

Optimal interventions for increasing healthy food consumption among low-income populations

Retsef Levi

MIT Sloan School of Management, retsef@mit.edu

Elisabeth Paulson

MIT Operations Research Center, epaulson@mit.edu

Georgia Perakis

MIT Sloan School of Management, georgiap@mit.edu

The federal government currently spends over \$100 billion per year on policies aimed to increase fruit and vegetable (FV) consumption among low income households. These include price-, nutrition education-, and access-related interventions. Currently, the government allocates funds to each type of intervention in an ad-hoc fashion, in some cases resulting in surprisingly disappointing outcomes. For example, access-related interventions have seen mixed results in many case studies, resulting in debate about the importance of supply-side interventions. This paper introduces a novel consumer behavioral model for grocery shopping dynamics, which is nested into a bi-level model for optimizing the government's investments. The government's goal is to increase fruit and vegetable (FV) consumption among low income households by utilizing strategic portfolios of interventions. Based on primitives estimated from data, the model agrees with known empirical evidence and suggests several new policy insights, for example that access-related interventions are only beneficial under specific conditions which depend heavily on how much the consumer values eating nutritiously. In a setting where the government cannot provide completely personalized interventions, it is found that high levels of FV consumption by can be achieved by subsetting consumers based on key characteristics and deploying smart group-level strategies.

Key words: Optimal subsidies, bi-level optimization, public policy, food policy, central planner

1. Introduction

Americans consume far less fruits and vegetables (FVs) than the recommended daily amount. This gap is particularly large for underserved populations (Dong and Lin 2009). Furthermore, the

link between lack of sufficient FV consumption and life-threatening chronic conditions, such as cardiovascular diseases and obesity, is well established. It is estimated that over 2 million deaths per year can be attributed to inadequate consumption of FVs (Miller et al. 2019). Thus, establishing effective policy strategies for increasing the consumption of FVs, especially among low income households, is a national priority with a substantial federal budget investment of over \$100 billion each year.

There are three categories of policy levers that the government currently utilizes in order to increase food security and improve diets: 1) Financial/monetary interventions that increase affordability; 2) Access-based interventions that increase physical access to healthy food; and 3) Nutrition education interventions. However, not only is the effectiveness of individual policy interventions debated, but how they can be strategically utilized together or be targeted to specific households or neighborhoods is rarely considered. This highlights the importance of developing a consumer behavioral model that captures the respective effects of all three types of interventions, as well as a corresponding optimization framework.

In addition to the food policy sector, there are many public policy initiatives where multiple levers, affecting individuals through different mechanisms, are employed to achieve the same goal. For example, consider the number of different interventions aimed at reducing recidivism, increasing voter turnout, or decreasing homelessness. Many interventions related to these initiatives can be categorized as either a financial-, education/awareness-, or access-related levers. An individual-level model, capable of capturing the impact of each class of intervention on the behavior of interest, would allow for joint optimization across the various levers and a more targeted policy approach.

This paper introduces a new consumer behavioral model and bi-level optimization framework to jointly optimize the government's investment into price-, education- and access-related interventions, with the goal of increasing a consumer's FV consumption. This modeling framework is not only sufficiently flexible to capture existing interventions, but is in fact validated by existing empirical evidence in the food policy literature. Moreover, it provides important insights that can

directly inform policy design at the household-level. The modeling framework can also readily be tailored to a number of different policy initiatives with similar classes of interventions.

Price-related interventions are perhaps the most commonly used and highly studied type of demand-side lever in food policy. The Supplemental Nutrition Assistance Program (SNAP)—formerly the food stamp program—is the largest example of a price-related policy lever, with an annual budget of around \$60 billion. Many other types of price-related policy levers exist, including those that specifically subsidize or incentivize healthy food purchasing (e.g., The Special Supplemental Nutrition Program for Women, Infants, and Children, and The Healthy Incentives Program in Massachusetts). Price-related interventions are largely found to be successful at increasing FV consumption, but only to a certain extent (Mabli et al. 2013, Olsho et al. 2016, Dong and Lin 2009). There is limited literature on the degree to which complementary interventions, such as nutrition education, could amplify the effects of price interventions (Waterlander et al. 2013).

Access-based interventions have gained popularity in recent years. The concept of *food deserts*—areas where the majority of the population does not have access to grocery stores (USDA ERS 2017)—has recently brought to light many access-related issues that disproportionately affect low-income households. Since 2014, the Healthy Food Financing Initiative has distributed over \$220 million dollars to community-based projects supporting increased access to healthy food in low-income food deserts (The Reinvestment Fund 2019). Strategies include building new grocery stores, increasing affordable transportation to grocery stores, and increasing the stock of healthy food at convenience or corner stores. Despite their rise in popularity, the effect of access-related interventions on diet is not well established (Allcott et al. 2019, Ver Ploeg and Rahkovsky 2016, Cummins et al. 2014, 2005). A key issue is a lack of understanding of the mechanisms linking access to food purchasing behavior and how access is affected by other household covariates and interventions.

The problem of increasing consumption or adoption of socially beneficial goods has been studied extensively in the operations management and operations research literature, particularly in the domains of green technology and vaccines (Taylor and Xiao 2019, Chemama et al. 2018, Levi

et al. 2016, Alizamir et al. 2016, Cohen et al. 2015, Raz and Ovchinnikov 2015, Avci et al. 2014, Sierzchula et al. 2014, Perakis and Lobel 2011). While the problem of increasing consumption of healthy food has a similar sentiment, it presents its own unique challenges and opportunities in terms of both consumer choice modeling and the distinct policy levers available.

1.1. Contributions

Novel and tractable bi-level optimization problem. This paper presents a novel optimization framework to describe consumer grocery shopping decision-making. The consumer's decision regarding their food purchases and shopping cycle is modeled as the solution to a personalized optimization problem, with input primitives that can be practically estimated from data. The impact of three different classes of interventions—nutrition education, price incentives, and access—are modeled as changes to different aspects of the consumer's optimization problem, for example, the price lever impacts the consumer's budget constraint whereas the education lever impacts their *value of nutrition*—i.e., how much they care about eating healthy, a component of their personal objective function.

The consumer-level model is nested into a bi-level problem in order to optimally choose the government's investment across these three policy levers to maximize consumers' FV consumption. The upper-level problem can be written as a mathematical program with equilibrium constraints, which is a problem that is typically difficult to solve (Luo et al. 1996). However, under mild assumptions, it can be shown that this problem is unimodal and can thus be solved efficiently in order to determine the government's optimal strategy.

Policy insights. The analysis suggests findings that both broadly agree with the documented empirical evidence, providing validation for the proposed model, as well as provide new detailed insights. In particular, the analysis finds that the impact of access is largely dependent on the consumer's value of nutrition, contributing to the debate surrounding the efficacy of access-based interventions. The paper studies how the optimal combination of interventions varies depending on the consumer's characteristics, such as their idiosyncratic taste for different foods and their value

of nutrition. In addition, two different types of price interventions are considered. The analysis concludes that lump sum interventions, which includes SNAP, are strictly dominated by subsidy interventions in terms of cost-effectiveness.

Personalized versus population-level strategies. The analysis indicates that the government can achieve high levels of FV consumption by strategically segmenting populations based on their value of nutrition and other characteristics and deploying group-level strategies across these segments. This is a practical and robust alternative to completely individualized strategies.

1.2. Related literature

This work primarily builds on two streams of literature: 1) Public health and food policy literature focusing on impact analysis of strategies for increasing FV consumption; and 3) Operations research literature on optimal incentives and policies for achieving socially beneficial outcomes.

Public health and food policy. The public health and food policy literature has studied the effectiveness of price-, access-, and nutrition education-based interventions extensively. SNAP has been largely found to be successful at alleviating food insecurity (Mabli et al. 2013), however its effect on diet composition remains uncertain (Gregory et al. 2013). SNAP and other similar programs can be characterized as *untargeted* price interventions since the benefits can be used to purchase any type of food without restriction. This is contrasted with *targeted* price interventions that are limited to specific foods. The Special Supplemental Nutrition Program for Women, Infants, and Children (WIC), which provides women and children with an allowance for purchasing healthy foods, and the Healthy Incentives Program (HIP) in Massachusetts, which gives rebates for SNAP dollars spent on FVs, are two examples of successful targeted programs (Bitler and Currie 2005, Olsho et al. 2016). Due in part to the success of these programs, a large body of theoretical work studies the effects of different types of price interventions; for example, point-of-purchase strategies (Steenhuis et al. 2011), taxes on unhealthy foods (Epstein et al. 2010), and a simulated comparison of untargeted and targeted subsidies (Mozaffarian et al. 2018).

There have been numerous analyses of the impact of nutrition education on diet. Beydoun and Wang (2008) and Wardle et al. (2000) find that nutritional attitude/knowledge has a positive

relationship with FV consumption. Axelson et al. (1985) and Spronk et al. (2014) perform a meta-analysis of studies on the relationship between nutrition education (or nutritional attitudes, beliefs, etc.) on diet, and conclude that there is generally a positive relationship with FV consumption, but the association is typically weak, and the effect estimate can vary widely from study to study. Guthrie et al. (2015) point out that the impact of information campaigns is extremely complex and depends on the specific campaign as well as the characteristics of the individuals. For example, reading nutrition labels is most effective for consumers with high literacy.

Some studies consider interactions between price and education interventions by estimating heterogeneous effects, and a much smaller body of work has specifically considered complementary effects. Both Dong and Leibtag (2010) and Waterlander et al. (2013) find that FV consumption is likely to be higher when price discounts are paired with nutrition education. However, there remains a lack of understanding of optimal usage of each type of intervention and how it varies depending on other consumer characteristics, which this paper seeks to address through a more general framework.

Although a positive relationship between access to grocery stores and FV consumption has been identified in the literature (Rose and Richards 2004, Larson et al. 2009), recent studies argue that access likely does not have a strong *causal* effect on FV consumption, in part because the introduction of new grocery stores in food deserts have seen mixed results (Ver Ploeg and Rahkovsky 2016, Cummins et al. 2014, 2005, Wrigley et al. 2003, Elbel et al. 2015, Weatherspoon et al. 2013, Dubowitz et al. 2015). Allcott et al. (2019) find that access has a very small effect on FV consumption on average, and that demand-side factors play a much larger role in determining food choices. This paper explores the possibility that a lack of consideration of appropriate heterogeneous effects is one reason for the disparate findings in the literature. Through analysis of the consumer-level model, the impact of access on individual household food choices based on their characteristics is elucidated so that these interventions can be more targeted and effective.

The mechanisms linking access to shopping behavior and FV consumption have also been studied only to a limited extent. Liese et al. (2014) find that accessibility of stores may have a relationship

to FV spending but only through its relationship to shopping frequency. Levi et al. (2018) find that access does impact FV spending through shopping frequency, but only among households who have a low value of nutrition. This paper further elucidates the findings of Liese et al. (2014) and predicts the findings of Levi et al. (2018) by introducing a modeling framework that enhances the understanding of how access interventions are likely to be impacted by other consumer characteristics and interventions.

Optimal interventions for social good. In the operations management and operations research literature, a modeling approach has been used to study a wide variety of policy optimization problems for social good. Two streams of literature in this area focus on increasing availability of antimalarial drugs (Taylor and Xiao 2019, Levi et al. 2016), and promoting green technology adoption (Chemama et al. 2018, Alizamir et al. 2016, Raz and Ovchinnikov 2015, Cohen et al. 2015, Avci et al. 2014, Sierzchula et al. 2014, Perakis and Lobel 2011). Other areas that have been studied include subsidies to achieve socially beneficial outcomes for trade-in remanufacturing (Zhang and Zhang 2018), housing subsidies for the poor (Gilbert 2004), food subsidies to alleviate poverty (Besley and Kanbur 1988), agricultural crop subsidies and their impacts on farmer behavior (Alizamir et al. 2018), and more general models for increasing consumer welfare through subsidies (Yu et al. 2018). This paper differentiates itself from the aforementioned literature both in terms of the application and the menu of interventions considered. Namely, this paper considers not only financial incentives but interventions related to education and physical access, which impact demand in different ways.

With regard to demand modeling, this paper employs a bi-level (or sequential game) approach in order to directly model the mechanisms that influence consumer behavior. A subset of the aforementioned literature also adopt this approach (Levi et al. 2016, Zhang and Zhang 2018, Yu et al. 2018, Perakis and Lobel 2011, Avci et al. 2014, Chemama et al. 2018). For example, Levi et al. (2016) consider a government whose objective is to maximize market consumption of a good by providing uniform subsidies to firms, with the lower-level problem reflecting competition

between firms. Avci et al. (2014) consider adoption of electric vehicles and model the lower-level problem as a utility-maximizing consumer whose utility function consists of four different additive components: utility from driving, range inconvenience, green utility, and direct costs. This paper also models the consumers as utility-maximizers—an approach that has been used to model food purchasing decisions Allcott et al. (2019), Dubois et al. (2014)—although the model in this paper is specifically tailored to capture the impacts of price, education, and access-related interventions on the consumer’s decisions.

The remainder of this paper is organized as follows. Section 2 explains the modeling framework—both the consumer-level model and the upper-level government objective. Section 3 presents the analytical results, which includes results related to the tractability of the model, new policy insights, and a comparison of targeted versus untargeted price interventions. Section 4 explores the government’s problem of choosing strategies across populations in a realistic setting where completely personalized interventions are not feasible. Finally, Section 5 concludes and summarizes the policy implications.

2. Model Description

This paper considers a bi-level model that captures a sequential game, depicted visually in Figure 1, in which the government first decides on a budget allocation to three types of interventions, and a consumer or household then makes their food purchasing and shopping decisions given these interventions and personal objective. The term consumer and household are used interchangeably, and the model can easily be applied to a household with multiple members who shares groceries. The government’s goal is to increase the consumer’s time-averaged consumption of healthy food, and can influence the consumer’s food shopping behavior by employing three different levers; specifically, pricing, education, and access-based interventions. Given the government’s budget allocation, the consumer’s behavior is affected. Specifically, the consumer decides on a bundle of foods to buy each store visit which also impacts their shopping frequency. To further introduce the model, the consumer’s decision problem is discussed next.

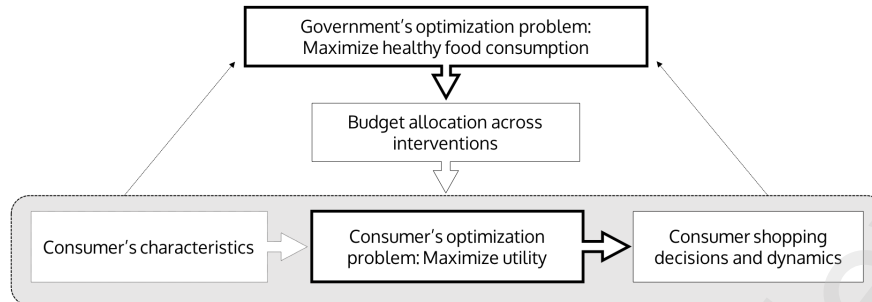


Figure 1 Illustration of bi-level model.

2.1. The consumer's shopping decisions and dynamics

The consumer-level model presented in this paper follows the modeling approach taken in Allcott et al. (2019) and Dubois et al. (2014), in which consumers are assumed to purchase quantities of various food groups in order to maximize their own personal utility, subject to a budget constraint. The utility gained from each food group depends on the characteristics of the food as well as the quantity consumed. This paper takes a similar modeling approach, however, it is specifically tailored to study the effects of access, value of nutrition, and price interventions on fruits and vegetable consumption. Although there are certainly many factors that impact consumers' food purchasing decisions and many modeling approaches that could be employed, a utility-maximization model has many benefits as noted by Dubois et al. (2014). For example, it is able to incorporate consumer-level heterogeneous preferences for different foods, connects to other commonly used choice models, and is easily generalizable to other settings.

Consider a consumer who makes recurring trips to the grocery store, and buys a bundle of food each visit. The decision regarding the bundle of food to buy, as well as the frequency of shopping trips, is assumed to depend on the consumers' characteristics. Three of these characteristics are of particular interest because of their relationship to the policy interventions that are considered in this paper. These characteristics are: 1) *Shopping disutility*, which includes the consumer's time, money, and opportunity cost required to make a shopping trip; 2) *Value of nutrition*, which can be thought of as how much the consumer cares about eating healthy; and 3) *Food budget*, which is the amount of money per unit time that the consumer puts towards food purchases. All

other characteristics of the consumer that impact their food purchasing decisions are nested into a consumer-specific idiosyncratic term, referred to as the consumer's *taste* for different food groups. This term can depend on the characteristics of the foods, as well as the consumers' characteristics.

For simplicity of exposition, food is categorized into two groups, which are referred to as “healthy” food (HF) and “regular” food (RF). These terms are used for simplicity, and with no intention of implying that all foods in the “regular” category are not healthy. In this paper, healthy food is considered to be fruits and vegetables (FVs), since the topic of interest is FV consumption. The modeling framework can naturally be extended to an arbitrary number of food groups with varying degrees of healthfulness. This extension is shown in Appendix EC.3.

