What Would We Eat if We Knew More: The Implications of a Large-Scale Change in Nutrition Labeling

Jason Abaluck*

December 31, 2012

Abstract

Abstract: This paper computes the welfare benefits of additional information about nutritional content in food by revealed preference and evaluates quantitatively whether the estimated behavioral response is consistent with information from experts on the relationship between diet and health. In doing so, it provides estimates of the impact of the law mandating nutrition labeling for all prepackaged foods in the US on nutrient consumption. Estimates derived from a structural model identified based on differential changes in information across foods are consistent with reduced form estimates comparing the change in calorie consumption among label users and non-label users. Taking the estimated willingness to pay for nutrient content as given implies that the labeling law led to an increase in consumer surplus of $25-40$ annually for each label user and that additional labeling regulations could generate a comparable benefit. Comparing the implicit value of nutrient information with a benchmark computed from medical evidence and the value of a statistical life (VSL) suggests that consumers are insufficiently responsive to health differences across foods. Taking this benchmark as the normative standard, revealed preference estimates understate the benefits of labeling by a factor of four, and thousands of dollars in additional per capita welfare gains could be realized by policies which would lead consumers to eat healthier foods.

*Thanks especially to my advisors Jon Gruber, Michael Greenstone, and Glenn Ellison, and to Karen Li for outstanding research assistance. Thanks also to Hunt Alcott, David Autor, Peter Diamond, Tatyanan Deryugina, Esther Duflo, Amy Finkelstein, Elena Harmon, Jerry Hausman, Panle Jia, Whitney Newey, Amanda Pallais, Michael Powell, Iuliana Pascu, Mar Reguant, Steven Ryan, Ashley Swanson and Joseph Shapiro for helpful comments and suggestions. Thanks also to Mary Brandt and Tomoko Shimikawa for assistance in obtaining data. Funding for this work was provided by NIA grant T32 AG000186-21.

1 Introduction

The World Health Organization estimates that individuals in the developed world could extend their life-span by a mean of 1.9-3.4 years through healthier dietary habits (World Health Organization 2002). Valuing these life-years at $\$100,000$ (Gruber and Koszegi 2001), this implies about a trillion dollars in life-years lost every year in the US alone by not eating the healthiest possible diet.[1] The aim of the paper is to determine the extent to which providing more information about nutrition shifts individuals towards healthier diets - and whether the response to this information, relative to the response to price changes, suggests that individuals are appropriately incorporating nutritional information into their food consumption decisions. Does one trillion dollars a year represent the willingness to pay for the taste and convenience of unhealthy foods, or does it represent gains that can be realized through policies which lead to healthier eating?

I investigate these questions by studying perhaps the largest case of mandated information provision in US history: the Nutrition Labeling and Education Act. This act mandated nutrition labeling of all prepackaged foods in the US beginning in 1994. I present evidence indicating that this law did impact consumption and develop a model of food demand as a function of nutrient characteristics which allows me to generate revealed preference estimates of its benefits. The model also allows me to evaluate the potential benefits from additional information about nutrient content such as the recent law mandating calorie labeling in all chain restaurants (Rosenbloom 2010). Finally, I compare the observed response to nutrient information to a benchmark response computed from medical evidence and the value of a statistical life (VSL).

The intuition is as follows: suppose we observe that consumers receive new information about nutrient content and negatively update their beliefs about the health consequences of cheeseburgers, but their consumption remains unchanged. If this occurs because they really love cheeseburgers, they should also be unresponsive to prices. If we observe that consumers readily substitute away from cheeseburgers when the price increases but not when they get bad news about health consequences, then they must place a low value on health. We can

[1] Details of these calculations are described in the online appendices.

compare the implicit value of health to the VSL estimated from other choices and ask whether consumer behavior appears consistent across settings. One can equivalently think of this exercise as starting with the value of the health benefits of eating different foods and asking what fraction of these benefits can be realized given the degree to which consumers already incorporate health information into their consumption decisions and given the willingness to substitute to foods which may be less desirable along other dimensions such as taste and price.

An important motivation for this project is to bridge the gap between two competing methodologies for analyzing policies which impact health. Many health policy analysts and some economists compute the benefits of such policies in terms of life-years saved, but do not consider whether consumers are made worse-off along other dimensions through substitution to otherwise less desirable alternatives (see e.g. World Health Organization 2002 and Variyam and Cawley 2006). Alternatively, many economists assume that individuals appreciate the full cost of their food consumption decisions and focus only on the benefits of health arising from externalities generated by the health insurance system (Bhattacharya and Sood 2006). My analysis follows O'Donoghue and Rabin (2006) and Gruber and Koszegi (2001) in accounting for revealed preference data, but also attempting to account for "internalities" from sub-optimal choices. While those analyses take a bottom-up approach to identifying internalities via a particular mechanism, my approach provides a top-down picture which connects directly with the health benefits of a particular policy by measuring the degree to which consumers incorporate those health benefits into their consumption decisions. To the extent that the VSL measured in other contexts reflects fully-informed and time-consistent decisions, one can think of my analysis as a kind of omnibus test which detects internalities generated by time inconsistent behavior as well as internalities generated by imperfect information about the relationship between diet and health.[2]

The main results are as follows: the structural model I estimate implies that the NLEA led to a reduction in calorie intake of 50-90 calories among label users. The model is identified

[2]O'Donoghue and Rabin (2006) note that the parameter β in their theoretical model could be interpreted to incorporate other sources of consumer error besides hyperbolic discounting and they simulate the implications of the model for different values of β. One can then think of this project as an attempt to operationalize their framework by calibrating a model allowing for multiple β's to reflect multiple dimensions of nutrient information.

based on the fact that consumption of calorie-intense foods fell within product groups and that calorie consumption fell most in those product groups which experienced the largest increase in labeling. This result is consistent with (the lower end) of several reduced form estimates identified by comparing calorie consumption among label users and non-label users, and with earlier studies of the impact of labeling on consumption. When I examine the response to particular nutrients, I find that the reduction in calories appears mainly to be due to a reduction in fat intake. The welfare benefits estimated via revealed preference imply that labeling led to a $25-40 annual gain, and that an additional gain of $40-$80 annually is possible with additional labeling. Reconsidering these benefits using benchmark preferences computed from the VSL gives answers four times as large. Taking into account the willingness to substitute to different foods and estimates of consumers current beliefs about health information, I find that the potential annual welfare gain from fully informed choices remains as high as 30-40% of the monetary value of the potential life-year gains from the healthiest possible diet.

Section 2 describes a simple example to explain how the structural model identifies the willingness to pay for each nutrient. Section 3 briefly reviews some related literature. Section 4 describes the available data on labeling, food consumption and prices. Section 5 reports reduced form estimates documenting that the NLEA appears to have led to a decrease in the consumption of high calorie foods relative to low calorie foods in product groups where labeling increased. Section 6 introduces the model of labeling and food demand, and reports estimates of the willingness to pay for nutrient content for several nutrients. In Section 7, I perform several welfare exercises evaluating the impact of the NLEA taking preferences as given, and in Section 8 I re-evaluate these gains given a normative benchmark derived from expert medical knowledge and assumptions about the value of additional life-years. Section 9 concludes.

2 A Simple Example

The main goal of this paper is to understand whether food consumption decisions adequately incorporate health information. My focus will be on the consumer response to nutrient information. I attempt to understand both how consumers respond to additional information about nutrient content and how they would respond if they were fully aware of how experts understand the relationship between nutrient content and health and behaved in a time-consistent way. There are two reasons I focus on nutrient information: first, because this explains a large amount of the variation in experts' assessment of the health consequences of different foods (Martin, Beshears, Milkman, Bazerman, and Sutherland 2009) and second, because I am able to use the nutrition labeling law as a source of variation to analyze the impact of nutrient information on consumption.

I define the willingness to pay for nutrient content as the marginal rate of substitution between nutrient content and prices. In this section, I present a simple example to make more transparent how I estimate the willingness to pay for nutrient content for a given nutrient. When nested in a broader system of food demand equations, these willingness to pay parameters allow me to compute the welfare gain from additional information about nutrient content and can be compared to benchmark normative values to determine if welfare gains exist from further informing consumers.

To identify the willingness to pay parameters, I ask the following question: what is the change in consumption induced by a change in information about nutrient content, and what is the magnitude of the change in price necessary to induce the same change in consumption? Table 1 presents a fictional product group consisting of 4 salad dressings. Of these four, three experience a change in labeling. Assume for now that consumer beliefs in all periods are directly observable and that calories are the only nutrient. The variation necessary to identify the parameters of interest is the ex post calorie content: after labeling, consumers learn that Dressing 2 is healthier than they thought (fewer calories / gram), Dressing 3 is as healthy as they thought, and Dressing 4 is less healthy than they thought. Further, there is a *relative* decline in the consumption of dressings for which individuals received bad news (despite the overall increase in dressing consumption): when the relative change in consumption of Dressings 2-4 is attributed entirely to the change in information, this implies that an increase of 1 calorie / gram in caloric intensity leads to a 1 gram reduction in consumption.

To move from this observation to an estimate of the willingness to pay for a reduction in calorie content, the magnitude of this decline must be compared to the marginal impact of price on consumption. Suppose that for each of the labeled products, $\frac{\partial N}{\partial p_j} = -1$ (where price is measured in cents). That is, a \$.01 increase in price/gram leads to a 1 gram reduction in demand. Then a \$.01 increase in price/gram has the same impact on consumption as a 1 calorie/gram increase in nutrient content. This implies a willingness to pay for a reduction in calorie content of \$.01 per calorie. Note that this identification method does not assume unhealthy foods are less desirable in cross-section, or even that individuals are equally as willing to substitute away from unhealthy foods as healthy foods. It may well be the case that individuals prefer the taste of foods with undesirable nutrient profiles and that they are less willing to give up these foods. To the extent that this is the case, the price elasticity for such foods will be smaller, so for a given change in information, the same observed change in consumption would translate into a larger willingness to pay for nutrient-content.

I assume in Table 1 that consumers' beliefs about calorie content are observable in all periods. Throughout, I maintain the assumption that label users know exactly the content of labeled foods. Below, I calibrate a model of beliefs about nutrient content for unlabeled foods using data from other studies which directly elicited beliefs about the nutrient content of unlabeled foods.

3 Literature Review

The theoretical literature examining the impact of information on health-related choices stretches back many years (e.g. Grossman 1972); the empirical literature is newer but growing. Alan Mathios and Pauline Ippolito conducted a series of studies in the early 1990s investigating the impact of nutritional information on dietary behavior after the removal of federal restrictions on health advertising. Ippolito and Mathios (1990) found that when the government lifted a regulation prohibiting the advertising of health claims in the mid-80s, consumption of high fiber cereals increase, and new cereal products higher in fiber were introduced. Ippolito and Mathios (1995) examine time series data on fat consumption during the same period and conclude that the removal of barriers to advertising on saturated fats led to a reduction in saturated fat consumption. A few recent studies have investigated the impact of calorie-labeling in restaurants: Elbel, Kersh, Brescoll, and Dixon (2009) collects consumer receipts from low-income fast food restaurants in New York and find that while consumers claim to use calorie information, there is no discernable impact of labeling on consumption. Bollinger, Leslie, and Sorensen (2010) finds a modest 6% reduction in calorie consumption at Starbucks after menus begin listing calorie information; calories consumed of food products fell by 14% per serving, while there was no change in the consumption of beverage products.

There is also a substantial literature outside of economics studying the impact of the NLEA. Moorman (1996) conducted detailed surveys of shoppers in supermarkets just before and just after the NLEA took effect to determine its impact on consumer information processing. She found that after the passage of the NLEA, consumers were more informed about the fat content of recently purchased products and spent more time comparing alternative products within food groups, especially for unhealthy products. She did not directly examine how this knowledge impacted food consumption. Balasubramanian and Cole (2002) also conduct in-store surveys, and they also find that post-NLEA consumers spend more time shopping, but attribute this to the increased presence of nutrition claims other than those appearing on the nutrition facts label. They find that after the passage of the NLEA, consumption of foods labeled "low fat" or "low sodium" increased.

None of the existing studies of label use attempt to evaluate the normative benefits of nutrition labeling via revealed preference. The three labeling studies most closely related to the analysis in this paper are Mathios (2000), Variyam and Cawley (2006) and Variyam (2008). Mathios (2000) studies the impact of the NLEA on the demand for salad dressings; he finds that prior to the NLEA only the lowest-fat salad dressings voluntarily labeled, and after the NLEA there was a significant decline in sales for the highest fat dressings. Variyam and Cawley (2006) and Variyam (2008) compare label users and non-label users using difference-in-difference methods to investigate the impact of the NLEA and label use on nutrient consumption and obesity (measured using body-mass index). They find decreases in obesity rates among non-hispanic white-women which they estimate leads to a $63 to $166 billion dollar reduction in life-years lost over a 10 year period.

The existing literature on nutrition labeling convincingly identifies the impact of labeling on consumption in specific settings or for a limited range of products (Bollinger et al. 2010, Mathios 2000, and Kiesel and Villas-Boas 2009). I extend this literature in three ways. First, using food diary data, I evaluate how labeling of particular products impacts overall food consumption given a model of satiation in which a reduction in the consumption of some foods leads to an increase in the consumption of other foods; this addresses the concern in earlier studies that a reduction in the consumption of labeled products may be offset by an increase in the consumption of other products. Second, the NLEA provided disaggregated information about the components of calorie-content from fats, proteins and carbohydrates, and so can be used to analyze whether individuals responded differently to calories of different sorts; e.g. whether consumers are more sensitive to calories from fats than to calories from proteins. Such analysis is especially urgent in light of the fact that the recently passed Patient Protection and Affordable Health Care Act includes provisions mandating that all large chain-restaurants post total calorie information. Third, the existing literature focuses on the positive impact of labeling on nutrient consumption and its potential impact on obesity: none of the existing studies evaluate its normative impact via revealed preference. The structural model I estimate allows me to consider the offsetting cost of being more nutritious: individuals are consuming potentially less desirable foods all else equal, so the potential health gains overstate the welfare increase from dietary changes.

The methodology in this paper relates to a broader literature in behavioral welfare economics. The analysis here can be thought of as an attempt to operationalize the theoretical framework developed in O'Donoghue and Rabin (2006) to analyze internalities in food consumption decisions in the context of a richer model of nutrient information and food demand. The method of determining a benchmark weight to attach to product characteristics based on values derived from decisions in other contexts has previously been applied in Abaluck and Gruber (2009) and Allcott and Wozny (2010). More generally, the analysis here is consistent with the choice-theoretic framework developed by Bernheim and Rangel (2008) given a set of assumptions about the welfare-relevant choice domain,[3] and directly follows the agenda

[3]The relationship between the exact normative assumption I make above - that the VSL from other contexts should be used to specify the normative-utility function - and the framework in Bernheim and Rangel (2008) depends on the positive explanation for why VSLs differ across settings. If the VSL from food consumption is low because consumers are imperfectly informed about the relationship between diet and health, then my assumption maps naturally into that framework. If the measured VSL is low due to self-control issues or due to a distrust of expert beliefs, then the relationship is potentially more complicated.

for behavioral welfare economics laid out by Beshears et al. (2008). I use a structural model to infer how estimated preferences would vary with more contextual information; ultimately, this model should be tested using data on how such contextual information impacts choices.

4 Data and Institutional Background

The Nutrition Labeling and Education Act was passed in 1991 and mandated the presence of nutrition labeling on all prepackaged foods beginning in 1994. I analyze the impact of this law on food consumption by combining labeling data from the FDA at the product group level, food diary data collected by the USDA and linked to information about individual label use behavior, a cross-section of (national-average) prices at the product level from the USDA, and price time-series at the product group level from the CPI.

4.1 Labeling Data

The standardized template mandated by the NLEA for nutrition labels is given in Figure 3. For a fixed serving size, the label reports the number of grams in a serving (or milliliters for beverages), the number of grams of total fat, saturated fat, carbohydrates (with sugars indicated when non-negligible), protein and fiber, and the number of milligrams of cholesterol and sodium. The label also reports these values as a percentage of the FDA's recommended daily value (RDV %). The label reports only RDV % for Vitamin A, Vitamin C, Calcium and Iron. I focus the analysis below on the nutrients for which the label provides exact quantities. Prior to the NLEA, any products which voluntarily reported nutritional information were legally required to use the format shown in Figure 4. Several differences are worth noting: the older label does not report the serving size in grams, it does not disaggregate fat into saturated fat or carbohydrates into sugar and fiber, and it does not report cholesterol. It reports the RDV % for only a subset of nutrients, and it does not report the total recommended intake for a 2,000 calorie diet or the number of calories per gram of fats, carbohydrates and proteins. These differences are taken into account when I specify consumers' information sets below.

The labeling data in this paper comes from two sources: the FDA's Food Labeling and Package Survey (FLAPS) produces a biannual estimate of the proportion of products labeled in 52 different product groups since 1979, and the Diet and Health Knowledge Survey provides information about individual label use behavior. The NLEA mandated the labeling of all prepackaged foods in the US, with a few exceptions for foods with negligible nutrient content (seasonings and spices). The law did not cover freshly prepared foods such as fruits and vegetables, or foods baked on site such as bakery products or restaurant products. The FLAPS survey indicates that 61% of prepackaged foods contained nutrition labels in 1988 (measured as a proportion of total food expenditure), 66% were labeled when the law was passed in 1991, and by 1995, 95% contained labels including all products mandated to contain labels by the NLEA. There was also substantial variation across product groups in the fraction of foods which were labeled prior to the NLEA. Figure 5 illustrates this variation for a sample of product groups.

