Information, Decision-Making and Health

by

Jason Abaluck

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of
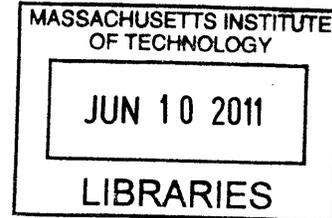
Doctor of Philosophy in Economics

at the

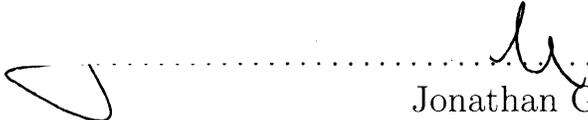MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2011

Author..............................
//   Department of Economics
May 15, 2011

Certified by ....
Jonathan Gruber
Professor of Economics
Thesis Supervisor

Certified by ....
Michaël Greenstone
3M Professor of Environmental Economics
Thesis Supervisor

Accepted by.....
Esther Duflo
Abdul Latif Jameel Professor of Poverty Alleviation and Development
Economics
Chairman, Department Committee on Graduate Theses

Information, Decision-Making and Health

by

Jason Abaluck

Abstract

This thesis consists of three essays on information, decision-making and health. All three concern the relationship between the choices consumers would make if they were "fully informed" in an appropriate sense and the choices we actually observe.

Chapter 1 considers how we can determine whether consumers are appropriately taking into account health information when they make their food consumption decisions. The fundamental idea is to determine the value of a statistical life (VSL) implicit in food consumption decisions and to compare this value with previous estimates of the VSL. The main positive result is that the VSL estimated from food consumption is about 1/10th as large as estimates from other contexts. I also consider the normative implications under the assumption that VSL estimates from other contexts indicate how individuals would behave if they were "fully informed" and discuss what additional evidence might support such an assumption.

Chapter 2, co-authored with Jonathan Gruber, performs an analogous exercise in the case of health care plans. Where Chapter 1 makes the normative assumption that consumers should value years of life equally regardless of where they come from (e.g. eating healthier foods or reducing risk of on-the-job death), Chapter 2 makes the normative assumption that consumers should value a dollar of cost savings equivalently whether it comes through premiums or out of pocket costs. This restriction can then be used to evaluate whether consumers are choosing appropriately. The chapter studies this question in the context of Medicare Part D Prescription Drug Plan, the most significant privatization of the delivery of a public insurance benefit in recent history.

Chapter 3 attempts to consider the circumstances in which the partial equilibrium welfare analyses performed in parts 1 and 2 extend to a general equilibrium setting in which prices and product characteristics respond endogenously to changes in demand. In particular, Chapter 3 derives conditions under which more information leads to welfare gains in general equilibrium taking into account the endogenous response of firms' pricing and product quality decisions.

3

Thesis Supervisor: Michael Greenstone
Title: 3M Professor of Environmental Economics

0.1 Acknowledgments

Thanks first to Jon Gruber for his invaluable guidance and advice and the 2,801 emails he sent me over the past 4 and a half years. More than anyone else, Jon has shaped my professional interests, how I think about what is worth studying, and how I go about doing research. Without his guidance about which papers to write, which conferences to attend, who to talk to, what to say, and what clothing to wear among many other things, I likely would have abandoned graduate school to pursue an unremarkable career as a journeyman NBA role-player or a very motion-sick astronaut. I plan to follow Jon's example by speaking quickly, carrying a big dataset, and attempting to revive indentured servitude among undergraduate students. His advice to "Be careful of co-authoring with more senior people... oh, by the way, I have 6 more projects we should work on" will no doubt continue to shape my research for many years to come.

