Essay n°1 : Paper Summary of *Food Deserts and the Causes of Nutritional Inequality*, by Allcott et al. (2019)

Romain FERNEX

May 2026

1 Introduction

A prominent strand of the public health literature holds that low-income Americans eat poorly because they lack access to healthy food. The areas where they live, often labelled "food deserts", tend to have fewer supermarkets, more convenience stores, and a narrower selection of fresh produce than wealthier neighbourhoods. Accordingly, federal and local governments have invested heavily in supply-side remedies. For instace, since 2011, the Healthy Food Financing Initiative alone has channelled nearly $200 million into opening new grocery retailers in underserved communities, on the premise that improving local supply will improve diets.

Combining two nationally representative Nielsen datasets on household grocery purchases and store-level sales with detailed nutritional information from the USDA, Allcott et al. (2019) first confirm that the nutrition gap is real, as the top income quartile purchases groceries that score 0.56 standard deviations higher on a standardised Health Index than the bottom quartile. Then the central question they answer is whether this gap indeed stems primarily from supply related factors (such as the density of supermarket presence) or from demand related ones.

Using two reduced-form event studies and a structural model of grocery demand, the authors show that supply-side factors, including differences in local store availability, product assortment and prices, account for only about 10% of the observed nutritional gap. The remaining 90% is driven by differences in demand, meaning that high-income and low-income households have fundamentally different preferences over dietary characteristics such as added sugar, whole grains and fresh produce. These findings suggest that building more supermarkets in food deserts would

do little to reduce nutritional inequality and that demand-oriented interventions such as means-tested subsidies for healthy foods are likely to be more cost-effective.

The paper contributes directly to the food desert debate by providing the first large-scale causal test of whether local grocery supply matters for dietary quality. It also advances the structural estimation of food demand, by building on the framework of Dubois, Griffith, and Nevo (2014) while introducing a novel price instrument to correct for food price endogeneity. Lastly it connects most closely to the literature on spatial variation in fodd prices and variety, particularly the work of Handbury and Weinstein (2015) and Handbury (2021). In the same spirit as the variety-adjusted price indices estimated by Handbury and Weinstein, they estimate an implicit price for each product group ,which is defined as the nominal price per calorie net of the utility value of its nutritional characteristics, and find it also differs based on income group.

2 Why Supply-Side Explanations Fall Short

The paper's first empirical contribution is to show, through two complementary event studies, that the supply-side channel, while plausible in theory, has no economically meaningful effect on dietary quality in practice.

2.1 The supermarket entry test

The first event study exploits the opening of 6,721 new supermarkets across the United States between 2004 and 2016. If food deserts cause unhealthy eating, then opening a supermarket in an underserved area should improve dietary quality for nearby residents. The authors track how household grocery purchases change before and after a new supermarket opens within a 10-minute drive, controlling for household fixed effects and time-varying demographics. By looking within the same household over time, they isolate the effect of the change in local supply from persistent differences in individual characteristics.

Their event-study estimates, presented in Figure IV and Table II, show that the entrant expenditure share increases by about 2 percentage points one year after entry in the full sample and by roughly 3 percentage points for households living in food deserts. In other words, when a retail chain adds a new store within a 15-minute drive, it earns an additional $2 to $3 for every

$100 in grocery expenditures by nearby residents. However, the middle panels of Figure IV and Panel A of Table II show that the combined expenditure share on grocery stores, supercentres and club stores rises by only a fraction of a percentage point, with no statistically detectable change in food deserts. Thus, they conclude that most of the additional spending at the entrant is merely displaced from other supermarkets rather than from smaller outlets with poorer assortments.

The effect on the Health Index is correspondingly small. The bottom panels of Figure IV and Panel B of Table II report that the Health Index does not increase in any economically meaningful way after supermarket entry, with point estimates close to zero and confidence intervals that exclude large improvements. Using the upper bound of the 95% confidence interval for the Health Index effect in the bottom income quartile and the observed difference in supermarket counts between high- and low-income households, the authors calculate that differential access to supermarkets can explain at most about 1.5% of the overall 0.56 standard deviation Health Index gap. An important reason for this limited impact is that, as the authors document using the National Household Travel Survey, even households living in ZIP codes with no supermarket still purchase about 85% of their groceries from supermarkets because they are willing to travel several miles to reach them.

2.2 The movers test for broader place effects

Even if supermarket proximity does not matter much, other place-related factors could still drive nutritional inequality. Local peer effects, cultural norms around eating, or systematic differences in the types of restaurants and advertising that surround low-income households might all shape food choices in ways that a supermarket entry study would miss. To test this broader hypothesis, Allcott et al. turn to a migration-based design where they follow the roughly 20,000 Homescan households that moved across ZIP codes or counties during the sample period and examine whether their Health Index converges toward the eating patterns of their new location. Combes and Gobillon (2015), for instance, use a similar strategy with movers across French cities to separate genuine agglomeration effects on productivity from the sorting of more productive workers into larger cities. The logic is comparable in both cases, as tracking the same individual before and after a change in environment makes it possible to distinguish true place effects from

selection, although the outcome of interest here is dietary quality rather than wages.