The labeling data from the DHKS indicates general label use behavior in 1989, 1994, 1995 and 1996 (e.g. use of nutritional information on the label, Often, Sometimes, Rarely, Never), and elicits nutrient specific label-use information in 1990, 1991, 1994, 1995 and 1996 (e.g. use of calorie / saturated fat / fiber information on the label, Often, Sometimes, Rarely, Never). Table 2 reports the proportion of female DHKS respondents in the age range I consider who report using the nutrition facts panel "Often" or "Sometimes" in each year the question was asked, as well as the proportion who report using the nutrition facts panel to examine the content of a particular nutrient "Often" or "Sometimes" in each year the question was asked. The proportion who report using the panel at all remains roughly constant at 75%. The most commonly used information concerned calorie content and total fats; the proportion reporting use of calorie content increased from 70% to 80% over the period of the sample, and the proportion reporting use of total fats increased from 75% prior to the NLEA to 80% in the years following the NLEA.

4.2 Food Consumption Data

The food consumption data comes from the Continuing Survey of Food Intake by Individuals (CSFII), the most complete data source recording individual dietary intake prior to 1999. The CSFII is linked to the Diet and Health Knowledge Survey (DHKS), which elicited information from CSFII respondents on their stated dietary goals and their understanding of the nutritional content and consequences of different dietary behaviors.

The CSFII is a repeated cross-section; food consumption data was collected in 1985, 1986, 1989, 1990, 1991, 1994, 1995 and 1996 for a total of 36,895 respondents. In each year, the first day of data, was elicited through an in-person interview in which the interviewer assisted the household member in developing accurate measurements of the quantity of food consumed. In subsequent days, the information comes from a food diary completed by the respondent. Data from both types of surveys are used in the analysis; I check that the main results are unaffected if I restrict to the first day of data in Appendix D.

There are 6400 unique food codes across all years. The level of aggregation extends to detailed generic descriptions of foods, but not to particular brands. A typical food name is "Fruit Punch Flavored Drink Powder". The database was designed so that foods would be coded separately if their nutrient content was distinctly different (e.g. there is a separate entry for "Low Calorie Fruit Punch Flavored Drink Powder"). For each individual, several days of consumption are reported spaced throughout the year. For each day of data, all foods consumed were reported along with the quantity of the food consumed, the place and time when it was consumed, how the food was prepared, where it was purchased, and which other household members were around when the food was consumed. The data on where it was purchased is especially important, since this allows me to distinguish between foods purchased in a given product group at a supermarket (which might be labeled) and foods purchased at a restaurant (which would not be labeled).

The CSFII provides several advantages over alternative data sources in analyzing the impact of nutrition labeling: first, unlike data from individual retail establishments, it attempts to give a complete record of food consumption for each individual so it is possible to study the impact of labeling on overall nutrient consumption and to understand how different substitution patterns would impact overall nutrient consumption. Second, it provides data on a representative sample of the US population. Third, it is conducted during the period when the NLEA was passed and implemented, so it is the only data source available which can be used to study the long-term impact of a large-scale change in nutrition labeling on food consumption.

Food diary data is known to understate total food consumption; I discuss earlier studies documenting the extent of this bias and several steps to investigate the robustness of the results to this bias in the Appendix D.

4.3 Price Data

The price data I employ comes from two sources: first, from the USDA, I have a cross-section of prices for almost all foods in the CSFII database in 2003. Unfortunately, this information was not collected during the sample period I study. To remedy this, I use CPI price series available for 24 of the 53 product groups to deflate the 2003 prices, and I deflate the remaining prices using a general food price index. Because it is difficult to find credible instruments for prices, I impute price elasticities from existing studies. All prices are expressed in 1990 dollars.

4.4 Sample Selection

I noted above that the full CSFII sample includes 36,895 respondents. This number includes both the main survey which was designed to give a representative sample of the US population in each year, and a separate sample which was collected prior to 1994 for only low income individuals. To avoid the complications associated with weighting the data, I drop all individuals from the separate low income sample. This leaves 28,965 respondents. The included demographic groups and the number of days of food diary data elicited also varied across years. The primary sample I use in the below analysis is the largest consistent demographic that can be constructed from 1985-1996; this includes two days of records for all women aged 19-50. This sample consists of 7,298 distinct individuals. The longer sample period allows for the inclusion of product-specific linear time-trends which absorb most of the variation if the shorter 1989-1996 window is used. The diet and health knowledge survey which is linked to the CSFII was completed only by individuals identified as the "main meal planner" in each household. Of the 7,298 women in the final sample, 6,436 - 88% - are identified as the main meal planner. These individuals form the primary sample used in the analysis which makes use of individual-level label use information. Additional details of sample creation are described in the online appendices.

5 Reduced Form Evidence

In this section, I present some reduced form evidence indicating that the NLEA led to a reduction in consumption of high-calorie foods relative to low-calorie foods. As discussed in Section 2, this variation - when combined with assumptions about prior information and price elasticities - can be used to identify the willingness to pay for a change in calorie content.

I aggregate the data to the product-year level and ask: does consumption of products with more calories/gram decline relative to products with fewer calories/gram when the proportion of products labeled in a product group increases? Because the proportion of products labeled may respond endogenously to demand, I also consider specifications in which I instrument for this proportion to isolate the variation induced by the NLEA. I begin by estimating the equation:

$$C_{jt} = \beta x_j L_{g(j)t} + d_j + d_t + e_{jt} \tag{1}$$

where C_{jt} gives the average per capita calories of product j consumed at time t, x_j gives the calories per gram of product j, $L_{g(j)t}$ gives the proportion of products labeled in product group g (to which j belongs) at time t, and d_j and d_t denote product and time fixed effects respectively. A separate instrument is constructed for each product group; these instruments are constructed by interacting x_j with a dummy variable which is 1 in product group g after the NLEA and 0 otherwise. I estimate models using both C_{jt} and $\ln(C_{jt})$ as the dependent variable.

The results are reported in Table 4. In all cases, the OLS and IV results are quite similar due to the fact that a large portion of the variation in labeling in the data is captured by comparing the pre- and post-NLEA periods. The coefficient in the linear models can be interpreted as the decline in consumption associated with an increase in caloric content of 1 calorie/gram when a product goes from unlabeled to labeled (or more precisely, when the proportion of products labeled in the product group in question goes from 0% to 100%). Note that 1 calorie / gram is a large increase; the mean food in the data has 2 calories / gram. To think about identification, it will be helpful to keep in mind a simple example with two product groups, A and B, in which all of the foods in product group A have more calories per gram than the foods in product group B.

The first specification includes only product and time fixed effects, so the coefficient of interest is identified based on whether consumption of high calorie foods declined relative to low calorie foods regardless of whether they are in the same product group; if consumption in product group A declines relative to product group B, this will lead us to estimate a negative coefficient. This specification implies that an increase of 1 calorie/gram is associated with a consumption decline of .09 calories (average consumption is 1 calorie; it is so small because most foods are not consumed by the vast majority of consumers). In the second specification, additional fixed effects are added for each group-year. This absorbs all across group variation, so the coefficient is identified by relative changes in the consumption of high and low calorie foods within product groups; a negative coefficient implies that individuals substituted away from the highest calorie foods within group A towards the lower calorie foods in that group (and the same for group B). The coefficient in this model is larger; it implies that an increase of 1 calorie/gram is associated with a relative decline in consumption of .13 calories when labeling changes. In the third specification, I include separate linear time trends for each product. The results in the first and second specifications are consistent with a story in which consumers develop a taste for low calorie foods over time in precisely those product groups where labeling is initially less common (which might occur because ex ante labeling is less common among less healthy product groups). The third specification shows that controlling for such time trends makes the coefficient even larger: 1 calorie/gram is associated with a .22 calorie decline in consumption.

The second panel in Table 4 reports the same three specifications using the log of calories as the dependent variable. The coefficient in the log models expresses the change in consumption associated with an increase in caloric content of 1 calorie/gram when labeling changes in percentage terms. The log models imply that each increase of 1 calorie/gram is associated with a decline from 4-39% depending on the specification. As above, the magnitude of the coefficient grows larger when we include group-year fixed effects, and larger still when we include product specific linear time trends. The last three specifications in the second panel rerun the linear specifications using only foods not covered by the NLEA as a falsification test. In all cases, the coefficient of interest is insignificantly different from zero or positive, which is consistent with a story in which labeling induced some substitution within product groups towards unlabeled foods.

Figure 6 replaces the independent variable with variables interacting the quintile of calories / gram with the proportion of products labeled and graphs the resulting coefficient in the linear model with group-year fixed effects. The figure suggests that labeling induced substitution throughout the distribution of calorie intensity, with consumption declining more for higher calorie foods.

In the next section, I develop a structural model which links these results to earlier studies of the impact of labeling on total calorie intake by explicitly modeling substitution across foods due to satiation. The structural model will also make explicit how labeling impacts consumers' information sets allowing us to determine which foods are likely to be most impacted by labeling and to separately identify the impact of information provided about different nutrients (the results reported so far could be due to the impact of information about calories or information about other nutrients correlated with calories). Finally, the structural model will specify consumers' prior information about nutrient content which will allow us to determine how much consumption changes in response to new information and, via a comparison to price elasticities, how much consumers are willing to pay for a change in nutrient content.

6 Nutrition Labeling and the Willingness to Pay for Nutrient Content

In this section, I develop a model of food demand which allows me to evaluate via revealed preference the welfare gains from information provision and ultimately, to assess quantitatively the magnitude of the response to information relative to that implied by expert medical knowledge combined with assumptions about the value of a statistical life.

The intuition behind the model is laid out in Section 2. The main idea is to capture how information about nutrient content impacts dietary choices relative to prices when consumers optimally choose their daily diets given their tastes for different food products, the relative prices of products, and the fact that there are diminishing returns to individual foods and overall satiation. I begin by laying out the formal model and deriving the estimating equation. Next, I specify the assumptions made about consumers' information sets. I then present estimates of the willingness to pay parameters. The meaning of these estimates in explored in the remainder of the paper by examining their welfare consequences.

6.1 Food Demand Equations

Let N_{ijt} denote the number of grams of product j consumed by individual i at time t. The utility of individual i of consuming diet $d_{it} = \{N_{i1t}, ..., N_{iJt}\}$ is given by:

$$\begin{aligned} U_{it} &= \sum_j (\gamma_j + \rho_j t + \epsilon_{ijt}) \frac{(K + N_{ijt})^{1-\frac{1}{\eta_j}}}{1 - \frac{1}{\eta_j}} \\ &\quad + \sum_n \alpha_n(X_{in}) \left(\sum_j N_{ijt} E_{ijt}(x_{nj}) \right) + \phi I_{it} \end{aligned} \tag{2}$$

where γ_j, ρ_j, ϵ_{ijt} parametrize individual i's taste for food at time t, η_j determines the elasticity of demand for food j, $E_{ijt}(x_{nj})$ gives the expected content of nutrient n in 100 grams of product j at time t as a function of the actual nutrient content x_{nj} (so $\sum_j N_{ijt} E_{ijt}(x_{nj})$ gives total expected consumption of nutrient n), and I_{it} gives individual i's income at time t. X_{in} gives average consumption of nutrient n for consumer i, computed by averaging over all reported days of consumption in the data. In the models reported in the main text, I assume that $\alpha_n(X_{in}) = \alpha_n \cdot \max(X_{in} - \bar{X}_{in}, \underline{X}_{in} - X_{in}, 0)$ where $\bar{X}_{in}$ and $\underline{X}_{in}$ are known constants so that the function mapping nutrient content to utility is piecewise linear as discussed below.

Note that in the model the actual nutrient content of a product j is fixed and does not change over time; as noted above, this is part of the definition of a product. Individuals have a well-defined maximization problem because they are aware of the number of grams of each product that they consume (or equivalently, the number of servings of 100 grams), but they may be uncertain about the nutrient content in each gram.

The key parameters of interest are the marginal utilities of nutrient consumption α_n (and in particular, α_n/ϕ, the marginal utility normalized by the marginal utility of income so that it is expressed in dollar terms). The model in principal allows the willingness to pay for nutrient content to vary with current total nutrient intake X_{nit}. The results reported in the text allow α_n to vary in a piecewise linear way depending on where current nutrient consumption falls relative to the FDA recommendation for someone of my age, gender and activity level (so for example, I have a willingness to pay to avoid sodium if I currently consume more than 2300 mg a day, but I am indifferent to sodium on the margin if my total intake is below that amount). There is some qualitative evidence from the DHKS that the marginal value of nutrients at least has the same sign for most individuals. When asked whether they consume "too much", "too little" or "about the right amount" of a series of nutrients "compared to what is healthy", the vast majority of respondents indicate that they consume either the right amount or too much calories, fats, saturated fats, sodium and cholesterol, and either the right amount or too little fiber and protein.[4] Table 3 gives the exact percentages.

Individuals maximize this subject to two constraints. First, the usual budget constraint:

$$\sum_j p_{jt} N_{ijt} + I_i \leq W_{it} \tag{3}$$

where p_{jt} is the price of product j at time t and W_{it} is wealth. And second, I assume that individuals always consume a constant weight of food in each day in expectation:

$$\sum_j E_{ij}(N_{ijt}) \leq N_i \tag{4}$$

The second constraint addresses a worry in many earlier studies of food labeling that if individuals consume less of some foods because they are labeled, they will just substitute

[4] To the extent that consumers care about nutrients for reasons other than health, this question may be insufficient to determine the sign of the marginal value. For example, if consumers believe their physical appearance would be improved by consuming more calories, they may desire consuming more calories on the margin even if they believe that doing so would be harmful to their health. I consider such concerns in more detail in Section 8.2.

to other foods (Bollinger et al. 2010). It is motivated by a stylized fact in the literature on consumer satiation: individuals tend to consume a constant weight of food (Rolls 2009); when given a pre-load consisting of a certain number of grams, they reduce their consumption later in the day by this number. The main consequence of this constraint in the model is that the impact of any shift in consumption on overall nutrient intake will be muted, since any decline in the grams consumed in some product groups must be offset by an increase in other product groups. The expectation in this equation is taken over the unobserved taste parameters ϵ_{ijt}; this capture the fact that the satiation constraint is not binding every single day. It is binding in expectation so that on average grams consumed do not change, but there may be fluctuations from day to day in the number of grams consumed.

An important simplifying assumption of the model is that substitution between foods occurs only through this satiation constraint. The available data is unsuited to the estimation of own and cross-price elasticities because I do not have exogenous pricing variation and because the information on how prices changed over time is only available at the product group level.[5] Instead, own-price elasticities are imputed for each product group from existing estimates through a procedure described in more detail below. In Appendix D, I confirm that the main results are not impacted if we allow for somewhat richer substitution patterns by allowing demand to vary with the number of low-fat substitutes in the same product group.

Because calories are a linear combination of fats, carbohydrates and proteins, I consider models with total calories included along with cholesterol and sodium, and with calories disaggregated into saturated fats, unsaturated fats, protein, (non-fiber) carbohydrates and fiber. Because the coefficients in the latter specification do not scale proportionately with their contribution to total calories, the model with aggregated calories cannot be exactly correct. There is no single willingness to pay for a change in calorie content; the willingness to pay for a change in calories depends on the underlying change in nutrients which results in the change in calories. Nonetheless, we can think of the estimated coefficient on calories

[5] Cross-price elasticities could in principle also be estimated from the changes in demand induced by labeling, but the estimated changes are too small to obtain estimates with any precision; the exercise is also made more difficult by the fact that labeling data is only available at the product group level, not the individual price level.

as the willingness to pay for a change in calories if the proportion of the change in calories due to each of the underlying nutrients is identical to the variation caused by the NLEA. ϕ gives the marginal utility of income, so α_n/ϕ gives the dollar willingness to play for a unit of nutrient n.