Thanks also to my other advisors, Michael Greenstone and Glenn Ellison. Both have given innumerable helpful comments regarding the content and presentation of my work. Michael has been especially helpful in thinking about how to present the results more clearly and succinctly - and this in turn has helped me develop a deeper understanding of what I am trying to do and what assumptions need to be made to do it. Michael also helpfully explained to me that he blames his wife whenever he has to turn down appealing job offers, so thanks also to Michael Greenstone's wife since I took his advice and blamed her when I was in a similar situation. Glenn has provided penetrating comments on several drafts, and has provided helpful guidance about which projects would be most interesting to pursue. Glenn's IO class is also one of the best classes I took at MIT and the lecture notes are my first point of reference whenever I start work on an IO-related project. Both Michael and Glenn have helped out with applications to fellowships and summer programs on extremely short notice (future students should reward their efforts by giving them ample time to prepare such documents in the future rather than giving them projects to do at the last minute because you know they will be helpful).

Many other faculty members have provided helpful comments on drafts of this work. Some of these contributions are mentioned in the acknowledgments to the individual papers,

but a few deserve special mention. Peter Diamond discussed early versions of this project with me and guided me away from several wrong turns. Jim Poterba has his finger on the pulse of the economics profession, and almost every comment or concern he has raised either in lunch presentations or in our mock interview helped me be better prepared when the same issue was raised on the Job Market. Amy Finkelstein, David Autor and Bill Wheaton also consistently provided helpful comments when I presented this work.

Given my rare genetic disorder which prevents me from dealing with logistical and administrative work until the very last minute, I am especially grateful to Gary King, John Arditi, Theresa Benvento, Deborah Jamiol, Katherine Swan, and Janet Stein who through their hard work have prevented this from causing serious problems. I am also grateful to the National Science Foundation Graduate Research Fellowships program and the NBER Health and Aging Pre-doctoral Fellowship for financial support.

My decision to pursue a PhD in economics was heavily influenced by my undergraduate advisor Richard Freeman, both through his advice and guidance, and his help in the admissions process. Having now worked with many undergraduate RAs of my own, I have a fuller appreciation of how unique the opportunity was to work for Richard both due to the amount of time we spent directly conversing and the nature of the tasks assigned (which typically required substantive economic and statistical thought as opposed to menial labor). The messiness of Richard's desk has also been a source of solace and inspiration.

I am of course grateful to my friends and classmates. I cannot list everything they have helped me with, but I can list everything important: Michael Powell helped me see more clearly at all stages of this project by recognizing and then eliminating the smudges from my glasses (Iuliana Pascu assisted by playing with Gustav the cat), Tatyana Deryugina lost $10 betting on Connect 4 and introduced me to her high-stakes friend Leonid who lost $100 betting on chess, Danielle Li was almost never racist, Amanda Pallais recently threatened to write using my blood if we forgot pens, David Simmons-Duffin read and occasionally commented on my twitter posts, Joanna Huey refused to provide legal advice, Matt Gline had an incident with fish immortalized in song by Alex Dahlen, and Lynn Wu took me to Trader Joe's. I'd also like to thank everyone who has spent time in the graduate computer lab, especially Joe Shapiro for his fast typing and Adam Sacarny who will buy and eat

anything if he is hungry. Anyone omitted from this list should take it as a personal insult and bide their time until an opportunity for vengeance arises.

This thesis is dedicated to my parents and brother. My parents conceived me, one of them birthed me and the other taught me how to multiply single-digit numbers then they paid for college all while feeding me and buying me juice. My brother was recently delegated the responsibility of helping me procure juice. I wouldn't trade my family for any family of complete strangers even if I lived in a dystopia where people had property rights over other people and such trades were legal and even commonplace. I love them all a lot.

Contents

List of Figures

13

List of Tables

Chapter 1

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

1.1 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,

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

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 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

18

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 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

---

[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.

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 1.2 describes a simple example to explain how the structural model identifies the willingness to pay for each nutrient. Section 3.2 briefly reviews some related literature. Section 1.4 describes the available data on labeling, food consumption and prices. Section 1.5 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 1.6 introduces the model of labeling and food demand, and reports estimates of the willingness to pay for nutrient content for several nutrients. In Section 1.7, I perform several welfare exercises evaluating the impact of the NLEA taking preferences as given, and in Section 1.8 I re-evaluate these gains given a normative benchmark derived from expert medical knowledge and assumptions about the value of additional life-years. Section 3.6 concludes.