Allcott et al. find that movers do exhibit a small convergence toward local eating patterns, as moving to a county where households eat more healthfully is associated with a modest improvement in the mover's own Health Index. The effect is statistically significant but economically small, explaining at most 2 to 3% of the nutrition gap. Moreover, the authors note that any endogeneity in the moving decision, such as unobserved lifestyle changes that simultaneously cause both the move and healthier eating, would bias these estimates upward, so even 2 to 3% is likely an upper bound on the true place effect over the time horizon they can observe. While Combes and Gobillon (2015) document that workers who move to denser cities experience significant and persistent wage gains, suggesting that agglomeration economies are a genuine place-based phenomenon rather than a pure sorting artefact, the convergence in the nutritional domain is far weaker and slower. Allcott et al. acknowledge that their panel only tracks households for a few years after a move, which limits their ability to estimate longer-run effects, but over the short-to-medium run, changing where someone lives does not appear to meaningfully change what they eat.

3 What Actually Drives Nutritional Inequality: The Structural Model

With the reduced-form evidence pointing away from supply-side explanations, the natural next question is whether the nutrition gap can instead be attributed to systematic differences in preferences across income groups. A structural model of grocery demand allows the authors to estimate how much households at each income level are willing to pay for different dietary characteristics and then to simulate their main counterfactual: what would low-income households eat if they had the same access as high-income households but kept their own preferences, and vice versa?

3.1 The model's intuition

Their model, adapted from Dubois, Griffith, and Nevo (2014), describes how a household allocates its grocery budget across product groups (milk, bread, candy, vegetables and so on) and how its demand for each group depends on the "implicit price" of that group. As mentioned in the introduction, the implicit price is not simply the monetary cost per calorie. Instead, it is the monetary price adjusted for the value the household places on the nutritional characteristics bundled into that product group, such as the amount of added sugar, sodium, whole grains or fresh fruit per calorie. A household that strongly values whole grains, for example, perceives a product group rich in whole grains as effectively cheaper and will then buy more of it.

The key parameters of interest are the household's willingness to pay for each of the 12 dietary components that make up the Health Index, which are estimated separately for each income quartile to capture how preferences for healthy eating vary systematically with income. The model assumes that within each product group, households substitute across individual products with a constant elasticity, while across product groups, expenditure shares follow a Cobb–Douglas form, meaning that a household allocates fixed proportions of its grocery budget to each category at given relative implicit prices. This structure is tractable and delivers clean identification, but it imposes homotheticity: the share of the grocery budget allocated to any product group does not change as household income rises, holding relative prices constant. In other words, the model attributes all income-related differences in dietary composition to differences in preference parameters across income groups rather than to non-linear income effects within a given preference structure. This assumption is worth scrutinising in light of Handbury (2021), who introduces a non-homothetic CES framework to study spatial price variation where the elasticity of the price index with respect to city income varies with household income. I shall, however, develop this point further in my second essay.

3.2 Identifying preferences: the price instrument

To estimate the preference parameters conditional on income, the model must confront a standard identification challenge, as households that like a particular product group may also seek out

lower prices for that group, or conversely invest in higher-quality items, which would make the observed price paid correlated with unobserved preferences and bias the estimates. Allcott et al. address this with a novel instrument that exploits the interaction of two sources of variation: the fact that different grocery chains have different comparative advantages in supplying different product groups and the fact that different chains are present in different geographic markets. To take their own illustrative example, suppose Safeway is especially efficient at sourcing pizza and Shaw's is especially efficient at sourcing apples. Then markets dominated by Safeway will have relatively low pizza prices and markets dominated by Shaw's will have relatively low apple prices, for reasons unrelated to local consumers' preferences for pizza versus apples. Concretely, the instrument is constructed by combining each chain's national price advantage in a given product group (measured using prices from all markets except the local one) with that chain's local market presence, and the identifying assumption is that geographic variation in prices driven by chains' cost advantages is uncorrelated with geographic variation in unobserved preferences, conditional on market and product group fixed effects.

This strategy is closely related, in its underlying logic, to the shift–share instrument proposed by Bartik (1991) and discussed by Combes and Gobillon (2015) in the context of urban productivity. In the Bartik framework, the instrument for local employment growth combines the local industry composition (shares) with the national growth rate of each industry (shifts), under the assumption that national industry trends are exogenous to local demand shocks. Allcott et al.'s instrument follows a parallel structure where the shares are the local presence of each retail chain and the shifts are each chain's national pricing advantage across product groups. Just as the Bartik instrument isolates demand-driven variation in local labour markets from supply-side structural differences across cities, Allcott et al.'s instrument aims to isolate cost-driven variation in local food prices from demand-driven variation in local food preferences.

The analogy also highlights a potential concern. In the Bartik literature, the validity of the instrument depends critically on whether the shares are exogenous. If certain chains choose to locate in markets where consumers happen to have strong unobserved preferences for the product groups in which those chains specialise, then the exclusion restriction would be violated.