Let $\theta_{ijt} = \mu_{it} + \phi p_{jt} - \sum_n \alpha_n E_{ijt}(x_{nj})$ where μ_{it} is the multiplier on the food amount constraint and $\gamma_{ijt} = \gamma_j + \rho_j t + v_{ij} + \epsilon_{ijt}$. We can rearrange the first order condition to give:

$$N_{ijt} = \max\{0, \left(\frac{\gamma_{ijt}}{\theta_{ijt}}\right)^{\eta_j} - K\} \tag{5}$$

6.2 Estimation

Let $Y_{ijt} = \left(\frac{\gamma_{ijt}}{\theta_{ijt}}\right)^{\eta_j} - K$. We can think of this as the latent demand for each good - if $Y_{ijt} < 0$ consumers will not consume any of product j, and if $Y_{ijt} > 0$, they will consume exactly Y_{ijt} grams. To make estimation of the model tractable, I Taylor-expend about z_0, the vector of parameter values in the first year when consumption is observed. Note that Y_{ijt} depends on the characteristics of other products $k \neq j$ only through total consumption which is captured by the μ_{it} term. Thus, because we are controlling for changes in μ_{it}, Y_{ijt} depends only on the characteristics of product j. In Appendix C, I show that Taylor-expanding about z_0 gives:

$$Y_{ijt} \quad \approx \quad w_{ij0}\left[-\phi p_{jt} + \sum_n \alpha_n(X_{in})E_{ijt}(x_{nj}) + t\hat{\rho}_j + \hat{\xi}_j + d\mu_{it}\right] + e_{ijt} \tag{6}$$

where $w_{ij0} \equiv \eta_{ij0}\frac{K+Y_{ij0}}{\theta_{ij0}}$, $\hat{\rho}_j$ and $\hat{\xi}_j$ are constants for each product and e_{ijt} is an error term which is independent of the included regressors. In Appendix C, I show that we can also express the weighting term as a function of observable prices and quantities: $w_{ijt} = \hat{\eta}_{ijt}\frac{E(N_{ijt}|N_{ijt}>0)}{\phi p_{jt}}$. Plugging this back into equation (6), gives:

$$Y_{ijt} \approx \hat{\eta}_{ij0}\frac{E(N_{ij0}|N_{ij0}>0)}{\phi p_{j0}}\left[-\phi p_{jt} + \sum_n \alpha_n(X_{in})E_{ijt}(x_{nj}) + t\hat{\rho}_j + \hat{\xi}_j + d\mu_{it}\right] + e_{ijt} \tag{7}$$

The scaling factor: $\hat{\eta}_{j0}\frac{E(N_{ij0}|N_{ij0}>0)}{p_{j0}}$ accounts for the fact that a change in nutrient content per gram is expected to result in a larger change in consumption for those products where consumption is more elastic (larger η_{j0}) and where a larger number of grams are consumed per serving (larger $E(N_{ij0}|N_{ij0} > 0)$). I assume that elasticities are constant across individuals within product groups so that $\hat{\eta}_{ij0} = \hat{\eta_{j0}}$ (in other words, I am assuming that observed price elasticities do not vary systematically with label use behavior, or with the total quantity of food consumed). The price elasticities for each product group are imputed based on the mean estimates from a recent survey article of price elasticity estimates, Andreyeva, Long, and Brownell (2010). For 33 of the 52 product group, no existing studies estimated a group-specific price-elasticity, so the average elasticity was used (the mean elasticity for all groups ranged from 0.34-0.79, with an interquartile range of .50-.75).[6] I compute $E(N_{ij0}|N_{ij0} > 0)$ in the data by dividing the population into ten cells based on deciles of total grams consumed and computing the average serving size for each food in each of those cells.

The full model is thus given by:

$$N_{ijt} = \max\{0, Y_{ijt}\} \tag{8}$$

where Y_{ijt} is given by equation (7). In Appendix C, I describe distributional assumptions on the primitives of the model so that e_{ijt} is normally distributed.[7] In this analysis, each individual-food is a separate observation. Foods which were not covered by the regulations in the NLEA as well as the bottom 5% of foods by total expenditure are included as an aggregated outside good.

I construct the set of instruments as the interaction of product-specific fixed effects with a dummy which is 1 after the NLEA and 0 prior to the NLEA. I estimate the model using the Smith-Blundell procedure (Wooldridge 2002, p. 530-533) which involves using the residuals from the first-stage regression as a control-function and correcting the standard errors for

[6] There is theoretical and empirical reason to believe that to the extent that identification is problematic, price elasticity estimates will be biased towards 0. The price elasticities reported in Chevalier, Kashyap, and Rossi (2003) for a limited number of foods using a plausibly exogenous instrument are 4-5 times larger than the price elasticities reported in Andreyeva et al. (2010). To the extent that the elasticities used in the model are biased downward, the willingness to pay parameters will be biased upward. Thus, using only the best identified elasticity parameters would only strengthen the conclusion that this willingness to pay is too low.

[7] The usual semiparametric estimators for censored regression models do not apply in this case because most foods are not consumed by the vast majority of consumers (Chay and Powell 2001). For example, the CLAD estimator would immediately trim all observations.

the variance in the first-stage estimates.

The inclusion of fixed effects raises a computational issue due to the large number of parameters as well as a conceptual issue due to the incidental parameters problem. The models reported in the main text treat μ_{it} as a random effect to avoid the incidental parameters problem; this assumption is problematic because it leads the error term of observation ijt to be correlated with the included variables of observation ikt for all k via μ_{it}. In Appendix D, I attempt to remedy this problem by estimating the full set of fixed effects using the computationally efficient procedure described in Greene (2001). While this procedure does generate bias due to the incidental parameters problem, the agreement between these results and the results reported in the main text suggests that the bias is not too severe. This agrees with the simulation evidence presented in Greene (2004), which suggests that the bias in the estimation of slope parameters in Tobit models from the inclusion of many fixed effects is less severe than the bias in binary choice models.

6.3 Specification of $E_{ijt}(x_{nj})$ and Identification of α_n

The willingness to pay for nutrient content α_n is identified using variation in perceived nutritional characteristics generated by nutrition labeling. In this section, I discuss the specification of $E_{ijt}(x_{nj})$, individual i's perceived content of nutrient n in 1 grams of product j at time t as a function of the actual content x_{nj}.

The key issue in the identification of the willingness to pay parameters is the degree to which beliefs about nutrient content in the absence of labels track actual nutrient content within product groups. Recall the example discussed in Section 2. In that hypothetical example, we observed in the data that consumption declined for foods with more calories per gram after a change in labeling. The results reported in Section 5 show that we observe this in the actual data as well. This observation needs to be combined with data on the change in beliefs about nutrient content to compute the elasticity of consumption with respect to a change in information (that elasticity normalized by price elasticities gives the willingness to pay for a change in nutrient content). If consumers had very accurate beliefs about nutrient content prior to the labeling law, then we would conclude that a small change in information led to the observed change in consumption which would imply a large willingness to pay; conversely, if consumers had inaccurate beliefs, then we would conclude that a large change in information led to the observed change in consumption which would imply a small willingness to pay.

Formally, I assume that the beliefs of non-label users and the beliefs of label users for unlabeled products can be written as an additive function of the average belief about nutrient content within product groups, the degree to which beliefs track actual nutrient content within product groups, and an idiosyncratic noise term. That is, for non-label users or unlabeled products:

$$E_{ijt}(x_{nj}) = E_g(x_{nj}) + a_g(x_{nj} - E_g(x_{nj})) + r_{ij} \tag{9}$$

The crucial issue in the identification of the willingness to pay parameters is the specification of the parameter a_g: this determines the degree to which consumers beliefs about nutrient content within product groups track the truth. I assume that label users (identified as individuals who "always" or "often" use nutrition labels) know the exact nutritional content of labeled foods. Ignoring the idiosyncratic error term, the change in information following a change in labeling is given by:

$$x_{nj} - E_{ijt}(x_{nj}) = (1 - a_{gn})(x_{nj} - E_g(x_{nj})) \tag{10}$$

As $a_{gn} \rightarrow 1$, consumers are fully informed about nutrient content prior to labeling, the change in information goes to 0, so the measured elasticity of consumption with respect to a change in information goes to infinity given the observed change in consumption.

Unfortunately, group-specific estimates of a_{gn} are not currently available, and for many nutrients, no estimates exist at all. In Appendix B, I estimate a_g based on a survey of consumers in Starbucks where consumers were explicitly asked their beliefs about the calorie content of food and drinks products they just purchased. In the estimates currently reported, I use the value $a_{gn} = 0.2$ calibrated from the Starbucks data for all product groups and nutrients; this reflects the fact that consumers have a very limited ability to distinguish the calorie content of food and drink products at Starbucks prior to labeling. In on-going work, I attempt to obtain separate estimates of a_{gn} for different product groups and nutrients via additional surveys.[8]

An important point to note is that the welfare benefits of the NLEA taking preferences as given do *not* depend on the parameter a_{gn}. The welfare benefits depend only on the elasticity of demand, the observed change in consumption and the degree of heterogeneity in information across individuals; the specification of prior information determines whether these benefits arise because consumers care a lot about nutrient content given a small change in information or whether consumers care just a little about nutrient content given a large change in information. The latter question is of interest here because we want to compare the implicit value of health in food consumption decisions with the value of health estimated in other contexts.

The specification of $E_g(x_{nj})$ does not impact the willingness to pay estimates directly due to the presence of product group level fixed effects, but it does impact the model's projections for the impact of labeling on consumption and the welfare analysis. In the specifications reported in the main text, $E_g(x_{nj})$ is estimated using dummy variables for each product group x label status x pre/post NLEA. This value is identified based on differential changes in the consumption of foods which experienced a change in labeling compared to those which did not.

6.4 Estimates of the Willingness to Pay for Nutrient Content

In this subsection, I report estimates of the willingness to pay for nutrient content. The main results are that I estimate a small but significant willingness to pay to avoid calories which appears to be due mostly to a willingness to pay to avoid fat; without more data on

[8] One particular worry is that a_{gn} might vary across product groups systematically depending on the proportion of products already labeled. For example, if the proportion of products labeled goes from 0% to 40%, then consumers may learn something about the nutrient content of currently unlabeled foods. Such contextual inferences would be unproblematic for the identification of the α_n parameters if they affected only the overall systematic bias in beliefs about unlabeled foods ($E_g(x_{nj})$), but they would be problematic if more labeling lead to improved discernment of the differences between unlabeled foods. For example, if a product group contains both yogurt and cream cheese, one might worry that labeling of some yogurt and some cream cheese product informs people more generally about the nutrient content of unlabeled yogurt and cream cheese products (making a_{gn} larger in those groups). To investigate this issue, I define sub-product groups based on the first two characters of the food code identifier and consider a specification with separate year specific fixed effects for each sub-product group. In this specification, reported in Appendix D, the willingness to pay parameters are identified only via within sub-group estimation, and the results do not qualitatively differ from the specification in the main text.

beliefs prior to the NLEA we cannot determine whether this is due to avoidance of saturated fat or unsaturated fat.

Before discussing the actual results, let us consider a back of the envelope calculation to see what the results discussed so far imply about the willingness to pay for a change in calorie content. The reduced form results suggest that a difference in actual nutrient content of 1 calorie/gram leads to a decline in consumption of 0.2 calories or about 20% when a product group goes from having no labels to having all foods labeled (1 calorie is the average consumption since most foods in the data are not consumed by the vast majority of consumers). The average price elasticity in the data is close to 0.5, so a 20% reduction in demand would be induced by a 40% increase in price. The average price of a gram of food in 1988 dollars is about 0.2 cents,[9] so a 40% increase in price corresponds to about .08 cents/gram. Given $a_g = 0.2$, when two foods differ by 1 calorie/gram in actual nutrient content, consumers will update their beliefs about this difference by 0.8 calories/gram after labeling. So if consumers respond to a change in information of 0.8 calories per gram the same way they respond to a price increase of .08 cents/gram, this implies a willingness to pay for a change in calorie content of about 0.1 cents/calorie.

Table 5 reports the estimates from the Tobit model described in equations (7) and (8). All of the models include fixed effects for each product as well as dummies for each (product group x year) and product-specific linear time trends. The reported models assume that the healthiest foods label within product groups when the proportion labeled is less than 100%. I assume further that a_{gn} in equation (10) is equal to 0.2, and I estimate $E_g(x_{nj})$ (the average prior for each product group) by comparing the change in consumption for foods which experience a change in labeling to foods which experience no change. Finally, I instrument for perceived nutrient content with group-specific instruments constructed by interacting group fixed effects with a post-NLEA dummy variable.

Model 1 from Table 5 reports estimates of the coefficients on calories, cholesterol and sodium. The coefficient on calories implies willingness to pay of 12 hundredths of a cent per calorie, very close to the value implied by the back of the envelope calculation above. In other

[9] Annual per capita expenditures in 1988 are about \$2000 according to the USDA, which translates to about \$5.50 / day for 2500 grams or about \$.002 per gram, consistent with the price data used.

words, when individuals learn that a food product has 100 more calories than they previously thought, their consumption decisions are impacted in the same way as an 12 cent increase in price. The model also estimates a marginally significant willingness to pay to avoid sodium. Model 2 disaggregates calories into total fats, protein, non-fiber carbohydrates, and fiber. This coefficient on total fats is large and significant, implying a willingness to pay of 67 hundredths of a cent per gram of fat and 52 hundredths of a cent per gram of carbohydrates. The point estimates also imply a negative willingness to pay for protein and a positive willingness to pay for fiber; these estimates are not significantly different from zero, but the standard errors are large. Note that a negative coefficient on proteins is not necessarily anomalous: if individuals are calorie conscious and use nutrition labels to avoid high-calorie foods, this would tend to produce a negative coefficient on proteins, carbohydrates and fiber even if individuals ignored information specifically about those nutrients.

Models 3 and 4 further disaggregate total fats into unsaturated fats and saturated fats. There is some uncertainty regarding the appropriate specification of this model since only total fats were listed on nutrition labels prior to the NLEA. It is therefore unclear what information - if any - we should assume label users possessed about the content of saturated or unsaturated fats in labeled foods prior to the NLEA. Model 3 assumes that the content of unsaturated and saturated fats was unknown even for labeled products prior to the NLEA while Model 4 assumes that both types of fat content were exactly known prior to the NLEA. The estimates in Model 3 imply that the willingness to pay to avoid fats comes mainly from unsaturated fats; the estimates in Model 4 implies that it comes from saturated fats, although the standard errors are larger. The results suggest that labeling lead to a reduction in fat consumption, but lacking information on consumers' beliefs regarding the specific content of unsaturated and saturated fats in labeled foods prior to the NLEA, we cannot determine whether this arose through avoidance of saturated or unsaturated fats.

I report a number of additional specifications as robustness checks in Appendix D. The main takeaway from these results is that structural model implies fairly small willingnesses to pay for nutrient information despite the fact that labeling appears to have had non-neglible impact on consumption.

7 Welfare Taking Preferences as Given

The results reported so far suggest that labeling did impact calorie consumption, but that the magnitude of the response to information about nutrient content can only be consistent with the medical evidence (given usual estimates of the VSL) if individuals have information from other sources which allows them to evaluate the health consequences of different diets. In this section, I use the estimated model to evaluate the welfare consequences of these claims taking estimated preferences as given. I ask what the estimated parameters imply about the partial equilibrium welfare gain from the NLEA and about the potential welfare gain from additional labeling regulations which would require mandatory labeling in restaurants and of fresh meats and vegetables. This exercise is of interest in its own right, and in the next section, it will serve as a baseline to which we can compare the results of the behavioral model which asks what these welfare benefits would be if consumers were fully informed.

As noted in section 6.3, the welfare benefits computed in this section do not depend on the specification of prior information. The welfare benefits depend only on the elasticity of demand and the observed change in consumption; the specification of prior information determines whether these benefits arise because consumers care a lot about nutrient content given a small change in information or whether consumers care just a little about nutrient content given a large change in information.[10]

7.1 Welfare Impact of the NLEA Taking Preferences as Given

The impact of labeling on welfare is most transparent if we consider a generalized price, $P_{jt} = \frac{\mu_{it}}{\phi} + p_{jt} - \sum_n \frac{\alpha_n}{\phi} E_{ijt}(x_{nj})$. This is the marginal cost of consumption in the structural model above (normalized by the marginal utility of money). The cost of an additional gram of product j depends on p_{jt}, the actual price, $\frac{\mu_{it}}{\phi}$, the cost of foregone satiation in dollar terms, and $\sum_n \frac{\alpha_n}{\phi} E_{ijt}(x_{nj})$, the cost of consuming the nutrients in that gram (for foods with desirable nutrients, this could be positive, reducing the generalized price). When foods are

[10]To the extent that the change in consumption cannot be separately estimated for each individual, the specification of prior information would impact the welfare calculation insofar as it is used to infer the degree to which the observed aggregate change in consumption varies across individuals; this point is discussed further below.

analysis in this section more transparent; that is, assume that $N_{ij} = a_{ij} - b_{ij}P_{ij}$. I will then repeat the analysis in the full structural model developed above. In the case of a linear demand curve, the welfare gain from labeling for a given product is given by: $W_{ij} = \frac{1}{2}|\Delta N_{ij}||\Delta P_{ij}| = \frac{1}{2}|b_{ij}|(\Delta P_{ij})^2$. Consider the simplest possible case in which only information about calorie content changes (so in particular, ignore any changes in information about other nutrients, changes in the multiplier, or changes in monetary price). In this case, $\Delta P_{ij} = \frac{\alpha_n}{\phi}\Delta c_{ij}$ where Δc_{ij} denotes the change in calorie content per gram. So in the simplest possible case, the welfare gain from labeling is given by:

$$\Delta W = \frac{1}{2}\left(\frac{\alpha_n}{\phi}\right)^2 \sum_j |b_{ij}|(\Delta c_{ij})^2 \tag{11}$$

In other words, the gain from labeling is proportional to the square of the price of a calorie (the previously estimated 7 hundredths of a cent) and the sum of squared deviations of perceived calories / gram from actual calories per gram weighted by the price responsiveness of each good (note that b_{ij} will tend to scale with the quantity of a food consumed, since the responsiveness of grams consumed to price per gram will be larger for foods with a larger serving size). In the more general case in which all nutrients are taken into account and in which the multiplier adjusts as well in response to labeling, $\Delta P_{ij} = \frac{\Delta\mu_i}{\phi} - \sum_n \frac{\alpha_n}{\phi}\Delta x_n$. This yields:

$$\Delta W = \frac{1}{2}\sum_j |b_j|(\frac{\Delta\mu_i}{\phi} - \sum_n \frac{\alpha_n}{\phi}\Delta x_n)^2 \tag{12}$$

The demand curve in the structural model is not a linear function of prices. I show in the online appendices that the welfare gain from labeling in this model is given by:

$$\Delta W = \sum_j \Big(CS(\hat{a}) - \hat{N}_{ij}(p^*_{jt} - \hat{p}_{jt})\Big) \tag{13}$$

where $CS(a) = -\frac{\sigma_j}{2\beta_j}\left(\Phi(a)a^2 + \Phi(a) - \phi(a)a\right)$, and $\hat{p_j}$ as the price at which consumption given full information would equal consumption when labels are not present, $\hat{N}_{ij}$ gives consumption in the world where labels are not present, $\hat{Y}_{jt}$ gives the predicted value of the latent variable in that world, and $\hat{a} = \hat{Y_{ijt}}/\sigma_j$.