1.2   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 1.1 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

21

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 1.1 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.

1.3 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 pro-

cessing. 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 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.

1.4 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

---

[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.

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.

1.4.1 Labeling Data

The standardized template mandated by the NLEA for nutrition labels is given in Figure 1-3. 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 1-4. 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 1-5 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 1.2 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.

1.4.2   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 A.4.

26

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 A.4.

1.4.3   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.

1.4.4   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.

1.5   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 1.2, 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:

$$C_{jt} = \beta x_j L_{g(j)t} + d_j + d_t + e_{jt} \tag{1.1}$$

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 1.4. 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 1.4 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 1-6 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

30

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.

1.6 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 1.2. 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.

1.6.1 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:

$$
\begin{aligned}
U_{it} &= \sum_{j} (\gamma_j + \rho_j t + \epsilon_{ijt}) \frac{(K + N_{ijt})^{1 - \frac{1}{\eta_j}}}{1 - \frac{1}{\eta_j}} \\
&+ \sum_{n} \alpha_n(X_{in}) \left(\sum_{j} N_{ijt} E_{ijt}(x_{nj}) \right) + \phi I_{it}
\end{aligned} \tag{1.2}
$$

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

32

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 1.3 gives the exact percentages.

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

$$\sum_j p_{jt} N_{ijt} + I_i \leq W_{it} \tag{1.3}$$

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:

$$\sum_j E_{ij}(N_{ijt}) \leq N_i \tag{1.4}$$

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 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.

---

[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 1.8.2.

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 A.4, 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 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:

$$N_{ijt} = \max\{0, \left(\frac{\gamma_{ijt}}{\theta_{ijt}}\right)^{n_j} - K\} \tag{1.5}$$

---

[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.

1.6.2 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 A.3, I show that Taylor-expanding about z_0 gives:

$$Y_{ijt} \approx w_{ij0}\left[-\phi p_{jt} + \sum_n \alpha_n(X_{in})E_{ijt}(x_{nj}) + t\hat{\rho}_j + \hat{\xi}_j + d\mu_{it}\right] + e_{ijt} \tag{1.6}$$

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 A.3, 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 (1.6), gives:

$$Y_{ijt} \approx \hat{\eta}_{ij0}\frac{E(N_{ij0}|N_{ij0} > 0)}{\phi p_{j0}}\left[-\phi p_{jt} + \sum_n \alpha_n(X_{in})E_{ijt}(x_{nj}) + t\hat{\rho}_j + \hat{\xi}_j + d\mu_{it}\right] + e_{ijt} \tag{1.7}$$

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:

$$N_{ijt} = \max\{0, Y_{ijt}\} \tag{1.8}$$

where Y_{ijt} is given by equation (1.7). In Appendix A.3, 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 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 A.4, 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

---

[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.

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.

1.6.3 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 1.2. 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 1.5 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:

$$E_{ijt}(x_{nj}) = E_g(x_{nj}) + a_g(x_{nj} - E_g(x_{nj})) + r_{ij} \qquad (1.9)$$

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:

$$x_{nj} - E_{ijt}(x_{nj}) = (1 - a_{gn})(x_{nj} - E_g(x_{nj})) \qquad (1.10)$$

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 A.2, 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]

---

[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 a_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

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.

1.6.4 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 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%

---

year specific fixed effects for each sub-product group. In this specification, reported in Appendix A.4, 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.

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 1.5 reports the estimates from the Tobit model described in equations (1.7) and (1.8). 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 (1.10) 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 1.5 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 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.

---

[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.

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 A.4. 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.

