated in regressions of the log household scores against an indicator denoting whether the household has a college-educated head, year-month fixed effects, and store dummies. Our "base case" excludes outliers (more than twice the distance between the 50th and 90th percentiles from the median), uses a P50 cutoff, and includes WIC households, SNAP/CNP households, and households with kids.

We interpret our results under the assumption that households shopping in the same store have equivalent access. It is possible, however, that socioeconomic differences in mobility could threaten the credibility of this assumption. For example, differences in car ownership may allow households with higher levels of income to shop more frequently, in which case they may purchase more healthful bundles because they can purchase perishable goods such as fresh produce and dairy. To examine whether differences in shopping behaviors alter the interpretation of our results, we replicate our analysis controlling for the number of shopping trips a household makes in each month. The final row of Panel A shows that the association between household education and purchase quality is robust to controlling for differences in shopping frequency.

In Panel B, we verify that our results are robust to the use of alternative samples. As in our time-series analysis, we verify that our results are not sensitive to our definition of outliers and that our results are not driven by households that either report or are eligible for various types of food assistance that many alter shopping behaviors. Our results are very similar regardless of the household sample used.

Finally, Panel C confirms that our results are robust to the use of alternative measures of nutritional quality. Our results are qualitatively robust to measuring nutritional quality using an absolute loss nutrient score, our expenditure score as defined in Appendix C, or the expenditure share on fruits and vegetables or soda. Since the calories purchased by a household within a given store-month conflates the total monthly calories purchased by the household with how the household divides its calorie purchases across multiple stores, we do not replicate our within-store analysis on total calories.

6 Discussion and Conclusion

Despite an absence of evidence drawing a causal link between disparities in retail access and disparities in nutritional consumption, much of the discussion surrounding food deserts assumes that equalizing access will eliminate nutritional disparities across different socioeconomic groups. Such an assumption underlies policies that aim to improve the quality of food purchases by increasing the availability of healthful products in areas with unhealthful consumption. On the contrary, our analysis indicates that the large socioeconomic disparities in nutritional consumption that we document across households are not driven by the relatively limited differences in access to healthy foods that we observe across neighborhoods with different socioeconomic compositions. Using two different sources of identifying variation, we find that access-improving policies alone will eliminate less than one fifth of existing socioeconomic disparities in nutritional consumption.

The results of our time-series analysis are smaller than our cross-sectional results. We find this unsurprising for two reasons. First, our time-series analysis only captures short-run effects of improvements in access on household consumption, while our cross-sectional results may also capture long-run changes in consumption behaviors that result from living in an environment with greater access to healthful products. Our results therefore indicate that access-improving policies will have larger effects in the long run than in the short run, although even in the long run such policies will leave over 82% of current socioeconomic disparities in nutritional consumption unresolved. Second, we expect differences in demand within a household over time to be more limited than differences in demand across households living in the same location. Therefore, even if household consumption responds fully in the short run, the bound that we estimate using within-household variation should be tighter than the bound that we estimate using across-household variation.⁴⁰

Since differences in demand across socioeconomic groups yield empirically relevant disparities, policy makers cannot expect access-improving policies alone to eliminate disparities in nutritional consumption. We note that there a range of other reasons why such policies may still be desirable, as efforts to improve access to nutritious foods may also improve the broader economic and social health of a neighborhood. Furthermore, as access to healthy foods is clearly a necessary—albeit insufficient—condition for healthy consumption, supply-side policies may be more successful if implemented alongside policies targeting demand-side determinants of nutritional consumption.

If disparities in retail access do not generate the consumption disparities that we observe, then something else is to blame. There are a range of other explanations for disparities in purchases, including differences in tastes or social norms, price sensitivities, and budget constraints. In order to successfully improve the consumption of low-SES households, we must first understand which factors are most important in explaining why demand varies across socioeconomic groups with equal access. We aim to identify these factors in future work.

⁴⁰Our within-store results yield a larger bound (18%) than our within-location results (13%). We find this unsurprising for a similar reason: differences in demand across households shopping in the same store should be more limited than differences in demand across households living in the same census tract.

References

- Aizer, A., J. Currie, P. Simon, and P. Vivier**, “Do low levels of blood lead reduce children’s future test scores?” *NBER Working Paper 22558*, 2016.
- Alcott, H., R. Diamond, and J.P. Dube**, “The geography of poverty and nutrition: Food deserts and food choices in the United States,” *Conference draft*, 2015.
- Anderson, M.L. and D.A. Matsa**, “Are restaurants really supersizing America?,” *American Economic Journal: Applied Economics*, 2011, pp. 152–188.
- Aussenberg, R.A.**, “SNAP and related nutrition provisions of the 2014 Farm Bill (P.L. 113-79),” *Congressional Research Service Report*, 2014.
- Bertail, Patrice and France Caillavet**, “Fruit and vegetable consumption patterns: A segmentation approach,” *American Journal of Agricultural Economics*, 2008, 90 (3), 827–842.
- Bitler, M. and S.J. Haider**, “An economic view of food deserts in the United States,” *Journal of Policy Analysis and Management*, 2011, 30 (1), 153–176.
- Carlson, A., M. Lino, and T.V. Fungwe**, “The low-cost, moderate-cost, and liberal food plans,” *United States Department of Agriculture, Center for Nutrition Policy and Promotion*, 2007.
- CDC**, “State initiatives supporting healthier food retail: An overview of the national landscape,” 2011.
- Chen, S., R.J. Florax, S. Snyder, and C.C. Miller**, “Obesity and access to chain grocers,” *Economic Geography*, 2010, 86 (4), 431–452.
- Chetty, R. and Lawrence Hendren N. Katz**, “The effects of exposure to better neighborhoods on children: New evidence from the Moving to Opportunity experiment,” *American Economic Review*, 2016, 106 (4).
- , **N. Hilger, E. Saez, D. Schanzenbach, and D. Yagan**, “How does your kindergarten classroom affect your earnings? Evidence from Project STAR,” *Quarterly Journal of Economics*, 2011, 126 (4).
- Courtemanche, C. and A. Carden**, “Supersizing supercenters? The impact of Wal-Mart Supercenters on body mass index and obesity,” *Journal of Urban Economics*, 2011, 69 (2), 165–181.
- Currie, J. and R. Walker**, “Traffic congestion and infant health: Evidence from E-ZPass,” *American Economic Journal: Applied Economics*, 2011, 3 (1), 65–90.
- , **L. Davis, M. Greenstone, and R. Walker**, “Environmental health risks and housing values: Evidence from 1,600 toxic plant openings and closings,” *American Economic Review*, 2015, 105 (2), 678–709.
- , **S. DellaVigna, E. Moretti, and V. Pathania**, “The effect of fast food restaurants on obesity and weight gain,” *American Economic Journal: Economic Policy*, 2010, pp. 32–63.
- Curtis, C.**, “First Lady Michelle Obama on making a difference in cities with food deserts,” *The White House*, 2011.

- Cutler, D., E. Glaeser, and J. Shapiro**, “Why have Americans become more obese?,” *Journal of Economic Perspectives*, 2003, pp. 93–118.
- Cutler, D.M. and A. Lleras-Muney**, “Understanding differences in health behaviors by education,” *Journal of health economics*, 2010, 29 (1), 1–28.
- Dingel, Jonathan I.**, “The determinants of quality specialization,” *Mimeograph, Columbia University*, 2014.
- Dubois, P., R. Griffith, and A. Nevo**, “Do prices and attributes explain international differences in food purchases?,” *American Economic Review*, 2014, 104 (3), 832–67.
- Eid, J., H.G. Overman, D. Puga, and M.A. Turner**, “Fat city: Questioning the relationship between urban sprawl and obesity,” *Journal of Urban Economics*, 2008, 63 (2), 385–404.
- Einav, L., E.S. Liebtag, and A. Nevo**, “On the accuracy of Nielsen Homescan data,” *USDA Economics Research Report*, 2008, (69).
- Elbel, B., A. Moran, L.B. Dixon, K. Kiszko, J. Cantor, C. Abrams, and T. Mijanovich**, “Assessment of a government-subsidized supermarket in a high-need area on household food availability and children’s dietary intakes,” *Public Health Nutrition*, 2015, *FirstView*, 1–10.
- Ellickson, Paul B.**, “Quality competition in retailing: A structural analysis,” *International Journal of Industrial Organization*, 2006, 24 (3), 521–540.
- Freedman, M. and A. Kuhns**, “Supply-side subsidies to improve food access and dietary outcomes: Evidence from the New Markets Tax Credit,” *Working Paper*, 2016.
- Grossman, M.**, “The relationship between health and schooling: What’s new?,” *NBER Working Paper 21609*, 2015.
- Harding, M. and M. Lovenheim**, “The effect of prices on nutrition: Comparing the impact of product- and nutrient-specific taxes,” *NBER Working Paper 19781*, 2014.
- Helpman, E. and P.R. Krugman**, *Market structure and foreign trade: Increasing returns, imperfect competition, and the international economy*, MIT press, 1985.
- Hottman, Colin**, “Retail Markups, Misallocation, and Store Variety in the US,” *Mimeograph, Columbia University*, 2014.
- Jones, E.**, “An analysis of consumer food shopping behavior using supermarket scanner data: differences by income and location,” *American Journal of Agricultural Economics*, 1997, pp. 1437–1443.
- Larson, N.I., M.T. Story, and M.C. Nelson**, “Neighborhood environments: Disparities in access to healthy foods in the US,” *American Journal of Preventive Medicine*, 2009, 36 (1), 74–81.
- 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**, “Neighborhoods, obesity, and diabetes—A randomized social experiment,” *New England Journal of Medicine*, 2011, 365, 1509–1519.