Based on the three characteristics above, as well as their idiosyncratic taste, the consumer decides how much HF and RF to buy each shopping trip and how often to go shopping. The shopping dynamic described is supported by the Nielsen consumer panel dataset. This public dataset contains grocery shopping visits and purchases for over 60,000 households. The authors classify each item purchased as either HF or RF based on whether they fall into the USDA category of fruits and vegetables. Therefore, for each grocery store visit, each households' HF and RF spending is calculated. Figure 2 illustrates four households' shopping decisions and dynamics over a 300 day period. Among these households, there is significant variation in shopping frequency, FV (or HF) purchasing, and non-FV (or RF) purchasing.

It is assumed that these shopping decisions and dynamics are the result of individual consumers solving a personalized optimization problem. Namely, the consumer chooses the food bundle, (h, u) —a quantity of HF (h) and RF (u)—by maximizing a combination of their *taste*, *value of nutrition*, and time-averaged *disutility* for grocery shopping. This decision also induces a shopping cycle, \mathcal{T} , which is the time between shopping trips. In order to determine (h, u) and subsequently



Figure 2 Shopping decisions and dynamics for four households contained in the 2016 Nielsen consumer dataset. FV and non-FV spending are standardized to account for household size.

\mathcal{T} , the consumer is assumed to solve the following optimization problem:

$$\begin{aligned}
 & \max_{u,h} f(u/\mathcal{T}, h/\mathcal{T}) + v \cdot g(h/\mathcal{T}) - d/\mathcal{T} \\
 & \text{s.t. } p_h h + p_u u \leq B \cdot \mathcal{T} \\
 & \mathcal{T} = T(h, u; d, v) \\
 & u, h \geq 0
 \end{aligned}
 \tag{M^c(B, d, v)}$$

The consumer's objective function is comprised of three components:

1. *The taste term, $f(h/\mathcal{T}, u/\mathcal{T})$:* The consumer's idiosyncratic utility for HF and RF, which is assumed to be concave and separable with positive partial third derivatives with respect to each argument. Note that this requirement is not restrictive as most concave functions, especially the most widely used utility functions, exhibit this property.
2. *The nutrition value term, $v \cdot g(h/\mathcal{T})$:* If the consumer has a positive value for nutrition ($v > 0$), additional utility is gained by purchasing HF. This is determined by the function $g(\cdot)$, which is assumed to be increasing and concave.
3. *The time-averaged disutility term, d/\mathcal{T} :* Upon each visit to the supermarket, the consumer incurs a shopping disutility cost of d , and thus the time-averaged shopping disutility is d/\mathcal{T} .

Households who grocery shop online, for example, would have a disutility close to zero and their shopping decisions would not be significantly impacted by this term. Households who travel far to grocery stores will be more impacted by this term.

The consumer's budget constraint is derived by assuming that the consumer has a fixed food budget per unit time, B , and accumulates an aggregate budget to spend on each store visit, $B \cdot \mathcal{T}$. The unit prices for HF and RF are p_h and p_u , respectively. In Section 3, Proposition 1 shows that Problem $M^c(B, d, v)$ is equivalent to a concave optimization problem by applying a transformation of variables.

Throughout the analysis it is assumed that the consumer's cycle time is given by $T(h, u; d, v) = t_1(h, u, v) + t_2(v, d)$. Intuitively, t_1 represents the time it takes for a critical mass of food to run out (*depletion time*), and t_2 is the additional time it takes to return to the store once this critical mass is depleted (*waiting time*). Assumption 1 considers more specific attributes of the functions $t_1(\cdot)$ and $t_2(\cdot)$.

ASSUMPTION 1 (Functional form of T). $t_1(\cdot)$ is of the form $t_1(\phi(v)h + (1 - \phi(v))u)$, where $\phi(v) : \mathbb{R}_+ \rightarrow \{0, 1\}$ and is non-decreasing in v , and $t_1(\cdot)$ is concave increasing in its argument. Additionally, $t_2(\cdot)$ is of the form $t_2(\phi(v) \cdot v, d)$ and is increasing in d , decreasing in $\phi(v) \cdot v$, and satisfies $t_2(\phi(v) \cdot v, d) > 0$.

Assumption 1 implies that $t_1(\cdot)$ is either a function of h or a function of u . The logic behind this assumption is the following. The depletion time, t_1 , can be thought of as the time required for a "critical food item"—an item that the consumer cares about having in stock—to deplete by some amount. For consumers who consider certain healthy, perishable items to be "critical," one of these food items will likely be the determining factor and thus $t_1(\cdot)$ will depend on h , implying that $\phi(v) = 1$. Otherwise, the limiting factor will be a regular food item and $t_1(\cdot)$ will depend on u , implying that $\phi(v) = 0$. Households with higher value of nutrition are more likely to have $t_1(\cdot)$ depend on h , and thus $\phi(v)$ is increasing in v .

The second component of $T(\cdot)$ represents the waiting time, given by $t_2(\phi(v) \cdot v, d)$. Once a food from the consumer's critical bundle is depleted to a certain level, the consumer has the desire to




Figure 3 Median average cycle time (in weeks), versus binned average healthy (FV) on the left, and binned average regular (non-FV) on the right. The transparency of the points represents the number of households in each bin.

return to the store. However, because of shopping disutility, the consumer may not return to the store immediately. The larger the shopping disutility, the longer the delay. The assumption is that if the consumer's critical food bundle contains HF, the increased urgency to return to the store may also depend on her value of nutrition. Thus, $t_2(\cdot)$ can depend on v when $\phi(v) = 1$.

Informed by the Nielsen consumer panel data, a specific functional form of $T(h, u; v, d)$ is introduced. Using the grocery shopping data, each household's average HF and RF spending per visit, standardized for household size, as well as their average time between store visits, is calculated. Figure 3 shows a coarsened version of cycle time versus FV spending (left) and non-FV spending (right) for households contained in the 2016 Nielsen consumer panel dataset. Based on these plots, cycle time is approximately linearly related to HF and RF spending for small- to medium-sized expenditures. However, in both cases there is a point after which cycle time is no longer as strongly correlated with spending. Motivated by Figure 3, Assumption 2 implies that cycle time is linearly increasing in either h or u ; however, due to perishability, the depletion time is bounded above by constants M^h and M^u .

ASSUMPTION 2 (Piece-wise linear T). t_1 is given by $t_1(h, u, v) = \min\{h\rho_h, M^h\}$ when $\phi(v) = 1$ and $t_1(h, u, v) = \min\{u\rho_u, M^u\}$ when $\phi(v) = 0$, where $\rho_h, \rho_u \in \mathbb{R}_+$. Additionally, t_2 is of the form $t_2(d, v) = \frac{w(d(\delta))}{r(v(\nu))}$, where $w(d(\delta))$ is non-decreasing in $d(\delta)$ and $r(v(\nu)) = \phi \cdot \hat{r}(v(\nu)) + (1 - \phi)r_0$, where $\hat{r}(\cdot)$ is concave non-decreasing in $v(\nu)$.

When Assumption 2 is employed, the analysis is focused on the case when (h^*, u^*) is such that $\rho_h h^* \leq M^h$ and $\rho_u u^* \leq M^u$. This means that the consumer is not buying more food than they could consume before it perishes (in other words, they are not wasting food)—a realistic assumption for those with limited financial means (Yu and Jaenicke 2020).

Waiting time is given by $t_2(d, v) = \frac{w(d(\delta))}{r(v(\nu))}$, where $w(d(\delta))$ is an increasing function of the consumer's disutility, and $r(v(\nu))$ can be thought of as the consumer's urgency to return to the store which is a function of her value of nutrition when $\phi(v(\nu)) = 0$. This aligns with Assumption 1.

2.2. The government's decision

The proposed model can capture a wide range of government interventions that impact the consumer's decision. This paper considers three types of interventions:

1. The financial support subsidies (β). Most of this paper considers an untargeted price intervention that increases the consumer's food budget (written as $B(\beta)$), modeled after the SNAP program. Section 3.3 considers a targeted price intervention that decreases the cost of HF.
2. The value intervention (ν). Investing in the value intervention, for example, by investing in nutrition education, increases the consumer's value of nutrition. This is captured by modeling the value of nutrition as a function $v(\nu)$.
3. The access intervention (δ). Investing in the access intervention decreases the consumer's shopping disutility. This is captured by modeling the disutility as a function $d(\delta)$. This could result from interventions such as building a new grocery store or providing easier transportation for grocery shopping.

The government's goal is to optimally choose a funding level for each lever in order to increase the consumer's average HF consumption. Therefore, the government's problem can be written as

$$\begin{aligned}
 & \max_{\beta, \delta, \nu} h^*/T(h^*, u^*; v(\nu), d(\delta)) \\
 & \text{s.t. } \beta + \delta + \nu \leq U \\
 & \beta, \delta, \nu \geq 0 \\
 & h^*, u^* = \operatorname{argmax} M^c(B(\beta), d(\delta), v(\nu))
 \end{aligned} \tag{M^g}$$

where U is the government's budget per consumer. The functions $v(\nu)$, $d(\delta)$, and $B(\beta)$ link the government's monetary investments in each intervention to value of nutrition, disutility, and food budget, respectively. In order to analyze the government's upper-level problem, the following assumptions are made.

ASSUMPTION 3.

- *The function $B(\beta)$ is increasing, concave, and satisfies $B(0) = B_0 > 0$*
- *The function $d(\delta)$ is decreasing, convex and satisfies $\lim_{\delta \rightarrow \infty} d(\delta) = 0$ and $d(0) = d_0 > 0$*
- *The function $v(\nu)$ is increasing, concave and satisfies $\lim_{\nu \rightarrow \infty} v(\nu) = v_{max}$ and $v(0) = v_0 \geq 0$*

The function $B(\beta)$ is assumed to be concave in order to represent the fact that a \$1 increase in SNAP benefits does not necessarily translate to a \$1 increase in the consumer's food budget since the consumer might divert some of her original food budget towards other uses. The function $d(\delta)$ is assumed to be convex and decreasing in order to capture diminishing marginal returns. Similarly, the concavity of $v(\nu)$ also captures diminishing returns.

3. Analytical results

This section presents results related to the tractability of the bi-level model as well as insights that can inform policy. Figure 4 summarizes the main findings and provides a mapping between the assumptions and results. Certain results hold under Assumption 1, the generic functional form of T , while others require the more specific assumption that T is a piece-wise linear function. Section 3.1 presents the results related to tractability, Section 3.2 presents policy insights, and Section 3.3 compares targeted and untargeted price subsidies.

3.1. Tractability of the bi-level model

First, consider the consumer's decision regarding the bundle of food purchased at each visit to the grocery store, which can be described by the total quantity of HF and RF bought at each visit, (h, u) . However, it is more convenient to equivalently describe the consumer's food bundle by the time averaged quantity of HF and RF, respectively, defined as $\tilde{h} := h/T(h, u; d, v)$ and $\tilde{u} := u/T(h, u; d, v)$. Lemma 1 (proof in Appendix EC.1), says that there is a one-to-one correspondence between h and \tilde{h} and between u and \tilde{u} .




Figure 4 Summary and mapping of analytical results.

LEMMA 1. Under Assumption 1, $\tilde{h}(h)$ and $\tilde{u}(u)$ are increasing functions.

Intuitively, Lemma 1 states that all else equal, an increase in the amount of HF (or RF) bought at the store will always result in an increase in the time-averaged quantity of HF (or RF). This is a consequence of the concavity of T with respect to h (or u). With this change of variables, t_1 can be written as $t_1(\phi\tilde{h} + (1 - \phi)\tilde{u}; v)$.

Therefore, Problem $M^c(B, d, v)$ can be re-written in the new variables as

$$\begin{aligned} \max_{\tilde{u}, \tilde{h}} \quad & f(\tilde{u}, \tilde{h}) + v(\nu) \cdot g(\tilde{h}) - d/T(\tilde{h}, \tilde{u}; d(\delta), v(\nu)) \\ \text{s.t.} \quad & p_h \tilde{h} + p_u \tilde{u} \leq B(\beta) \\ & \tilde{u}, \tilde{h} \geq 0 \end{aligned} \tag{M^c(\beta, \delta, \nu)}$$

The following proposition (proof in Appendix EC.1), asserts that Problem $\tilde{M}^c(\beta, \delta, \nu)$ is concave with linear constraints and is thus tractable.

PROPOSITION 1. Under Assumption 1, Problem $\tilde{M}^c(\beta, \delta, \nu)$ is concave.

By Proposition 1, the Karush-Kuhn-Tucker (KKT) conditions of Problem $\tilde{M}^c(\beta, \delta, \nu)$ can be thought of as a system of equations linking the government's strategy to the consumer's food bundle decision. The most interesting case is when both \tilde{h}^* and \tilde{u}^* are positive, meaning that the consumer purchases positive quantities of both HF and RF.

To assist in the analysis of the upper-level problem, the government's feasible space, $\mathcal{D} = \{(\beta, \delta, \nu) : \beta \geq 0, \delta \geq 0, \nu \geq 0, \beta + \delta + \nu \leq U\}$, is subsetted into two disjoint continuous domains \mathcal{D}_0 and \mathcal{D}_1 such that $\mathcal{D}_0 \cup \mathcal{D}_1 = \mathcal{D}$, given by

$$\mathcal{D}_0 = \{(\beta, \delta, \nu) \in \mathcal{D} : \phi(v(\nu)) = 0\}, \quad \mathcal{D}_1 = \{(\beta, \delta, \nu) \in \mathcal{D} : \phi(v(\nu)) = 1\}.$$

The fact that \mathcal{D}_0 and \mathcal{D}_1 are continuous is clear by noting $\phi(v(\nu))$ is non-decreasing in ν . The main tractability result is presented next, specifically, that determining the government's optimal budget allocation across access, price, and education-related interventions is computationally tractable. The following theorem asserts that the function $\tilde{h}^*(\nu, \beta, \delta)$ —average HF consumption as a function of the government's investment—is unimodal over both domains \mathcal{D}_0 or \mathcal{D}_1 when either investment in the access intervention, δ , is fixed, or under additional mild conditions on the modeling functions.

THEOREM 1. *For fixed investment in the access intervention ($\delta = \delta_0$), the consumer's average HF consumption as a function of the investment in the price and education intervention — $\tilde{h}^*(\nu, \beta, \delta_0)$ — is unimodal on both domains \mathcal{D}_0 and \mathcal{D}_1 . For non-fixed δ , $\tilde{h}^*(\nu, \beta, \delta)$ is unimodal on both \mathcal{D}_0 and \mathcal{D}_1 provided that the following additional conditions are met:*

1. $d''(\delta)w(d(\delta)) - 2d'(\delta)^2w'(d(\delta)) \geq 0$, and (Condition 1)
2. The education intervention has a non-negative budget elasticity. (Condition 2)

Condition 1 is a statement about the convexity of the function $d(\delta)$. In particular, for Condition 1 to be met it is necessary (but not sufficient) that the function $1/d(\delta)$ is strictly concave. Condition 2 implies that as the government's budget *shrinks*, the government will not allocate *more* funds to the education intervention, and vice versa. This is a regularity condition on the impact of the interventions, and can equivalently be written in terms of the modeling parameters and functions. Details on these conditions can be found in Appendix EC.1 along with the proof of Theorem 1.

3.2. Policy insights

This section presents the main policy insights obtained by analyzing the proposed model. Proposition 2 characterizes the impact of the untargeted price and value intervention— β and ν —on average

HF and RF consumption, under the most general form of T given by Assumption 1. Corollary 1 characterizes the effect of access on average HF consumption under the piece-wise linear version of T given in Assumption 2. Finally, Proposition 3 explores how the government's optimal budget allocation between the three interventions changes depending on the consumer's characteristics.

PROPOSITION 2. *Under Assumption 1, the following characterizes the impact of the untargeted price and value intervention (β and ν) on average HF and RF consumption (\tilde{h}^* and \tilde{u}^*) when both are positive ($\tilde{h}^*, \tilde{u}^* > 0$):*

1. *HF and RF consumption are increasing in the untargeted price intervention, β . However, the proportion of healthy to RF consumed, $(\tilde{h}^*/\tilde{u}^*)$, is not necessarily increasing in β .*
2. *HF consumption is increasing in the value intervention, ν , and RF consumption is decreasing in ν .*

The proof of Proposition 2, which utilizes the KKT conditions from Problem $\tilde{M}^c(\beta, \delta, \nu)$, is relegated to Appendix EC.1. Appendix EC.1 also includes a characterization of the impact of the access intervention on average HF consumption under Assumption 1, which is quite complex. Therefore, for brevity, the impact of access is characterized under the piece-wise linear form of T , given in Assumption 2, in the following corollary.

COROLLARY 1. *Under Assumption 2, the following characterizes the impact of access on average HF consumption when $\tilde{h}^*, \tilde{u}^* > 0$: When cycle time depends on HF purchasing ($\phi(v) = 1$), average HF consumption, \tilde{h}^* , is increasing in the access intervention, δ , and average RF consumption, \tilde{u}^* , is decreasing in δ if and only if $w(d) - dw'(d) < 0$. When cycle time depends on RF purchasing ($\phi(v) = 0$), \tilde{h}^* is increasing in δ and \tilde{u}^* is decreasing in δ if and only if $w(d) - dw'(d) > 0$.*

Parts 1 and 2 of Proposition 2 align with the empirical evidence regarding the impact of food stamps and nutrition education on FV consumption. Namely, these results capture the fact that although food stamps are found to increase food security and total FV consumption, at the same time they can have mixed effects on overall diet composition (Mabli et al. 2013, Gregory et al.