Table 6 gives the results of this analysis. In each year, I compute the welfare gain from all of the labeling which has occurred since 1985 in Model 1 (with calories, sodium and cholesterol). In the years following the NLEA, labeling leads to a welfare gain of $0.07-$0.11 per day, or about $28-$40 annually in the structural model and about $50-$60 annually in the linear model. The table also decomposes the welfare benefits into the direct benefits - the benefits from foods which experienced a change in labeling - and the indirect benefits from foods which experienced no change in labeling but for which consumption changed due to substitution. About 5% of the estimated welfare benefits come from the latter type of foods.

The FDA estimated the total cost of NLEA implementation to be between 1.4 billion and 2.3 billion dollars (Food and Administration 1993). Industry estimates were slightly higher; the National Food Processors Association estimated compliance costs from 3.3-4 billion dollars (Van Wagner 1992).[11] Taking just the sample population studied in this paper - females aged 19-50 - implies 46.2 million label users in 1993 (62 million women times 75% who use labels), which in turn implies an annual welfare gain of roughly 1-2 billion dollars. This estimate is not a complete welfare analysis since it ignores general equilibrium factors, as well as other impacts of the NLEA such as nutrition claims legislation. Nonetheless, it suggests that the information provision element of the NLEA was successful judged from a revealed preference standpoint; it paid for itself in the first few years of the program.

7.2 Potential Welfare Gains from Additional Labeling

We can also use the model to evaluate the welfare gains from additional nutrition labeling. The NLEA exempted all fresh foods from labeling regulations, including all restaurants and fresh meat and vegetables sold in grocery stores. Recently, the Affordable Health Care Act mandated calorie posting in all chain restaurants (Rosenbloom 2010). In this section, I ask

[11] This estimate does not take into account the opportunity cost of "package real estate" taken up by nutrition labeling to the exclusion of other advertising or information. This cost is difficult to quantify. An informal survey suggests that most packages have space in the back or side where additional information could be included if it were valuable, which suggests that this opportunity cost is not too large (it still may not be zero because even information with positive value would not be provided if the aesthetic costs were too large).

question: what are the potential welfare gains from labeling all foods which are currently unlabeled?

This exercise requires some additional assumptions. I gloss over distinctions between different types of labeling: it may be that individuals respond differently to calories posted on restaurant menus than to nutrition facts on the back of packaging; the rough agreement between my results and the results in Bollinger et al. (2010) suggest that the responses may not be too different, but that is far from convincingly established (this agreement is discussed in Appendix A). I also do not have information about the proportion of restaurants which voluntarily provide calorie information. Many restaurants provide nutritional information either in a booklet or on their website, but surveys suggest that this information is rarely used and that despite its presence, individuals do not have accurate beliefs about the nutritional content of alternative products (Bollinger et al. 2010). The majority of consumers do report using calorie information when it is posted as prominently as prices (Elbel, Kersh, Brescoll, and Dixon 2009).

The FLAPS survey indicates that 60% of fresh meats and vegetables currently carry nutrition labeling. Data does not appear to exist on either the proportion of restaurants which currently use prominent nutrition labeling of any kind; I assume that it is 0%. Given these assumptions, the analytical framework is otherwise identical to that in the Section 7.1: the demand curve with labeling is computed, and the welfare gain from labeling comes from the fact that individuals purchase the wrong quantity of each food if they are not properly informed about nutrient content.

Table 7 gives the results of this analysis and compares the potential welfare gain from more labeling to the welfare gain from the NLEA for each of the four models reported in Section 6.4. The welfare gain ranges from $40-$80 annually per person and is typically larger than the welfare gain from the NLEA; this occurs despite the fact that the NLEA applies to a larger fraction of overall consumption because of the assumption that labeling in restaurants increases from 0% to 100%.

8 Welfare Gains with Consumer Errors

The welfare computations in the previous sections take the estimated willingnesses to pay for nutrient content at face value. In this section, I attempt to determine if the estimated parameters are consistent with the available medical evidence and VSLs estimated in other contexts given what consumers already know about health.

The normative implications of this comparison are not completely clear-cut. At the very least, to the extent that VSLs differ across settings, there is a positive puzzle as to why the marginal willingness to trade-off money and expected life-years differs.[12] I will also consider the implications of a strong normative assumption: what gains from policy are possible if we assume that the VSL estimated in other settings is a better guide to consumers' best interests than the VSL estimated from food consumption? The plausibility of this normative assumption will depend on the answer to the positive question as to why VSLs differ across settings, as I discuss at greater length below.

With those caveats in mind, I proceed as follows. I begin by considering a thought-experiment which suggests that nutrient information plays little role in people's consumption decisions; if consumers are responsive to health in their food consumption decisions, it must be from other sources of information. I then lay out a theoretical model which makes explicit the role that assumptions about health information play in computing a normative benchmark for the estimated willingness to pay parameters. Finally, I operationalize that model and I use the benchmark parameters to reconsider the welfare exercises in the previous section and to evaluate the potential additional welfare gains if consumers' response to nutrient information were consistent with the medical knowledge of experts and the VSL from other contexts.

8.1 Implications of the Estimated Parameters for Healthier Diets

One way to assess the magnitude of the estimated willingness to pay parameters is to ask what they imply about the potential gross benefits of healthier diets ("gross" in the sense

[12] For clarification, I am defining the VSL here as whatever marginal value of a statistical life is implicit in consumers' food consumption decisions given the observed response to nutrient content.

of ignoring the off-setting cost from such diets being otherwise less desirable). If we take preferences as given, we are assuming that choices are optimal if consumers use nutrition labels so any potential health gains would be offset by the fact that consumers would enjoy healthier foods less all things considered; the net benefits of eating different foods must be negative. We can however still ask the following question: what is the welfare gain implied by the estimated parameters if consumers could continue to eat exactly the same foods, but the *nutrient profile* of those foods were altered so that it matched that of a much healthier diet? In other words, what would be the welfare gain implied by the estimated parameters if consumers could continue to eat pizza and cheeseburgers, but the nutrient intake were as if they consumed tofu and broccoli?

The main result in this section is that this gain is very small relative to the magnitude implied if we start with pre-existing estimates of the impact of diet on health and the value of life-years. This does not necessarily imply that consumers are undervaluing nutrient information: it may be that nutrient information is redundant given everything else that consumers know about the health consequences of diets.[13] I investigate this question in more detail in section 8.3 and by attempting to characterize consumers' beliefs about the health consequences of foods from all sources relative to those of experts. The exercise in this section demonstrates only that if individuals are adequately incorporating health information into their food consumption decisions, they are not doing it via information about nutrient content.

To conduct this exercise, I construct a profile of nutrients for the healthiest possible diet based on FDA recommendations and the relationship between nutrient intake and the theoretical minimum baseline dietary risk factors described in Ezzati et al. (2003) which is used in the introduction to compute the welfare gains from longer life; on a per capita basis, these gains range from \$2,500-\$4,500 year if we start with standard assumptions about the VSL and discount factor (or \$2000-\$3500 in 1990 dollars, in which prices in the data are

[13]In the context of the model estimated above, it may be that the product fixed effects incorporate judgments about health which are omitted from this exercise. That is, if we want to compute the true revealed preference gain from the counterfactual world in which cheeseburgers were as healthy as broccoli, we would first need to decompose the product fixed effect into the portion driven by health considerations and the portion driven by everything else. The model I consider in the next section attempts to do something like this.

expressed). In this section, I use the value of \$3,000 in 1990 dollars as a benchmark.

The benchmark healthy diet is described in Table 8. The recommendation for calorie intake is computed separately for each individual based on their height and self-reported exercise habits, and the rule that a BMI of 18.5 kg/m^2 minimizes health risk as in Ezzati et al. (2003).[14] The remainder of the recommendations are taken from the National Academy of Sciences Food and Nutrition Board based on their survey of the medical literature (Trumbo, Schlicker, Yates, and Poos 2002); they characterize for each nutrient a consumption range which minimizes health risk. The value of moving to this nutrient profile holding fixed all other aspects of foods can be computed straightforwardly from the willingness to pay estimates by multiplying the coefficient for each nutrient by the difference in nutrient intake between nutrients actually consumed and nutrients consumed in the healthy diet profile. The assumptions I make are conservative in the sense that this is a very extreme nutrient profile (e.g. zero consumption of saturated fats is not a realistic goal). It may be that a less extreme nutrient profile would yield most of the possible health gains, but the implied value of a less extreme nutrient profile would be even smaller, strengthening the result that the apparent willingness to pay for nutrients is insufficient to account for the value of the health consequences of alternative diets according to experts.

Since this calculation depends heavily on the change in the number of calories consumed, it is important to directly address the issue of missing data in the CSFII. In the numbers reported in this section, I scale total grams consumed so that mean calorie consumption matches a population average goal as estimated from the Behavioral Risk Surveillance Survey over the time period in question. I consider two alternative ways of performing this scaling: in the first case, I scale the consumption of all individuals in the data; in the second case, I select a sampling of individuals whose total consumption is consistent with the estimated BMI given the average energy requirements of a person of that BMI given their reported age and sex according to the USDA (Trumbo, Schlicker, Yates, and Poos 2002).

The results for Models 1-4 from the previous section for each of these cases are reported in Table 9. Depending on the model used, the welfare gains range from \$40-\$130 annually. This is 1-2 orders of magnitude smaller than the gains of \$2000-\$3500 annually computed

[14]For a sedentary female of average height, this translates to 1816 calories/day.

from the VSL and medical evidence. This suggests that if individuals are sensitive to the health consequences of their dietary behavior, nutrient information plays little direct role.

8.2 A Normative Theory of Food Consumption and Health

In this section, I develop a theory describing how we can compute a benchmark responsiveness to new health-relevant information given what consumers already know, assumptions about discount factors and the value of a statistical life, and medical evidence about the relationship between nutrient content and health.

Define a diet as a vector of grams consumed $\{N_1, ..., N_J\}$ for each of J foods in an individual's choice set, and let $h(d)$ denote the best prediction of the life-years gained from consuming diet d relative to a benchmark diet d_0. Consumers choose $\{N_{i1}, ..., N_{iJ}\}$ so that:

$$N_{ij} = \arg\max_{j} v(N_{i1}, ..., N_{iJ}) + \beta_i \tilde{h}_i(N_{i1}, ..., N_{iJ}) + \sum_n \alpha_{in} \left(E_{it}(X_{in})\right) - \sum_j p_j N_{ij} \qquad (14)$$

where $\tilde{h}_i(d)$ denotes consumer i's beliefs about the life expectancy consequences of consuming diet d relative to a benchmark, and β_i gives consumer i's value of a statistical life, defined as the marginal rate of substitution between expected life years and income (in this case, the marginal utility of income is normalized to 1). As in the structural model, I assume that α_{in} is piecewise linear, so that $\alpha_n(X_{in}) = \alpha_n \cdot \max(X_{in} - \bar{X}_{in}, \underline{X}_{in} - X_{in}, 0)$.

This is a more general version version of the structural model estimated above with one exception. In this specification, utility from taste and utility from all health considerations are written as additively separable; in the model above, both of these terms were encompassed in the "taste" term which indicated utility from all sources other than nutrient information. The benchmark responsiveness to new health information for each nutrient, the $\{\alpha^*_{in}\}$, are defined as the parameters which solve:

$$\arg\max_{\alpha_{in}} v(N^*_{i1}, ..., N^*_{iJ}) + \beta^*_i h(N^*_{i1}, ..., N^*_{iJ}) - \sum_j p_j N^*_{ij} \qquad (15)$$

where N^*_{ij} are solutions to equation 14 and β^*_i gives the normatively appropriate value of a statistical life (I discuss the relationship between β^*_i and β_i below).

In words, the benchmark parameters are defined as the additional responsiveness to nutrient information which would maximize utility given the true health consequences of alternative dietary behaviors and given the degree to which health is already taken into account. Think of a consumer with some existing knowledge of the health consequences of different diets ($\tilde{h}$). One way of conceptualizing the α_{in}^* parameters is to imagine that consumers have access to h at some point, but not at the time when they make their purchasing decisions. Instead, they can carry around in their memories a few parameters telling them how to best approximate the decisions they would make if they knew h given only the nutrient information available at the point of purchase. The α_{in}^* are the parameters they would carry around with them. Alternatively, we can think of these parameters as defined by a social planner: given what consumers know about health and given the true health consequences of alternative foods, the benchmark parameters give the optimal person-specific tax on nutrient content (assuming that the objective function in equation 15 is the appropriate normative standard).[15]

Provided total intake for a given nutrient lies in the costly range,[16] we can compute the solution to equation 15 by implicitly differentiating the objective function in equation 15 and then substituting in for $\frac{\partial v}{\partial N_{ij}^*}$ and $\frac{\partial N_{ij}}{\partial \alpha_{in}}$ from the first order conditions for equation 14. The benchmark parameters are then characterized by the system of linear equations:

$$(Sx)' q = 0 \tag{16}$$

where S is the J x J matrix of marginal price effects defined by $s_{kj} = \frac{\partial N_{ik}}{\partial p_j}$, x is the J x N matrix of nutrient contents where x_{jn} gives the content of nutrient n in one gram of product j and q is the J x 1 vector whose jth element is given by: $q_{ij} = \frac{\partial h}{\partial N_{ij}} - \frac{\partial \tilde{h}}{\partial N_{ij}} - \sum_n \alpha_{in}^* x_{nj}$.

[15] A more realistic formulation of the problem of optimal sin taxes might assume that the government is restricted to a single tax rate for all consumers, and that this tax rate cannot be perfectly conditioned on current overall nutrient intake. This is the subject of a separate investigation.

[16] If nutrient intake lies in the range with zero costs, I assume that individuals are unresponsive to nutrient information and the benchmark parameters are also 0. If optimal nutrient intake falls exactly on the cut-off point, the benchmark parameter for each nutrient is given by the value which leads total intake of that nutrient to exactly equal the cut-off value given the remainder of the structural parameters. Given the estimated utility functions, one can thus solve for the benchmark parameters in each of these three cases (optimal nutrient content in the costly range, on the cut-off point and in the range with zero costs) and then directly check which case maximizes utility.

I show in Appendix E that in this case we can write the benchmark parameters as a function of $\text{Cov}(\beta_i^* \frac{\partial h}{\partial N_{ij}} - \beta_i \frac{\partial \tilde{h}}{\partial N_{ij}}, \tilde{x}_{nj})$, and $E(\tilde{x}_{nj})E(\beta_i^* \frac{\partial h}{\partial N_{ij}} - \beta_i \frac{\partial \tilde{h}}{\partial N_{ij}})$ for each nutrient n where $\tilde{x}_{nj} = \sum_k \frac{\partial N_j}{\partial p_k} x_{nk}$ tells us how much a change in α_n will impact consumption of product j. Intuitively, suppose nutrient n is a "bad"; if consumers understate marginal health costs a lot for foods whose consumption is very sensitive to α_n, then being more sensitive to variation in nutrient n (α_n further from 0) will help correct for this understatement. In a world with no cross-price elasticities, this gives the intuitive condition that the benchmark parameters will be large in magnitude if the bias in marginal health costs is highly correlated with nutrient content.

In the following subsections, I compute these sufficient statistics and use them to compute the benchmark parameters. In the remainder of this section, I discuss the interpretation of the benchmark parameters, and in particular, the relationship between these parameters and the observed responsiveness to nutrient information. The model above assumes that nutrients only enter utility for health reasons, but this may be incorrect. For example, consumers may care about calorie consumption due to its impact on physical appearance regardless of the related health consequences. In this case, we could rewrite equation 15 as:

$$\arg\max_{\alpha_{in}} v(N_{i1}^*, ..., N_{iJ}^*) + \beta_i^* h(N_{i1}^*, ..., N_{iJ}^*) + \sum_n \delta_{in} X_{in} - \sum_j p_j N_{ij}^* \tag{17}$$

and the appropriate benchmark weights on X_n will be given by $\alpha_{in}^* + \delta_{in}$. I do not attempt to directly estimate δ_n below; instead, I argue using survey evidence that δ_{in} will typically have the same sign as α_{in}^*. In the case of calories, α_{in}^* will understate the degree to which individuals want to avoid calories because the vast majority believe their physical appearance would also be improved by consuming fewer calories. The value I compute is therefore a conservative benchmark; to the extent that the estimated willingness to pay for calorie content is still smaller in magnitude than α_{in}^*, I am potentially understating the degree to which individuals are less responsive than they should be to calorie information.