- Ogden, C.L., M.M. Lamb, M.D. Carroll, and K.M. Flegal**, “Obesity and socioeconomic status in adults: United States, 2005–2008,” *NCHS Data Brief No. 50*, 2010.
- Oster, E.**, “Diabetes and diet: Purchasing behavior change in response to health,” *AEJ: Applied Economics (Forthcoming)*, 2017.
- Park, J.L., R.B. Holcomb, K.C. Raper, and O. Capps**, “A demand systems analysis of food commodities by US households segmented by income,” *American Journal of Agricultural Economics*, 1996, 78 (2), 290–300.
- Ploeg, M. Ver, L. Mancino, J.E. Todd, D.M. Clay, and B. Scharadin**, “Where do Americans usually shop for food and how do they travel to get there? Initial findings from the National Household Food Acquisition and Purchase Survey,” *USDA Economic Information Bulletin No. (EIB-138)*, 2015.
- Powell, L.M., S. Slater, D. Mirtcheva, Y. Bao, and F.J. Chaloupka**, “Food store availability and neighborhood characteristics in the United States,” *Preventive Medicine*, 2007, 44 (3), 189–195.
- Rahkovsky, I. and S. Snyder**, “Food choices and store proximity,” 2015, (ERR-195).
- Song, H.-J., J. Gittelsohn, M. Kim, S. Suratkar, S. Sharma, and J. Anliker**, “A corner store intervention in a low-income urban community is associated with increased availability and sales of some healthy foods,” *Public Health Nutrition*, 2009, 12 (11), 2060–2067.
- Todd, J.E., L. Mancino, E.S. Leibtag, and C. Tripodo**, “Methodology behind the quarterly food-at-home price database,” *United States Department of Agriculture, Economic Research Service*, 2010.
- Volpe, R., A. Okrent, and E. Leibtag**, “The effect of supercenter-format stores on the healthfulness of consumers’ grocery purchases,” *American Journal of Agricultural Economics*, 2013, 95 (3), 568–589.
- and — , *Assessing the healthfulness of consumers’ grocery purchases*, BiblioGov, 2013.
- Waldfogel, J.**, “Preference externalities: An empirical study of who benefits whom in differentiated-product markets,” *The RAND Journal of Economics*, 2003, 34 (3), 557–568.
- Weatherspoon, D., J. Oehmke, A. Dembélé, M. Coleman, T. Satimanon, and L. Weatherspoon**, “Price and expenditure elasticities for fresh fruits in an urban food desert,” *Urban Studies*, 2013, 50 (1), 88–106.
- Zenk, S.N., A.M. Odoms-Young, C. Dallas, E. Hardy, A. Watkins, J. Hoskins-Wroten, and L. Holland**, “‘You have to hunt for the fruits, the vegetables’: Environmental barriers and adaptive strategies to acquire food in a low-income African American neighborhood,” *Health Education & Behavior*, 2011, p. 1090198110372877.

For Online Publication

A Data Appendix

Household Consumption

We use the Nielsen Homescan data to examine the grocery purchases made by households. This dataset is collected by the National Consumer Panel (NCP), a joint venture between Nielsen and IRI, and provided by Nielsen through the USDA. As mentioned in Section 2, the Homescan data contains transaction-level purchase information for a representative panel of households across the entire US. See Harding and Lovenheim (2014) for a detailed description of how households are recruited and encouraged to report purchases on a weekly basis. While the number of participating households varies from year to year, we observe 114,286 unique households over our sample period (2006 through 2011).

The Homescan data includes information on demographics and residential location for each household in the panel. Households are asked to update their demographic information every year that they are in the sample, so the reported demographics should be relevant for the household's consumption decisions in that year. Households record each household head's education in one of six categories: grade school, some high school, high school graduate, some college, college graduate, or post-college graduate. To insure that household heads have completed their education, we only consider households in which at least one household head is older than 25. In most of our analyses, we consider an indicator denoting whether either household head has a college degree. When we instead work with education continuously, we assign each household head a number of years of education assuming that some high school corresponds to 10 years, some college corresponds to 14 years, and post college corresponds to 18 years. For households with two household heads, we use their average years of education.

Households record their income in one of 19 categories: under 5,000; 5,000-7,999; 8,000-9,999; 10,000-11,999; 12,000-14,999; 15,000-19,999; 20,000-24,999; 25,000-29,999; 30,000-34,999; 35,000-39,999; 40,000-44,999; 45,000-49,999; 50,000-59,999; 60,000-69,999; 70,000-99,999; 100,000-124,999; 125,000-149,999; 150,000-199,999; 200,000+. We assign households an income equal to the midpoint of their income category for each bounded category and an income of \$260,000 for the "\$200,000 and above" category. Where noted, we adjust the resulting household income for household size using the OECD equivalence scale. According to this scale, the first adult in the household receives a weight of 1, all other adults receive weights of 0.5, and each child

receives a weight of 0.3 (<http://www.oecd.org/eco/growth/OECD-Note-EquivalenceScales.pdf>).

To avoid adding noise by including inconsistent reporting, we only keep household-level observations where at least one purchase with non-zero calories is reported in each of three weeks in a given month (87% of households, 67% of household-month observations, 52% of household-store-month observations). Next, as nutrient scores that are more than twice the distance between the 50th and 90th percentiles from the median likely reflect measurement error, we further exclude these outliers from our preferred sample (less than 1% of remaining households, 7% of remaining household-month observations, 8% of remaining household-store-month observations). Finally, as households who qualify for either WIC, SNAP, or reduced school lunch pricing may have different consumption patterns as a result of institutional details, where noted we exclude households who qualify for various forms of food assistance (13% of remaining households, 48% of remaining household-month observations, 31% of remaining household-store-month observations). In particular, we exclude households who report WIC participation to Nielsen, households with monthly income below the threshold for SNAP based on household size (refer to <http://www.fns.usda.gov/snap/eligibility> for the cut-offs used), and households with annual income below the threshold for free and reduced school lunch based on household size (refer to <http://www.fns.usda.gov/school-meals/income-eligibility-guidelines> for the cut-offs used). As shown in Tables 8 and 10, our results are robust to an alternative definition of outliers and to the exclusion of households who qualify for food assistance. Summary statistics for the main household sample used in our analysis are provided in Table A.1.

One concern with using the Homescan data to examine socioeconomic disparities in consumption is that reporting diligence may vary systematically with household SES. Einav et al. (2008) study the credibility of the self-recorded data in the 2004 Homescan sample. They find that reporting errors in the Homescan data are on the same order of magnitude as those commonly found in earnings and employment-status data, although the reporting errors found in the Homescan sample are more pronounced for higher income and more educated households. Across all households, however, Einav et al. (2008) find that purchase locations and quantities are reported more accurately than prices.⁴¹ Our results rely primarily on purchase locations and quantities, although our results are qualitatively consistent when we replicate our analyses using measures based on prices (see Appendix C for results based on recommended expenditure shares).

⁴¹To reduce measurement error, Nielsen replaces many of the prices recorded by households with prices reported in store-level data.

Retail Environments

The Homescan data only provides a limited picture of the retail environments in which households are making their purchase decisions. There are two problems with using the Homescan data to characterize retail environments: First, if no household in the Homescan sample shops at a given store, then we do not observe from the data that this store exists. Second, even if we do observe households shopping in a given store, we only observe the products that they actually purchase, not the full variety of products offered. Because of these limitations, we use two additional datasets, both maintained by Nielsen, to obtain a more comprehensive picture of the retail environments that households face.

In order to observe the full set of stores available to households, we use a sample of the Nielsen TDLinX data provided to us through the USDA. The TDLinX data contains the names and geo-coded locations of all food stores in the US.⁴² Our sample contains information on all stores in the grocery, convenience, drug, wholesale club, and mass merchandiser categories.⁴³ There are 284,050 stores across these five categories (Table A.8). As dollar stores may attract customers from different socioeconomic profiles than other mass merchandisers, we separate dollar stores from mass merchandisers by making a new channel code for mass merchandisers with a sub-channel code description containing “Dollar Store.”

While the TDLinX data tells us about the number and types of stores that households have access to, it provides us with no direct information about product offerings within these stores. To see the full set of food products available at a subset of stores, we use the Nielsen Scantrack data. This data is provided by the Kilts-Nielsen Data Center at the University of Chicago Booth School of Business; refer to <http://research.ChicagoBooth.edu/nielsen> for information on availability and access. The Scantrack data contains weekly sales and quantities at the UPC level. This information is collected by point-of-sale systems that are located in over 30,000 retailers across the US. Stores are divided into five categories in the Scantrack data: grocery, convenience, drug, mass merchandiser, and liquor. Unlike the Homescan data, the Scantrack data does not track random weight purchases.

Despite this detailed information on prices and product offerings, the Scantrack data covers a more limited range of retail outlets than the TDLinX data and only provides us with the county, not the precise geo-coded location, of each store. Where possible, we obtain the geo-coded loca-

⁴²TDLinX materials state that the data provides “universal coverage and unique codes for every store in retail trade channels” (see http://www.nielsen.com/content/dam/nielsen/en_us/documents/pdf/Fact%20Sheets%20III/Nielsen%20TDLinX.pdf).

⁴³Our sample does not include information on stores in the Superette, gas station, liquor store, or cigarette outlet categories.

tions of Scantrack stores by using a concordance between the Homescan, Scantrack, and TDLinx data provided by the USDA. For every store in which a Homescan panelist is observed shopping, the concordance provides both the Scantrack and the TDLinx store bearing the reported name that is closest to the household's residence. To first verify the quality of the concordance, we merge in store-level information provided in both the Scantrack and the TDLinx data. We drop TDLinx-Scantrack matches for which the 3-digit zip codes or county codes do not match across the identified TDLinx and Scantrack stores (17.15% of store-year level observations) or for which the retailer or parent codes do not match the reported store name (40.07% of remaining store-year level observations). Finally, since the methodology used to create the concordance results in multiple TDLinx stores being matched to the same Scantrack store, we select the TDLinx store that is most frequently linked to a given Scantrack store over all years from the remaining TDLinx-Scantrack matches. This methodology allows us to extract the geo-coded locations of 62.2% of Scantrack stores.