2013). Furthermore, the model captures the consensus in the literature that nutrition education has a positive impact on FV consumption (Wardle et al. 2000, Axelson et al. 1985). These results help validate the proposed model in this paper.

Corollary 1 illustrates the intricate relationship between access and FV consumption which has also been well documented in the literature. Simply stated, Corollary 1 asserts that the access intervention does not always increase average FV consumption. In what follows, the focus is on explaining this phenomenon intuitively.

All else equal, the consumer would like to shop less frequently in order to minimize her time-averaged shopping disutility. When cycle time depends on HF purchasing (i.e., when $\phi(v) = 1$), the consumer can increase cycle time by increasing HF spending. Therefore, when the magnitude of the disutility term increases, the consumer will purchase more HF per visit (corresponding to more HF on average by Lemma 1), and when the magnitude of the disutility term decreases, the consumer will purchase less HF per visit (corresponding to less HF on average by Lemma 1). Therefore, the impact of the access intervention on HF consumption depends on its effect on the magnitude of the disutility term.

Consider the disutility term, given by $d(\delta)/T(\tilde{h}; d(\delta))$. Under the piece-wise linear form of T given by Assumption 2, and by applying the change of variables $h \rightarrow T \cdot \tilde{h}$, the disutility term can be written as $\frac{d(\delta)(1-\rho_h \tilde{h})r}{w(d(\delta))}$. Note that $1 - \rho_h \tilde{h} > 0$ because $T > \rho_h h$ and $\tilde{h} = h/T$. Therefore, when $\frac{\partial d(\delta)/w(d(\delta))}{\partial \delta} > 0$, increasing the access intervention will increase the magnitude of the disutility term and thus cause an increase HF spending when $\phi(v) = 1$, or cause an increase in RF spending when $\phi(v) = 0$. Noting that $d'(\delta) < 0$, the condition $\frac{\partial d(\delta)/w(d(\delta))}{\partial \delta} > 0$ is equivalent to $w(d) - dw'(d) < 0$, which is the condition stated in Corollary 1 for the case that $\phi(v) = 1$.

In the case that $\phi(v) = 0$, similar logic shows that HF consumption is increasing in δ if and only if $w(d) - dw'(d) > 0$. Notice that because $w(d)$ is concave, it is always true that $w(d) - dw'(d) > 0$. Therefore, when $\phi(v) = 0$, average HF consumption always increases as access increases. In other words, when $w(d)$ is concave, access interventions are generally effective for consumers with a low

value of nutrition (who are more likely to have $\phi(v) = 0$). This agrees with the empirical findings of Levi et al. (2018).

In conclusion, since the access intervention can either increase or decrease the magnitude of the disutility term, its impact on HF is also mixed. The impact of access on HF depends on the relationship between disutility and cycle time, as well as consumer's value of $\phi(v)$.

This complex relationship between access and food purchasing helps explain why the impact of access-related interventions has been so highly debated in the literature (Ver Ploeg and Rahkovsky 2016, Cummins et al. 2014, 2005, Wrigley et al. 2003, Elbel et al. 2015, Weatherspoon et al. 2013, Dubowitz et al. 2015). This framework and corresponding insights can assist in predicting the impact of future access-related interventions on consumers based on their individual characteristics. Using this approach, more targeted access-related interventions can be deployed, increasing their cost-effectiveness.

Corollary 2 (below) highlights the fact that the government will only invest in access when it increases average HF consumption, namely, when the conditions of Corollary 1 are met. When the conditions are not met, the government can fix $\delta^* = 0$, which makes the government's optimization problem simpler as stated in Theorem 1.

COROLLARY 2. *Under Assumption 2, the government's optimal solution, $(\beta^*, \nu^*, \delta^*)$, will only have $\delta^* > 0$ when the conditions of Corollary 1 are met.*

It is also interesting to explore how the government's optimal budget allocation changes with respect to certain characteristics of the consumer. For simplicity, the case where $\delta^* = 0$ is discussed. Equivalently, the following results apply when δ (investment into the access intervention) is fixed and the government only considers jointly investing between the price and education levers.

PROPOSITION 3. *For a consumer with fixed disutility, $d(\delta)$, the government's optimal allocation between the price and education levers satisfies:*

1. *Allocation to the education lever is decreasing (and thus allocation to the price lever is increasing) as the consumer's initial value of nutrition increases.*

2. *Allocation to the price lever is decreasing (and thus allocation to the education lever is increasing) as the consumer's utility for RF has less diminishing returns (i.e., as $f(\tilde{h}, \tilde{u})$ becomes more linear in \tilde{u})*
3. *The government's allocation could change in either direction as the consumer's budget changes.*

Part 1 of Proposition 3 provides the intuitive result that consumers with higher value of nutrition are better off with more investments in price-related interventions. Part 2 says that when the consumer's utility for RF becomes less concave, i.e., when $\frac{\partial f^2(\tilde{h}, \tilde{u})}{\partial \tilde{u}^2}$ is closer to zero, consumers are better off with education-related interventions. Intuitively this result makes sense. Specifically, consumers with a less nonlinear utility in terms of RF are more likely to spend their food budget on RF. Therefore, increasing their food budget will not be effective at increasing HF consumption since most of the money will go towards RF.

Finally, Part 3 in Proposition 3 tells us that households with larger food budgets are not necessarily better off with value-related interventions. In other words, simply because a consumer has been given food stamps does not mean that they would be benefited by nutrition education more than they would be benefited by an even larger food budget. These insights reveal the importance of tailoring interventions to each individuals' characteristics. However, complete personalization of interventions is often not realistic in public policy settings. Therefore, the Section 4 focuses on determining optimal group-level strategies across populations. In what follows, the effectiveness of targeted versus untargeted price interventions is discussed.

3.3. Targeted price subsidies

Thus far, the analysis has considered an untargeted price intervention, which corresponds to the national food stamp program, SNAP. In this section, another common type of price intervention—a targeted price subsidy—is considered that decreases the price of HF (which is, in effect, what many targeted price interventions do). This is modeled by modifying the consumer's budget constraint in Problem $\tilde{M}^c(\beta, \delta, \nu)$. Let β_{unt} denote the untargeted price subsidy level (equivalent to β in the previous sections) and β_{tar} denote the targeted price subsidy lever. The untargeted budget

constraint is $p_h \tilde{h} + p_u \tilde{u} = B(\beta_{unt})$, whereas the targeted budget constraint is given by $(p_h - \beta_{tar}) \tilde{h} + p_u \tilde{u} = B$. In order to fairly compare the two interventions, it is assumed that β_{unt} linearly affects the food budget, meaning that $B(\beta_{unt}) = B + \beta_{unt}$. In the targeted case, the government's budget constraint is given by $\tilde{h}^* \beta_{tar} + \nu + \delta \leq U$ with the additional constraint that $\beta_{tar} \leq p_h$.

The main result of this section reveals that targeted subsidies are more cost-effective than untargeted subsidies. Consider fixing the total amount of money that the government spends on price subsidies to k , (so that $\beta_{unt} = k$ in the untargeted case and $\beta_{tar} \tilde{h}^* = k$ in the targeted case). Let $\tilde{h}_{tar}^*(k)$ be the consumer's average HF consumption under the targeted price subsidy scheme, and $\tilde{h}_{unt}^*(k)$ be the average HF consumption under the untargeted price subsidy scheme.

PROPOSITION 4. *Targeted subsidies strictly dominate untargeted subsidies, in the sense that $\tilde{h}_{tar}^*(k, \nu_0, \delta_0) \geq \tilde{h}_{unt}^*(k, \nu_0, \delta_0)$, for all k and fixed ν_0, δ_0 .*

The proof of Proposition 4 can be found in Appendix EC.1, and relies on the fact that the consumer's feasible set of food bundles are identical in both the targeted and untargeted case when the government spends k dollars on either type of price subsidy.

4. From individuals to populations

The model presented thus far concerns individual consumers. However, in reality the government often has to select and implement strategies over populations of consumers. It may be neither practical nor realistic for the government to provide personalized interventions to each individual in the population. This section presents a method for assigning households to groups and choosing group-level strategies.

When designing strategies across populations, the location of households plays a critical role for certain interventions. Many access-related interventions are location-dependent. For example, a new grocery store will only impact households living within a certain distance to the new store. However, other types of access interventions are largely location-independent—meaning that they can be given to individual households without much consideration of the household's location—for example, subsidizing transit to grocery stores or grocery delivery services. This section considers

both types of access interventions. Price and education interventions are both assumed to be location-independent, which is realistic in practice. Intervention bundles consisting of location-independent interventions can be given to groups of households regardless of their locations. On the other hand, location-dependent interventions (also referred to as simply location interventions) can only be applied to groups of households within the same neighborhood.

The population of households considered for this analysis is described in the next section. The government's optimization problem for choosing the group-level strategies and location interventions is then described, along with the proposed method for determining the groups. Finally, the government's optimal group-level strategy is computed and each household's FV consumption under their group-level strategy is compared to their FV consumption under their individual optimal strategy.

4.1. Population data

In order to use the proposed model to inform strategies across populations and within neighborhoods, granular shopping and geographic data is required for a sample of households within the neighborhoods and populations of interest. For SNAP households, which are often considered the population of interest in food policy, this granular data is accessible. SNAP benefits are administered to households through a personal electronic benefits transfer (EBT) card, which are used at retailer outlets like a debit card. Therefore, every transaction involving an EBT card is recorded. In states like Massachusetts, which use a targeted price incentive in conjunction with SNAP (called the Healthy Incentives Program), certain FV purchases are recorded separately. Thus, SNAP spending on FVs can be roughly calculated for every household that participates in this supplementary program. This method would only capture SNAP transactions, however, the literature regarding food spending patterns of SNAP households could be used to estimate households' non-SNAP expenditures (e.g., Tiehen et al. (2017), Dorfman et al. (2018)).

Disutility and value of nutrition can also be readily estimated from data. Disutility can be estimated through demographic and geographic information, which are both inquired about when

a household signs up for SNAP. The proximity of households to grocery stores—a large component of disutility—is easily determined from this geographic information. Furthermore, information such as employment and household composition can be used to estimate the opportunity cost of grocery shopping—another potential contributor to disutility.

Value of nutrition can be estimated a number of ways. The USDA administers many surveys that contain questions regarding the households' beliefs and attitudes towards nutrition. These questions could be widely disseminated to SNAP households as part of the sign-up process or during the re-certification process. Value of nutrition could also be estimated at the neighborhood level by using Google search data to assess the frequency of nutrition-related searches, as in Levi et al. (2018).

Because much of the data described above is not publicly or currently available, this analysis considers a subset of the population of households contained in the 2016 Nielsen consumer panel dataset to be the population of interest. Certain attributes of the households—specifically, their value of nutrition and disutility—are synthetically generated (according to a process described in Appendix EC.4.2) in order to demonstrate the functionality of the model. The subset of households from the panel used in this analysis are those who satisfy a set of criteria: including (i) households below the federal poverty guideline, (ii) households that have at least 20 grocery store trips recorded in 2016, and (iii) households that have an average shopping cycle time of less than 20 days. The resulting dataset consists of 1,193 households. Using this semi-synthetic data, the modeling parameters and functions are estimated for each household according to the procedure described in Appendix EC.4.

Because the Nielsen dataset does not contain many households within a single neighborhood, the households contained in the dataset are assigned to four synthetic “neighborhoods.” The neighborhood assigned to household i is denoted n_i . In the case of location interventions, the government's strategy will depend on the characteristics of the neighborhoods. For illustrative purposes it is assumed that neighborhoods are correlated in terms of their value of nutrition and shopping disutility. The characteristics of the four neighborhoods are the following.

- Neighborhood 1. Low value of nutrition, low shopping disutility.
- Neighborhood 2. High value of nutrition, low shopping disutility.
- Neighborhood 3. Low value of nutrition, high shopping disutility.
- Neighborhood 4. High value of nutrition, high shopping disutility.

4.2. Group-level objectives

The optimal intervention bundle for individual households varies widely depending on their characteristics, and are quite complex in the sense that these optimal strategies cannot be predicted well using linear models (more details on the optimal individual strategies for the Nielsen population of households are given in Appendix EC.5.1). Therefore, it may be unrealistic for the government to provide completely personalized interventions. Suppose that the government instead wishes to deploy interventions across groups as well as within neighborhoods.

Each household is assigned to a group, independent of their neighborhood. The government's goal is to choose a single strategy consisting of an allocation to the price, education, and location-independent access intervention, denoted $(\beta_k, \delta_k, \nu_k)$ to deploy to every household in \mathcal{P}_k , for $k = 1, \dots, M$, where M is the number of groups and \mathcal{P}_k is the population of households in group k . In addition, the government must choose in which neighborhoods, if any, to implement a location intervention. Therefore, every household receives a group-level intervention bundle as well as potentially a location intervention, depending on their neighborhood.

There are many possible objective functions that the government could employ in order to choose the group-level interventions. Different objective functions can be thought of as corresponding to different notions of fairness (Bertsimas et al. 2011). This paper considers two approaches. The first objective considered is the utilitarian solution of maximizing the sum of HF consumption among all households. The second approach is a max-min objective, where the government wishes to employ the strategy that achieves the highest minimum outcome among households in each group. An outcome in this case is defined as the ratio of the percentage increase in FV consumption under the chosen strategy to the percentage increase in FV consumption under each household's optimal

individual strategy. The problem formulations, as well as solution technique, are given in Appendix EC.5.2. The budget constraint for both problems is modeled as $\beta_k + \nu_k + \delta_k \leq U - \frac{cL}{M}$, where L is the number of location interventions chosen and c is the cost of a single location intervention. Let M_{util}^P represent the utilitarian optimization problem and $M_{max-min}^P$ represent the max-min optimization problem.

The performance of an arbitrary strategy \mathbf{s} is evaluated based on its *performance ratio*. For each strategy \mathbf{s} , its optimality percentage for household i is defined as

$$\text{perf_ratio}_i(\mathbf{s}) := \frac{\tilde{h}_i^*(\mathbf{s}) - \tilde{h}_i^*(\mathbf{0})}{\tilde{h}_i^*(\beta_i^*, \nu_i^*, \delta_i^*) - \tilde{h}_i^*(\mathbf{0})},$$

which is the ratio of the percent increase in FV spending under strategy \mathbf{s} to the percent increase in FV spending under the optimal individual strategy for household i , denoted $(\beta_i^*, \nu_i^*, \delta_i^*)$. The term $\tilde{h}_i^*(\mathbf{0})$ denotes household i 's HF spending without any interventions. Without location interventions, the performance ratio is always between zero and one. However, in neighborhoods with a location intervention it is possible that the performance ratio could exceed one, depending on the effectiveness of the location intervention.

4.3. Determining groups

The goal of this section is to develop simple methods for producing “smart” groups. The key question is: *How should households be segmented into groups in order to maximize performance of population-level strategies?*

In order to choose the groups, a multi-output regression tree is trained to predict the optimal individual intervention bundles, $(\beta_i^*, \nu_i^*, \delta_i^*)$. Households are classified into groups according to their corresponding leaves of the tree. To produce m groups, a tree with m leaves is trained. Details are given in Appendix EC.5.3. This approach is employed because it produces interpretable groups based on a small number of covariates, which is desirable in practice and simple to implement.

When choosing the number of groups, there is a trade-off between performance and complexity. This analysis focuses on determining three groups. A regression tree with three leaves is thus trained

to predict the households' individual optimal intervention bundles. The trained regression tree can be seen in Figure EC.5. The covariates chosen for the splits of the tree are *value of nutrition* and *taste* for RF. Once the households are partitioned into three groups, Problems M_{util}^P and $M_{max-min}^P$ are solved.

4.4. Group-level strategy performance

The solutions to Problems M_{util}^P and $M_{max-min}^P$ prescribe neighborhoods in which to implement location interventions (e.g., which neighborhoods to build a new grocery store in), as well as funding levels for the education, price, and location-independent access interventions, which are unique to each group. This section assesses the performance of these group-level strategies.

The government's decision will depend on the cost of the location intervention, denoted c . First consider the case when the location interventions are free. In this case, the government's optimal decision is to implement a location intervention in Neighborhoods 1 and 3, which are the neighborhoods where most households have a low value of nutrition. This is line with Proposition 2, which says that access interventions are only effective at increasing HF consumption for households with a low value of nutrition, under a certain range of modeling parameters. In this case, the modeling parameters (in particular, the parameters of the function $w(d(\delta))$) are estimated to fall in this range. Therefore, even when location interventions are free, the government chooses to implement them only in neighborhoods with a low value of nutrition.

As c increases, the government's choice regarding the number of location interventions changes. In the scenario presented in this section, the utilitarian strategy will only fund location interventions in both Neighborhoods 1 and 3 if $c \leq .09U$, whereas the max-min strategy will fund both location interventions when $c \leq .51U$. There is also a range of c for both the max-min and utilitarian objectives such that a location intervention is implemented only in Neighborhood 3. This neighborhood contains households with a low value of nutrition and a high shopping disutility, and is thus the neighborhood where access interventions are most effective.