Allcott et al. provide evidence against this concern, showing that the instrument is uncorrelated with predicted preferences based on observable demographics and that the results are robust to adding region-by-product-group fixed effects that absorb broad geographic patterns in tastes. These checks support the validity of the instrument, though , as acknowledged by the authors, they cannot rule out more granular correlations between chain location choices and local taste heterogeneity.

3.3 What the estimates reveal

The structural estimates reveal that higher-income households place significantly greater value on six out of eight "healthy" dietary components, including whole fruit, other fruit, whole grains, vegetables and dairy, and significantly greater distaste for two of the four "unhealthy" ones, sodium and added sugar. The case of added sugar is notable, as low-income households are estimated to have a weakly positive willingness to pay for added sugar, meaning they view it as a desirable product feature, while high-income households are willing to pay to avoid it. Summarising these preferences into a single willingness-to-pay measure for overall dietary quality, the top income quartile values a maximally healthy diet at \$1.14 per 1,000 calories, nearly three times the \$0.43 that the bottom quartile is willing to pay.

To decompose the nutrition gap based on these estimates, the authors run a series of counterfactuals. First, they give all income groups the prices observed for the top quartile, which barely changes the gap. Next, they also equalise the product characteristics available and the gap shrinks by only about 9%. Finally, they equalise the preference parameters themselves and the gap closes almost entirely, leading to the conclusion that roughly 90% of nutritional inequality is demand-driven and only about 10% is supply-driven. Robustness checks using alternative specifications and additional geographic fixed effects yield consistent results, with supply accounting for 7 to 12% across specifications.

The authors also investigate which observable characteristics mediate this income–preference relationship using the omitted variable decomposition of Gelbach (2016). They find that education explains about 20% of the correlation between income and healthy grocery demand, and that nutrition knowledge, measured through a detailed quiz administered to Homescan panelists,

explains another 14%. These are correlations rather than causal estimates, but they suggest that improving nutritional literacy and educational attainment might shift demand toward healthier foods more effectively than expanding supply.

4 Policy Implications and Limitations

To compare the cost-effectiveness of supply-side policies with demand-side alternatives, the authors simulate a hypothetical means-tested subsidy for healthy foods. Using the structural model to simulate how bottom-quartile households would respond to a proportional subsidy on above-median Health Index products, the authors find that an annual expenditure of $84 million could match the dietary improvement that would result from one additional supermarket entry within a 10-minute drive for all bottom-quartile households. A more ambitious subsidy of about $830 million per year, roughly 1% of the existing SNAP budget, could match the improvement from equalising supply conditions across income groups. Finally, a subsidy of about $11 billion per year, roughly 15% of the SNAP budget, could close the nutrition gap entirely by raising the bottom quartile's Health Index to that of the top quartile.

These estimates rest on a partial equilibrium model that holds store locations, product assortments and prices fixed, ignoring how retailers might adapt to changes in consumer demand. If a subsidy shifted demand toward healthier products, retailers could respond by expanding their healthy offerings in low-income areas, entering new markets or adjusting prices. The current simulations therefore provide a useful benchmark but may understate the longer-run contribution of supply responses to the effectiveness of demand-side policy. The reliance on a homothetic demand structure also pushes the analysis towards finding a dominant role for demand. As discussed above, if part of the income-related variation in dietary composition reflects budget constraints interacting with non-homothetic preferences rather than fixed taste differences, the 90% demand figure partly reflects this modelling choice.

Taken together, the reduced-form and structural results suggest that a well-designed means-tested subsidy for healthy foods could be more cost-effective than further supply-side subsidies to store entry, at least within the partial-equilibrium setting that the authors model. At the same

time, the assumptions of a homothetic demand structure and of no retailer response to policy changes both tend to favour demand-side explanations over supply-side ones. In my second essay, I therefore revisit these assumptions and explore how allowing for non-homothetic preferences and richer supply responses might change the relative importance of demand and supply in explaining nutritional inequality, and what this implies for the effectiveness and cost of policies such as SNAP-based healthy-food subsidies.

References

Bartik, Timothy J. (1991). *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute for Employment Research.

Combes, Pierre-Philippe and Laurent Gobillon (2015). "The Empirics of Agglomeration Economies". In: *Handbook of Urban and Regional Economics*. Ed. by Gilles Duranton, Vernon Henderson, and Will Strange. Vol. 5A. Amsterdam: North-Holland.

Dubois, Pierre, Rachel Griffith, and Aviv Nevo (2014). "Do Prices and Attributes Explain International Differences in Food Purchases?" In: *American Economic Review* 104.3, pp. 832–867.

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

Handbury, Jessie (2021). "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across US Cities". In: *Econometrica* 89.6, pp. 1679–2715.

Handbury, Jessie and David Weinstein (2015). "Goods Prices and Availability in Cities". In: *Review of Economic Studies* 82.1, pp. 258–296.