There are several additional reasons we might expect the estimated willingness to pay for nutrient content to differ from α_{in}^*. Consumers may be unsure how to map nutrient information into health consequences. When I attempt to estimate α_{in}^* below, I consider some survey evidence suggesting that consumers' beliefs about what constitutes a healthy diet differ from expert beliefs. One explanation for this discrepancy is that consumers are ignorant of expert beliefs about the relationship between diet and life expectancy and would change their beliefs to match expert beliefs if the latter could be communicated in a meaningful way. To the extent that this explanation is right, it would seem to support taking the benchmark parameters as the appropriate normative standard.

An alternative explanation is that even if consumers were fully informed about expert beliefs they would still respond differently from the benchmark that I calculate; that is, β_i, the full information VSL measured in the context of food consumption may differ from β_i^*, the full information VSL measured in other settings. This may occur for several reasons. Among others: consumers may distrust expert beliefs either for good reason (informed skepticism about the methodologies of nutritional epidemiologists or taking into account factors not considered below such as technological progress in treating diet related illness) or for bad reasons ("I know an old woman who ate a jar of lard every day and lived to be 120"); self-control issues may be especially relevant in the setting of food consumption; individuals may respond differently to many small decisions which lead to large health consequences than one large decision; and choices may depend on the entire distribution of mortality risk as opposed to just the expected number of life years. The normative implications of many of these explanations are uncertain (e.g. Bernheim and Rangel (2008) discusses several possible normative criteria in the case of time-inconsistency).

Below, I make the strong normative assumption that the VSL estimated in other settings is the normatively correct VSL, and that consumers err to the extent that the full-information VSL implicit in food consumptions decisions differs from this benchmark. Future work clarifying the positive explanation for the discrepancy between the estimated response to nutrient information and the benchmark value computed will shed light on the plausibility of this normative assumption.

8.3 Calculation of Benchmark Parameters

To calculate the benchmark willingnesses to pay for nutrient content, I begin by combining evidence from a survey of experts about the relative health ratings of different foods with evidence about the long-term health consequences of different diets in order to compute an estimate of the life-expectancy consequences of consuming a unit of each food which I express in dollar terms using the estimates of the value of a life-year. Next, I compute estimates of the parameters characterizing current consumer beliefs by minimizing the distance to expert beliefs subject to the constraint that consumer beliefs rationalize the judgment that the diet reported by consumers in a survey to be a healthy diet is as healthy as the benchmark healthy diet given by experts. I then use the estimated expert beliefs and consumer beliefs to compute the normative benchmark for the willingness to pay estimates by solving the system of linear equations characterized by equation 30.

My calculation of the marginal life expectancy consequences of consuming each food implied by medical evidence proceeds in four steps. First, I characterize a range of nutrient intakes which minimize health risks based on expert recommendations. Second, for all diets outside this range, I compute the distance from the benchmark healthy diet weighting each nutrient based on weights derived from a survey of experts. Third, I scale the difference between a given diet and the benchmark healthy diet into life years based on the assumption that the average diet leads to a loss of .04 life-years annually (this is justified by the calculation in online Appendix based on (World Health Organization 2002)). Fourth, I convert this into a dollar amount based on estimates of the distribution of the value of a statistical life. I start with an average VSL of \$6.4 million (Viscusi and Aldy 2003) and compute the value of each life year by assuming that there is a constant value of a life year and that the VSL is the present discounted value of all additional life years. That is, I solve, $E_i(\sum_{t=0}^{T_i(a_i)-a_i} \delta^t V_{year}) = V^*$ where the expectation is taken over all individuals in the data, a_i indicates age and $T_i(a_i)$ indicates life expectancy conditional on age a_i (this procedure is similar to that used in Gruber and Koszegi (2001). The value of the marginal life-year is given by $\delta^{T_i-a_i}V_{year}$. I assume $\delta = .96$. This value will vary across consumers based on their age, but it implies that the average consumer loses about \$3,000 worth of life-years by consuming their current diet rather than the healthiest possible diet (I consider below the impact of allowing for some heterogeneity in the VSL). The result of this calculation is a function which expresses the health cost of all diets in dollar terms which I can use to determine the marginal health cost of all foods. The benchmark healthy diet is again summarized in Table 8. More details of this calculation are given in Appendix E.

The next step in the calculation is to determine what individuals already know about the life expectancy consequences of consuming different diets. To perform this calculation, I make use of consumers' answer to the following question: "how many servings would you say a person of your age and sex should eat each day for good health from food group [X]?" Given a characteristic serving from each food group, we can use consumers' answer to this question to characterize their beliefs about an optimal diet. As a robustness check, I investigate whether consumers who list a greater proportion of servings from a given food group tend to consume healthier foods from that group; I find that this is not the case. Details are again given in Appendix E. In future work, I hope to improve the characterization of consumers current beliefs about life expectancy consequences by drawing on additional survey and choice evidence.[17]

The characterization of the dollar-equivalent true life expectancy consequences and the dollar-equivalent of consumers' beliefs about these consequences along with the matrix of price elasticities derived from the structural model allows me to compute the benchmark parameters from equation 30 separately for each consumer. The average of these benchmark weights across all individuals for each of the four specifications in Table 5 are given alongside the estimated willingnesses to pay for each nutrient in Table 10. The benchmark responsiveness to calorie information is about 9 times the estimated responsiveness. In the disaggregated models, the benchmark responsiveness to total fat content is again about 8-9 times the estimated coefficient, and the benchmark responsiveness to fiber, sodium and cholesterol are 2-7 times what is estimated depending on the specification.

One worry with this calculation is that the least healthy consumers may eat the foods they do precisely because they have a lower VSL. (Kniesner, Viscusi, and Ziliak 2010) estimates the degree of heterogeneity in the VSL and estimates a 10th percentile of $3.5 million, a median of $7.5 million, and a 90th percentile of $22 million. So if we conservatively assume

[17]While we cannot expect consumers to reliably report their current beliefs about health consequences of different foods expressed in life expectancy units, we can elicit beliefs about health consequences from all sources and examine how movement along this range of beliefs induced via exogenous information provision impacts consumption decisions relative to prices. That is, the same methodology used in this paper to compute a dollar-equivalent of nutrient information can be used to compute a dollar equivalent of health information from all sources.

the healthiest possible diet. What fraction of this $3,000 represent gains to consumer surplus that consumers can achieve by eating healthier foods? This depends on two factors: the willingness to substitute across foods (as captured by price elasticities) and the degree to which consumers already incorporate health information into their food consumption decisions (as captured by the disparity between α_n^* and α_n). Because consumption is not infinitely elastic, consumers will not immediately switch to the healthiest possible diet even if they understand its impact on life expectancy. The welfare gains from healthier eating are determined in part by the elasticity of substitution: the larger the elasticity, the more readily consumers will substitute towards healthier foods. The welfare gains are also determined partly by the degree to which consumers already incorporate nutrient information into their food consumption decisions; this is given by α_n. Finally, the welfare gains depend on the degree to which consumers already incorporate health information from other sources. This was taken into account in the calculation of α_n^* through the function $\tilde{h}$. If consumers already accurately appraise the health consequences of different foods, then α_n^* will be small. $3,000 dollars annually represents the welfare gain that would be achievable if consumption were infinitely elastic, if $\alpha_n = 0$, and if α_n^* were computed assuming $\tilde{h} = 0$ (that is, assuming that beliefs about health currently play no role in consumer demand). The analysis in this section considers the welfare gain when α_n, $\tilde{h}$ and price elasticities are estimated from data.

The analysis here follows closely the welfare analysis laid out in Section 7. The estimated model indicates the choices individuals actually make given their perceived generalized price. For each food, the benchmark willingness to pay parameters determine a true generalized price based on the nutrient profile, and we can compute the potential welfare gain from better choices by asking what consumption would be for each food if consumers faced the true generalized price rather than the price they actually face. When we reevaluate the benefits of the NLEA, three generalized prices are relevant: the price consumers face in the counterfactual world with 1985 labeling, the price consumers face given labeling and the weights they actually attach to nutrients, and the price consumers would face given labeling if they weighted nutrients correctly according to the benchmark parameters. Figure 2 describes how consumer surplus is computed for each product given these prices.

Table 11 reports the results of these additional analyses. Using the benchmark preference parameters increases the estimated welfare gain from labeling by a factor of 4-5; the gains from eating healthier foods are substantially larger when assessed using these benchmark parameters. The model implies that consumers could gain over a thousand dollars in additional surplus per year if their assessment of the marginal cost of consumption for each food were perfectly in accordance with what is implied by the estimated benchmark preferences. Given the assumption that the gross health benefit of consuming the healthiest possible diet is $3000 per year, this number suggests that 30-40% of the possible gross health gains from consuming a healthier diet could be internalized as consumer surplus.

9 Conclusion

Diet has important consequences for health; this paper analyzes whether individuals incorporate these consequences into their food consumption decisions using information about nutrient content from nutrition labels. I find that food consumption decisions do appear to respond to labeling laws and that while the estimated responses are modest, such laws easily pass a cost-benefit test and additional labeling would likely lead to welfare gains. The magnitude of the response is nonetheless far too small given the health consequences of different diets and the fact that nutrient information should be valuable in computing those health consequences given what consumers already know about the value of those health consequences from other contexts. Accounting for the facts that healthy foods are otherwise less desirable and that consumers already have some information about health, the net benefit to consumers possible from consuming healthier foods is 30-40% of the value of the gross health benefit from switching to the healthiest possible diet.

The value reported here for the potential welfare gains from consuming healthier foods is by no means definitive: instead, it is a first pass attempt to apply the model developed in Section 8.2 to evaluate the magnitude of the welfare gains from more informed choices. The normative assumptions made could be better assessed given an explanation of why the estimated willingnesses to pay for nutrient information differ from the normative benchmark. There are many possible explanations ranging from incomplete understanding of the relationship between nutrient content and health to distrust of expert information, time inconsistency, and other contextual and framing effects which impact the way that consumers respond to health information in this setting relative to other settings. Ideally, choice evidence from experiments could be used to replace any reliance on survey evidence in forecasting how individuals would choose under ideal circumstances. The magnitude of the disparity identified in this paper between the observed responsiveness to nutrient information and the normative benchmark provides a standard which can be used to assess the results of such experiments: how much of the disparity is explained by each of these factors alone and in concert?

Beyond provision of nutrient information, there are several additional instruments a social planner might consider to alter dietary behavior; these include nutritional education programs, “nudges” such as listing healthier foods earlier on menus (Downs, Loewenstein, and Wisdom 2009), or more paternalistic measures such as taxes or subsidies on certain types of foods or even outright bans. One can view the model in this paper as an extension of the model developed in O’Donoghue and Rabin (2006) to evaluate the welfare consequences of such policies. In ongoing work, I use the model to evaluate Pigouvian taxes on negative nutrients and subsidies for positive nutrients designed to correct for the apparent underresponsiveness of consumers to this information.

More generally, the approach developed in this paper can be used to analyze any choices where some desired characteristic of alternatives is observable and can be independently priced. Food consumption decisions depend in large part on unobservable taste parameters, but they also depend on perceptions of healthiness, and we have estimates of the value of health from other settings. We can then ask: are consumers aware of the variation in this desirable characteristic when they make their choices? To the extent that this is not the case, if expert knowledge can be used to determine the amount of the desired characteristic in the available alternatives, we can determine the scope for welfare improvements from policies designed to lead to better choices. Whether the policy under consideration is a relatively innocuous information provision, a “nudge” from alternative choice architecture, a tax or subsidy or a more stringent restriction on choice, it is of interest to determine quantitatively the magnitude of the potential benefits from better choices so that these can be properly weighed against the costs.

References

Abaluck, J. and J. Gruber (2009). Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program. *NBER Working Paper*.

Allcott, H. and N. Wozny (2010). Gasoline Prices, Fuel Economy, and the Energy Paradox. *Working Papers*.

Andreyeva, T., M. Long, and K. Brownell (2010). The impact of food prices on consumption: a systematic review of research on the price elasticity of demand for food. *American Journal of Public Health 100*(2), 216–22.

Balasubramanian, S. and C. Cole (2002). Consumers' Search and Use of Nutrition Information: the Challenge and Promise of the Nutrition Labeling and Education Act. *The Journal of Marketing*, 112–127.

Bernheim, B. and A. Rangel (2008). Choice-theoretic foundations for behavioral welfare economics. *The Methodologies of Modern Economics. Oxford University Press, forthcoming*.

Beshears, J., J. Choi, D. Laibson, and B. Madrian (2008). How are preferences revealed? *Journal of Public Economics 92*(8-9), 1787–1794.

Bhattacharya, J. and N. Sood (2006). Health insurance and the obesity externality. *Advances in Health Economics and Health Services Research 17*(1), 279–318.

Bollinger, B., P. Leslie, and A. Sorensen (2010). Calorie Posting in Chain Restaurants. *NBER Working Paper*.

Chay, K. and J. Powell (2001). Semiparametric censored regression models. *Journal of Economic Perspectives 15*(4), 29–42.

Chevalier, J., A. Kashyap, and P. Rossi (2003). Why don't prices rise during periods of peak demand? Evidence from scanner data. *American Economic Review 93*(1), 15–37.

Cutler, D., E. Glaeser, and J. Shapiro (2003). Why have Americans become more obese? *Journal of Economic Perspectives 17*(3), 93–118.

Downs, J., G. Loewenstein, and J. Wisdom (2009). Strategies for promoting healthier food choices. *American Economic Review 99*(2), 159–164.

Elbel, B., R. Kersh, V. Brescoll, and L. Dixon (2009). Calorie labeling and food choices: a first look at the effects on low-income people in New York City. *Health Affairs 28*(6), w1110.

Ezzati, M., S. V. Hoorn, A. Rodgers, A. D. Lopez, C. D. Mathers, and C. J. Murray (2003). Estimates of global and regional potentil health gains from reducing muliple major risk factors. *The Lancet 362*(9380), 271 – 280.

Food and D. Administration (1993). Regulatory impact analysis of the final rules to amend the food labeling regulations. *Federal Register 50*(3).

Fulgoni III, V., D. Keast, and A. Drewnowski (2009). Development and Validation of the Nutrient-Rich Foods Index: A Tool to Measure Nutritional Quality of Foods. *Journal of Nutrition 139*(8), 1549.

Greene, W. (2001). Estimating econometric models with fixed effects. *NYU Leonard N. Stern School of Business Department of Economics Working Paper*.

Greene, W. (2004). Fixed effects and bias due to the incidental parameters problem in the Tobit model. *Econometric Reviews 23*(2), 125–147.

Grossman, M. (1972). On the concept of health capital and the demand for health. *Journal of Political economy 80*(2).

Gruber, J. and B. Koszegi (2001). Is Addiction Rational? Theory and Evidence.. *Technology 1261*.

Honoré, B. E. (2008). On marginal effects in semiparametric censored regression models. Princeton University.

Ippolito, P. and A. Mathios (1990). Information, advertising and health choices: a study of the cereal market. *The RAND Journal of Economics 21*(3), 459–480.

Ippolito, P. and A. Mathios (1993). New food labeling regulations and the flow of nutrition information to consumers. *Journal of Public Policy & Marketing 12*(2), 188–205.

Ippolito, P. and A. Mathios (1995). Information and Advertising: The Case of Fat Consumption in the United States. *The American Economic Review*, 91–95.

Kiesel, K. and S. Villas-Boas (2009). Another Nutritional Label-Experimenting with Grocery Store Shelf Labels and Consumer Choice. *UC Berkeley: Department of Agricultural and Resource Economics, UCB. Retrieved from: http://www. escholarship. org/uc/item/6st6d0rr*.

Kniesner, T., W. Viscusi, and J. Ziliak (2010). Policy relevant heterogeneity in the value of statistical life: New evidence from panel data quantile regressions. *Journal of Risk and Uncertainty 40*(1), 15–31.

Livingstone, M. and A. Black (2003). Markers of the validity of reported energy intake. *Journal of Nutrition 133*(3), 895S.

Martin, J., J. Beshears, K. Milkman, M. Bazerman, and L. Sutherland (2009). Modeling Expert Opinions on Food Healthfulness: A Nutrition Metric. *Journal of the American Dietetic Association 109*(6), 1088–1091.

Mathios, A. (2000). The Impact of Mandatory Disclosure Laws on Product Choices: An Analysis of the Salad Dressing Market*. *The Journal of Law and Economics 43*(2), 651–678.

Mojduszka, E. and J. Caswell (2000). A test of nutritional quality signaling in food markets prior to implementation of mandatory labeling. *American Journal of Agricultural Economics 82*(2), 298.

Moorman, C. (1996). A Quasi Experiment to Assess the Consumer and Informational Determinants of Nutrition Information Processing Activities: The Case of the Nutrition Labeling and Education Act. *Journal of Public Policy & Marketing*, 28–44.

O'Donoghue, T. and M. Rabin (2006). Optimal sin taxes. *Journal of Public Economics 90*(10-11), 1825–1849.

Rolls, B. J. (2009). The relationship between dietary energy density and energy intake. *Physiology & Behavior 97*(5), 609 – 615. Proceedings from the 2008 Meeting of the Society for the Study of Ingestive Behavior - SSIB 2008, Proceedings from the 2008

Meeting of the Society for the Study of Ingestive Behavior.

Rosenbloom, S. (2010). Calorie data to be posted at most chains. *The New York Times*.