Even without location interventions, the group-level strategies achieve high performance ratios. Under the utilitarian strategy, without any location interventions the average performance ratio is




Figure 5 Performance of population strategies without a location intervention (black) compared, and difference performance (shown in green/red) when a location intervention is placed in Neighborhoods 1 and 3.

84% (with a median of 94% and minimum of 16%), and under the max-min strategy the average performance ratio is 81% (with a median of 84% and minimum of 34%).

Figure 5 shows how these performance ratios change at the household level when location interventions are implemented in Neighborhoods 1 and 3 at a cost of $c = .09U$. As expected, most households in Neighborhoods 1 and 3 are benefited by the location interventions, and most households in Neighborhoods 2 and 4 would be better off without it, since the location interventions decrease the amount of funds left for other interventions. The benefit of the location interventions is drastic for many households in Neighborhoods 1 and 3, and the downside for is small for many households in Neighborhoods 2 and 4. Furthermore, the households with the lowest performance ratios to begin with are the ones who are most benefited.

It is also worth noting that some households in Neighborhoods 2 and 4 are benefited by the location interventions. These appear to be the households with the lowest performance ratios without the location intervention. The group-level strategies for these households contain a positive investment in the access intervention, however these particular households are not benefited by the access intervention and hence have low performance ratios. When the location interventions are implemented in Neighborhoods 1 and 3, the group-level strategy responds by investing less in the location-independent access intervention, and thus these households are benefited. Therefore, when neighborhoods are correlated in terms of their value of nutrition, location interventions can effectively improve the performance of group-level strategies.

The above results suggest a number of policy insights. First, group-level strategies based on “smart” groups formed using value of nutrition and taste for RF can achieve near optimal levels of HF consumption for many households. This implies that completely personalized interventions may not be necessary, especially in cases where providing personalized interventions is difficult or costly.

Second, deploying access-related interventions at the population level should be done with caution. Access interventions are only effective among households with certain characteristics (described in Proposition 2), and therefore the characteristics of neighborhoods should be carefully assessed before implementation location-based access interventions. If value of nutrition can effectively be measured at the neighborhood level, then location interventions could be strategically deployed in a manner that is beneficial to most households.

5. Discussion and Policy Implications

The model proposed in this paper provides both a new tool for informing future policymaking, as well as new insights into current policies. First, the paper finds that targeted price subsidies are strictly better than untargeted price subsidies at increasing FV consumption. This suggests that the SNAP program could be more effective at increasing HF consumption if it also provided a targeted supplemental program, such as the Healthy Incentives Program in Massachusetts.

Many of the results elucidate the importance of the consumer’s value of nutrition, and the need for standardized and precise methods for its measurement. Currently, USDA surveys often inquire about nutrition-related behaviors and knowledge of nutrition. However, how the responses actually connect to the consumer’s true value of nutrition is not well established. Developing a specific series of questions aimed at understanding the extent to which nutrition affects the consumers’ thought process regarding their food purchasing choices would be extremely helpful in determining value of nutrition, in addition to other techniques such as analysis of Google search data.

Value of nutrition plays an important role in determining which portfolio of interventions will be the most effective for different consumers. The impact of access on consumer behavior is complex

and creates two competing incentives for the consumer. The direction of each incentive is highly influenced by the consumer's value of nutrition. In particular, investing in access is likely to be effective among consumers with low value of nutrition. Therefore, access interventions should be deployed *after* estimating individual-level or neighborhood-level value of nutrition. This could help preemptively identify situations where access interventions are unlikely to be effective, thus saving time and money. The complex relationship between access and value of nutrition helps explain why the results of access interventions have been mixed, and without considering the appropriate effect heterogeneity, the impact of access cannot be elucidated.

The analysis of group-level strategies shows that it is not necessary to provide completely individualized interventions in order to achieve large increases in FV consumption. It appears that by subsetting households based on value of nutrition and idiosyncratic taste, effective group-level strategies can be deployed. This is a simpler and more practical alternative to completely personalized interventions. Location-dependent access interventions can also be useful but only when they target neighborhoods with a low average value of nutrition.

Outside of food policy, the model proposed in this paper can be used to optimize investment allocation in a variety of different public policy settings. For example, consider a non-governmental organization's efforts to increase voter turnout. These may include education campaigns (about various policy initiatives, candidates, where to vote, etc.), voter registration initiatives, scheduled rides to polling places, and many more. These (primarily education- and access-related) interventions impact an individual's decision about whether to vote in very different ways. A model that captures the impact of these interventions on individual decision-making could be used to determine optimal personalized strategies. Furthermore, group-level strategies could be constructed by first developing consumer-level models on a sample of the population for which rich data is available. The methods of Section 4 can then be used to determine group assignments and group-level strategies. A group classifier (such as the decision tree in Section 4) could be constructed that is based solely on a couple covariates that are widely observable for the general public in order to classify all potential voters into groups.

Models such as those developed in this paper and described above are especially relevant in policy settings where data is largely observational and experiments are often costly or infeasible. Furthermore, most public policy initiatives have multiple types of levers that impact individuals through various mechanisms. In these cases, strategic policy design relies on the development of realistic data-driven models. Individual-level models that are able to simultaneously capture the effects of different levers provide many benefits, including an increased understanding of 1) how different interventions interact with one another, 2) how interventions interact with characteristics of the individual, and 3) how to centrally allocate funds across the various interventions. Use of these methods in policy design can lead to more personalized and cost-effective strategies, and ultimately better outcomes.

Acknowledgments

This research is supported in part by the National Science Foundation Graduate Research Fellowship. Researcher(s) own analyses calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at the University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher(s) and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

References

- Alizamir S, de Véricourt F, Sun P (2016) Efficient feed-in-tariff policies for renewable energy technologies. *Operations Research* 64(1):52–66.
- Alizamir S, Iravani F, Mamani H (2018) An analysis of price vs. revenue protection: Government subsidies in the agriculture industry. *Management Science* 65(1):32–49.
- Allcott H, Diamond R, Dubé JP, Handbury J, Rahkovsky I, Schnell M (2019) Food deserts and the causes of nutritional inequality. *The Quarterly Journal of Economics* 134(4):1793–1844.
- Avcı B, Girotra K, Netessine S (2014) Electric vehicles with a battery switching station: Adoption and environmental impact. *Management Science* 61(4):772–794.
- Axelson ML, Federline TL, Brinberg D (1985) A meta-analysis of food- and nutrition-related research. *Journal of Nutrition Education* 17(2):51 – 54, ISSN 0022-3182, URL <http://www.sciencedirect.com/science/article/pii/S0022318285801758>.

- Bertsimas D, Farias VF, Trichakis N (2011) The price of fairness. *Operations Research* 59(1):17–31.
- Besley T, Kanbur R (1988) Food subsidies and poverty alleviation. *The Economic Journal* 98(392):701–719.
- Beydoun MA, Wang Y (2008) How do socio-economic status, perceived economic barriers and nutritional benefits affect quality of dietary intake among us adults? *European Journal of Clinical Nutrition* 62(3):303.
- Bitler MP, Currie J (2005) Does WIC work? The effects of WIC on pregnancy and birth outcomes. *Journal of Policy Analysis and Management* 24(1):73–91.
- Chemama J, Cohen MC, Lobel R, Perakis G (2018) Consumer subsidies with a strategic supplier: Commitment vs. flexibility. *Management Science* 65(2):681–713.
- Cohen MC, Lobel R, Perakis G (2015) The impact of demand uncertainty on consumer subsidies for green technology adoption. *Management Science* 62(5):1235–1258.
- Cummins S, Flint E, Matthews SA (2014) New neighborhood grocery store increasing awareness of food access but did not alter dietary habits or obesity. *Health Affairs* 33(2):283–291.
- Cummins S, Petticrew M, Higgings C, Findlay A, Sparks L (2005) Large scale food retailing as an intervention for diet and health: quasi-experimental evaluation of a natural experiment. *J Epidemiol Community Health* 59(1035-1040).
- Delyon B, Lavielle M, Moulines E (1999) Convergence of a stochastic approximation version of the em algorithm. *Annals of statistics* 94–128.
- Dong D, Leibtag E (2010) Promoting fruit and vegetable consumption: are coupons more effective than pure price discounts? Economic Research Report 96, U.S. Department of Agriculture, Economic Research Service.
- Dong D, Lin BH (2009) Fruit and vegetable consumption by lowincome americans: Would a price reduction make a difference? Economic Research Report 70, U.S. Department of Agriculture, Economic Research Service.
- Dorfman JH, Gregory C, Liu Z, Huo R (2018) Re-Examining the SNAP Benefit Cycle Allowing for Heterogeneity. *Applied Economic Perspectives and Policy* 41(3):404–433, ISSN 2040-5790, URL <http://dx.doi.org/10.1093/aep/ppy013>.
- Dubois P, Griffith R, Nevo A (2014) Do prices and attributes explain international differences in food purchases? *American Economic Review* 104(3):832–67.
- Dubowitz T, Ghosh-Dastidar M, Cohen DA, Beckman R, Steiner ED, Hunter GP, Flórez KR, Huang C, Vaughan CA, Sloan JC, et al. (2015) Changes in diet after introduction of a full service supermarket in a food desert. *Health Affairs* 34(11):1858.
- Elbel B, Moran A, Dixon LB, Kiszko K, Cantor J, Abrams C, Mijanovich T (2015) Assessment of a government-subsidized supermarket in a high-need area on household food availability and children's dietary intakes. *Public Health Nutrition* 18(15):2881–2890.

- Epstein LH, Dearing KK, Roba LG, Finkelstein E (2010) The influence of taxes and subsidies on energy purchased in an experimental purchasing study. *Psychological Science* 21(3):406–414.
- Gilbert A (2004) Helping the poor through housing subsidies: lessons from chile, colombia and south africa. *Habitat international* 28(1):13–40.
- Gregory CA, Ver Ploeg M, Andrews M, Coleman-Jensen A (2013) Supplemental nutrition assistance program (SNAP) participation leads to modest changes in diet quality. Economic Research Report 147, U.S. Department of Agriculture, Economic Research Service.
- Guthrie J, Mancino L, Lin CTJ (2015) Nudging consumers toward better food choices: Policy approaches to changing food consumption behaviors. *Psychology & Marketing* 32(5):501–511.
- Larson N, Story M, Nelson M (2009) Neighborhood environments: Disparities in access to healthy foods in the U.S. *American Journal of Preventative Medicine* .
- Levi R, Paulson E, Perakis G (2018) Fresh fruit and vegetable consumption: The impact of access and value. *Working Paper* .
- Levi R, Perakis G, Romero G (2016) On the effectiveness of uniform subsidies in increasing market consumption. *Management Science* 63(1):40–57.
- Liese AD, Bell BA, Barnes TL, Colabianchi N, Hibbert JD, Blake CE, Freedman DA (2014) Environmental influences on fruit and vegetable intake: results from a path analytic model. *Public Health Nutrition* 17(11):2595–2604.
- Luo ZQ, Pang JS, Ralph D (1996) *Mathematical programs with equilibrium constraints* (Cambridge University Press).
- Mabli J, Ohls J, Dragoset L, Castner L, Santos B, et al. (2013) Measuring the effect of Supplemental Nutrition Assistance Program (SNAP) participation on food security. Technical report, Mathematica Policy Research.
- Miller V, Cudhea F, Singh G, Micha R, Shi P, Zhang J, Onopa J, Karageorgou D, Webb P, Mozaffarian D (2019) Estimated global, regional, and national cardiovascular disease burdens related to fruit and vegetable consumption: An analysis from the global dietary database. *Proc. Nutrition 2019*.
- Mozaffarian D, Liu J, Sy S, Huang Y, Rehm C, Lee Y, Wilde P, Abrahams-Gessel S, Jardim TdSV, Gaziano T, et al. (2018) Cost-effectiveness of financial incentives and disincentives for improving food purchases and health through the US Supplemental Nutrition Assistance Program (SNAP): A microsimulation study. *PLoS medicine* 15(10):e1002661.
- Nielsen SF, et al. (2000) The stochastic em algorithm: estimation and asymptotic results. *Bernoulli* 6(3):457–489.
- Olsho LE, Klerman JA, Wilde PE, Bartlett S (2016) Financial incentives increase fruit and vegetable intake among Supplemental Nutrition Assistance Program participants: a randomized controlled trial of the USDA Healthy Incentives Pilot. *The American Journal of Clinical Nutrition* 104(2):423–435.

- Perakis G, Lobel R (2011) Consumer choice model for forecasting demand and designing incentives for solar technology. Technical Report 4872-11, MIT Sloan Research Paper.
- Raz G, Ovchinnikov A (2015) Coordinating pricing and supply of public interest goods using government rebates and subsidies. *IEEE Transactions on Engineering Management* 62(1):65–79.
- Rose D, Richards R (2004) Food store access and household fruit and vegetable use among participants in the US Food Stamp Program. *Public Health Nutrition* .
- Sierzchula W, Bakker S, Maat K, Van Wee B (2014) The influence of financial incentives and other socio-economic factors on electric vehicle adoption. *Energy Policy* 68:183–194.
- Spronk I, Kullen C, Burdon C, O'Connor H (2014) Relationship between nutrition knowledge and dietary intake. *British Journal of Nutrition* 111(10):1713–1726.
- Steenhuis IH, Waterlander WE, De Mul A (2011) Consumer food choices: the role of price and pricing strategies. *Public Health Nutrition* 14(12):2220–2226.
- Taylor TA, Xiao W (2019) Donor product-subsidies to increase consumption: Implications of consumer awareness and profit-maximizing intermediaries. *Production and Operations Management* .
- The Reinvestment Fund (2019) Online, URL <https://www.reinvestment.com/initiatives/hffi/>.
- Tiehen L, Newman C, Kirilin JA (2017) The food-spending patterns of households participating in the supplemental nutrition assistance program: Findings from USDA's FoodAPS. Economic Research Report, Economic Information Bulletin 176, U.S. Department of Agriculture, Economic Research Service.
- USDA ERS (2017) URL <https://www.ers.usda.gov/data-products/food-access-research-atlas/>.
- Ver Ploeg M, Rahkovsky I (2016) Recent evidence on the effects of food store access on food choice and diet quality. *USDA Economic Research Service* .
- Wardle J, Parmenter K, Waller J (2000) Nutrition knowledge and food intake. *Appetite* 34(269-275).
- Waterlander WE, de Boer MR, Schuit AJ, Seidell JC, Steenhuis IH (2013) Price discounts significantly enhance fruit and vegetable purchases when combined with nutrition education: a randomized controlled supermarket trial. *The American Journal of Clinical Nutrition* 97(4):886–895.
- Weatherspoon D, Oehmke J, Dembélé A, Coleman M, Satimanon T, Weatherspoon L (2013) Price and expenditure elasticities for fresh fruits in an urban food desert. *Urban Studies* 50(1):88–106.
- Wrigley N, Warm D, Margetts B (2003) Deprivation, diet, and food-retail access: findings from the Leeds 'food deserts' study. *Environmental and Planning A* 35:151–181.
- Yu JJ, Tang CS, Shen ZJM (2018) Improving consumer welfare and manufacturer profit via government subsidy programs: Subsidizing consumers or manufacturers? *Manufacturing & Service Operations Management* 20(4):752–766.
- Yu Y, Jaenicke EC (2020) Estimating food waste as household production inefficiency. *American Journal of Agricultural Economics* 102(2):525–547.

Zhang F, Zhang R (2018) Trade-in remanufacturing, customer purchasing behavior, and government policy.
Manufacturing & Service Operations Management 20(4):601–616.

Preprint not peer reviewed

Electronic Companion

EC.1. Proofs of statements

LEMMA 1. *Under Assumption 1, $\tilde{h}(h)$ and $\tilde{u}(u)$ are increasing functions.*

Proof of Lemma 1.

We prove this lemma for \tilde{h} and h , and note that the proof for \tilde{u} and u is analogous. Recall that the function $T(h, u; d, v)$ can be written as $T(\phi h + (1 - \phi)u; d, v)$ where $\phi \in \{0, 1\}$. In other words, T depends on h when $\phi = 1$, and T depends on u when $\phi = 0$. It is easy to see that $\tilde{h}(h)$ is a one-to-one function when $\phi = 0$ since in this case, T has no dependency on h and thus $\tilde{h}(h)$ is linearly increasing in h .

We now consider the case when T depends on h . To simplify notation we will write $T(h, u; d, v)$ as simply $T(h)$ or T . First note that $T(0) = t_1(0) + t_2(d, v) > 0$ since $t_2(d, v) > 0$ and $t_1(0) \geq 0$. We can write $\frac{\partial \tilde{h}}{\partial h} = \frac{\partial(h/T(h))}{\partial h} = \frac{T - h \frac{\partial T}{\partial h}}{T^2}$ which is non-negative whenever $\frac{T(h)}{h} \geq \frac{\partial T}{\partial h}$. Notice that the average slope of $T(h)$ from 0 to $h = l$ is equal to $(T(l) - T(0))/l$. Since $T(h)$ is concave increasing, $\frac{\partial T}{\partial h} \Big|_l$ is always smaller than the average slope of $T(h)$ from 0 to l , so $\frac{\partial T}{\partial h} \Big|_l \leq (T(l) - T(0))/l < T(l)/l$ for all l , where the last inequality follows from the fact that $T(0) > 0$. Thus, $T(h)/h > \frac{\partial T}{\partial h}$ for all h and we see that $\tilde{h}(h)$ is an increasing function. Furthermore, because $\tilde{h}(0) = 0$ and $\tilde{h}'(h)$ is finite, the function $\tilde{h}(h)$ is also one-to-one. ■

PROPOSITION 1. *Under Assumption 1, Problem $\tilde{M}^c(\beta, \delta, \nu)$ is concave.*

Proof of Proposition 1.