1.7   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 1.6.3, 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]

1.7.1 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 not labeled, individuals act as if they face a generalized price $\hat{P}_{jt}$ which is different from the actual price P_{jt}^* obtained by substituting in the true nutrient content x_{nj} for the perceived nutrient content (ignore for the moment the fact that labeling also impacts the multiplier μ_{it}). If labeling conveys good news, individuals will underconsume prior to labeling, and if labeling conveys bad news, individuals will overconsume prior to labeling. The lost consumer surplus is given by the area between the true generalized price if nutrient content were known

---

[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.
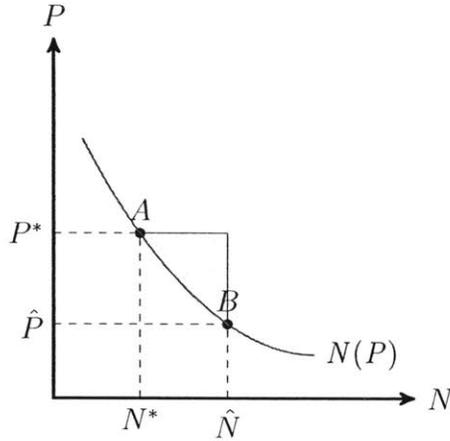
Figure 1-1: $N(P)$ gives the demand curve for a sample product j as a function of the generalized price P. $\hat{P}$ gives the apparent price prior to nutrition labeling and P^* gives the price after nutrition labeling which is also the appropriate normative benchmark in this section. The shaded region gives the welfare gain from labeling for this product: individuals receive bad news about the nutrient content so the welfare gain comes from realizing that the marginal cost of consumption P^* is higher than they thought absent labeling.

P^*_{jt} and the demand curve as shown in Figure 1-1. The total change in consumer surplus is computed by summing across all products.

The welfare benefits computed in this section are a lower bound on the partial equilibrium consumer surplus because the current calculation ignores heterogeneity across individuals in prior information for a given food; that is, even if we observe no change in aggregate demand for a given food in response to labeling, this may mask the fact that some individuals received bad news and consumed less while others received good news and consumed more. Future drafts will take this factor into account as well by calibrating the degree of heterogeneity across individuals from survey evidence.

I will start by considering a simplified model with a linear demand curve to make the 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:

$$\Delta W = \frac{1}{2}\left(\frac{\alpha_n}{\phi}\right)^2 \sum_j |b_{ij}|(\Delta c_{ij})^2 \tag{1.11}$$

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:

$$\Delta W = \frac{1}{2}\sum_j |b_j|(\frac{\Delta \mu_i}{\phi} - \sum_n \frac{\alpha_n}{\phi}\Delta x_n)^2 \tag{1.12}$$

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:

$$\Delta W = \sum_j \left(CS(\hat{a}) - \hat{N}_{ij}(p_{jt}^* - \hat{p}_{jt})\right) \tag{1.13}$$

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 1.6 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 -

44

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.

1.7.2 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 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

---

[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).

not be too different, but that is far from convincingly established (this agreement is discussed in Appendix A.1). 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 1.7.1: 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 1.7 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 1.6.4. 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%.

1.8 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.

1.8.1 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 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

---

[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.

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 1.8.3 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 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 1.8. 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

---

[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.

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

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 1.9. 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 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.

1.8.2 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.

49

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:

$$N_{ij} = \arg\max_j v(N_{i1}, ..., N_{iJ}) + \beta_i \tilde{h}_i(N_{i1}, ..., N_{iJ}) + \sum_n \alpha_{in}\left(E_{it}(X_{in})\right) - \sum_j p_j N_{ij} \quad (1.14)$$

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:

$$\arg\max_{\alpha_{in}} v(N_{i1}^*, ..., N_{iJ}^*) + \beta_i^* h(N_{i1}^*, ..., N_{iJ}^*) - \sum_j p_j N_{ij}^* \quad (1.15)$$

where N_{ij}^* are solutions to equation 1.14 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 infor-

50

mation 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 1.15 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 1.15 by implicitly differentiating the objective function in equation 1.15 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 1.14. The benchmark parameters are then characterized by the system of linear equations:

$$(Sx)' q = 0 \qquad\qquad (1.16)$$

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: $q_{ij} = \frac{\partial h}{\partial N_{ij}} - \frac{\partial \tilde{h}}{\partial N_{ij}} - \sum_n \alpha_{in}^* x_{nj}$.