One concern with the Scantrack data is that participation of retailers may systematically vary either across neighborhoods or across store types. Reassuringly, the average share of TDLinx stores appearing in our geo-coded Scantrack sample is not statistically different across tracts with different socioeconomic profiles. However, there are differences in the percent of TDLinx stores appearing in our geo-coded Scantrack sample across store types. This is in part due to better coverage for certain store types in the Scantrack data and in part due to our geo-coding procedure. While the Scantrack data contains 75% fewer grocery stores than the TDLinx data, we are able to extract the geo-coded location of approximately 90% of Scantrack grocery stores. The numbers are similar for drug stores: the Scantrack data contains 65% fewer drug stores than the TDLinx data, but we are able to extract the geo-coded locations of 70% of drug stores in the Scantrack sample. The coverage for mass merchandisers and convenience stores is weaker. For convenience stores this is primary due to the fact that the Scantrack sample only contains 2% as many convenience stores as the TDLinx data, whereas for mass merchandisers it is primarily because we are only able to extract the geo-coded locations of 25% of mass merchandisers in the Scantrack sample.

Nutritional Information

To obtain nutritional information for the products purchased by Homescan panelists and sold in Scantrack stores, we use IRI's nutritional database. The IRI Nutrition Database provides nutritional information for over 700,000 unique UPCs throughout the entire length of our sample. As described in Section 2, the database contains information on the quantity of macro-nutrients and

vitamins per serving, serving size in weight, and the number of servings per container at the UPC level. IRI collects this information directly from product labels. Since product characteristics can change without a change in the product's UPC, IRI revises its database when an updated version of an existing product is received and includes a time stamp of when the change was made. We use a version of the database that includes a snapshot of the market as of July 30th each year. We assume that these product characteristics are relevant for that calendar year.

We merge the IRI database with the Homescan and Scantrack data to uncover the full nutritional profiles of products we observe being purchased by households and sold in stores. These merges are not perfect: only 45% of the UPCs in the Homescan data and 57% of the UPCs in the Scantrack data are in the IRI nutrition database. We impute nutritional information for products not in the IRI data using the average nutritional information for UPCs in the same product module and product group with the same values for all other relevant characteristics, including brand, flavor, form, formula, style, and type. This same procedure is used to impute the nutritional information for random weight purchases in 2006 in the Homescan data. In addition, the nutritional profiles for 87% of UPCs are not available in every year from 2006 to 2011. For such UPCs, we impute the nutritional information in each missing year using the nutritional information for the same UPC in the prior year. If the nutrient profile of a UPC is missing in 2006, we impute the nutritional information using the information from the first year the UPC's nutritional information is available.

To assess the quality of our nutritional imputation, we compare the nutrient scores of bundles with no missing nutritional information to counterfactual nutrient scores in which some of the nutritional information is imputed. While we can also compare true nutritional information to counterfactual, imputed nutritional information on a nutrient by nutrient basis, comparing true and imputed nutrient scores is a more parsimonious way to evaluate the strength of the imputation. To do so, we set the nutritional information to missing for a random sample of 20% of the UPCs in the IRI database. We then use the nutritional information for the remaining 80% of UPCs to impute the nutritional information for the UPCs with "missing" nutritional information. Finally, we randomly draw with replacement 1,000 bundles of 400 products each and compute the nutrient score of the bundle using either the true or the partially-imputed nutritional information. The results of this exercise demonstrate that the imputation works quite well: the average absolute difference between the nutrient scores for a given bundle is only 6.8% of the average nutrient score across all bundles with a standard deviation of 8.7%.

We are not concerned that the imputed nutritional information biases our results for two reasons. First, the percent of purchased UPCs with imputed nutritional information does not vary systematically with household characteristics. Therefore, even if the imputed nutritional informa-

tion introduces noise, it does so equally across household socioeconomic profiles. Second, all of our results are robust to using measures of nutritional quality that do not rely on the precise nutritional information of each UPC. As shown in Section C, we obtain the same results if we use a measure of nutritional quality based on recommended expenditure shares for food categories as opposed to recommendations for particular nutrients.

Driving and Transit Times

To estimate how long it would take a household residing in a given census tract to access local stores, we scrape driving and transit times from Google Maps. In particular, we pull the driving and transit times between the centroid of each census tract and all stores in our TDLinx sample within 40km (these times are current as of April 2015). While coverage on Google Maps is good, it is not perfect: we are able to obtain driving (transit) times for 67% (61%) of store-tract pairs under 40km. Across census tracts, the average share of stores within 40km for which driving (transit) times are non-missing is 63% (37%). Reassuringly, there is no statistically significant difference in the share of non-missing driving and transit times across tracts with different socioeconomic profiles. We further note that our results are robust to the use of distance weights in place of driving or transit time weights when calculating access kernel densities, and our coverage of distance for store-tract pairs is 100%.

Neighborhood Demographics

While the Homescan data contains demographics for sample households, it only provides us with a limited picture of tract-level demographics. To measure the distribution of education and income in the neighborhoods in which Nielsen households reside and Nielsen stores are located, we use the five-year pooled (2007-2011) American Community Survey (ACS).

While the Nielsen datasets use 2000 census tract boundaries, the 2007-2011 ACS uses census tract boundaries as defined in 2010. We construct demographics from the ACS for the census tract boundaries used by Nielsen as follows. First, using the crosswalk provided by the Census at https://www.census.gov/geo/maps-data/data/tract_rel_download.html, we compute population shares that represent the share of the population from a given 2000 census tract residing in every overlapping 2010 census tract boundary in 2010. That is, letting $pop_{i(y)}^{10}$ denote the population in 2010 in census tract i whose boundary was defined in year y , we compute $w_{i(00),j(10)} = \frac{pop_{i(00)}^{10} \cap pop_{j(10)}^{10}}{pop_{i(00)}^{10}}$ for all tract boundaries such that $pop_{i(00)}^{10} \cap pop_{j(10)}^{10} \neq 0$. For a given 2000 census tract, we then compute the share of college-educated residents as reported in the 2007-

2011 ACS by taking a population-weighted average of college-educated shares across 2010 tract boundaries that overlap with the 2000 tract boundary in question. If, for example, a 2000 census tract was split across two census tracts in 2010, the share of residents with at least a college degree in the 2000 census tract would be given by $educ_{i(00)}^{ACS} = w_{i(00),j(10)} \cdot educ_{j(10)}^{ACS} + w_{i(00),j'(10)} \cdot educ_{j'(10)}^{ACS}$, where $educ_{i(y)}^{ACS}$ denotes the share of college-educated residents as measured in the 2007-2011 ACS in census tract i defined in year y .

We compute tract-level median income using a similar procedure. Using the share of residents in each of the 16 binned categories provided in the ACS for every 2010 census tract boundary, we take a population-weighted average to compute analogous shares for 2000 tract boundaries. Median income is then set to the midpoint of the income bin for which at least 50% of the population in the 2000 census tract boundary has income either in or below.

B Supplementary Tables and Figures

B.1 Supplementary Tables

Table A.1: Summary Statistics: Household Demographics

	All [1]	By Socioeconomic Status				
		College [2]		Non-College [3]		
<i>Household level data</i>						
Number of households	99,524					
Average months in sample (std dev)	25.7 (19.3)					
<i>Household-year level data</i>						
Number of households						
2006	36,591	17,863	18,728			
2007	61,357	30,677	30,680			
2008	59,568	30,182	29,386			
2009	57,704	29,623	28,081			
2010	57,827	30,311	27,516			
2011	45,558	23,975	21,583			
Mean (std dev) of:						
Income (1000s)						
2006	45.4 (29.5)	55.0 (33.0)	36.3 (22.0)			
2007	45.1 (28.1)	54.0 (31.1)	36.2 (21.4)			
2008	44.9 (27.8)	53.5 (30.5)	36.1 (21.5)			
2009	45.5 (29.2)	54.0 (32.0)	36.5 (22.7)			
2010	44.1 (23.5)	50.9 (23.9)	36.7 (20.6)			
2011	44.3 (23.9)	51.1 (24.4)	36.8 (20.8)			
Education (years)						
2006	14.2 (1.8)	15.7 (1.2)	12.7 (1.1)			
2007	14.2 (1.8)	15.6 (1.2)	12.7 (1.0)			
2008	14.2 (1.8)	15.6 (1.2)	12.7 (1.0)			
2009	14.2 (1.8)	15.6 (1.2)	12.7 (1.0)			
2010	14.2 (1.8)	15.5 (1.2)	12.7 (1.0)			
2011	14.2 (1.8)	15.5 (1.2)	12.7 (1.0)			

Notes: The above figure presents summary statistics of demographics for Homescan panelists included in our main estimation sample. Income is residualized from household size fixed effects. The base sample of households represented here excludes outliers with regards to nutritional consumption (more than twice the distance between the 50th and 90th percentiles from the median of household nutrient scores) and includes WIC households, SNAP/CNP households, and households with kids.