Since $f(\tilde{u}, \tilde{h})$ and $v \cdot g(\tilde{h})$ are concave, it is only necessary to prove concavity of the disutility term, which is equivalent to convexity of $1/T(\phi \tilde{h} + (1 - \phi)\tilde{u}; v, d)$ (since the disutility term has a minus sign in front). Assume that $\phi = 1$ (although the same argument holds when $\phi = 0$). In this case we simplify $T(\phi \tilde{h} + (1 - \phi)\tilde{u}; v, d)$ by writing $T(\tilde{h})$. The function $1/T(\tilde{h})$ is convex if and only if $-T(\tilde{h})T''(\tilde{h}) + 2T'(\tilde{h})^2 \geq 0$.

Notice that Assumption 1 is an assumption about the effect of h and u on T , not the variables \tilde{h} and \tilde{u} . Thus, in order to make sense of the condition $-T(\tilde{h})T''(\tilde{h}) + 2T'(\tilde{h})^2 \geq 0$, it will be converted to a condition in terms of h instead of \tilde{h} . By writing $\tilde{h} = h(\tilde{h})/T(h(\tilde{h}); d, v)$ and taking the implicit derivatives $h'(\tilde{h})$ and $h''(\tilde{h})$, we obtain

$$h'(\tilde{h}) = \frac{T^2}{T - h \frac{\partial T}{\partial h}}, \quad h''(\tilde{h}) = \frac{T^3(-2h \frac{\partial^2 T}{\partial h^2} + T(2 \frac{\partial T}{\partial h} + h \frac{\partial^2 T}{\partial h^2}))}{(T - h \frac{\partial T}{\partial h})^3}$$

Using the chain rule,

$$2 \left(\frac{\partial T}{\partial \tilde{h}} \right)^2 - T \frac{\partial^2 T}{\partial \tilde{h}^2} = 2 \left(\frac{\partial T}{\partial h} \frac{\partial h}{\partial \tilde{h}} \right)^2 - T \left(\frac{\partial^2 T}{\partial h^2} \left(\frac{\partial h}{\partial \tilde{h}} \right)^2 + \frac{\partial T}{\partial h} \frac{\partial^2 x}{\partial \tilde{h}^2} \right) = - \frac{T(h)^6 \frac{\partial^2 T}{\partial h^2}}{(T - h \frac{\partial T}{\partial h})^3}$$

where the last equality follows by substituting in $h'(\tilde{h})$ and $h''(\tilde{h})$. This quantity is positive when T is concave increasing in h and $T(0) > 0$, both of which hold under Assumption 1.

■

THEOREM 1. *For fixed investment in the access intervention ($\delta = \delta_0$), the consumer's average HF consumption as a function of the investment in the price and education intervention — $\tilde{h}^*(\nu, \beta, \delta_0)$ — is unimodal on both domains \mathcal{D}_0 and \mathcal{D}_1 . For non-fixed δ , $\tilde{h}^*(\nu, \beta, \delta)$ is unimodal on both \mathcal{D}_0 and \mathcal{D}_1 provided that the following additional conditions are met:*

1. $d''(\delta)w(d(\delta)) - 2d'(\delta)^2w'(d(\delta)) \geq 0$, and (Condition 1)
2. The education intervention has a non-negative budget elasticity. (Condition 2)

Proof of Theorem 1.

When δ is fixed at δ_0 , ν can be written as a function of β by taking advantage of the fact that the government's budget constraint is always tight at optimality. In particular, $\nu^*(\beta) = U - \delta_0 - \beta$. The government's problem can therefore be considered univariate. The goal of the government is to maximize $\tilde{h}^*(\beta, \nu^*(\beta), \delta_0)$. We will first show that this function is unimodal in β on \mathcal{D}_1 . Through implicit differentiation of the consumer's *KKT* conditions (found in Appendix EC.2), we can write $\frac{\partial \tilde{h}^*(\nu(\beta), \beta, \delta_0)}{\partial \beta}$ as

$$\frac{\partial \tilde{h}^*(\beta, \nu^*(\beta), \delta_0)}{\partial \beta} = \frac{p_u^2 v'(\nu^*) g'(\tilde{h}^*) + p_h B'(\beta) f^{(0,2)}(\tilde{h}^*, \tilde{u}^*) + \frac{p_h p_u^2 d_0 v'(\nu^*) r'(v(\nu^*))}{w(d_0)}}{p_u^2 \nu^* g''(\tilde{h}^*) + p_h^2 f^{(0,2)}(\tilde{h}^*, \tilde{u}^*) + p_u^2 f^{(2,0)}(\tilde{h}^*, \tilde{u}^*)} \quad (\text{EC.1a})$$

$$\frac{\partial \tilde{u}^*(\beta, \nu^*(\beta), \delta_0)}{\partial \beta} = \frac{-p_h p_u g'(\tilde{h}^*) v'(\nu^*) + p_u B'(\beta) (v(\nu^*) g''(\tilde{h}^*) + f^{(2,0)}(\tilde{h}^*, \tilde{u}^*)) - \frac{p_u p_h \rho_h d_0 v'(\nu^*) r'(v(\beta))}{w(d_0)}}{p_u^2 \nu^* g''(\tilde{h}^*) + p_h^2 f^{(0,2)}(\tilde{h}^*, \tilde{u}^*) + p_u^2 f^{(2,0)}(\tilde{h}^*, \tilde{u}^*)} \quad (\text{EC.1b})$$

Where \tilde{h}^* and \tilde{u}^* are shorthand for $\tilde{h}^*(\beta, \nu^*(\beta), \delta_0)$ and $\tilde{u}^*(\beta, \nu^*(\beta), \delta_0)$, and ν^* is shorthand for $\nu^*(\beta) := U - \delta_0 - \beta$. Notice that $\frac{\partial \tilde{u}^*(\beta, \nu^*(\beta), \delta_0)}{\partial \beta}$ is always positive (both the numerator and denominator of Equation EC.1b are negative), meaning that as the government allocates more money towards the consumer's food budget and less towards education, RF consumption consistently increases. In order to prove that $\tilde{h}^*(\beta, \nu^*(\beta), \delta_0)$ is unimodal in β , we will show if $\left. \frac{\partial \tilde{h}^*(\beta, \nu^*(\beta), \delta_0)}{\partial \beta} \right|_{\beta_1} < 0$ for some β_1 , then $\left. \frac{\partial \tilde{h}^*(\beta, \nu^*(\beta), \delta_0)}{\partial \beta} \right|_{\beta} < 0$ for all $\beta > \beta_1$. This implies that once the function $\tilde{h}^*(\beta, \nu^*(\beta), \delta_0)$ begins to decrease, it will remain decreasing, implying that there can be at most one local maximum.

Notice that the denominator of Equation EC.1a is always negative. Therefore, the sign of the derivative is determined by the numerator of Equation EC.1a. Suppose that $\left. \frac{\partial \tilde{h}^*(\beta, \nu^*(\beta), \delta_0)}{\partial \beta} \right|_{\beta_1} < 0$ for some β_1 . Consider the derivative of the numerator with respect to β , evaluated at β_1 . This is given by

$$p_u^2 v'(\nu) \frac{\partial \tilde{h}^*}{\partial \beta} \Big|_{\beta_1} g''(\tilde{h}^*) - p_u^2 g'(\tilde{h}^*) v''(\nu) + p_h (B'(\beta) \frac{\partial \tilde{u}^*}{\partial \beta} \Big|_{\beta_1} f^{(0,3)}(\tilde{h}^*, \tilde{u}^*) + B''(\beta) f^{(0,2)}(\tilde{h}^*, \tilde{u}^*)) - \frac{p_u^2 \rho_h d(r'(v) v''(\nu) + v'(\nu)^2 r''(v))}{w(d)}$$

We will show that for all $\beta > \beta_1$, the numerator is positive. Evaluated at β_1 , this quantity is positive under Assumptions 2 and 3, in addition to the assumption that $f(\tilde{h}, \tilde{u})$ has a positive third derivative. Note that this holds for any β_1 such that $\left. \frac{\partial \tilde{h}^*(\beta, \nu^*(\beta), \delta_0)}{\partial \beta} \right|_{\beta_1} < 0$. Therefore, $\frac{\partial \tilde{h}^*(\beta, \nu^*(\beta), \delta_0)}{\partial \beta}$ is decreasing on the domain $\beta \in (\beta_c, U]$ where β_c is such that $\frac{\partial \tilde{h}^*(\beta, \nu^*(\beta), \delta_0)}{\partial \beta} = 0$. Therefore, $\tilde{h}^*(\beta, \nu^*(\beta), \delta_0)$ can have at most one local maximum, and can have no local minimum.

On the domain \mathcal{D}_0 , the expression for $\frac{\partial \tilde{h}^*(\beta, \nu^*(\beta), \delta_0)}{\partial \beta}$ is still given by Equation EC.1a except that the term $\frac{\rho_h p_u^2 d_0 v'(\nu(\beta)) r'(v(\beta))}{w(d_0)}$ disappears from the numerator since r does not depend on v when $\phi = 0$. Thus this case is even simpler and we obtain the same result as before.

We now prove the second part of the theorem, where δ is not fixed. Consider the function $\tilde{h}^*(\beta^*(\delta), \nu^*(\delta), \delta)$ where $\nu^*(\delta)$ and $\beta^*(\delta)$ are the optimal values of ν and β for a given δ , which we now know are unique. We will show that the function $\tilde{h}^*(\beta^*(\delta), \nu^*(\delta), \delta)$ is unimodal in δ , implying that $\tilde{h}^*(\nu, \beta, \delta)$ is unimodal. We can write

$$\left. \frac{\partial \tilde{h}^*(\beta(\delta), \nu(\delta), \delta)}{\partial \delta} \right|_{(\beta^*(\delta), \nu^*(\delta), \delta)} = \left(\frac{\partial \tilde{h}^*(\beta, \nu, \delta)}{\partial \delta} + \nu^*(\delta) \frac{\partial \tilde{h}^*(\beta, \nu, \delta)}{\partial \nu} + \beta^*(\delta) \frac{\partial \tilde{h}^*(\beta, \nu, \delta)}{\partial \beta} \right) \Big|_{(\beta^*(\delta), \nu^*(\delta), \delta)}$$

By noting that $U = \delta + \nu + \beta$, we see that $\nu^*(\delta) + \beta^*(\delta) = -1$. Furthermore, at optimal ν and β , we must have

$$\left. \frac{\partial \tilde{h}^*(\beta, \nu, \delta)}{\partial \nu} \right|_{(\beta^*(\delta), \nu^*(\delta), \delta)} = \left. \frac{\partial \tilde{h}^*(\beta, \nu, \delta)}{\partial \beta} \right|_{(\beta^*(\delta), \nu^*(\delta), \delta)}$$

for every δ . Combining these two observations, we see that

$$\left. \frac{\partial \tilde{h}^*(\beta(\delta), \nu(\delta), \delta)}{\partial \delta} \right|_{(\beta^*(\delta), \nu^*(\delta), \delta)} = \left(\frac{\partial \tilde{h}^*(\beta, \nu, \delta)}{\partial \delta} - \frac{\partial \tilde{h}^*(\beta, \nu, \delta)}{\partial \nu} \right) \Big|_{(\beta^*(\delta), \nu^*(\delta), \delta)}$$

On the domain \mathcal{D}_1 this can be written as

$$\frac{\partial \tilde{h}^*(\beta^*(\delta), \nu^*(\delta), \delta)}{\partial \delta} = \frac{-\rho_h p_u^2 r(v(\nu^*)) d'(\delta) (w(\delta) - d(\delta) w'(\delta)) + p_u^2 v'(\nu^*) w(\delta) (w(\delta) g'(\tilde{h}^*) + \rho_h d r'(v(\nu^*)))}{w(\delta)^2 (p_u^2 f^{(2,0)}(\tilde{h}^*) + p_h^2 f^{(0,2)}(\tilde{u}^*) + p_u^2 v(\nu^*) g''(\tilde{h}^*))} \quad (\text{EC.2})$$

where $w(\delta)$, ν^* , \tilde{h}^* , and \tilde{u}^* are shorthand for $w(d(\delta))$, $\nu^*(\delta)$, $\tilde{h}^*(\beta^*(\delta), \nu^*(\delta), \delta)$ and $\tilde{u}^*(\beta^*(\delta), \nu^*(\delta), \delta)$, respectively. On the domain \mathcal{D}_0 we obtain the expression

$$\frac{\partial \tilde{h}^*(\beta^*(\delta), \nu^*(\delta), \delta)}{\partial \delta} = \frac{\rho_u p_h p_u r d'(\delta) (w(\delta) - d(\delta) w'(\delta)) + w(\delta)^2 p_u^2 v'(\nu^*) g'(\tilde{h}^*)}{w(\delta)^2 (p_u^2 f^{(2,0)}(\tilde{h}^*) + p_h^2 f^{(0,2)}(\tilde{u}^*) + p_u^2 v(\nu^*) g''(\tilde{h}^*))} \quad (\text{EC.3})$$

In what follows, we use the same proof technique as above and show that on each domain, once the numerators of Equations EC.2 and EC.3 are negative, they must remain negative as δ increases. This is true so long as $\nu^{*\prime}(\delta) \leq 0$ and $d''(\delta)w(d(\delta)) - 2d'(\delta)^2 w'(d(\delta)) \geq 0$, which are the conditions of the theorem (the constraint that $\nu^{*\prime}(\delta) \leq 0$ is what is referred to as a *non-negative budget elasticity*, and details on this condition are provided after the proof).

First consider the derivative over \mathcal{D}_0 . Let $\mathcal{L}(\delta)$ be defined as

$$\mathcal{L}(\delta) = \rho_u p_h p_u r \frac{d'(\delta)}{w(d(\delta))^2} (w(d(\delta)) - d(\delta) w'(d(\delta))) + p_u^2 v'(\nu^*(\delta)) g'(\tilde{h}^*(\delta))$$

Notice that $\mathcal{L}(\delta) \geq 0 \implies \frac{\partial \tilde{h}^*(\beta^*(\delta), \nu^*(\delta), \delta)}{\partial \delta} \leq 0$. Our goal is to show that if there exists a δ_c such that $\left. \frac{\partial \tilde{h}^*(\beta^*(\delta), \nu^*(\delta), \delta)}{\partial \delta} \right|_{\delta_c}$ is a local maximum (implying that $\mathcal{L}(\delta_c) = 0$ and $\mathcal{L}(\delta_c + \epsilon) > 0$ for ϵ small enough), that $\mathcal{L}(\delta_c + \Delta) > 0$ for all $\Delta > 0$.

Consider the two terms in $\mathcal{L}(\delta)$ separately. The second term, $p_u^2 v'(\nu^*(\delta)) g'(\tilde{h}^*(\delta))$ is always increasing with respect to δ for δ such that $\frac{\partial \tilde{h}^*(\beta^*(\delta), \nu^*(\delta), \delta)}{\partial \delta} \leq 0$ and $\nu^{*\prime}(\delta) \leq 0$. We will also show that the first term, $p_h p_u \frac{d'(\delta)}{w(d(\delta))^2} (w(d(\delta)) - d(\delta) w'(d(\delta))) := L(\delta)$, is increasing under the assumptions of the theorem. We can compute $\frac{\partial L}{\partial \delta}$ as

$$\frac{\partial L}{\partial \delta} = \frac{(w(d(\delta)) - d(\delta) w'(d(\delta))) (d''(\delta) w(d(\delta)) - 2d'(\delta)^2 w'(d(\delta))) - w(d(\delta)) d'(\delta)^2 d(\delta) w''(d(\delta))}{w(d(\delta))^3}$$

Note that $w(d(\delta)) - d(\delta) w'(d(\delta))$ is assumed to be positive since \tilde{h}^* is only increasing in δ if this is true. Therefore, $\frac{\partial L}{\partial \delta}$ is necessarily positive when $d''(\delta) w(d(\delta)) - 2d'(\delta)^2 w'(d(\delta)) \geq 0$, which is the condition of the theorem.