Sawaya, A., K. Tucker, R. Tsay, W. Willett, E. Saltzman, G. Dallal, and S. Roberts (1996). Evaluation of four methods for determining energy intake in young and older women: comparison with doubly labeled water measurements of total energy expenditure. *American Journal of Clinical Nutrition 63*(4), 491.

Trumbo, P., S. Schlicker, A. Yates, and M. Poos (2002). Dietary reference intakes for energy, carbohydrate, fiber, fat, fatty acids, cholesterol, protein and amino acids. *Journal of the American Dietetic Association 102*(11), 1621–1630.

Van Wagner, L. R. (1992). On the road to nlea. *Food Processing*.

Variyam, J. (2008). Do nutrition labels improve dietary outcomes? *Health Economics 17*(6), 695–708.

Variyam, J. and J. Cawley (2006). Nutrition Labels and Obesity. *NBER Working paper*.

Viscusi, W. and J. Aldy (2003). The value of a statistical life: a critical review of market estimates throughout the world. *Journal of Risk and Uncertainty 27*(1), 5–76.

Willett, W. (1998). *Nutritional epidemiology.* Oxford University Press, USA.

Wooldridge, J. (2002). *Econometric analysis of cross section and panel data.* The MIT press.

World Health Organization, W. (2002). The World Health Report 2002: Reducing Risks, Promoting Healthy Life.

Figure 3: The Sample Nutrition Facts Panel Provided by the USDA (post-NLEA)

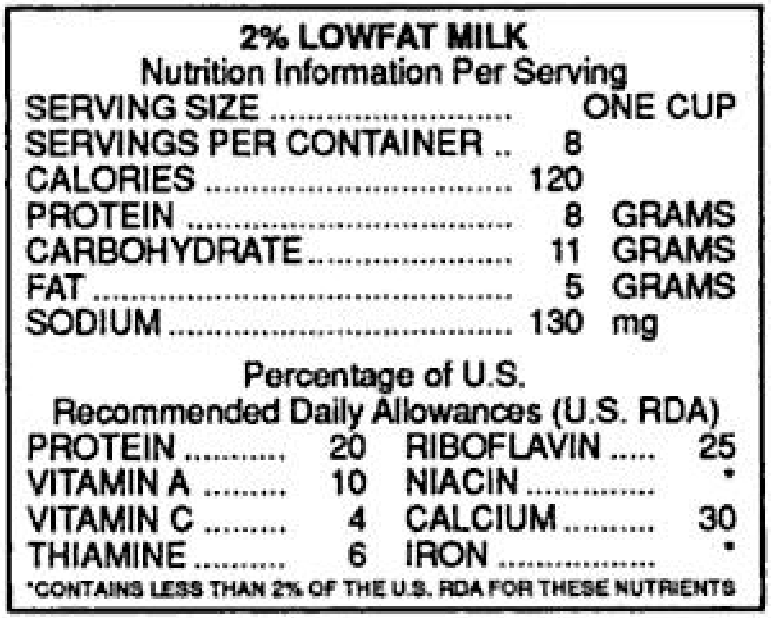

2% LOWFAT MILK
Nutrition Information Per Serving

SERVING SIZE		ONE CUP
SERVINGS PER CONTAINER	8	
CALORIES	120	
PROTEIN	8	GRAMS
CARBOHYDRATE	11	GRAMS
FAT	5	GRAMS
SODIUM	130	mg

Percentage of U.S.
Recommended Daily Allowances (U.S. RDA)

PROTEIN	20	RIBOFLAVIN	25
VITAMIN A	10	NIACIN	*
VITAMIN C	4	CALCIUM	30
THIAMINE	6	IRON	*

*CONTAINS LESS THAN 2% OF THE U.S. RDA FOR THESE NUTRIENTS

Figure 4: A Sample Pre-NLEA Nutrition Label

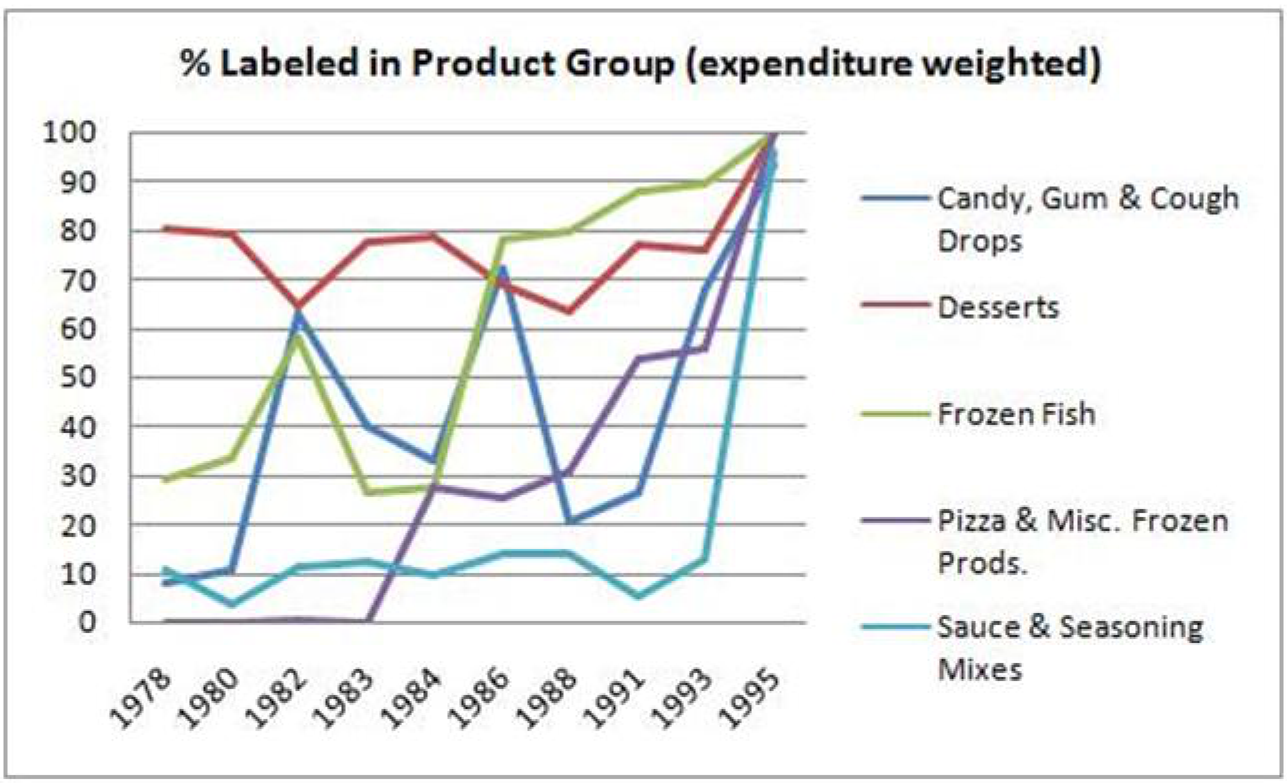

Figure 5: Proportion of Products Labeled by Year for a Sample of Product Groups

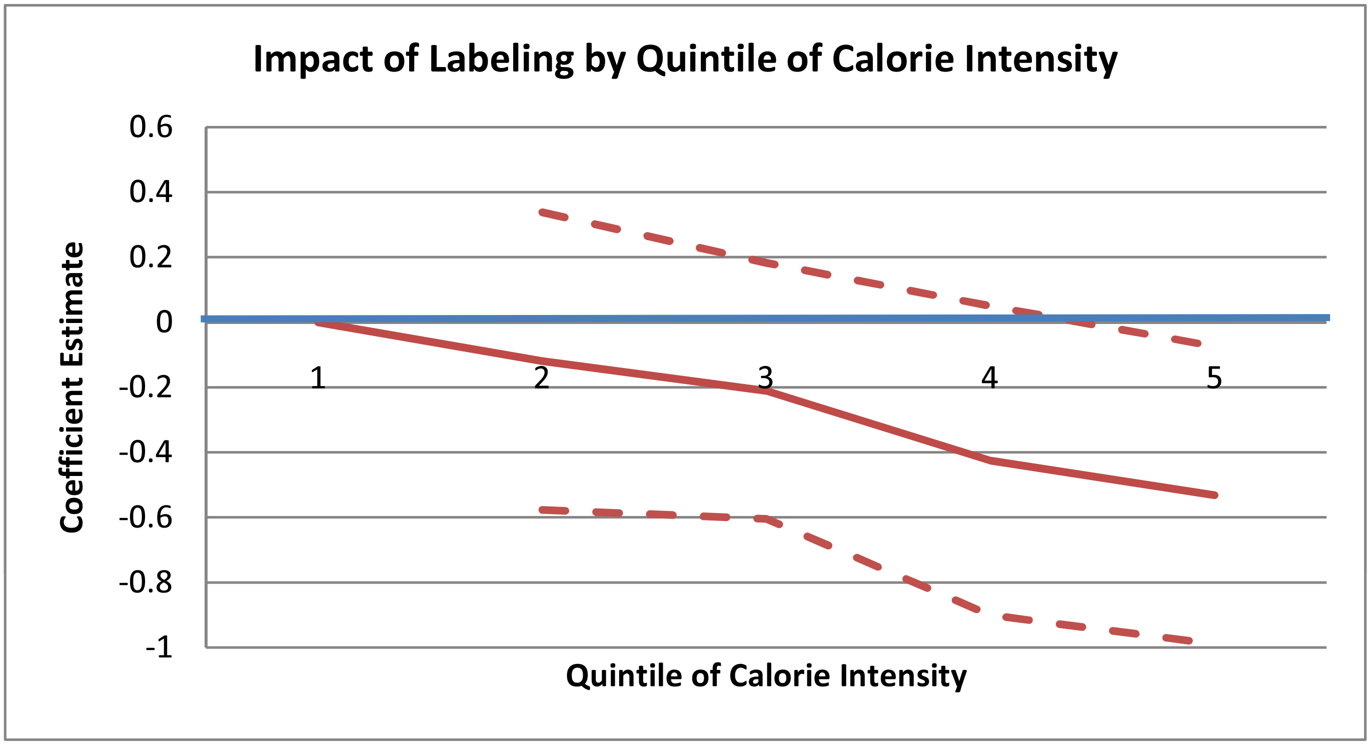

Figure 6: This graph shows the coefficients from estimating equation 1, but substituting quantiles of calorie intensity interacted with the proportion of products labeled for the independent variable. The dotted lines indicate 95% confidence intervals. The coefficient can be interpreted as the change in consumption in the nth quintile relative to the first quintile when the proportion of products labeled in a product group increases from 0% to 100%. A coefficient of -0.5 indicates that consumption falls by 0.5 calories in the quintile of interest relative to the first quintile when labeling increases by 0% to 100%.

Table 1: Sample Product Group

	Label Status		Calories / gram		Consumption (grams)	
Food	Period 0	Period 1	Prior Belief	Actual Content	Period 0	Period 1
Dressing 1	Labeled	Labeled	1	1	20	24
Dressing 2	Not Labeled	Labeled	3	1	20	26
Dressing 3	Not Labeled	Labeled	3	3	20	24
Dressing 4	Not Labeled	Labeled	3	5	20	22

Table 2: Percentage of DHKS Respondents Indicating Use of Nutrient on Label

	Nfacts Panel	Calories	Total Fats	Sat. Fats	Fiber	Sugars	Sodium	Cholesterol	Vitamins
1989	74	-	-	-	-	-	-	-	-
1990	-	68	72	-	53	68	66	71	62
1991	-	68	74	-	53	66	60	68	61
1994	77	79	81	70	54	67	67	69	67
1995	79	81	82	70	54	67	66	66	68
1996	73	81	82	71	50	64	63	63	66
Total	75	75	78	70	53	66	64	68	64

Each value gives the percentage of DHKS respondents in the indicated year who indicated that they "Often" or "Sometimes" used the information on the indicated nutrient on the nutrition facts panel out of all respondents indicating "Often", "Sometimes", "Rarely" or "Never". In 1989, the question asking respondents about use of the Nutrition Facts panel (the first column) replaced the "Often" choice with "Always".

Table 3: Beliefs about Current Nutrient Consumption Relative to What is Healthy

	Calories	Total Fats	Sat. Fats	Fiber	Protein	Sodium	Cholesterol
Too Much	51	60	50	3	8	32	41
Too Little	6	4	5	47	19	6	4
Right Amount	43	36	46	50	72	62	55

Each value gives the percentage of DHKS respondents among those who indicated whether they consumed "too much", "too little" or "about the right amount" of the nutrient in question compared to what is healthy.

Table 4: Reduced Form Evidence: Labeling and Calorie Consumption

	Linear OLS	Linear OLS	Linear OLS	Linear 2SLS	Linear 2SLS	Linear 2SLS
Calories · Labeling	-.0924**	-.1349**	-.2214*	-.1032**	-.1289**	-.2271**
	(.0215)	(.0459)	(.1027)	(.0204)	(.0428)	(.0859)
Product F.E.	YES	YES	YES	YES	YES	YES
Year F.E.	YES	-	-	YES	-	-
Group-year F.E.	NO	YES	YES	NO	YES	YES
Product Time Trends	NO	NO	YES	NO	NO	YES
Observations	8419	8419	8419	8419	8419	8419
Sample	NLEA	NLEA	NLEA	NLEA	NLEA	NLEA
	Log 2SLS	Log 2SLS	Log 2SLS	Linear 2SLS	Linear 2SLS	Linear 2SLS
Calories · Labeling	-.0393	-.2159**	-.3875**	-.0026	.1790**	-.0228
	(.0262)	(.0551)	(.1157)	(.0166)	(.0363)	(.0758)
Product F.E.	YES	YES	YES	YES	YES	YES
Year F.E.	YES	-	-	YES	-	-
Group-year F.E.	NO	YES	YES	NO	YES	YES
Product Time Trends	NO	NO	YES	NO	NO	YES
Observations	8419	8419	8419	7534	7534	7534
Sample	NLEA	NLEA	NLEA	No NLEA	No NLEA	No NLEA

* indicates significance at the 5% level and ** indicates significance at the 1% level. Each observation is a food-year. The dependent variable is average consumption of food j in calories at time t (including 0's) in the linear specification, and the natural log of that in the log specifications. The dependent variable is the interaction between calories per gram (a constant for each product) and the proportion of products which are labeled. The IV specifications construct a separate instrumental variable in each product group which is the interaction of calories per gram and a dummy which is 1 after the NLEA in the product group in question and 0 otherwise. Specifications with group-year fixed effects include a separate fixed effect for each group-year, rendering the year fixed effects redundant. Specifications with product time trends include a separate linear time trend for each product. Specifications with the "NLEA" sample include all prepackaged foods, while specifications with the "No NLEA" sample include all fresh foods and foods consumed at restaurants as a falsification test.

Table 5: Tobit Results: w/ and w/o Time Fixed Effects, w/ and w/o Heteroskedasticity

	(1)	(2)	(3)	(4)
Calories	-.1214**	-	-	-
(kcal)	(.0310)			
Total Fats	-	-.6713**	-	-
(grams)		(.2014)		
Saturated Fats	-	-	-.0016	-2.713
(grams)			(.2761)	(1.555)
Unsaturated Fats	-	-	-1.214**	-.1221
(grams)			(.2788)	(.3815)
Protein	-	-.4276	1.211	-1.590*
(grams)		(1.011)	(.9781)	(.7031)
Carbohydrates	-	-.5230*	-.3100	-.2911
(grams)		(.2586)	(.1960)	(.1431)
Fiber	-	2.911	1.568	2.120
(grams)		(2.717)	(1.088)	(1.711)
Cholesterol	-.0849	-.0181	-.0416	-.0120
(mg)	(.0629)	(.0431)	(.0773)	(.0723)
Sodium	-.0144*	.0201	.0075	.0120
(mg)	(.0091)	(.0211)	(.0041)	(.0090)
# of Consumer-Days	14596	14596	14596	14596
# of Observations	10671750	10671750	10671750	10671750

* indicates significance at the 10% level and ** indicates significance at the 5% level. Each observation is a patient day-food. Estimation is by maximum likelihood. All specifications include fixed effects for each food, group-year fixed effects, product-group specific dummies for unlabeled foods, prices with a fixed coefficient determined as outlined in the text, and dummy variables for missing prices and deflators. Model 2 disaggregates calories into fats, proteins, carbohydrates and fibers, and models 3 and 4 further disaggregate fats into saturated fats and unsaturated fats. Model 3 assumes perfect knowledge of saturated and unsaturated fat content prior to the NLEA for labeled foods, while Model 4 assumes no knowledge prior to the NLEA for labeled foods (i.e. they are treated just like unlabeled foods). All values are expressed in 1990 dollars.

Table 6: Average Annual Welfare Gain in Dollars

	Linear			Structural		
	Total	Direct	Indirect	Total	Direct	Indirect
1985	0	0	0	0	0	0
1986	1	1	0	0	2	2
1989	4	4	2	8	8	0
1990	8	9	2	8	9	0
1991	7	9	1	8	8	2
1994	48	45	6	38	38	4
1995	52	48	4	37	37	4
1996	60	55	4	41	37	4

Estimated welfare gain from additional labeling since 1985 in the linear and structural models. The direct column gives the welfare gain from foods which experienced a change in labeling. The indirect column gives the gain from substitution for foods which experienced no change in labeling. All values are expressed in 1990 dollars.