Table A.2: Summary Statistics: Nutritional Quality of Household Purchases

	All [1]	By Socioeconomic Status				
		College [2]		Non-College [3]		
<i>Household-month level data</i>						
Number of observations (millions)	2.55	1.28		1.27		
Mean (std dev) of:						
Nutrient score	1.16 (1.4)	1.29 (1.5)	1.04 (1.2)			
Expenditure score	6.90 (2.9)	7.22 (3.1)	6.58 (2.6)			
Calories (1000s)	99.63 (66.5)	97.10 (65.7)	102.17 (67.2)			
Total fat per 100 calories	4.00 (1.1)	3.94 (1.1)	4.07 (1.0)			
Expenditure share:						
Fruits and vegetables	8.53 (7.5)	9.18 (7.9)	7.88 (7.0)			
Soda	5.94 (7.8)	5.54 (7.5)	6.34 (8.2)			
<i>Household-month-store level data</i>						
Number of observations (millions)	5.52	2.88		2.64		
Mean (std dev) of:						
Nutrient score	0.68 (0.9)	0.72 (1.0)	0.63 (0.8)			
Expenditure score	4.76 (3.7)	4.95 (3.9)	4.55 (3.5)			
Calories (1000s)	27.08 (37.2)	26.50 (36.6)	27.72 (37.9)			
Total fat per 100 calories	3.70 (2.0)	3.65 (2.0)	3.76 (2.0)			
Expenditure share:						
Fruits and vegetables	7.66 (14.6)	8.06 (15.0)	7.23 (14.1)			
Soda	6.44 (16.5)	5.99 (15.8)	6.92 (17.1)			

Notes: The above figure presents summary statistics for the nutritional quality of purchases made by Homescan panelists in our main estimation sample. The base sample of households represented here excludes outliers with regards to nutritional consumption (more than twice the distance between the 50th and 90th percentiles from the median of household nutrient scores) and includes WIC households, SNAP/CNP households, and households with kids.

Table A.3: Correlations between Measures of Nutritional Quality

<i>Household-Month Purchases</i>	Nutrient score	Exp. score	Total calories	Fat per calorie	Share: soda	Share: fruit/veg.
Nutrient score	1	0.20	-0.04	-0.36	-0.06	0.17
Expenditure score	0.20	1	0.14	-0.03	-0.16	0.56
Total calories	-0.04	0.14	1	0.08	0.03	-0.09
Fat per calorie	-0.36	-0.03	0.08	1	-0.22	-0.09
Exp. share: soda	-0.06	-0.16	0.03	-0.22	1	-0.19
Exp. share: fruit/veg.	0.17	0.56	-0.09	-0.09	-0.19	1

<i>Household-Store-Month Purchases</i>	Nutrient score	Exp. score	Total calories	Fat per calorie	Share: soda	Share: fruit/veg.
Nutrient score	1	0.26	0.19	-0.22	-0.02	0.15
Expenditure score	0.26	1	0.29	-0.06	-0.14	0.52
Total calories	0.19	0.29	1	0.12	-0.02	0.00
Fat per calorie	-0.22	-0.06	0.12	1	-0.28	-0.15
Exp. share: soda	-0.02	-0.14	-0.02	-0.28	1	-0.10
Exp. share: fruit/veg.	0.15	0.52	0.00	-0.15	-0.10	1

Notes: The above figure presents correlations across measures of the nutritional quality of household purchases in our main estimation sample.

Table A.4: Three Sample Bundles

Sample Bundle:	Amount (oz)			Sample Bundle:	Amount (oz)		
	Healthy	Mixed	Unhealthy		Healthy	Mixed	Unhealthy
Cereal	12.25	6.125	0	Potato Chips	0	5.5	11
Russet Potatoes	160	80	0	Milk - 2% Fat	0	64	128
Broccoli Florets	12	6	0	American Cheese	0	6	12
Carrots	16	8	0	Bacon	0	8	16
Kidney Beans	30	15	0	Breakfast Scramble	0	12	24
Iceberg Lettuce	16	8	0	Butter Grade AA	0	4	8
Strawberries	16	8	0	Coca Cola	0	72	144
Orange Juice	64	32	0	Oreo Cookies	0	1.125	2.25
Low-fat Yogurt	36	18	0	Mayo	0	1.875	3.75
Chicken Breast	48	24	0	Frozen Pizza	0	56.60	113.20
Tuna - Chunk Light	20	10	0				
Peanut Butter	18	9	0				
Egg - Grade A Large	24	12	0				

Notes: The above table shows the composition of three sample bundles. To determine the food products included in each of these bundles, we select among the most widely purchased UPCs in each TFP food category.

Table A.5: TFP Healthful and Unhealthful Food Categories

Healthful	Unhealthful
Whole grain products	Non-whole grain breads, cereals, rice,
Potato products	pasta, pies, pastries, snacks, and flours
Dark green vegetables	Whole milk products
Orange vegetables	Cheese
Canned and dry beans, lentils, and peas	Beef, pork, veal, lamb, and game
Other vegetables	Bacon, sausage, and luncheon meats
Whole fruits	Fats and condiments
Fruit juices	Soft drinks, sodas, fruit drinks, and ades
Reduced fat, skim milk, and low-fat yogurt	Sugars, sweets, and candies
Chicken, turkey, and game birds	Soups
Eggs and egg mixtures	Frozen or refrigerated entrées
Fish and fish products	
Nuts, nut butters, and seeds	

Notes: We determine which TFP food categories are healthful and unhealthful using the recommendations from the Quarterly Food-at-Home Price Database (QFAHPD) indicators for which of 52 food groups are healthful and unhealthful. The QFAHPD categories were created by the USDA (see Todd et al. (2010) for more details).

Table A.6: Household Characteristics and Nutritional Quality of Purchases: Individual Nutrients

	Healthful Nutrients				
	Fiber	Iron	Calcium	Vitamin A	Vitamin C
College-Educated	0.096*** (0.002)	0.042*** (0.002)	0.063*** (0.002)	0.109*** (0.004)	0.298*** (0.008)
Observations	2,553,494	2,553,494	2,553,494	2,553,494	2,553,494
R^2	0.015	0.009	0.013	0.006	0.015

	Unhealthful Nutrients			
	Total Fat	Sat. Fat	Sodium	Cholesterol
College-Educated	-0.038*** (0.002)	-0.030*** (0.003)	-0.037*** (0.002)	-0.045*** (0.003)
Observations	2,553,494	2,553,494	2,553,494	2,553,494
R^2	0.011	0.007	0.011	0.002

Notes: Standard errors are in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The dependent variable in each regression is the normalized deviation of a household's per calorie consumption of a given nutrient in a given month from the recommended consumption. Standard errors are clustered by household. All variables are standardized by the variable's standard deviation.

Table A.7: Summary Stats: Tract Demographics, Store Density, Product Availability, and Pricing

	By Socioeconomic Status					
	All [1]	High Educ. [2]		Low Educ. [3]		
<i>Tract level data</i>						
Number of tracts with Homescan panelists	47,012	24,639		22,373		
Mean (std dev) of:						
Median income	55,286 (25,585)	68,411 (27,148)		40,833 (12,854)		
College-educated share	0.27 (0.17)	0.39 (0.15)		0.13 (0.05)		
Number of tracts with stores within 40km:						
All tracts	46,990	24,623		22,367		
Tracts with driving times	36,951	19,316		17,635		
Tracts with transit times	24,895	15,463		9,432		
Mean (std dev) of:						
Store concentration (by bandwidth):						
Driving time, 10 min	89.9 (140)	102.2 (148)	76.4 (128)			
Driving time, 5 min	18.7 (30)	19.7 (30)	17.5 (29)			
Distance, 3km	25.0 (59)	28.0 (64)	21.8 (54)			
Distance, 5km	57.5 (134)	66.5 (147)	47.7 (118)			
Transit time, 30 min	112 (235)	113 (248)	110 (212)			
Transit time, 45 min	258 (477)	266 (498)	245 (441)			
Distance (km) to Nearest:						
Store	1.46 (2.2)	1.19 (1.7)	1.75 (2.6)			
Grocery store	2.64 (3.7)	2.03 (2.8)	3.31 (4.5)			
Healthy grocery store	2.82 (4.0)	2.17 (2.9)	3.53 (4.8)			
Total number of stores within:						
0.5km	1.37 (4.3)	1.35 (4.5)	1.38 (4.1)			
0.5 to 1km	4.00 (10.0)	4.04 (10.6)	3.95 (9.3)			
1 to 2km	13.1 (32)	13.9 (33)	12.3 (30)			
2 to 4km	42.4 (104)	47.5 (111)	36.7 (95)			
4 to 8km	130 (317)	154 (354)	104 (268)			
8 to 16km	365 (817)	449 (924)	272 (669)			
16 to 32km	824 (1552)	1,043 (1719)	583 (1302)			
Average local store:						
Nutrient score (availability)	0.61 (0.1)	0.62 (0.1)	0.60 (0.1)			
Price index	1.06 (0.0)	1.07 (0.0)	1.05 (0.0)			
Healthy-to-unhealthy price ratio	1.00 (0.0)	1.00 (0.0)	1.00 (0.0)			

Notes: The above figure presents summary statistics for demographics, store density, product availability, and pricing for the tracts in which Nielsen Homescan panelists reside. Neighborhood demographics are taken from the 2007-2011 ACS; store-level information is taken from the Nielsen TDLinx and Scantrack data in 2011. See Footnote 18 for a description of how tracts are separated according to education levels. A grocery store is considered “healthy” if its store-level nutrient score is above the median across all stores.

Table A.8: Summary Statistics: Store-Level Product Availability and Pricing

	All [1]	By Socioeconomic Status				
		High Educ. [2]		Low Educ. [3]		
<i>Store level data</i>						
Number of stores						
TD Linx	284,050	125,144	158,906			
TD Linx-Kilts merged	21,744	13,327	8,383			
Mean (std dev) of:						
Nutrient score (availability)	0.61 (0.2)	0.62 (0.2)	0.59 (0.2)			
Nutrient score (sales)	0.79 (0.3)	0.84 (0.3)	0.72 (0.2)			
Expenditure score (availability)	6.77 (2.1)	6.90 (2.1)	6.57 (2.1)			
Expenditure score (sales)	6.87 (3.4)	7.13 (3.5)	6.46 (3.2)			
Price indexes:						
Aggregate	1.059 (0.07)	1.063 (0.07)	1.051 (0.06)			
Healthy-to-unhealthy (all)	1.001 (0.04)	1.003 (0.04)	0.999 (0.05)			
Healthy-to-unhealthy (storable)	1.002 (0.04)	1.002 (0.04)	1.001 (0.04)			

Notes: The above figure presents summary statistics for the number of stores in the Nielsen TDLinx data and the store count, product availability, and pricing for the stores in the merged TDLinx-RMS dataset in January 2011. See Footnote 18 for a description of how tracts are separated according to education levels.