Now consider the domain \mathcal{D}_1 . Again we will treat each term in the numerator separately (after first dividing each term by $w(d(\delta))^2$ and $r(v(\nu^*(\delta)))$). It is straightforward to see that the term $\frac{p_u^2 v'(\nu^*(\delta)) g'(\tilde{h}^*(\delta))}{r(v(\nu^*(\delta)))}$ is increasing in δ when $\frac{\partial \tilde{h}^*}{\partial \delta} \leq 0$ and $\nu^{*\prime}(\delta) \leq 0$. The term $\frac{p_u^2 v'(\nu^*(\delta)) \rho_h d(\delta) r'(v(\nu^*(\delta)))}{w(d(\delta)) \cdot r(v(\nu^*(\delta)))}$ is also increasing when $\nu^{*\prime}(\delta) \leq 0$. Finally, by an analogous argument as to the one on domain \mathcal{D}_0 , it is clear that the term $-\frac{d'(\delta) (w(d(\delta)) - d(\delta) w'(d(\delta)))}{w(d(\delta))^2}$ is increasing when \tilde{h}^* is increasing in δ and $d''(\delta) w(d(\delta)) - 2d'(\delta)^2 w'(d(\delta)) \geq 0$. ■

The condition $d''(\delta)w(d(\delta)) - 2d'(\delta)^2w'(d(\delta)) \geq 0$ is a statement about the convexity of the function $d(\delta)$. By Corollary 2, in order for \tilde{h}^* to be increasing on \mathcal{D}_0 we require that $w(d(\delta))$ is concave on \mathcal{D}_0 , which implies that $w(d(\delta)) - d(\delta)w'(d(\delta)) \geq 0$, and in order for \tilde{h}^* to be increasing on \mathcal{D}_1 we require that $w(d(\delta))$ is convex on \mathcal{D}_1 , implying that $w(d(\delta)) - d(\delta)w'(d(\delta)) \leq 0$. Let $\epsilon_1(d)$ be such that $w(d(\delta)) - \epsilon_2(\delta)w'(d(\delta)) = 0$ on domain \mathcal{D}_1 . The only way that all three of these conditions can hold is if $\frac{d''(\delta)}{2d'(\delta)^2} \leq \epsilon_1(\delta) < d(\delta)$. Which is equivalent to upper bounding the second derivative of $1/d(\delta)$ by some $\epsilon_2(\delta) < 0$.

The condition that $\nu^{*'}(\delta) \leq 0$ intuitively means that as more of the government's budget is allocated to the access intervention (and thus the total available budget for the education and price intervention *decreases*), the allocation to the education intervention does not *increase*. This is an intuitive condition and when it is not met, the government's optimization problem need not be unimodal. This condition can be written as a (rather complicated) condition of the modeling parameters and functions by noting that $\nu^*(\delta)$ solves

$$p_u^2 v'(\nu^*(\delta)) g'(\tilde{h}^*) + p_h B'(\beta^*(\delta)) f^{(0,2)}(\tilde{h}^*, \tilde{u}^*) + \frac{\rho_h p_u^2 d(\delta) v'(\nu^*(\delta)) r'(v(\nu^*(\delta)))}{w(d(\delta))} = 0,$$

which is the numerator of Equation EC.1a, and that $\nu^*(\delta) = U - \delta - \beta^*(\delta)$. Taking the implicit derivative with respect to δ gives $\nu^*(\delta)$ as a function of the modeling constraints.

PROPOSITION 2. *Under Assumption 1, the following characterizes the impact of the untargeted price and value intervention (β and ν) on average HF and RF consumption (\tilde{h}^* and \tilde{u}^*) when both are positive ($\tilde{h}^*, \tilde{u}^* > 0$):*

1. *HF and RF consumption are increasing in the untargeted price intervention, β . However, the proportion of healthy to RF consumed, (\tilde{h}^*/\tilde{u}^*), is not necessary increasing in β .*
2. *HF consumption is increasing in the value intervention, ν , and RF consumption is decreasing in ν .*

In addition to Proposition 2 as stated in the main text, we include Part 3 which characterize the impact of access on HF spending.

3. *When $\phi(v) = 1$ (namely, when cycle time depends on the inventory of HF items), HF consumption, \tilde{h}^* , is increasing in the access intervention, δ , and RF consumption, \tilde{u}^* is decreasing in δ if and only if $G(h) \geq 0$, where*

$$G(x) := -\frac{\partial T}{\partial x} \left(T - x \frac{\partial T}{\partial x} \right) \left(T - x \frac{\partial T}{\partial x} - d(\delta) \frac{\partial T}{\partial d(\delta)} \right) - d(\delta) T x \frac{\partial^2 T}{\partial x^2}$$

When $\phi(v) = 0$ (namely, when cycle time depends on the inventory of RF items), HF consumption, \tilde{h}^* , is increasing in the access intervention, δ , and RF consumption, \tilde{u}^* is decreasing in δ if and only if $G(u) \leq 0$.

Before proving Proposition 2, we prove a lemma that will be useful.

LEMMA EC.1. Under Assumption 1, when $\phi = 1$, $\frac{\partial^2 1/T}{\partial h^2} \geq 0$, $\frac{\partial^2 1/T}{\partial \tilde{h} \partial v(\nu)} \leq 0$, and $\frac{\partial^2 1/T}{\partial \tilde{h} \partial d(\delta)} \geq 0$. When $\phi = 0$, $\frac{\partial^2 1/T}{\partial \tilde{u}^2} \geq 0$ and $\frac{\partial^2 1/T}{\partial \tilde{u} \partial d(\delta)} \geq 0$.

[Proof of Lemma EC.1] First we focus on the second derivative of $1/T$ with respect to \tilde{h} and \tilde{u} . Notice that $\frac{\partial^2 1/T}{\partial h^2} = \frac{2 \frac{\partial T}{\partial h}^2 - T \frac{\partial^2 T}{\partial h^2}}{T^3}$. Furthermore, $\frac{\partial T}{\partial h} = \frac{\partial h}{\partial h} \frac{\partial T}{\partial h}$ and $\frac{\partial^2 T}{\partial h^2} = \frac{\partial^2 h}{\partial h^2} \frac{\partial T}{\partial h} + \frac{\partial h}{\partial h}^2 \frac{\partial^2 T}{\partial h^2}$. Finally, $\frac{\partial h}{\partial \tilde{h}}$ and $\frac{\partial^2 h}{\partial \tilde{h}^2}$ can be computed by implicitly differentiating $\tilde{h} - h(\tilde{h})/T(h(\tilde{h})) = 0$ as in the proof of Proposition 1. Combining these expressions and simplifying $\frac{\partial^2 1/T}{\partial h^2}$ yields

$$\frac{\partial^2 1/T}{\partial h^2} = \frac{-T^3 \frac{\partial^2 T}{\partial h^2}}{(T - h \frac{\partial T}{\partial h})^3}$$

When T is concave in h , this quantity is positive. An analogous argument proves this result when $\phi = 0$ for $\frac{\partial^2 1/T}{\partial \tilde{u}^2}$.

Now consider the mixed derivatives of $1/T$. By following a similar method—rewriting $T(\tilde{h}; d(\delta), v(\nu))$ as a function of h instead of \tilde{h} —and utilizing the separability of $T(\cdot)$ with respect to v and h , we have

$$\frac{\partial^2 1/T}{\partial \tilde{h} \partial v(\nu)} = \frac{2 \frac{\partial T}{\partial d(\delta)} \frac{\partial T}{\partial h} - T \frac{\partial^2 T}{\partial h \partial v(\nu)}}{T^3} = \frac{\frac{\partial T}{\partial v} (\frac{\partial T}{\partial h} (T - h \frac{\partial T}{\partial h}) - T h \frac{\partial^2 T}{\partial h^2})}{(T - h \frac{\partial T}{\partial h})^3} \quad (\text{EC.4})$$

The RHS of EC.4 has the same sign as $\frac{\partial T}{\partial v}$, which is negative. If we had taken the derivative with respect to $d(\delta)$ instead, the RHS would be positive since $\frac{\partial T}{\partial d(\delta)}$ is negative. The same method can be applied when $\phi = 0$ and T depends on \tilde{u} . Therefore, $\frac{\partial^2 1/T}{\partial \tilde{h} \partial v(\nu)} \leq 0$, $\frac{\partial^2 1/T}{\partial \tilde{h} \partial d(\delta)} \geq 0$, and $\frac{\partial^2 1/T}{\partial \tilde{u} \partial d(\delta)} \geq 0$. ■

Proof of Proposition 2.

The proof of each of these statements relies on considering the derivatives $\frac{\partial \tilde{h}^*}{\partial x}$ and $\frac{\partial \tilde{u}^*}{\partial x}$, where $x \in \{\beta, \nu, \delta\}$. These derivatives are obtained by implicitly differentiating KKT_{cons} , and can be found in Appendix EC.2.

Part 1. The fact that \tilde{h}^* and \tilde{u}^* are increasing in β follows by considering the derivatives $\frac{\partial \tilde{h}^*}{\partial \beta}$ and $\frac{\partial \tilde{u}^*}{\partial \beta}$ in Appendix EC.2 and noting that, by Assumption 3 and Lemma EC.1, they are both always positive.

The fact that the FV ratio, \tilde{h}^*/\tilde{u}^* , is not necessarily increasing in β is evident by considering $\frac{\partial \tilde{h}^*/\tilde{u}^*}{\partial \beta}$. In the case when $\phi = 1$, which is positive if and only if

$$\frac{\tilde{u}^* p_h f^{(0,2)}(\tilde{h}^*, \tilde{u}^*)}{\tilde{h}^* p_u (f^{(2,0)}(\tilde{h}^*, \tilde{u}^*) + \nu g''(\tilde{h}^*) - d(\delta) \frac{\partial^2 1/T}{\partial \tilde{h}^2})} \geq 1$$

In general his condition depends on the current values of HF and RF consumption, as well as their relative prices. A special case where this condition does *not* hold is when the consumer's preferences are nearly linear in \tilde{u} , meaning that there are no diminishing returns for RF consumption. A similar analysis reveals that the ratio of FVs can also be increasing or decreasing in the case that $\phi = 0$.

Part 2. This is proven by considering the equations in Appendix EC.2. When $\phi = 1$, $\frac{\partial \tilde{h}^*}{\partial \nu}$ is positive when $\phi = 1$ if $\frac{\partial^2 1/T}{\partial \tilde{h} \partial \nu} \leq 0$. By Lemma EC.1, this condition holds.

The derivative $\frac{\partial \tilde{h}^*}{\partial \nu}$ when $\phi = 0$ is more straightforward since there is no dependency of T on ν when $\phi = 0$. An argument similar to the proof of Part 1 shows that in this case, $\frac{\partial \tilde{h}^*}{\partial \nu} \geq 0$. Finally, since the food budget remains unchanged as ν changes, \tilde{u}^* must decrease if \tilde{h}^* increases in order for the consumer's budget constraint to remain tight.

Part 3. By again considering the equations in Appendix EC.2, when $\phi = 1$, \tilde{h}^* is increasing in δ and \tilde{u}^* is decreasing in δ if and only if $\frac{\partial 1/T}{\partial \tilde{h}} + d(\delta) \frac{\partial^2 1/T}{\partial \tilde{h} \partial d(\delta)} \geq 0$. Similarly, when $\phi = 0$, \tilde{h}^* is increasing in δ and \tilde{u}^* is decreasing in δ if and only if $\frac{\partial 1/T}{\partial \tilde{u}} + d(\delta) \frac{\partial^2 1/T}{\partial \tilde{u} \partial d(\delta)} \leq 0$. We will focus on the case when $\phi = 1$. First, note that

$$\frac{\partial 1/T}{\partial \tilde{h}} = -\frac{\frac{\partial T}{\partial \tilde{h}}}{T^2} = \frac{-\frac{\partial T}{\partial \tilde{h}}}{T - h \frac{\partial T}{\partial \tilde{h}}} \leq 0$$

Where the inequality comes from the fact that $T - h \frac{\partial T}{\partial \tilde{h}} \geq 0$ by the concavity of T with respect to h . By following a similar method to the proof of Lemma EC.1—rewriting $T(\tilde{h}; d(\delta), v(\nu))$ as a function of h instead of \tilde{h} —and utilizing the separability of $T(\cdot)$ with respect to d and h , By Lemma EC.1, $\frac{\partial^2 1/T}{\partial \tilde{h} \partial d(\delta)} \geq 0$. Simplifying $\frac{\partial 1/T}{\partial \tilde{h}} + d(\delta) \frac{\partial^2 1/T}{\partial \tilde{h} \partial d(\delta)}$ results in the condition given in the Theorem. The same logic can be applied when $\phi = 0$. ■

COROLLARY 1. *Under Assumption 2, the following characterizes the impact of access on average HF consumption when $\tilde{h}^*, \tilde{u}^* > 0$: When cycle time depends on HF purchasing ($\phi(v) = 1$), average HF consumption, \tilde{h}^* , is increasing in the access intervention, δ , and average RF consumption, \tilde{u}^* , is decreasing in δ if and only if $w(d) - dw'(d) < 0$. When cycle time depends on RF purchasing ($\phi(v) = 0$), \tilde{h}^* is increasing in δ and \tilde{u}^* is decreasing in δ if and only if $w(d) - dw'(d) > 0$.*

Proof of Corollary 1.

This is a direct consequence of Proposition 2 in the case when T is linear in h or u . The result can be seen by examining the derivative $\frac{\partial \tilde{h}^*}{\partial \delta}$ under Assumption 2. ■

COROLLARY 2. *Under Assumption 2, the government's optimal solution, $(\beta^*, \nu^*, \delta^*)$, will only have $\delta^* > 0$ when the conditions of Corollary 1 are met.*

Proof of Corollary 2.

This is a direct consequence of Corollary 1 ■.

PROPOSITION 3. *For a consumer with fixed disutility, $d(\delta)$, the government's optimal allocation between the price and education levers satisfies:*

1. *Allocation to the education lever is decreasing (and thus allocation to the price lever is increasing) as the consumer's initial value of nutrition increases.*
2. *Allocation to the price lever is decreasing (and thus allocation to the education lever is increasing) as the consumer's utility for RF has less diminishing returns (i.e., as $f(\tilde{h}, \tilde{u})$ becomes more linear in \tilde{u})*
3. *The government's allocation could change in either direction as the consumer's budget changes.*

Proof of Proposition 3

The proof of Proposition 3 follows by considering Equation EC.1a—the derivative of \tilde{h}^* with respect to β and $\nu(\beta)$ —and whether it increases or decreases as certain parameters change. As one example, consider the marginal concavity of $f(\tilde{h}, \tilde{u})$ with respect to \tilde{u} . Suppose that the numerator of Equation EC.1a is equal to zero at some β_c . If $f^{(0,2)}(\tilde{h}, \tilde{u})$ increases (corresponding to $f_u(\tilde{u})$ becoming less concave), the numerator of Equation EC.1a increases at every point β . This means that at β_c , the numerator of Equation EC.1a is positive, hence $\frac{\partial \tilde{h}^*(\beta, \nu(\beta), \delta_0)}{\partial \beta}$ is negative (since the denominator is negative). Thus, since we know $\tilde{h}^*(\beta, \nu(\beta), \delta_0)$ is unimodal in β , it must be the case that Equation EC.1a is equal to zero at a β less than β_c , implying that a smaller amount of funding to the price intervention than before is now optimal. In conclusion, decreasing the concavity of $f_u(\tilde{u})$ results in an optimal allocation with less funding to the price intervention. A similar logic can be followed to prove the other statements of Proposition 3. ■

PROPOSITION 4. *Targeted subsidies strictly dominate untargeted subsidies, in the sense that $\tilde{h}_{tar}^*(k, \nu_0, \delta_0) \geq \tilde{h}_{unt}^*(k, \nu_0, \delta_0)$, for all k and fixed ν_0, δ_0 .*

Proof of Proposition 4

Consider the consumer's KKT conditions in both the targeted and untargeted subsidy case. In the untargeted case we have

$$\begin{cases} \left. \frac{\partial f}{\partial \tilde{h}} \right|_{\tilde{h}=\tilde{h}^*} + c_\nu^{-1}(B-k) \left. \frac{\partial g}{\partial \tilde{h}} \right|_{\tilde{h}=\tilde{h}^*} + (d-\delta) \frac{\phi \frac{\partial T}{\partial \tilde{h}} - (\phi-1) \frac{\partial T}{\partial \tilde{u}} \frac{p_h}{p_u}}{T^2} = \left. \frac{\partial f}{\partial \tilde{u}} \right|_{\tilde{u}=\tilde{u}^*} \frac{p_h}{p_u} \\ p_h \tilde{h}^* + p_u \tilde{u}^* - (B+k) = 0 \end{cases} \quad (\text{EC.5})$$

And in the targeted case, by substituting $\beta_{pu} = k/\tilde{h}^*$, we obtain

$$\begin{cases} \left. \frac{\partial f}{\partial \tilde{h}} \right|_{\tilde{h}=\tilde{h}^*} + c_\nu^{-1}(B-k) \left. \frac{\partial g}{\partial \tilde{h}} \right|_{\tilde{h}=\tilde{h}^*} + (d-\delta) \frac{\phi \frac{\partial T}{\partial \tilde{h}} - (\phi-1) \frac{\partial T}{\partial \tilde{u}} \frac{(p_h-k/\tilde{h}^*)}{p_u}}{T^2} = \left. \frac{\partial f}{\partial \tilde{u}} \right|_{\tilde{u}=\tilde{u}^*} \frac{p_h-k/\tilde{h}^*}{p_u} \\ p_h \tilde{h}^* + p_u \tilde{u}^* - (B+k) = 0 \end{cases} \quad (\text{EC.6})$$

Notice that the budget constraints in both cases are identical. In the targeted case, the budget constraint is written as $(p_h - k/\tilde{h}^*)\tilde{h}^* + p_u \tilde{u}^* = B$ which simplifies to $p_h \tilde{h}^* + p_u \tilde{u}^* - (B+k) = 0$.