Table 7: Average Welfare Gain from NLEA and New Labeling

NLEA	Structural	Linear
Model 1	40.6	60.1
Model 2	33.4	49.2
Model 3	32.1	51.0
Model 4	28.3	56.0
More Labeling	Structural	Linear
Model 1	68.9	79.2
Model 2	78.5	72.1
Model 3	58.1	74.8
Model 4	40.1	55.1

Estimated welfare gain in dollars per year in Models 1-4 from the change in labeling from 1985-1996. The first panel gives the observed welfare gain from the NLEA. The second panel gives the additional counterfactual welfare gains that would have occurred if more products had been labeled over this period. All values are expressed in 1990 dollars.

Table 8: Benchmark Healthy Diet

Nutrient	Recommendation	Penalty
Calorie Intake	BMI of $18.5 kg/m^2$	X
Saturated Fat	As low as possible	-0.4708
Unsaturated Fat	20-35% of calorie intake	-0.0538
Protein	>10% of calorie intake	0.123
Carbohydrates	45-65% of calorie intake	-0.03
Fiber	>25g	0.561
Cholesterol	As low as possible	-0.00398
Sodium	<2300 mg	-0.00254

The first column reports the nutrient intake recommended by the National Academy of Sciences Food and Nutrition Board to minimize health risk. The second column gives the weights derived from a survey of experts by regressing health ratings for each food on nutrient content. All values are expressed in 1990 dollars. In models with saturated and unsaturated fats aggregated into total fats, I assume that any fat intake greater than 0 imposes a health risk.

Table 9: Annual Consumer Surplus from Nutrient Profile of Healthiest Diet (dollars)

Scaled	Scaled to Match BMI	Representative BMI
Model 1	100.2	123.5
Model 2	49.3	62.1
Model 3	45.0	51.3
Model 4	94.8	131.0

Table 10: Estimated vs. Benchmark Preferences

	Model 1		Model 2		Model 3	Model 4	
	Estimated	Benchmark	Estimated	Benchmark	Estimated	Estimated	Benchmark
Calories	-0.12	-0.95					
Tfat			-0.67	-5.49			
Sfat					0.00	-2.64	-11.13
Ufat					-1.21	-.12	-1.41
Protein			-0.43	2.90	1.21	-1.59	3.12
Carbo			-0.52	-0.34	-0.31	-0.29	-0.41
Fiber			2.91	15.51	1.57	2.12	14.59
Choles	-0.09	-0.24	-0.02	-0.08	-0.04	-0.01	-0.09
Sodium	-0.01	-0.07	0.02	-0.07	0.01	0.01	-0.06

For each of the four models in Table 5, this Table compares the estimated willingness to pay parameters with the benchmark parameters computed from medical evidence and VSL estimates given the characterization of consumers' current beliefs about the health consequences of different foods. All values are expressed in 1990 dollars.

Table 11: Annual Welfare Gains Re-evaluated Given Scaled Preferences

	Model 1	Model 2
Original		
NLEA	41	33
More Labeling	69	79
Re-evaluated		
NLEA	159	172
More Labeling	281	315
Additional Welfare Gains	941	1151

Estimated welfare gain in dollars per year. The original estimates are the estimates reported in Tables 6 and 7 computed using the willingness to pay parameters estimated via revealed preferences. The re-evaluated welfare gains recompute the gains from changes in consumption due to additional labeling using the scaled preferences as the normative benchmark. "Additional Welfare Gains" are the welfare gains possible if the marginal price of consumption were perfectly in line with the estimated benchmark preferences. All values are expressed in 1990 dollars.

A The Impact of the NLEA on Daily Caloric Intake

In this section, I use the structural model to generate the predicted change in consumption of label users relative to non-label users and I relate this change to the existing literature and to additional reduced form analyses. Depending on the specification used, the structural model implies a 50-90 calorie decline in consumption among label users relative to non-label users; this range is consistent with earlier studies of the impact of labeling on consumption.

From equations (7) and (8), we can compute the predicted consumption of each food as follows:

$$E(N_{ijt}) = \Phi(\hat{Y}_{ijt}/\sigma_j)(\hat{Y}_{ijt} + \sigma_j \frac{\phi(\hat{Y}_{ijt}/\sigma_j)}{\Phi(\hat{Y}_{ijt}/\sigma_j)}) \tag{18}$$

where $\phi(\cdot)$ is the standard normal density function and $\Phi(\cdot)$ is the standard normal distribution function, σ_j is the standard deviation of e_{ijt} and $\hat{Y}_{ijt}$ is the predicted value of Y_{ijt}.

I assume that label users and non-label users differ only in the specification of expected nutrient content for labeled foods. Label users are assumed to know the exact content of these foods, while non-label users know only their prior belief; given the estimated willingnesses to pay, non-label users are less able to substitute towards foods with desirable nutrient profiles.

Given simulated consumption, it is straightforward to compute the expected value of total caloric intake. This is given by:

$$E(C_{it}) = \sum_j x_{jt} E(N_{ijt}) \tag{19}$$

where x_{jt} gives the actual calories per gram of product j at time t. Given these projections, I compute a difference in difference estimate of the impact of the NLEA on label users relative to non-label users.[19]

Appendix Table 1 presents the results for the four models reported in Table 5 which

[19] In particular, I compute the change in the consumption of label users as the average simulated calorie consumption in 1994-1996 minus the average simulated calorie consumption in 1989-1991. I compute the change in consumption of non-label users via the same method. The difference in difference estimate is the difference between the change in consumption for label users and the change in consumption for non-label users.

estimate willingness to pay parameters at varying levels of aggregation. Depending on the level of aggregation, the projected change in calorie consumption ranges from 45 to 96 calories.

Given reported label use behavior, we can also compute the difference in difference estimator directly. This estimator is unfortunately confounded by selection due to the fact that the pool of label-users is changing over time. To correct for this, I have considered both triple difference estimates using changes in consumption of foods not impacted by the NLEA as a control group and estimates using a pseudo-panel constructed based on predicted label use behavior by demographic cell which allows me to control directly for the change in the proportion of label users within cells. The best identified of these specifications suggest that the NLEA led to a decline in calorie consumption of 50-100 calories among label users relative to non-label users, consistent with the projections of the structural model. Details of these estimates are available upon request.

The online version of this appendix contains a more detailed comparison of my results with the results of two earlier studies of nutrition labeling: Bollinger et al. (2010) and Variyam and Cawley (2006). The magnitude of the observed response to labeling is consistent with the findings of these earlier studies.

B Specification of $E_{ijt}(x_{nj})$

The willingness to pay for nutrient content α_n is identified using variation in perceived nutritional characteristics generated by nutrition labeling. In this section, I discuss the specification of $E_{ijt}(x_{nj})$, the expected content of nutrient n in 100 grams of product j at time t as a function of the actual content x_{nj}. The following questions need to addressed to specify this variable:

- Who uses labels?
- Which products are labeled within product groups?
- What beliefs do consumers have about the nutrient content of labeled and unlabeled foods?

The data available indicate for all "main meal planners" whether they use labels and how frequently on an "Often", "Sometimes", "Rarely" or "Never" scale in 1990, 1991, 1994, 1995, 1996. Because labeling data is only available from 1990 onward and to avoid selection issues generated by the fact that the data is a repeated cross-section, I specify the model as if the only information available on label-use were the proportion of individuals using labels for each nutrient in each year.[20] This information is sufficient to identify the impact on consumption of having more products labeled (the inframarginal impact of labeling) while controlling for the impact of an increase in label-use (the marginal effect).

The data give the proportion of products labeled in each year in each of 52 product groups; I do not know at the product level whether a food product is labeled if fewer than 100% of products are labeled. I consider two alternative assumptions: either that the healthiest products within each group voluntarily label, or that all products label randomly with probability equal to the proportion of products in each product group which are labeled. The index used to compute the health of each product is taken from Fulgoni III et al. (2009).[21] There is evidence that healthier products are more likely to voluntarily label, although the effect seems to vary by product group (Mathios (2000) finds that only the healthiest salad dressing label while Mojduszka and Caswell (2000) finds mixed results across product groups). The main text assumes that the healthiest products label, while random labeling is considered as a robustness check in Appendix D.

The specification of prior information is discussed in the main text. The parameter a_g is estimated based on a survey of consumers in Starbucks conducted by the authors of Bollinger, Leslie, and Sorensen (2010). Consumers are asked to estimate the calorie content of food and drink products that they purchased, and this value can be compared to the actual nutrient content. I restrict to those consumers who purchased a single food item. The estimate of a_g for food items is .19 with a standard error of (.20). The estimate for beverages varies depending on whether we include caffeinated beverages with close to 0 calories. If these drinks are excluded by restricting to beverages with at least 20 calories, the estimate is .11

[20] Prior to 1990, I assume that the proportion of label users for each nutrient remains constant at the 1990 level.

[21] The health index is computed by first finding the nutrient content in a fixed serving (I use 100 grams) and then computing: $protein/50 + fiber/50 + VitaminA/5000 + VitaminC/60 + Calcium/1000 + Iron/18 - SaturatedFat/20 - Sodium/2400$.

(.23). If these drinks are included, the estimate is .66 (.05). In other words, consumers appear to recognize that some drinks have almost know calories, but they are unable to distinguish between the calorie content of foods or drinks with a non-negligible number of calories. In the models in the main text, I use the estimate $a_g = 0.2$ for all product groups.

C Estimating Equation for Structural Model

Let $Y_{ijt} = \left(\frac{\gamma_{ijt}}{\theta_{ijt}}\right)^{\eta_j} - K$ where, as in the text, $\theta_{ijt} = c + \mu_{it} + \phi p_{jt} - \sum_n \alpha_n E_{ijt}(x_{nj})$ and $\gamma_{ijt} = \gamma_j + \rho_j t + v_{ij} + \epsilon_{ijt}$. I Taylor-expend about z_0, the vector of parameter values in the first year when consumption is observed. Note that Y_{ijt} depends on the characteristics of other products $k \neq j$ only through total consumption which is captured by the μ_{it} term. Thus, because we are controlling for changes in μ_{it}, Y_{jt} depends only on the characteristics of product j. Taylor-expanding about z_0 gives:

$$\begin{aligned} Y_{ijt} \approx{}& Y_{ij0} + \frac{\partial Y_{ijt}}{\partial p_{jt}}(z_0)dp_{jt} + \frac{\partial Y_{ijt}}{\partial \mu_{it}}(z_0)d\mu_{it} \\ &+ \sum_n \frac{\partial Y_{ijt}}{\partial E_{ijt}(x_{nj})}(z_0)dE_{it}(x_{nj}) + \frac{\partial Y_{ijt}}{\partial \gamma_{ijt}}(z_0)(\rho_j t + d\epsilon_{ijt}) \end{aligned} \tag{20}$$

Note that I have implicitly assumed that $\alpha_n(X_{nit})$ does not change over time. This assumption ignores the complications that arise at the boundary for the small fraction of individuals predicted to cross the threshold points where α_n changes from 0 to a non-zero magnitude for each nutrient. Instead, all individuals are treated as if $\alpha_n(X_{nit}) = \alpha_n(X_{ni})$, a constant over time determined by the value of X_{ni} observed for each individual (recall that we observe only a single time period for each individual since the data are a repeated cross-section).

To evaluate equation (20), note that for any parameter z_{ijt} in θ_{ijt} we have:

$$\frac{\partial Y_{ijt}}{\partial z_{ijt}}(z^*) = \frac{\partial \theta_{ijt}}{\partial z_{ijt}}(z^*) \cdot -\frac{\eta_{ijt}(Y^*_{ijt} + K)}{\theta^*_{ijt}} \tag{21}$$

Note also that for parameters q_{ijt} in γ_{ijt}, we have:

$$\frac{\partial Y_{ijt}}{\partial q_{ijt}}(z^*) \;\;=\;\; \frac{\partial \gamma_{ijt}}{\partial q_{ijt}}(z^*) \cdot -\frac{\eta_{ijt}(Y^*_{ijt}+K)}{\gamma^*_{ijt}} \tag{22}$$

This allows us to rewrite equation (20) as:

$$Y_{ijt} \approx Y_{ij0} + \eta_{ij0}\frac{K+Y_{ij0}}{\theta_{ij0}}\left[-\phi dp_{jt} + \sum_n \alpha_n(X_{in})dE_{ijt}(x_{nj}) + d\mu_{it} + \frac{t\rho_j + du_{ijt}}{(K+Y_{ij0})^{1/\eta_j}}\right] \tag{23}$$

I write $E_{ijt}(x_{nj}) = L_i E^L_{ijt}(x_{nj}) + (1-L_i)E^U_{ijt}(x_{nj})$ where L_i is a dummy variable for label use, E^L represents the beliefs of label users and E^U represents the beliefs of non-label users. We can rewrite this as: $E_{ijt}(x_{nj}) = E_t(L_i)E^L_j(x_{nj}) + (1-E_t(L_i))E^U_j(x_{nj}) + (L_i - E_t(L_i))(E^L_j(x_{nj}) - E^U_j(x_{nj}))$.

Define $w_{ij0} \equiv \eta_{ij0}\frac{K+Y_{ij0}}{\theta_{ij0}}$, and $e_{ijt} = Y_{ij0} + w_{ij0}\sum_n \alpha_n(X_{in})\left[(L_i - E_t(L_i))(E^L_j(x_{nj}) - E^U_j(x_{nj}))\right] + \frac{w_{ij0}t\rho_j}{(K+Y_{ij0})^{1/\eta_j}} - \frac{w_{ij0}t\rho_j}{E_j(K+Y_{ij0})^{1/\eta_j}} + \frac{w_{ij0}du_{ijt}}{(K+Y_{ij0})^{1/\eta_j}}$. This implies that we can rewrite equation (23) as:

$$\begin{aligned} Y_{ijt} &\approx w_{ij0}\left[-\phi(p_{jt}-p_{j0}) + \sum_n \alpha_n(X_{in})(E_{ijt}(x_{nj}) - E_{ij0}(x_{nj})) + \frac{t\rho_j}{E(K+Y_{ij0})^{1/\eta_j}}\right] + e_{ijt} \\ &= w_{ij0}\left[-\phi p_{jt} + \sum_n \alpha_n(X_{in})E_{ijt}(x_{nj}) + t\hat{\rho}_j + \hat{\xi}_j + d\mu_{it}\right] + e_{ijt} \end{aligned} \tag{24}$$

where $E_{ijt}(x_{nj}) = \left[E_t(L_i)E^L_j(x_{nj}) + (1-E_t(L_i))E^U_j(x_{nj})\right]$, $\hat{\rho}_j = \frac{\rho_j}{E(K+Y_{ij0})^{1/\eta_j}}$ and $\hat{\xi}_j = -\phi p_0 - \sum_n \alpha_n(X_{in})\left[E_0(L_i)E^L_j(x_{nj}) - (1-E_0(L_i))E^U_j(x_{nj})\right]$.

I now attempt to rewrite w_{ijt} as a function of the prices elasticity, prices and quantities. In particular, define the marginal price elasticity as:

$$\hat{\eta}_{ijt} \equiv -\frac{\partial E(N_{ijt})}{\partial p_{jt}}\frac{p_{jt}}{E(N_{ijt})} \tag{25}$$

where the expectation is taken over all individuals for each j and t. Note that we can write: $\frac{\partial E(N_{ijt})}{\partial p_{jt}} = \frac{\partial Y_{ijt}}{\partial p_{jt}}P(Y_{ijt}>0) = -\phi w_{ijt}P(Y_{ijt}>0)$. The first equality follows even if we allow for heteroskedastic errors provided we use a marginal effect defined in (Honoré 2008). In particular, consider $\lim_{\delta\to 0} E\left[\frac{\max\{0,Y(x+\delta)+\epsilon\} - \max\{0,Y(x)+\epsilon\}}{\delta}\right] = \frac{\partial Y}{\partial x}P(N>0)$. This marginal effect corresponds to the thought experiment: what happens to N if we change x by a small amount holding ϵ constant. In the more standard case, if ϵ_{ij} is dependent on x due to heteroskedasticity, we would allow ϵ to change as well when we perturbed x; this leads to a much more complicated expression. Note further that $E(N_{ijt}) = E(Y_{ijt}|Y_{ijt} > 0)P(Y_{ijt} > 0)$. Thus, equation (25) simplifies to:

$$\hat{\eta}_{ijt} = \phi w_{ijt} \frac{p_{jt}}{E(N_{ijt}|N_{ijt} > 0)} \tag{26}$$

where I have also used the fact that $E(Y_{ijt}|Y_{ijt} > 0) = E(N_{ijt}|N_{ijt} > 0)$. We can rearrange this to solve for w_{ijt} (and thus w_{ij0}). Substituting the resulting expression back into equation (24), gives:

$$Y_{ijt} \approx \hat{\eta}_{ij0} \frac{E(N_{ij0}|N_{ij0} > 0)}{\phi p_{j0}} \left[-\phi p_{jt} + \sum_n \alpha_n(X_{in}) E_{ijt}(x_{nj}) + t\hat{\rho}_j + \hat{\xi}_j + d\mu_{it} \right] + e_{ijt} \tag{27}$$

The full model is thus given by:

$$N_{ijt} = \max\{0, Y_{ijt}\} \tag{28}$$

where Y_{ijt} is given by equation (27) where

$$\begin{aligned} e_{ijt} &= Y_{ij0} + w_{ij0} q_{ijt} \text{ and} \\ q_{ijt} &= \sum_n \alpha_n(X_{in}) \left[(L_i - E_t(L_i))(E_j^L(x_{nj}) - E_j^U(x_{nj})) \right] + \\ & \quad \frac{t\rho_j}{(K + Y_{ij0})^{1/\eta_j}} - \frac{t\rho_j}{E_j(K + Y_{ij0})^{1/\eta_j}} + \frac{du_{ijt}}{(K + Y_{ij0})^{1/\eta_j}} \end{aligned} \tag{29}$$

The usual semiparametric estimators for censored regression models do not apply in this case because most foods are not consumed by the vast majority of consumers (Chay and Powell 2001). For example, the CLAD estimator would immediately trim all observations. For this reason, I parametrically specify the distribution of the error term. I assume that $Y_{ij0} \sim_{i.i.d.} N(0, \sigma^2)$ and $q_{ijt} \sim_{i.i.d.} N(0, \tau^2)$ and that they have constant correlation ρ. This implies that $e_{ijt} \sim N(0, \sigma^2 + w_{ij0}^2 \tau^2 + 2w_{ij0}^2 \rho\sigma^2\tau^2) = N(0, \sigma^2 + \hat{\eta}_j^2 \frac{E(N_{ijt}|N_{ijt}>0)^2}{\phi p_{jt}} \tau^2 (1 + 2\rho\sigma^2))$.