Table A.9: Correlations between Measures of Access to Stores and Nutrition

	Concentration index			Distance to nearest			Avg. score
	Dist.	Tran.	Drive	Store	Groc.	HGS	
Concentration indexes							
Distance, 10km	1	0.95	0.95	-0.25	-0.26	-0.26	0.81
Transit time, 30 min	0.95	1	0.92	-0.24	-0.25	-0.25	0.79
Driving time, 10 min	0.95	0.92	1	-0.30	-0.31	-0.32	0.75
Distance to nearest							
Store	-0.25	-0.24	-0.30	1	0.73	0.70	-0.14
Grocery store	-0.26	-0.25	-0.31	0.73	1	0.95	-0.15
Healthy grocery store (HGS)	-0.26	-0.25	-0.32	0.70	0.95	1	-0.17
Average local nutrient score	0.87	0.81	0.79	-0.16	-0.16	-0.18	1

Notes: The above figure presents correlations across measures of access across tracts. A grocery store is considered "healthy" if its store-level nutrient score is above the median across all stores.

Table A.10: Neighborhood Characteristics and Store Concentration

<i>Panel A</i>	(1) All	(2) Grocery	(3) Conven.	(4) Drug	(5) Club	(6) Dollar	(7) Mass
Ln(Med. Income)	0.072*** (0.005)	0.093*** (0.005)	0.054*** (0.005)	0.091*** (0.005)	-0.06*** (0.005)	-0.14*** (0.007)	0.072*** (0.006)
Observations	36,951	36,470	36,489	36,398	36,640	19,860	31,175
R^2	0.005	0.009	0.003	0.008	0.004	0.020	0.005
<i>Panel B</i>	(1) All	(2) Grocery	(3) Conven.	(4) Drug	(5) Club	(6) Dollar	(7) Mass
Indicator for HIHE	0.28*** (0.01)	0.32*** (0.01)	0.25*** (0.01)	0.32*** (0.01)	0.022* (0.01)	-0.20*** (0.01)	0.28*** (0.01)
Observations	36,951	36,470	36,489	36,398	36,640	19,860	31,175
R^2	0.020	0.025	0.015	0.024	0.000	0.010	0.019

Notes: Standard errors are in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Observations are at the tract-year level. All variables are standardized by the variable's standard deviation. These results are for 2010; they are representative of other years in the TDLinx sample. "HIHE" denotes tracts with above median income and share of college-educated residents; 43% of tracts are HIHE. This table replicates Table 4 using median income or the interaction of indicators denoting tracts with above median income and above median college-educated shares in place of college-educated shares.

Table A.12: Nutritional Quality of Food at Home vs. Food Away from Home in FoodAPS

	Nutrient Scores	
	College	Non-College
<i>Food at home</i>		
Average	1.014	0.854
Standard error	(0.059)	(0.027)
Observations	629	2,928
<i>Food away from home</i>		
Average	0.664	0.600
Standard error	(0.026)	(0.009)
Observations	629	2,928

Notes: The above table presents raw averages of household nutrient scores in the FoodAPS data for households with and without a college-educated household head.

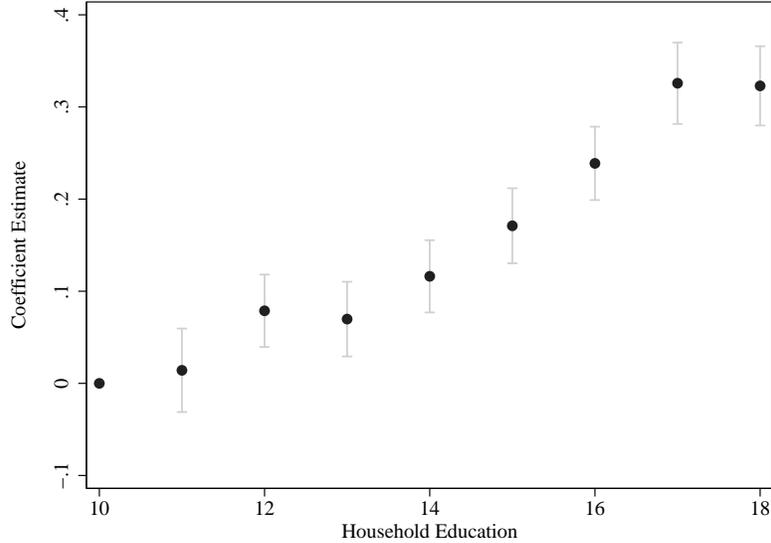
Table A.11: Socioeconomic Disparities in Nutritional Consumption within Tracts: Robustness

	N (1)	College Coefficient (SE)		Diff (%) (4)	T-stat of Raw Diff (5)
		No FEs (2)	Tract FEs (3)		
Full sample	2,553,494	0.169***	0.122***	-28%	-7.13
Fixed effect identification sample	1,638,121	0.139***	0.122***	-13%	-2.38
A. Different Controls					
Non-socioeconomic demographics	1,638,118	0.104***	0.093***	-10%	-1.44
Ln(Exp. per 100 cal)	1,638,118	0.113***	0.101***	-10%	-1.59
Ln(Trips per month)	1,638,121	0.139***	0.122***	-13%	-2.38
B. Different Household Samples					
Alternative outliers (dropping >P99 & <P1)	1,729,639	0.154***	0.137***	-11%	-1.91
Excluding WIC	1,621,537	0.139***	0.121***	-13%	-2.36
Excluding WIC & below SNAP inc.	1,409,195	0.135***	0.119***	-11%	-1.95
Excluding WIC & below SNAP/CNP inc.	1,336,742	0.134***	0.120***	-11%	-1.80
Excluding WIC, SNAP/CNP, & with kids	999,215	0.103***	0.096***	-8%	-0.83
C. Alternative Nutrition Measures					
Ln(Random weight nutrient score)	190,979	0.104***	0.113***	9%	0.51
Ln(Expenditure score)	1,638,000	0.068***	0.063***	-7%	-1.26
Expenditure share on fruits & vegetables	1,638,003	0.010***	0.008***	-19%	-2.91
Expenditure share on soda	1,638,000	-0.007***	-0.006***	-19%	1.89
Ln(Total calories (thousands))	1,638,121	-0.045***	-0.032***	-30%	1.85

Notes: The above table presents average raw and residualized household-store-level scores across households with different socioeconomic profiles. Residualized scores are obtained by subtracting census tract fixed effects estimated in regressions of the log household scores against an indicator denoting whether the household has a college-educated head, year-month fixed effects, and census tract dummies. Our “base case” excludes outliers (more than twice the distance between the 50th and 90th percentiles from the median), uses a P50 cutoff, and includes WIC households, SNAP/CNP households, and households with kids.

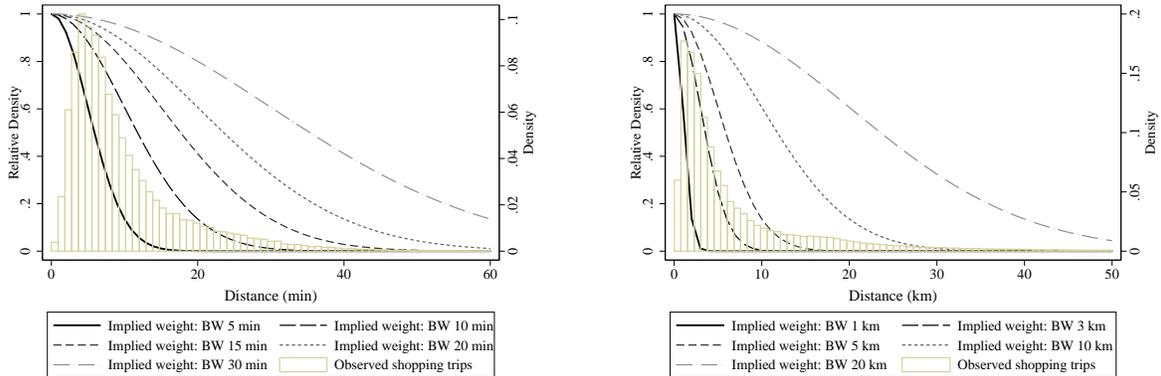
B.2 Supplementary Figures

Figure A.1: Household Nutrient Scores by Education



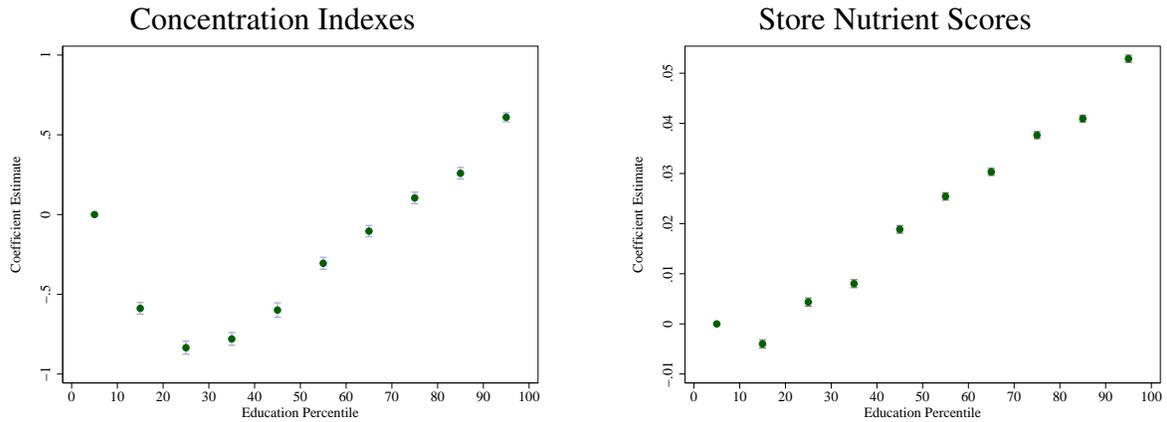
Notes: The above plot depicts the association between household education and the nutritional quality of household purchases. Observations are at the household-month level. The dots are the coefficient estimates on education dummies from an expenditure-weighted regression of log household nutrient scores on education dummies controlling for year-month fixed effects. The bars depict 95% confidence intervals.