Let $LHS_{unt}(\tilde{h}, \tilde{u})$ and $RHS_{unt}(\tilde{h}, \tilde{u})$ be the quantities on the LHS and RHS of the first equation in System EC.5, respectively, and similarly for $LHS_{pu}(\tilde{h}, \tilde{u})$ and $RHS_{tar}(\tilde{h}, \tilde{u})$ for System EC.6. Suppose that (\tilde{h}', \tilde{u}') satisfies System EC.5 for a fixed k , so $LHS_{unt}(\tilde{h}', \tilde{u}') = RHS_{unt}(\tilde{h}', \tilde{u}')$. For the same k , (\tilde{h}', \tilde{u}') will not satisfy Equation EC.6, because $RHS_{tar}(\tilde{h}', \tilde{u}') < RHS_{unt}(\tilde{h}', \tilde{u}')$, and thus $RHS_{tar}(\tilde{h}', \tilde{u}') < LHS_{tar}(\tilde{h}', \tilde{u}')$ since $LHS_{tar}(\tilde{h}', \tilde{u}') = LHS_{unt}(\tilde{h}', \tilde{u}') = RHS_{unt}(\tilde{h}', \tilde{u}')$.

Now consider some point $(\tilde{h}'', \tilde{u}'')$ such that $\tilde{h}'' < \tilde{h}'$, $\tilde{u}'' > \tilde{u}'$. We know that $LHS_{tar}(\tilde{h}, \tilde{u})$ is increasing in \tilde{h} , so $LHS_{tar}(\tilde{h}'', \tilde{u}'') > LHS_{tar}(\tilde{h}', \tilde{u}')$. We also know that $RHS_{tar}(\tilde{h}, \tilde{u})$ is decreasing as \tilde{u} increases and as \tilde{h} increase, so $RHS_{tar}(\tilde{h}'', \tilde{u}'') < RHS_{tar}(\tilde{h}', \tilde{u}')$. Therefore, $(\tilde{h}'', \tilde{u}'')$ cannot satisfy System EC.6 since $RHS_{tar}(\tilde{h}'', \tilde{u}'') < RHS_{tar}(\tilde{h}', \tilde{u}') < LHS_{tar}(\tilde{h}', \tilde{u}') < LHS_{tar}(\tilde{h}'', \tilde{u}'')$. Therefore, there must be a solution to System EC.6 $(\tilde{h}''', \tilde{u}''')$ such that $\tilde{h}''' > \tilde{h}'$, $\tilde{u}''' < \tilde{u}'$. (Note that \tilde{h} and \tilde{u} are required to move in opposite directions to ensure that the budget constraint is still satisfied). Therefore, under the same budget allocation to targeted and untargeted price subsidies, the consumer's optimal HF spending is always higher in the targeted subsidy case. ■

EC.2. Derivatives

The derivatives shown below are derived by differentiating the KKT equations of Problem $\tilde{M}^c(\beta, \delta, \nu)$, which are given by

$$\begin{cases} \left. \begin{aligned} \frac{\partial f}{\partial h} \Big|_{\tilde{h}=\tilde{h}^*} + v(\nu) \frac{\partial g}{\partial h} \Big|_{\tilde{h}=\tilde{h}^*} + d(\delta) \frac{\phi \frac{\partial T}{\partial h} - (\phi-1) \frac{\partial T}{\partial \tilde{u}} \frac{p_h}{p_u}}{T^2} = \frac{\partial f}{\partial \tilde{u}} \Big|_{\tilde{u}=\tilde{u}^*} \frac{p_h}{p_u} \\ p_h \tilde{h}^* + p_u \tilde{u}^* - B(\beta) = 0 \end{aligned} \right\} & (KKT_{cons}) \end{cases}$$

To find the derivative of \tilde{h}^* and \tilde{u}^* with respect to x , for example, we take the implicit derivatives of the KKT equations from Problem $\tilde{M}^c(\beta, \delta, \nu)$ with respect to x and then solve for $\frac{\partial \tilde{h}^*}{\partial x}$ and $\frac{\partial \tilde{u}^*}{\partial x}$. In the equations below, all derivatives are evaluated at $(\tilde{h}^*, \tilde{u}^*)$ (for example, the term $\frac{\partial^2 1/T}{\partial \tilde{u} \partial \nu}$ is shorthand for $\frac{\partial^2 1/T}{\partial \tilde{u} \partial \nu} \Big|_{(\tilde{h}^*, \tilde{u}^*)}$).

$$\frac{\partial \tilde{h}^*}{\partial \nu} = \frac{-p_u v'(\nu)(p_u g'(\tilde{h}) - dp_u \phi \frac{\partial^2 1/T}{\partial \tilde{h} \partial \nu})}{p_h^2 f^{(0,2)} + p_u^2 f^{(2,0)} + p_u^2 v(\nu) g''(\tilde{h}) - \phi d(\delta) p_u^2 \frac{\partial^2 1/T}{\partial \tilde{h}^2} - (1-\phi) d(\delta) p_h^2 \frac{\partial^2 1/T}{\partial \tilde{u}^2}} \quad (\text{EC.7a})$$

$$\frac{\partial \tilde{u}^*}{\partial \nu} = \frac{p_h v'(\nu)(p_u g'(\tilde{h}) - dp_u \phi \frac{\partial^2 1/T}{\partial \tilde{h} \partial \nu})}{p_h^2 f^{(0,2)} + p_u^2 f^{(2,0)} + p_u^2 v(\nu) g''(\tilde{h}) - \phi d(\delta) p_u^2 \frac{\partial^2 1/T}{\partial \tilde{h}^2} - (1-\phi) d(\delta) p_h^2 \frac{\partial^2 1/T}{\partial \tilde{u}^2}} \quad (\text{EC.7b})$$

$$\frac{\partial \tilde{h}^*}{\partial \beta} = \frac{B'(\beta) \left(p_h f^{(0,2)} + d(\delta) (-p_h (1-\phi) \frac{\partial^2 1/T}{\partial \tilde{u}^2}) \right)}{p_h^2 f^{(0,2)} + p_u^2 f^{(2,0)} + p_u^2 v(\nu) g''(\tilde{h}) - \phi d(\delta) p_u^2 \frac{\partial^2 1/T}{\partial \tilde{h}^2} - (1-\phi) d(\delta) p_h^2 \frac{\partial^2 1/T}{\partial \tilde{u}^2}} \quad (\text{EC.7c})$$

$$\frac{\partial \tilde{u}^*}{\partial \beta} = \frac{B'(\beta) \left(p_u f^{(2,0)} + p_u \nu g''(\tilde{h}) - \phi d(\delta) p_u \frac{\partial^2 1/T}{\partial \tilde{h}^2} \right)}{p_h^2 f^{(0,2)} + p_u^2 f^{(2,0)} + p_u^2 v(\nu) g''(\tilde{h}) - \phi d(\delta) p_u^2 \frac{\partial^2 1/T}{\partial \tilde{h}^2} - (1-\phi) d(\delta) p_h^2 \frac{\partial^2 1/T}{\partial \tilde{u}^2}} \quad (\text{EC.7d})$$

$$\frac{\partial \tilde{h}^*}{\partial \delta} = \frac{p_u d'(\delta) \left(-p_h (1-\phi) \frac{\partial 1/T}{\partial \tilde{u}} + p_u \phi \frac{\partial 1/T}{\partial \tilde{h}} + d(\delta) (-p_h (1-\phi) \frac{\partial^2 1/T}{\partial \tilde{u} \partial \delta} + p_u \phi \frac{\partial^2 1/T}{\partial \tilde{h} \partial \delta}) \right)}{p_h^2 f^{(0,2)} + p_u^2 f^{(2,0)} + p_u^2 v(\nu) g''(\tilde{h}) - \phi d(\delta) p_u^2 \frac{\partial^2 1/T}{\partial \tilde{h}^2} - (1-\phi) d(\delta) p_h^2 \frac{\partial^2 1/T}{\partial \tilde{u}^2}} \quad (\text{EC.7e})$$

$$\frac{\partial \tilde{u}^*}{\partial \delta} = \frac{p_h d'(\delta) \left((1-\phi) p_h \frac{\partial 1/T}{\partial \tilde{u}} - \phi p_u \frac{\partial 1/T}{\partial \tilde{h}} + d(\delta) (p_h (1-\phi) \frac{\partial^2 1/T}{\partial \tilde{u} \partial \delta} - p_u \phi \frac{\partial^2 1/T}{\partial \tilde{h} \partial \delta}) \right)}{p_h^2 f^{(0,2)} + p_u^2 f^{(2,0)} + p_u^2 v(\nu) g''(\tilde{h}) - \phi d(\delta) p_u^2 \frac{\partial^2 1/T}{\partial \tilde{h}^2} - (1-\phi) d(\delta) p_h^2 \frac{\partial^2 1/T}{\partial \tilde{u}^2}} \quad (\text{EC.7f})$$

EC.3. Model extension: N food groups

This section shows how the original consumer-level model can be extended to N food groups. For clarity of exposition, the model in the paper focuses on two food groups—*healthy food* and *regular food*. In reality, of course, foods have varying degrees of healthfulness. In this section we extended the proposed model to consider an arbitrary number of food groups, with different degrees of healthfulness. Since the focus of the paper is on fruit and vegetable consumption, fruits and vegetables should always remain one of the food groups considered.

Let x_i , $i = 1, \dots, N$, be the quantity of food group i purchased in a single shopping trip. For example, x_1 could be fruits and vegetables, x_2 could be grains, x_3 could be protein, etc. Each food group has its own nutritional benefits. Let $v(\nu) g_i(x_i/T)$ be the “nutrition value term” of the consumer’s objective function for food group i . Intuitively, this is the boost in utility that the consumer realizes by purchasing food group i , based on its nutritional benefits. The functions $g_i(\cdot)$ can be adjusted to reflect the food group’s nutrient level and

the consumer's utility for having this nutrient. The only requirement is that $g_i(\cdot)$ must be concave for all i . A food group that is not nutritious (e.g., junk food) would have $g_i(x_i/T) = 0$. The consumer's optimization problem is written as

$$\begin{aligned} \max_{\mathbf{x}} \quad & f\left(\frac{x_1}{T}, \dots, \frac{x_N}{T}\right) + v(\nu) \cdot \sum_{i=1}^N g_i\left(\frac{x_i}{T}\right) - \frac{d(\delta)}{T} \\ \text{s.t.} \quad & \sum_{i=1}^N p_i x_i \leq B(\beta) \cdot T \\ & T = T(x_1, \dots, x_n; v(\nu), d(\delta)) \\ & x_i \geq 0 \quad \forall i \in \{1, \dots, N\} \end{aligned} \tag{EC.8}$$

Analogous to Problem $M^c(B, d, v)$, we assume that the function $f\left(\frac{x_1}{T}, \dots, \frac{x_N}{T}\right)$ is concave and separable. Additionally, we assume that $T(\mathbf{x}; v, d)$ depends on only one x_i . The intuition for this assumption, similar to the explanation in Section 2.1, is the following: A consumer will feel a sense of urgency to return to the store once a "critical food item" (as explain in Section 2.1) is depleted to a certain level. Therefore, $T(\cdot)$ depends only on the limiting food (i.e., the one that runs out first) from within the consumer's set of critical food items.

Using the same change of variables as in the original formulation (i.e., $\tilde{x}_i := x_i/T$), we see that Proposition 1 extends to Problem EC.8. Namely, it can be shown that Problem EC.8 is also concave. Therefore, the same techniques can be used to analyze the bi-level model (albeit, with more complex algebraic equations). Additionally, Lemma EC.1 easily extends since T only depends on one x_i . Using these facts, the analysis of the extended model can be accomplished using similar methods to the analysis of the original model.

To give an example of how the results can be extended, we consider the impact of the education intervention, the untargeted price intervention, and the access intervention on fruit and vegetable spending in the case of N food groups. Food group 1 will be used to represent fruits and vegetables, but the other two food groups may also have nutritional value.

We find that, in this extended model, the impact of the price intervention and the access intervention on FV spending is the same as in the original model. Analogously to Proposition 2 and Corollary 1, the price intervention always increases FV spending and the access intervention only increases FV spending in certain cases, depending on the functional form of T .

The impact of the nutrition education intervention, on the other hand, is a bit more complex in this extended model. The impact of the education intervention on fruit and vegetable consumption now depends on the relative slopes of the functions $g_i(\cdot)$, for $i = 1, \dots, N$, compared to the cost of each food group. In other words, now that all food groups could have nutritional benefits, the nutrient value per dollar of each food group plays a role. We will assume that fruits and vegetables have the highest nutrient benefit per dollar, given by the following condition:

$$\frac{g_1(y)}{p_1} \geq \frac{g_i(y)}{p_i} \text{ for all } i, y > 0$$

When T depends on x_1 , and the condition above holds, fruit and vegetable spending will always increase as the education intervention increases. When T depends on $x_{i'}$ for $i' \neq 1$, however, it is possible that fruit and vegetable spending may *not* increase as the education intervention increases. The intuition is the following: Food group i' has some nutritional value, and also contains the “limiting food.” Increasing value of nutrition decreases T , meaning that households will end up going to the store more frequently. In order to go to the store *less* frequently (which is desirable), the household needs to buy more of food group i' , which contains the limiting food. In order to do so, they may divert money away from fruits and vegetables (or other food groups).

On the other hand, the education intervention also increases the attractiveness of fruits and vegetables to the consumer. Therefore, the overall effect of the education intervention on fruit and vegetable spending will depend on the relative magnitude of these two competing effects. Mathematically, the latter effect (i.e., the effect of nutrition education on the attractiveness of FVs) will outweigh the former effect when value of nutrition has a small impact on $T(x_{i'}; v, d)$.

In this version of the extended model, the nutrition education intervention impacts every *all* food group i where $g_i(x_i) \neq 0$. In this sense, we are modeling a very broad nutrition education program that encourages the consumption of all foods with nutritional benefits. While this is a realistic model for many programs, there are also nutrition education programs that specifically encourage the consumption of fruits and vegetables (e.g. VeggieRx Nutrition and Cooking classes, the Fresh Fruit and Vegetable Program, etc.). If we focus on only these programs, and write the *nutrition value term* as

$$v(\nu)g_1(\tilde{x}_1) + \sum_{i=2}^N g_i(\tilde{x}_i),$$

the impact of the nutrition education program on FV spending is identical to the original model. Namely, it will always increase FV spending.

In conclusion, the consumer-level model can readily be extended to include an arbitrary number of food groups. The nutritional value of each food group will play a role in determining the effect of the nutrition education intervention; however, if we restrict ourselves to fruit and vegetable-specific education programs, the results remain consistent with the original model.

EC.4. Computational set-up

EC.4.1. Parameter estimation

This section presents a method for estimating the parameters of the consumer's optimization problem, given a dataset containing:

- Information on grocery shopping trips over a period of time for a population of consumers, including the types of foods purchased (or at least a flag for healthy versus regular food) and the prices paid.
- Enough information to calculate relative disutility measures for each household for getting to the grocery store (e.g., access information, employment information, household composition information)
- Information on the household's dietary beliefs and attitudes, such that relative measure of each household's "value of nutrition" could be elucidated.
- A subsample of households whose nutritional value of effectively zero.

Note that a dataset containing this type of information is not unrealistic—in fact, two datasets currently exist which each contain components of this information, and the union of these two datasets would result in the complete set of required information. The USDA's recent FoodAPS dataset contains detailed information on over 4,000 household's shopping behavior over a one week period, as well as information regarding their access to grocery stores and value of nutrition. Unfortunately, one week is not enough data to elucidate household's shopping behavior. The Nielsen consumer panel dataset, on the other hand, contains decades worth of data on household's shopping patterns. Unfortunately, it does not contain the necessary information on access and value of nutrition. Therefore, we use this dataset's shopping information but synthetically generate each household's access and value of nutrition, using a process described in the following section.

EC.4.2. Nielsen data and synthetic covariates

The Nielsen consumer panel dataset does not contain information related to the households' value of nutrition or access to grocery stores. Thus, these covariates were synthetically generated for each household. First, we generate N random "values of nutrition," where N is the number of households, according to a lognormal distribution with arbitrarily chosen mean of -1 and variance of 1. Because it has been found in the literature that households with a higher value of nutrition will consume a larger portion of fruits and vegetables compared to other foods (Levi et al. 2018), we sorted the households according to the ratio of their average FV spending to non-FV spending per grocery store visit and assigned the highest value of nutrition to the household with the highest average FV spending, etc. After this assignment, we add random noise (drawn from a lognormal distribution with mean zero and variance 0.3) to each household's value of nutrition. This way, value of nutrition is correlated, but not perfectly correlated, with the ratio of average FV spending to non-FV spending.

A similar procedure is used to construct each household's shopping disutility covariate. In this case, we use the fact that access—and thus likely disutility—has been found to be associated with shopping cycle time (Liese et al. 2014, Levi et al. 2018). Each household's disutility is drawn from a lognormal distribution with mean equal to $T_i^{0.1}$, where T_i is household i 's average cycle time, and variance 0.01.

Plots of the synthetically generated value of nutrition versus the ratio of average FV spending to non-FV spending, and synthetically generated disutility versus average cycle time, are shown in Figure EC.1.