This is a heteroskedastic Tobit model where the variance $\hat{\sigma^2} = \sigma^2 + \hat{\eta}_j^2 \frac{E(N_{ijt}|N_{ijt}>0)^2}{\phi p_{jt}} \tau^2 (1 + 2\rho\sigma^2)$ varies across foods based on the elasticity of demand and the average serving size. To implement this, I compute the index $\hat{\eta}_j^2 \frac{E(N_{ijt}|N_{ijt}>0)^2}{p_{jt}^2}$ for each individual and food and estimate a separate variance for each of 20 quantiles of this index. Further details of estimation are discussed in the main text.

D Robustness of Specification of Structural Model

In this section, I discuss several estimates designed to check the robustness of the willingness to pay estimates to the assumptions made in the main text. These estimates are reported in Appendix Table 2. All specifications include the same control variables as specification (1) in Table 5. All specifications in this section use the linear model rather than the piecewise linear model used in the main text (in future drafts these checks will be repeated with the piecewise linear results). Column 1 repeats specification (1) from Table 5, except with the linear rather than piecewise linear model.

The first issue I consider is alternative assignment of labels to products within product groups when the proportion labeled is less than 100%. The estimates in the main text assume that only the healthiest foods labeled. The estimates in column 2 of Table 2 assume that the probability that a food is labeled in a given year is equal to the proportion of products in its product group which label (so that labeling is random within product group). These estimates are an informal bootstrap, in that they average point estimates and standard errors from 5 alternative estimates (the standard errors are not bootstrapped, since they are the average of the standard errors computed for each individual estimate rather than the standard error of the 5 specifications I have run). These estimates suggest that assignment of labels does not change the main results.

The second issue I consider is whether the introduction of new products in response to the labeling law may be biasing the willingness to pay estimates. In the model, the introduction of new products impacts demand for existing products only through the constraint on the total amount individuals can eat. Nonetheless, in a more realistic model the introduction of similar products would be more likely to induce substitution than the introduction of a random food product. To deal with this, for each product, I create a variable which indicates the number of low fat (or otherwise nutrient enriched) versions available in each year. Column 3 of Table 2 reports the model with this variable included; it has little impact on the willingness to pay estimates.

A third issue I consider is attrition. Studies which attempt to validate food intake from food diaries or 24-hour recall in person interviews find that these methods understate food consumption by roughly 200 calories or 10% of total intake, with the degree of understatement greater in food diary data (Sawaya, Tucker, Tsay, Willett, Saltzman, Dallal, and Roberts 1996). In the CSFII, on days when in person interviews are conducted, the reported average nutrient calorie for females aged 19-50 is 1640 calories, and on days when food diaries are used the average is 1520 calories. Combined with some assumptions about average energy expenditure, these numbers imply steady state weights below those measured in the same population (Livingstone and Black 2003 and Cutler et al. 2003). The degree of bias also appears to vary across individuals, with larger understatement of total calorie intake for more obese individuals and to vary across food groups, with understatement especially common for side dishes such as cooked vegetables and eggs (Willett 1998). Because this is a potentially serious problem, I consider three alternative specifications to deal with attrition. Column 4 restricts only to the first day of data for each individual in which attrition is less severe. Column 5 restricts to those product groups which Willett (1998) finds that attrition is least problematic. Finally, Column 6 scales all estimates of nutrients and grams consumed so that the estimated caloric intake is consistent with the mean bmi reported over this period. These estimates are similar in magnitude to the estimates reported in main text.

Column 7 reports estimates of the model with the individual-fixed effects included. The estimates with individual-fixed effects are comparable to those obtained in the model in which these are treated as a random effect.

Column 8 reports estimates of the model with the additional sub-product group / year fixed effects discussed in footnote 8. Once again, the estimates are comparable to the baseline case.

One important issue I have not yet addressed are impacts of the NLEA through avenues other than nutrition labeling. In addition to mandating nutrition labeling of prepackaged foods and altering the format of nutrition labels, the NLEA standardized the language allowed for nutrient content claims elsewhere on the packaging. The standardization rules apply to absolute nutrient claims (e.g. "low fat" requires 3 g of fat or less per serving), relative nutrient claims (e.g. "Reduced Fat" requires 25% less fat than the reference food), and health claims (only an existing list of health claims are allowed, and foods touting the health benefits of a particular nutrient must meet the requirements for absolute health claims for that nutrient) (Ippolito and Mathios 1993). After 1991, the FLAPS survey collected information on the proportion of products in each product group making nutrition claims in several different categories. This data is currently being processed, and once it is made available it will be possible to control for nutrition claims as well.

E Details of Behavioral Welfare Calculation

Provided $X_{in} \notin \{\bar{X}_{in}, \underline{X}_{in}\}$, we can compute the solution to equation 15 by implicitly differentiating the objective function in equation 15 and then substituting in for $\frac{\partial v}{\partial N_{ij}}$ and $\frac{\partial N_{ij}}{\partial \alpha_{in}}$ from the first order conditions for equation 14. The benchmark parameters are then characterized by the system of linear equations:

$$(Sx)' q = 0 \tag{30}$$

where S is the J x J matrix of marginal price effects defined by $s_{kj} = \frac{\partial N_{ik}}{\partial p_j}$, x is the J x N matrix of nutrient contents where x_{jn} gives the content of nutrient n in one gram of product j and q is the J x 1 vector whose jth element is given by: $q_{ij} = \frac{\partial h}{\partial N_{ij}} - \frac{\partial \tilde{h}}{\partial N_{ij}} - \sum_n \alpha^*_{in} x_{nj}$. Define $W = (Sx)'$. Equation 30 defines a system of N equations, one for each nutrient, given by: $\sum_j w_{nj}(\frac{\partial h}{\partial N_{ij}} - \frac{\partial \tilde{h}}{\partial N_{ij}} - \sum_n \alpha^*_{in} x_{nj}) = 0$. If we divide these equations by J and take the limit as J $\to \infty$, then we can write the benchmark parameters as a function $E\left(w_{nj}(\frac{\partial h_j}{\partial N_{ij}} - \frac{\partial \tilde{h}_j}{\partial N_{ij}})\right)$ for each nutrient, which in turn depends on $\text{Cov}(\beta^*_i \frac{\partial h}{\partial N_{ij}} - \beta_i \frac{\partial \tilde{h}}{\partial N_{ij}}, \tilde{x}_{nj})$, and $E(\tilde{x}_{nj})E(\beta^*_i \frac{\partial h}{\partial N_{ij}} - \beta_i \frac{\partial \tilde{h}}{\partial N_{ij}})$ for each nutrient n where $\tilde{x}_{nj} = \sum_k \frac{\partial N_j}{\partial p_k} x_{nk}$ tells us how much a change in α_n will impact consumption of product j.

I begin by describing the metric used to calculate the dollar cost of alternative diets compared to diets which minimize health risk. Martin, Beshears, Milkman, Bazerman, and

Sutherland (2009) survey nutritional experts and elicit a health rating for each of 205 different foods in light of their nutritional characteristics on a scale of -5 to 5. This rating is then regressed on the underlying characteristics to recover the relative weight attached to different nutrients by experts in evaluating food healthiness. The authors perform several additional checks which suggest agreement among experts regarding the relative weights attached to different nutrients. I use the weights recovered from this regression to evaluate the relative importance of different nutrients in computing the distance of a given diet from the range of benchmark healthy diets. Let $\underline{X}_n$ denote the minimum recommended consumption of nutrient n in the benchmark diet and $\bar{X}_n$ the maximal recommended consumption (for protein and fiber, $\bar{X}_n = \infty$). I compute the distance from a given diet d to the benchmark healthy diet as:

$$w(d) = \sum_n \alpha_n \max(X_n^d - \bar{X}_n, \underline{X}_n - X_n^d, 0) \tag{31}$$

where δ_n is the negative of the absolute value of the coefficient from the Martin et al. (2009) regression (this appropriately accounts for the fact that Fiber and Protein consumption below the recommended level negatively impacts health). Thus, $w(d) = 0$ for any diet in the benchmark range, and $w(d) < 0$ for diets outside the benchmark range. I scale $w(d)$ into life years by choosing α such that $l(d) = \alpha w(d)$ and $E(l(d))$ across all consumers is .04 life-years. Finally, I convert this into a dollar amount based on estimates of the distribution of the value of a statistical life. I start with an average VSL of \$6.4 million (Viscusi and Aldy 2003) and compute the value of each life year by assuming that there is a constant value of a life year and that the VSL is the present discounted value of all additional life years. That is, I solve, $E_i(\sum_{t=0}^{T_i(a_i)-a_i} \delta^t V_{year}) = V^*$ where the expectation is taken over all individuals in the data, a_i indicates age and $T_i(a_i)$ indicates life expectancy conditional on age a_i (this procedure is similar to that used in Gruber and Koszegi (2001). The value of the marginal life-year is given by $\delta^{T_i - a_i} V_{year}$. I assume $\delta = .96$. This value will vary across consumers based on their age, but it implies that the average consumer loses about \$3,000 worth of life-years by consuming their current diet rather than the healthiest possible diet (I consider below the impact of allowing for some heterogeneity in the VSL). The result of this calculation is a function which expresses the health cost of all diets in dollar terms which I can use to determine the marginal health cost of all foods.

I next describe how I characterize consumers current beliefs about the life expectancy consequences of alternative diets. I make use of consumers' answer to the following question: "how many servings would you say a person of your age and sex should eat each day for good health from food group [X]?" for each of grains, fruits, vegetables, dairy and meat and poultry products. Because consumers may understand the word "serving" differently from the official definition, I rescale all of their responses so that the total number of servings indicated would match their food consumption in grams (which I assumed in the structural model was fixed). This is a conservative assumption, in the sense that the only deviations detectable between consumers' beliefs and the benchmark diet recommended by experts are deviations in the relative consumption of different food groups. I consider two strategies for characterizing a typical serving. In the first case, I assume that the nutrient profile of a serving from each food group is the average profile constructed by weighting all foods in that group by their proportion of group-consumption in grams. In the second case, I use the nutrient profile constructed using only those individuals who rank in the top 5% in terms of the health of their consumption in the group in question. I report results from the first case; the results in the second case are not appreciably different.

Define a set of indicator variables A_n which indicate the range into which X_n falls ($A_n = 0$ if $X_n <= \underline{X}_n$, $A_n = 1$ if $\underline{X}_n < X_n < \bar{X}_n$, $A_n = 2$ if $\bar{X}_n <= X_n$). We can write $h(d) = \sum_n [S\delta_n(A_n)X_n + Q_n(A_n)]$ where $\delta_n(0) = -\delta_n$, $\delta_n(1) = 0$ and $\delta_n(2) = \delta_n$, and $Q_n(0) = S\delta_n\underline{X}_n$, $Q_n(1) = 0$ and $Q_n(2) = -S\delta_n\bar{X}_n$. For each food we can define $h_{ij}(d) = \sum_n \left[S\delta_n(A_n)x_n + \frac{Q_n(A_n)}{N}\right]$ with the property that $h_i(d) = \sum_j N_{ij}h_{ij}(d)$. I normalize the $h_i(d)$ by subtracting a constant so that the life expectancy consequences of the benchmark diet are normalized to 0.

Define $\tilde{h}_j(A)$ as consumers beliefs about the marginal life expectancy consequence of consuming a unit of food j given their current nutrient consumption (which is summarized by A, as defined in Section 8.2). Let $d^* = (N_1^*, ..., N_J^*)$ denote the benchmark diet recommended by experts and $d' = (N_1', ..., N_J')$ denote the diet consumers believe is healthiest. Let $A' = \{A_n\}$ evaluated at diet d' and $A^* = \{A_n\}$ evaluated at diet d^*. To characterize consumer beliefs over the entire range of possible nutrient intakes (that is, all possible A rather than just A'), I assume that the $\tilde{h}_j(A_1) - \tilde{h}_j(A_2) = h_j(A_1) - h_j(A_2)$. This is again conservative in the sense that it assumes that consumers correctly evaluate changes in the marginal value of food with respect to their overall nutrient intake.

I compute $\tilde{h}_j(A')$ for each food by minimizing the distance from $\tilde{h}_j(A')$ to $h_j(A')$ while nonetheless rationalizing the judgment that the diet given by consumers is healthier than the benchmark healthy diet recommended by experts. That is, I solve:

$$\begin{aligned} (\tilde{h}_1, ..., \tilde{h}_J) \quad &\equiv \quad \underset{(\tilde{h}_1,...,\tilde{h}_J)}{\operatorname{argmin}} \sum_n \left[Cov(h_j(A') - \tilde{h}_j(A')), \tilde{x}_{nj})\right]^2 + \sum_j E(h_j(A') - \tilde{h}_j(A')) \\ s.t. \quad &\quad \sum_j N'_j \tilde{h}_j(A') \geq \sum_j N^*_j \tilde{h}_j(A^*) \end{aligned} \tag{32}$$

Because the benchmark is actually a range of risk-minimizing nutrient intakes rather than a particular profile of food consumption, many different diets are consistent with this range. For this reason, I also maximize equation 32 over the set of diets $(N^*_1, ..., N^*_J)$ which are at least as healthy as the benchmark diet (i.e. $\sum_j N^*_j h_j \geq 0$, the normalized health of the benchmark diet). That is, I solve:

$$\begin{aligned} (\tilde{h}_1, ..., \tilde{h}_J) \equiv \underset{(\tilde{h}_1,...,\tilde{h}_J,N^*_1,...,N^*_J)}{\operatorname{argmin}} \sum_n \left[Cov(h_j(A') - \tilde{h}_j(A')), \tilde{x}_{nj})\right]^2 + \sum_j E(h_j(A') - \tilde{h}_j(A')) \; s.t. \\ \sum_j N'_j \tilde{h}_j(A') \geq \sum_j N^*_j \tilde{h}_j(A^*), \sum_j N^*_j h_j(A^*) \geq 0, \sum_j N^*_j = \bar{N} \end{aligned} \tag{33}$$

The second constraint, $\sum_j N^*_j h_j \geq 0$ states that candidate benchmark diet is consistent with the range of values given in Table 8 (whose health value is normalized to 0). The third constraint, $\sum_j N^*_j = \bar{N}$, states that candidate benchmark diets must also have total consumption in grams equal to a fixed constant.

This calculation results in estimate of $\tilde{h}_j(A)$, consumers' beliefs about the life expectancy consequences of consuming a unit of each product j, expressed in dollar equivalents as a function of current nutrient consumption from which I can directly compute the covariance and expectation parameters of interest.

Appendix Table 1: Impact of NLEA on Label Users

Specification	Calories/Day
Model 1	-96.3
Model 2	-61.1
Model 3	-62.3
Model 4	-45.1

The reported value is the models projection of the changed in calories consumed per day induced by the NLEA for each model from Table 5.

Appendix Table 2: Robustness of Willingness to Pay Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Calories	-.1078**	-.0812**	-.1085**	-.1251*	-.1415**	-.1283**	-.1211**	-.0895**
	(.0260)	(.0211)	(.0241)	(.0526)	(.0401)	(.0309)	(.0114)	(.0210)
Cholesterol	-.0864	.0123	-.0710	-.0322	-.1011*	-.1028	-.0913**	-.0312
	(.0648)	(.0070)	(.0510)	(.0782)	(.0415)	(.0771)	(.0315)	(.0514)
Sodium	-.0145*	-.0110	-.0209*	.0252	.0215*	-.0173*	-.0155**	-.0212*
	(.0090)	(.0140)	(.0102)	(.0442)	(.0100)	(.0107)	(.0050)	(.0108)

Column 1 replicates specification (1) in Table 5. Column 2 assigns labels randomly within product group rather than by health. Column 3 includes an additional control for the number of low fat substitutes. Column 4 restricts just to food consumption data from the 1st day, reported based on an in-person interview (the 24-hour recall data). Column 5 restricts to those product groups where attrition is least problematic in this type of data. Column 6 scales all consumption to match a BMI benchmark. Column 7 includes individual fixed effects. Column 8 includes sub-product group-year fixed effects.