Figure A.2: Implied Bandwidth Weights vs. Distribution of Household Shopping Trips



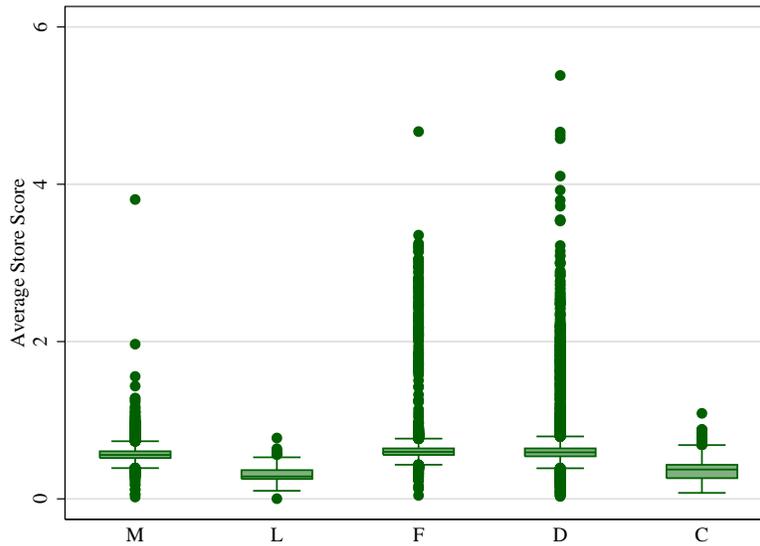
Notes: The above plots depict the relative distribution of observed shopping trips and the implied weights from various Gaussian kernel density bandwidths. The left (right) subplot uses driving times (distances). The bars depict the propensity of households to visit stores at various driving times/distances from their census tract centroid relative to stores at their census tract centroid (i.e. a driving time of 0 minutes and a distance of 0km). Similarly, the dashed lines represent the relative propensities implied by a Gaussian kernel with various bandwidths.

Figure A.3: Store Concentration Indexes by Tract Education



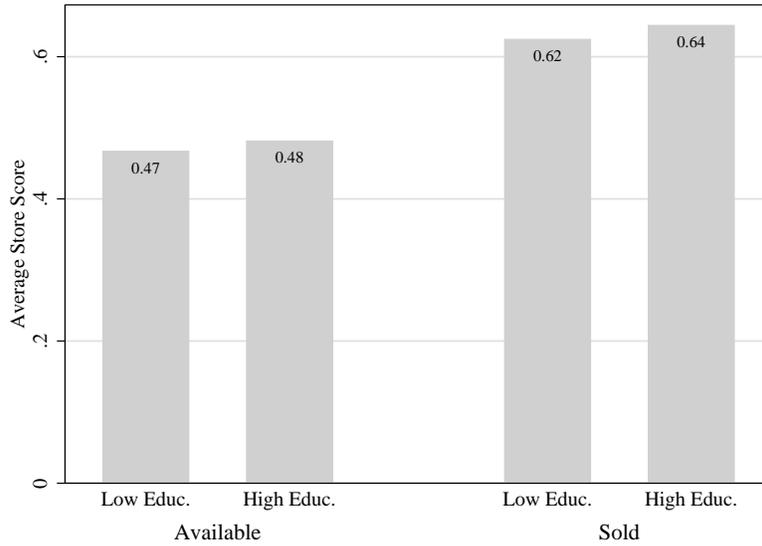
Notes: The above plots depict the association between neighborhood demographics and the concentration of stores (left subplot) and store nutrient scores (right subplot). Observations are at the tract-year level. The dots are the coefficient estimates on education dummies from a regression of log store concentration indexes or log kernel densities of the nutrient scores of stores surrounding a tract on education dummies controlling for year-month fixed effects. Education dummies are constructed using tract-level shares of college-educated residents from the ACS. The bars depict 95% confidence intervals.

Figure A.4: Store Nutrient Scores Across Channels



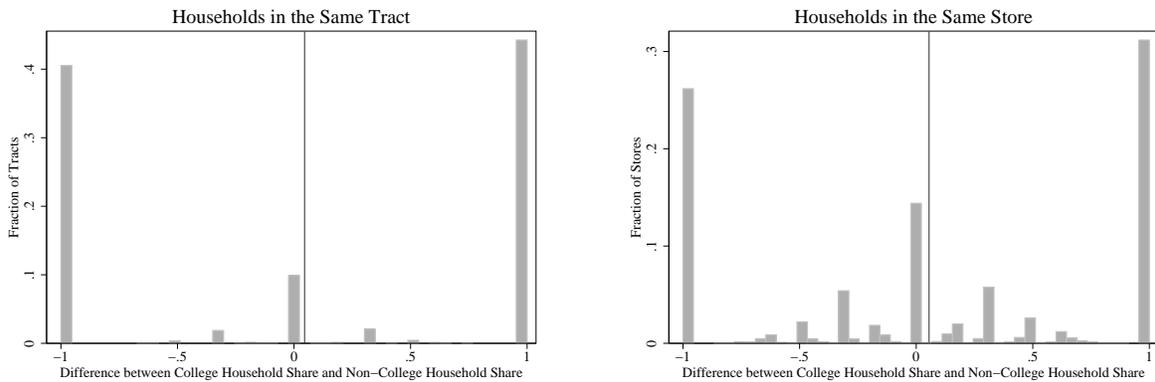
Notes: The above figure presents distributions of store-level nutrient scores by channel. Stores in the Scantrack data are divided into five channels: mass merchandiser (M), liquor (L), food (F), drug (D), and convenience (C). These results are for January 2010; they are representative of the other months in the Scantrack sample.

Figure A.5: Imputed Nutrient Scores Across Census Tracts: Available vs. Sold



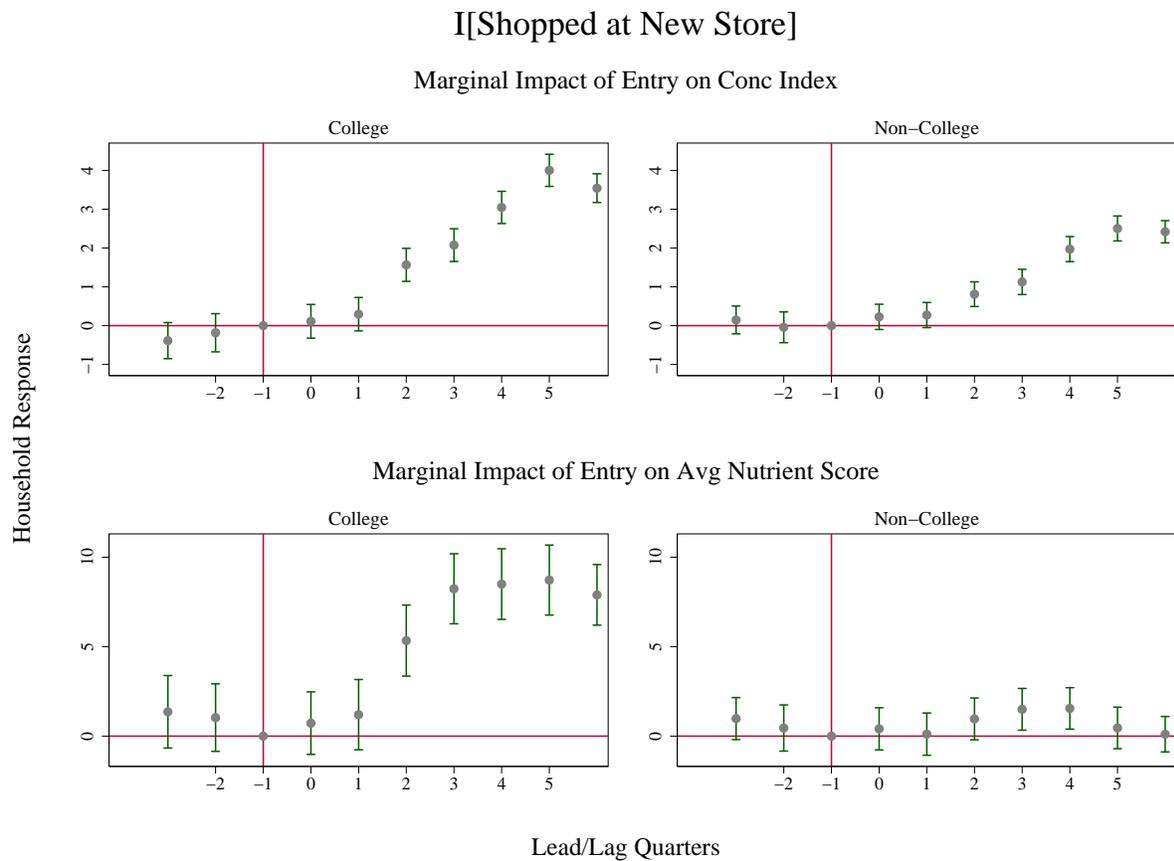
Notes: The above figure presents average of imputed store-level nutrient score across census tracts with different socioeconomic compositions. The “available” bars use the averages of store-level nutrient scores computed using national-sales weights for each store type in the Scantrack data to impute store-level nutrient scores (availability) for all stores in the TDLinx data. Similarly, the “sold” bars use the averages of store-level nutrient scores computed using store-sales weights for each store type in the Scantrack data to impute store-level nutrient scores (sold) for all stores in the TDLinx data. See Footnote 18 for a description of how tracts are separated according to education levels. These results are for January 2010; they are representative of other months in the Scantrack sample. A meticulous reader may wonder whether it is possible for the nutrient scores of a nationally representative consumer to be lower than the nutrient scores of bundles actually sold across all neighborhoods. This is not an error but rather an artifact of a skewed distribution of store-level nutrient purchases combined with an index that does not reward healthy deviations.