EC.4.3. Identification

Consider a dataset containing the outcome variables h_i^* , u_i^* and T_i^* along with the relative measures of disutility, d_i , and value of nutrition, v_i for a population of consumers, $i \in \mathcal{P}$. Since we have multiple shopping trips per households, the values of h_i^* , u_i^* and T_i^* can be considered to be the average amount of HF and RF bought each visit, and the average time between store visits. We assume that these outcomes are realizations from a population of consumers who are solving Problem $M^c(B, d, v)$.

We will assume that T_i^* is of the form $T_i^*(h^*, u^*; d, v) = \min\{(\phi h^* \rho_h + (1 - \phi)u^* \rho_u), \phi T_m^h + (1 - \phi)T_m^u\} + \frac{w(d)}{r(v)}$. The functions $f(\cdot)$, $g(\cdot)$, $w(\cdot)$, and $r(\cdot)$ are mixed-effects functions, meaning that they contain both individual-level parameters as well as population-level




Figure EC.1 The covariates value of nutrition and disutility resulting from the synthetic covariate generation process.




Figure EC.2 Causal graph when $\phi = 0$. Red denotes the outcome variables and blue denotes the observed covariates.

parameters. For example, we will estimate $r_i(v_i)$ for household i by assuming the structure $r_i(v_i) = \max\{r_{0_i} - \bar{r} \cdot v_i, r_{min}\}$ where \bar{r} is fixed across all households, and r_{0_i} can vary between households but is drawn from a common distribution, with density function $f_r(r_{0_i}; \bar{r}_0, \sigma^2)$ with mean \bar{r}_0 and variance σ^2 . Therefore, we will say that \mathbf{r} is parametrized by the population parameter vector $\mathbf{r} = (\bar{r}, \bar{r}_0, \sigma^2)$ with random (or individual-level) component \mathbf{r}_0 .

Let $f(\cdot)$, $g(\cdot)$, $w(\cdot)$, and $r(\cdot)$ be parametrized, respectively, by $\boldsymbol{\eta}$, $\boldsymbol{\gamma}$, $\boldsymbol{\omega}$, and \mathbf{r} . We will assume that $f(\cdot)$ has random component f_0 , and that the functions $g(\cdot)$ and $w(\cdot)$ do not contain any random components. We seek a method for determining the unknown population- and individual-level parameters and using the observational data described above.

Figure EC.2 demonstrates the causal relationship between the covariates, parameters, functions, and outcome variables in the model, which will be useful for identification.

First consider the task of estimating \bar{r}_0 and ω . By conditioning on u^* and h^* , we remove the causal pathways between $w(d)$ and u^* or h^* as well as between $r(v)$ and u^* or h^* . After this conditioning and assuming independence between d and v , marginal changes in d and v , and the observed impacts on T^* , provides identification for $r(v)$ and $w(d)$. Notice that this implies that $\mathbf{r}(\mathbf{v})$ and $w(\mathbf{d})$ can be determined by observing T^* and conditioning on h^* and/or u^* . Practically speaking, with a finite number of households it may not be possible to condition exactly on h^* and u^* and still retain enough data. Therefore, in practice we approximate this conditioning by clustering households together with similar values of h^* and u^* . After determining \bar{r}_0 and ω , \bar{r}_0 and the latent components \mathbf{r}_0 are estimated for each household by penalizing large values of the variance, σ^2 .

Now consider estimation of ρ_h, ρ_u , and ϕ_i (or more generally, the marginal effect of h^* or u^* on T^*). Since $\phi_i \in \{0, 1\}$, this can be accomplished by conditioning on $w(d)$ and $r(v)$, since there is a causal pathway between $w(d)$ and h^*/u^* as well as between $r(v)$ and h^*/u^* . Once ρ_h, ρ_u , and ϕ_i are determined, variations in u_i^* and h_i^* within individual households and the subsequent impact on T_i^* allows for identification of ϕ_i .

In order to determine $f(\cdot, \cdot)$, we first condition on $v = 0$ so that the value term in the government's objective function is equal to zero. Intuitively this means that we consider only households whose purchasing decisions are not impacted by nutritional considerations. Variations in these consumers' budgets allows for identification on η (which contains the mean and variance of \mathbf{f}_0). Finally, γ and \mathbf{f}_0 can be identified through variations in value and budget for all households.

In general the process that we follow is:

1. Condition on u^* and h^* to determine \bar{r}_0 and ω using observations of T^*
2. Estimate \mathbf{r}_0 and \bar{r} using the previous estimates of \bar{r}_0 and ω , penalizing larger values of σ^2
3. Estimate ρ_h, ρ_u and ϕ using observations of h^*, u^* , and T^*
4. Condition on $v = 0$ and utilize variation in consumers' budgets to determine η .
5. Estimate γ and \mathbf{f}_0 using observations of T^*, h^* and u^*

EC.4.4. Estimation procedure

The method above concerns identification, but does not provide a method for producing estimates of the parameters at each of the three steps above after correctly conditioning. This section presents an estimation method and computational results for the entire parameter estimation problem given specific functional forms of the modeling functions.

After appropriate conditioning to ensure that the parameters of interest can be identified, estimating the parameters is a maximum likelihood problem that would be straightforward to solve if the functions relating the parameters and covariates to the outputs (T^* , h^* , and u^*) were simple. However, because h^* and u^* are implicitly defined (by the consumer's nonlinear KKT conditions), and due to the inclusion of both individual- and population-level parameters, the maximum likelihood problem is not straightforward. In general, with outcome variables \mathbf{y} , covariates \mathbf{x} , population-level parameters $\boldsymbol{\theta}$ and individual-level (random or latent) effects \mathbf{z} , maximum likelihood estimation seeks to maximize the marginal data likelihood, namely $f(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})$, with respect to $\boldsymbol{\theta}$. This is a natural quantity to wish to maximize since we know the ground truth of the outcome variables. In non-linear mixed effects (or latent variable) models, maximizing the marginal data likelihood is often done using stochastic expectation-maximization (EM) algorithms, which includes an ‘‘S-step’’ to approximate taking an expectation over $f(\mathbf{z}|\boldsymbol{\theta}, \mathbf{y}, \mathbf{x})$ which is often difficult in non-linear mixed-effects models. Instead of taking this expectation, the S-step samples from the distribution $f(\mathbf{z}|\boldsymbol{\theta}, \mathbf{y}, \mathbf{x})$, which is often accomplished using MCMC methods. For more information on stochastic EM, see Nielsen et al. (2000) and Delyon et al. (1999).

However, for the specific problem presented in this paper, each iteration of stochastic EM requires solving a system of nonlinear equations many times, and is not guaranteed to converge. Therefore, we choose an alternative approach which appears to work well in practice. Instead of seeking to maximize the marginal data likelihood, we maximize the complete likelihood, $f(\mathbf{y}, \mathbf{z}|\mathbf{x}, \boldsymbol{\theta})$ with respect to both \mathbf{z} and $\boldsymbol{\theta}$. The advantage of this approach is that it can be solved by one non-linear program, and produces estimates of both \mathbf{z} and $\boldsymbol{\theta}$ simultaneously. This method is known to produce similar estimates to those produced by maximizing the marginal data likelihood, and is much simpler and faster in this context.

We demonstrate the success of the proposed identification and estimation strategy method by generating synthetic ‘‘observational data’’ according to the proposed model with known parameters, and then recovering the underlying parameters. We assume that the functions $f(\cdot, \cdot)$, $g(\cdot)$, $r(\cdot)$ and $w(\cdot)$ are of the forms listed in Table EC.1.

We follow the identification procedure and estimation technique described above in order to estimate the parameters of the functional forms described in Table EC.1. Figure EC.3 (left) shows the bootstrapped joint density of the estimates for (\bar{r}) and ω . The straight

Table EC.1 Functional forms for parameter estimation

Function	Form
$f_i(\tilde{h}, \tilde{u})$	$\frac{(1+\tilde{h})^{1-\eta_{h_i}}-1}{1-\eta_{h_i}} + \frac{(1+\tilde{u})^{1-\eta_{u_i}}-1}{1-\eta_{u_i}}$
$g(\tilde{h})$	$c_\gamma \cdot \frac{\tilde{h}^{1-\eta_\gamma}-1}{1-\eta_\gamma}$
$r_i(v_i)$	$\frac{1}{r_{0_i}-\bar{r}^*v}$
$w(d)$	d^ω
$T(h, u, v, d)$	$\min\{(\phi h^* \rho_h + (1 - \phi)u^* \rho_u), \phi T_m^h + (1 - \phi)T_m^u\} + \frac{w(d)}{r(v)}$




Figure EC.3 Bootstrapped joint density for the estimation of ω and δ (left), $\bar{\eta}_h$ and $\bar{\eta}_u$ (middle) and η_γ and c_γ (right). Each bootstrapped sample contains 30 consumers and 20 observations per consumer.

lines show the true value of these parameters, and lighter blue indicates a higher density. The middle of Figure EC.3 shows the joint density of the estimation of $\bar{\eta}_h$ and $\bar{\eta}_u$. Finally, the rightmost figure shows the density of the joint estimation of η_γ and c_γ .

EC.4.5. Parameter estimates

The procedure described in the preceding section, and functional forms shown in Table EC.1, are used to estimate the parameters for the partially synthetic Nielsen consumer panel data. The estimated parameters are shown in Table EC.2.

Table EC.2 Estimates of modeling parameters from Nielsen data

Parameter	w	c_γ	η_γ	\bar{r}	ρ_h	ρ_u	$\mathbb{E}[r_{0_i}]$	$\mathbb{E}[\eta_{h_i}]$	$\mathbb{E}[\eta_{u_i}]$	$\mathbb{E}(\phi_i)$
Estimate	0.3	2.34	0.48	0.0	0.03	0.023	0.57	2.44	0.22	0.62

EC.4.6. Government parameters and functions

The modeling parameters and functions used to determine the optimal individual and population-level strategies for the data are given in Table EC.3.

Table EC.3 Modeling parameters for the upper-level problem

Parameter/function	U	$B(\beta)$	$v(\nu)$	$d(\delta, z)$
Value/functional form	3	$B + \beta$	$\frac{v \cdot \nu}{\nu + U}$	$\frac{(.5zd + (1-z)d) \cdot \delta}{\delta + U}$

EC.5. Population-level strategies

EC.5.1. Optimal individual strategies

The government's optimal investment between value, price, and location-independent access interventions for each individual consumer in the population is found by solving the Problem (M^g), the results of which are shown in Figure EC.4 as a ternary plot. Each point on Figure EC.4 represents the optimal intervention for one household. Note that the optimal individual interventions do not depend on the households' locations. As seen in Figure EC.4, the optimal strategy for individual households can vary widely. Many households' optimal intervention bundles contain a positive investment in only two out of the three interventions. However, there are some households whose optimal strategy includes a positive investment in all three (those that lie in the interior of the triangle in Figure EC.4). Although the optimal individual intervention bundles are a function of the household's characteristics and original shopping behavior, it is quite difficult to directly predict the optimal individual intervention bundles using linear models. This illustrates the complexity of the relationship between the consumer's characteristics and the optimal intervention bundles. However, these prediction methods are employed in the next section to help inform population-level strategies.

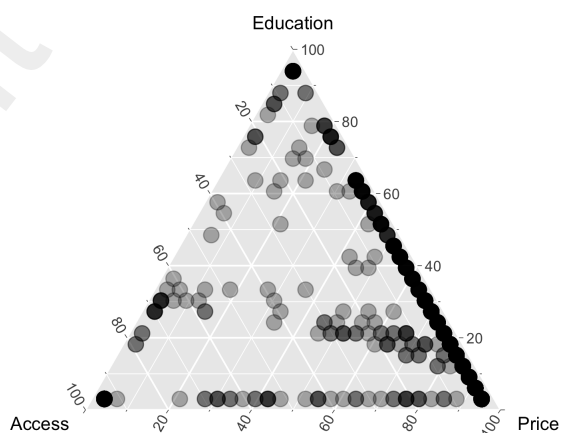


Figure EC.4 Optimal strategies for individual consumers.

EC.5.2. Group-level optimization problems

The utilitarian objective is given by Problem M_{util}^P .

$$\begin{aligned}
 \mathbf{s}_{util}^P &= \max_{\substack{\beta_k, \nu_k, \delta_k, z_n \\ k \in \{1, \dots, M\} \\ n \in \{1, \dots, N\}}} \sum_{i \in \mathcal{P}} \tilde{h}_i^*(\beta_{k_i}, \nu_{k_i}, d(\delta_{k_i}, z_{n_i})) & (M_{util}^P) \\
 \text{s.t. } & \beta_k + \nu_k + \delta_k \leq U - \frac{c}{M} \sum_{n=1}^4 z_n \quad \forall k \\
 & z_n \in \{0, 1\} \quad \forall n \\
 & \beta_k, \nu_k, \delta_k \geq 0
 \end{aligned}$$

where c is the cost of a single location intervention (e.g., the cost of building or subsidizing a new grocery store), and thus the government's remaining budget each group-level strategy is $U - \frac{c}{M} \sum_n z_n$. The budget constraint enforces that each household receives the same level of government funding, regardless of where they are located or into which group they are classified. The term $d(\delta_{k_i}, z_{n_i})$ is household i 's shopping disutility as a function of both the location-dependent and location-independent access interventions. The functional form of $d(\cdot)$ can be found in Appendix EC.3. It is assumed that if an access intervention is implemented in a neighborhood, each household's shopping disutility in that neighborhood is impacted equally. However, Problem M_{util}^P can easily be extended to the case where households within a neighborhood are impacted by the access intervention differently.

The max-min objective is given by Problem $M_{max-min}^P$.

$$\begin{aligned}
 \mathbf{s}_{max-min}^P &= \max_{\substack{\beta_k, \nu_k, \delta_k, z_n \\ k \in \{1, \dots, M\} \\ n \in \{1, \dots, N\}}} \sum_{k=1}^M \min_{i \in \mathcal{P}_k} \frac{\tilde{h}_i^*(\beta_k, \nu_k, d(\delta_k, z_{n_i})) - \tilde{h}_i^*(\mathbf{0})}{\tilde{h}_i^*(\beta_i^*, \nu_i^*, d_i^*) - \tilde{h}_i^*(\mathbf{0})} & (M_{max-min}^P) \\
 \text{s.t. } & \beta_k + \nu_k + \delta_k \leq U - \frac{c}{M} \sum_n z_n \quad \forall k
 \end{aligned}$$

where $(\beta_i^*, \nu_i^*, \delta_i^*)$ denotes the optimal individual allocation for household i (shown in Figure EC.4), and $\tilde{h}_i^*(\mathbf{0})$ is household i 's average HF consumption with no government interventions. The objective function is thus the ratio of the percentage increase in FV spending obtained through a population-level strategy and the percentage increase obtained through the optimal individual-level strategy.

Problems M_{util}^P and $M_{max-min}^P$ are not necessarily unimodal and thus may be quite difficult (at least theoretically) to solve computationally. In practice, since the government chooses funding levels from a discrete set of values (i.e., funding allocations are typically not determined down to the cent), a grid search approach is practical and fast, and is employed in this paper. Specifically, the following procedure is followed. First, each household's FV spending under each possible intervention bundle is determined. For each given intervention bundle, determining each household's FV spending is equivalent to solving Problem $\tilde{M}^c(\beta, \delta, \nu)$, which is concave (Proposition 1) and can be solved quickly. Second, the best intervention bundle for each group can be easily determined given the information from the first step.

EC.5.3. Multi-output regression trees

Figure EC.5 shows the multi-output regression tree that was used to determine the groups discussed in Section 4.3. Note that the goal of this tree is *not* prediction; the goal is to subset the population strategically. The parameter v is the household's value of nutrition, and the parameter η_u controls the concavity of the function $f(\tilde{h}, \tilde{u})$ with respect to \tilde{u} (the specific functional form can be seen in Table EC.1). The predicted individual strategy is denoted in the tree as `value=` (ν, β, δ) , meaning that the investment in education is the first number, investment in the price intervention is second, and access is third.

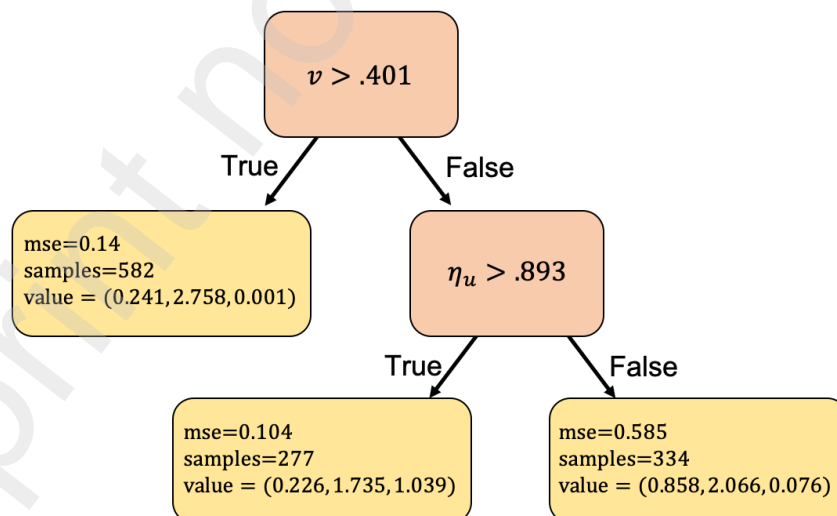


Figure EC.5 Multi-output regression tree used to choose groups in Section 4.3.