Figure A.8: Overlap in Household SES in the Same Retail Environment



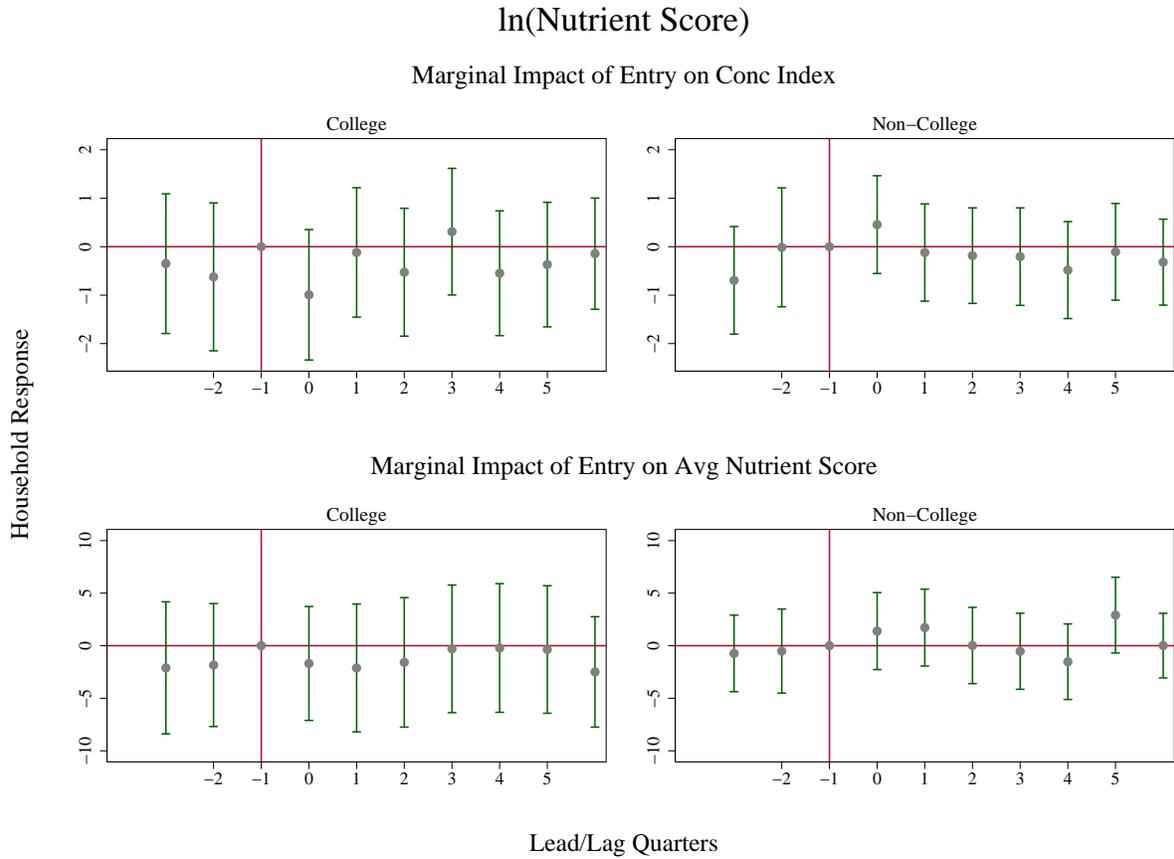
Notes: The plot on the left (right) is a household-weighted histogram of the difference in the shares of sample households with and without a college-educated household head across tracts (stores) in January 2011. Each bar therefore reflects the fraction of sample households represented by tracts (stores) whose college vs non-college divide is equal to the value on the x-axis.

Figure A.6: Event Study Analysis of Store Entry: Household Shopping Behavior



Notes: The above plots display the results from an event study analysis of store entry on household shopping behavior. Each plot depicts coefficient estimates on event time indicators from a regression of an indicator denoting whether a household shopped at the entering store in a given quarter on household fixed effects, quarter-year fixed effects, and a full set of event time dummies. The unmarked first and last points denote the pre-period and post-period coefficients, respectively. The first (second) panel depicts the coefficient estimates on event time indicators interacted with the entering store's marginal impact on the tract-level concentration index (the tract-level kernel density of store-level nutrient scores); the first (second) column is estimated on the sample of households with (without) a household head with a college education. An entry is considered an event for a household if a single store entered over the sample period within 4km of the household's census tract centroid and the household is in the sample for all lead and lag periods.

Figure A.7: Event Study Analysis of Store Entry: Nutritional Quality of Purchases



Notes: The above plots display the results from an event study analysis of store entry on the nutritional quality of household purchases. Each plot depicts coefficient estimates on event time indicators from a regression of log household-quarter nutrient scores on household fixed effects, quarter-year fixed effects, and a full set of event time dummies. The unmarked first and last points denote the pre-period and post-period coefficients, respectively. The first (second) panel depicts the coefficient estimates on event time indicators interacted with the entering store's marginal impact on the tract-level concentration index (the tract-level kernel density of store-level nutrient scores); the first (second) column is estimated on the sample of households with (without) a household head with a college education. An entry is considered an event for a household if a single store entered over the sample period within 4km of the household's census tract centroid and the household is in the sample for all lead and lag periods.

C Alternative Measure of Nutritional Quality: The Expenditure Score

In the following section, we reproduce our results using an alternative measure of nutritional quality. In particular, we introduce a second index, the “expenditure score,” that examines how products purchased by households and offered in stores deviate from recommendations for food group expenditure shares. After defining the expenditure score both for households and stores in Section C.1, we reproduce our main figures and tables using the expenditure score in place of the nutrient score in Section C.2. Note that results for our time-series and cross-section analyses using the expenditure score in place of the nutrient score are provided in robustness Tables 8 and 10, respectively.

C.1 Definition

Household Purchase Quality The USDA Center for Nutrition Policy and Promotion (CNPP) designs food plans for consumers based on recommendations from the Dietary Guidelines for Americans. Our second index, the “expenditure score,” examines how a household’s grocery purchases on each food group deviate from the CNPP’s recommended expenditure shares in the “thrifty food plan” (TFP). The expenditure index follows the measure used by Volpe et al. (2013) and Oster (2017).

The expenditure score for the grocery purchases recorded by household h in month t is defined as

$$Expenditure\ Score_{ht} = \left[\sum_{c \in C_{Healthful}} (sh_{cht} - sh_{ch}^{TFP})^2 \mid sh_{cht} < sh_{ch}^{TFP} + \sum_{c \in C_{Unhealthful}} (sh_{cht} - sh_{ch}^{TFP})^2 \mid sh_{cht} > sh_{ch}^{TFP} \right]^{-1}$$

where c indexes TFP food categories, sh_{cht} denotes the percent of household h ’s grocery expenditures in month t spent on products in category c , and sh_{ch}^{TFP} is the category c expenditure share, also in percent units, that the TFP recommends for a household with the same gender-age profile as household h .

We use the recommended individual expenditure shares from the TFP outlined in Carlson et al. (2007) to construct recommended household expenditure shares, sh_{ch}^{TFP} . The recommended category c expenditure share for each household member i , denoted by sh_{ci}^{TFP} , is determined by his/her

age and gender profile. We assign weights to each household member following the OECD equivalence scale and calculate the food expenditure weights as $w_{adult} = \frac{1+(n_{adult}-1)\times 0.5}{n_{adult}}$ and $w_{child} = \frac{0.3}{1+(n_{adult}-1)\times 0.5+n_{children}\times 0.3}$. The recommended category c expenditure share for household h is a weighted average of the recommended category c expenditure shares for each household member, i.e., $sh_{ch}^{TFP} = \sum_i w_i sh_{ci}^{TFP}$. Our results are robust to using the recommended individual expenditure shares from the low-cost, moderate-cost, or liberal food plans instead of those from the TFP.

The TFP provides recommendations for individual-level expenditure shares in 24 food categories. We matched the TFP food groups with Nielsen products using the Quarterly Food-at-Home Price Database (QFAHPD) developed by Todd et al. (2010). In particular, we aggregate the 52 QFAHPD food groups to the 24 TFP food categories using the correspondence created by Volpe and Okrent (2013). In doing so, we find that two TFP food categories, cheese and meat, contain both healthful and unhealthful food groups. Since the vast majority of cheese and meat purchases are of UPCs that fall into the unhealthful QFAHPD food groups, we assume that the aggregate TFP cheese and meat categories are unhealthful. All of our results are robust to assuming that these food categories are instead healthful. Refer to Table A.5 for the full list of healthful and unhealthful food categories that we use.

The expenditure score penalizes households for having a lower-than-recommended expenditure share for healthful food categories ($c \in C_{Healthful}$) and for having a higher-than-recommended expenditure share for unhealthful categories ($c \in C_{Unhealthful}$). As there are no clear guidelines for which food categories are most important for health, the index construction gives equal weight to all categories. For example, an underconsumption of whole fruits and an overconsumption of frozen or refrigerated entrees are treated the same. We follow Volpe et al. (2013) and take the inverse of the squared loss function so that higher scores are indicative of healthfulness.

The expenditure and nutrient scores consider the healthfulness of consumer purchases from two complementary perspectives, and each measure has its strengths and its weaknesses.⁴⁴ Since consumers choose foods rather than nutrients, the expenditure score is more closely related to consumer demand. Furthermore, expenditures on specific food groups, such as fruits and vegetables, are used by many other studies, and thus the expenditure score is more comparable to previous research.⁴⁵ Finally, the expenditure score takes into account other nutrients, such as zinc and potassium, which are not displayed on the nutritional facts panel and are therefore not included

⁴⁴The household-level expenditure and nutrient scores are positively correlated (correlation coefficient of 0.20; see Table A.2).

⁴⁵The correlation between our expenditure score and expenditure shares on fruits and vegetables is 0.56 (Table A.2).

in the nutrient score. The expenditure score, on the other hand, distinguishes between products in the same food category, e.g. frozen fish fillets versus fish sticks, that will be missed by the expenditure score. The nutrient score is also not sensitive to systematic variations in the price of foods purchased by different socioeconomic groups. If, for example, low-SES and high-SES consumers purchase identical quantities of cheese, but high-SES consumers purchase more expensive varieties, then for all else equal, expenditure scores will differ by household SES. The nutrient score, on the other hand, will reflect that both groups have similar diets.⁴⁶

Available Product Quality The expenditure score for store s in month t can be written as

$$Expenditure\ Score_{st} = \left[\sum_{c \in C_{Healthful}} (sh_{cst} - sh_{ch}^{TFP})^2 | sh_{cst} < sh_{ch}^{TFP} + \sum_{c \in C_{Unhealthful}} (sh_{cst} - sh_{ch}^{TFP})^2 | sh_{cst} > sh_{ch}^{TFP} \right]^{-1}$$

where c again indexes TFP food categories. sh_{cst} is the representative household's predicted category c expenditure share in store s in month t , calculated as

$$sh_{cst} = \sum_{u \in U_{cst}} \left(\frac{v_{ut}}{\sum_{u \in U_{st}} v_{ut}} \right)$$

Here, U_{cst} is the set of TFP-category c UPCs with positive sales in store s in month t , U_{st} is the set of all UPCs with positive sales in store s in month t , and v_{ut} is the total value of sales of UPC u across all stores in the national Scantrack sample in month t . We look at the distance of this representative household's category expenditure shares from the TFP's recommended category expenditure shares for a "typical" household, consisting of a male of age 19-50, a female of age 19-50, one child of age 6-8, and one child of age 9-11. We denote the recommended expenditure share in category c for this modal household by sh_{ch}^{TFP} .⁴⁷

⁴⁶To address the sensitivity of expenditure scores to prices, we recompute household food category expenditures using the average price per module instead of the actual price paid. Expenditure scores based on this alternative measure of expenditures are comparable to expenditure scores calculated using observed expenditures.

⁴⁷The store-level expenditure and nutrient scores are positively correlated (correlation coefficient of 0.49).

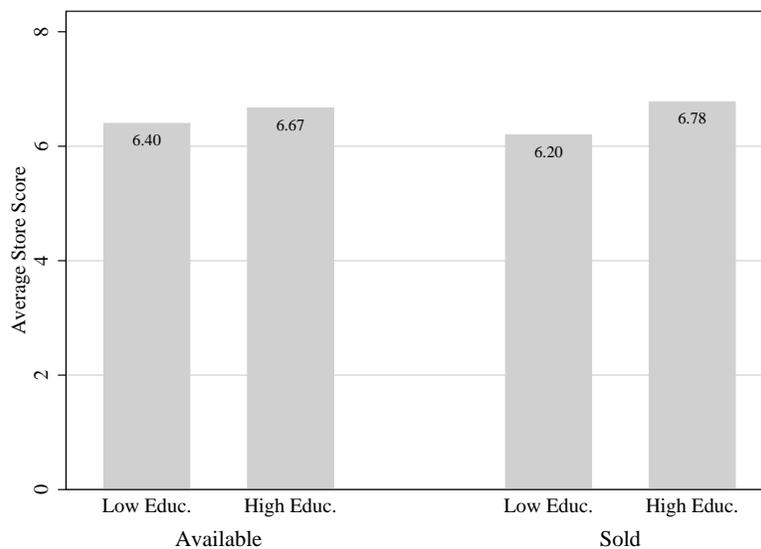
C.2 Results

Table A.13: Household Characteristics and Nutritional Quality of Purchases

	Ln(Expenditure Score)				
	(1)	(2)	(3)	(4)	(5)
College-Educated	0.196*** (0.0057)	0.161*** (0.0055)			0.110*** (0.0058)
Education			0.0833*** (0.0030)		
Ln(Income)				0.154*** (0.0030)	0.135*** (0.0031)
Observations	2,553,287	2,553,287	2,553,287	2,553,287	2,553,287
R^2	0.015	0.068	0.011	0.028	0.031
Demo. Controls	No	Yes	No	No	No

Notes: Standard errors are in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Observations are at the household-month level. All variables are standardized by the variable's standard deviation, standard errors are clustered by household, and year-month fixed effects are included. Column (2) includes controls for household demographics, including household size dummies, average head of household age, a dummy for marital status of household heads, dummies for households with either a female or male household head only, a dummy for the presence of children, and dummies for whether the household reports being white, black, Asian, or Hispanic. All specifications include expenditure weights. This table replicates Table 2 using the expenditure score in place of the nutrient score.

Figure A.9: Expenditure Scores Across Census Tracts: Available versus Sold



Notes: The above figure presents raw store-level expenditure score averages, computed using either national-sales weights (left) or store-sales weights (right), across census tracts with different socioeconomic compositions. See Footnote 18 for a description of how tracts are separated according to education levels. These results are for January 2010; they are representative of other months in the Scantrack sample. This figure replicates Figure 1 using the expenditure score in place of the nutrient score.

D Store Inventory and Pricing

D.1 Store-level Nutrient Scores

The nutrient score for store s in month t is given by

$$\begin{aligned} \text{Nutrient Score}_{st} = & \left[\sum_{j \in J_{\text{Healthful}}} \left(\frac{pc_{jst} - pc_j^{DGA}}{pc_j^{DGA}} \right)^2 \Big|_{pc_{jst} < pc_j^{DGA}} \right. \\ & \left. + \sum_{j \in J_{\text{Unhealthful}}} \left(\frac{pc_{jst} - pc_j^{DGA}}{pc_j^{DGA}} \right)^2 \Big|_{pc_{jst} > pc_j^{DGA}} \right]^{-1} \end{aligned}$$

where j again indexes nutrients, $J_{\text{Healthful}}$ and $J_{\text{Unhealthful}}$ are defined as in Section 3.1, and pc_j^{DGA} is the DGA's recommendation for the per calorie consumption of nutrient j . pc_{jst} is the per calorie amount of nutrient j that would be purchased by a representative household in store s in month t , calculated as

$$pc_{jst} = \sum_{u \in U_{st}} \left[\left(\frac{v_{ut}}{\sum_{u \in U_{st}} v_{ut}} \right) pc_{ju} \right]$$

where pc_{ju} is the per calorie amount of nutrient j in UPC u , U_{st} is the set of all UPCs with positive sales in store s in month t , and v_{ut} is the total value of sales of UPC u across all stores in the national Scantrack sample in month t .

D.2 Store-level Price Indexes

The aggregate price index for store s in month t is given by

$$P_{st} = \prod_{u \in U_{st}} \left(\frac{p_{ust}}{p_{ut}} \right)^{\frac{v_{ut}}{\sum_{u \in U_{st}} v_{ut}}}$$

where p_{ust} is the sales-weighted average price of UPC u in store s in month t , p_{ut} is the sales-weighted average price of UPC u across all stores in the Scantrack sample in month t , and U_{st} denotes the full set of UPCs sold in store s in month t . This price index summarizes how the average price of each UPC that the store offers compares to the national average price for the UPC.

To measure the spatial distribution of the cost of healthy and unhealthy eating, we further consider store-level price indexes for healthful and unhealthy products. For each store, the healthful (unhealthful) price index summarizes how the average price of each healthful (unhealthful) UPC that the store offers compares to the national average price for that UPC. The price index of health-

ful products offered in store s in month t is defined as

$$P_{st}^{Healthful} = \prod_{u \in U_{st}^{Healthful}} \left(\frac{p_{ust}}{p_{ut}} \right)^{\frac{v_{ut}}{\sum_{u \in U_{st}^{Healthful}} v_{ut}}}$$

where $U_{st}^{Healthful}$ is the set of all UPCs sold in store s in month t that are classified in a healthful TFP food category. Analogously, the price index of unhealthy products offered in store s in month t is given by

$$P_{st}^{Unhealthful} = \prod_{u \in U_{st}^{Unhealthful}} \left(\frac{p_{ust}}{p_{ut}} \right)^{\frac{v_{ut}}{\sum_{u \in U_{st}^{Unhealthful}} v_{ut}}}$$

where $U_{st}^{Unhealthful}$ is the set of all UPCs sold in store s in month t that are classified in an unhealthy TFP food category.

As our focus is on the accessibility of healthful versus unhealthy foods, we consider the ratio of a store's healthful-to-unhealthy price indexes, i.e. $\frac{P_{st}^{Healthful}}{P_{st}^{Unhealthful}}$. This ratio, which we refer to as the "relative price index" and denote by $P_{st}^{Relative}$, compares a store's average markup over national prices for the healthful products it offers to its average markup over national prices for the unhealthy products it offers. A store with a higher relative price index charges relatively more than average for its healthful versus its unhealthy products than a store with a lower relative price index.