I show in Appendix A.5 that in this case we can write the benchmark parameters as a function of $\mathrm{Cov}(\beta_i^* \frac{\partial h}{\partial N_{ij}} - \beta_i \frac{\partial \tilde{h}}{\partial N_{ij}}, \tilde{x}_{nj})$, 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

---

[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.

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 1.15 as:

$$\arg\max_{\alpha_{in}} v(N_{i1}^*, ..., N_{iJ}^*) + \beta_i^* h(N_{i1}^*, ..., N_{iJ}^*) + \sum_n \delta_{in} X_{in} - \sum_j p_j N_{ij}^* \qquad (1.17)$$

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.

1.8.3   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 A.13.

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 1.8. More details of this calculation are given in Appendix A.5.

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

54

consume healthier foods from that group; I find that this is not the case. Details are again given in Appendix A.5. 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 A.13 separately for each consumer. The average of these benchmark weights across all individuals for each of the four specifications in Table 1.5 are given alongside the estimated willingnesses to pay for each nutrient in Table 1.10. 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 that all of the health consequences of poor dietary behavior come from people in the 10th percentile of the VSL distribution, the scaling factor for the true health consequences and the associated benchmark parameters will be about 40% smaller than reported in Table 1.10.[18] This is still several times larger than the observed responsiveness.
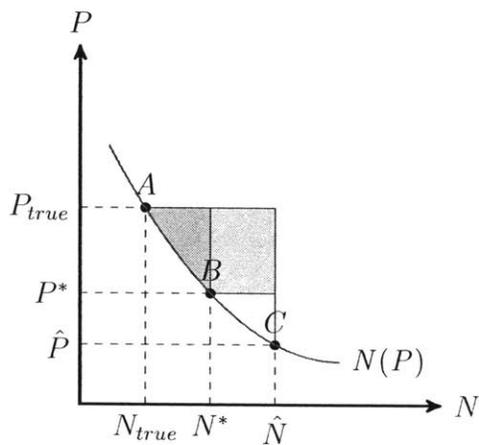
Figure 1-2: $N(P)$ gives the demand curve for a sample product j as a function of the generalized price P. $\hat{P}$ gives the apparent price prior to nutrition labeling, P^* gives the price after nutrition labeling, and P_{hat} gives the true price if individuals were fully aware of all health relevant factors. The most lightly shaded region gives the welfare gain from labeling judged from the benchmark of P^*, the value computed in Section 1.7.1. The medium-gray region gives the correction to this welfare gain when the change in consumption is judged using the fully informed marginal cost P_{true}. The most darkly shaded region gives the additional welfare gain that could be realized if individuals observed P_{true} and changed their consumption to N_{true}.

1.8.4   Welfare Gains Reconsidered

Given these normative benchmarks, I can then ask: supposing that the correct normative utility function values nutrients at α_n^* (the average value of the benchmark parameter) rather than the observed α_n, what is the welfare loss due to the fact that individuals behave as if nutrients are valued at α_n?

Suppose as above that consumers could gain \$3,000 in life-years annually by consuming 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.

---

[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.

[18]More generally, the benchmark parameters do not scale exactly with the function h holding $\tilde{h}$ fixed, but due to the way $\tilde{h}$ is estimated above, changing the scale of h likewise changes the scale of $\tilde{h}$, and thus the scale of the benchmark parameters by the same amount given equation A.13.

The analysis here follows closely the welfare analysis laid out in Section 1.7. 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 1-2 describes how consumer surplus is computed for each product given these prices.

Table 1.11 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.

1.9 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 1.8.2 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.

Nutrition Facts

Serving Size 0 cup (000g)
Servings Per Container 0

Amount Per Serving

Calories 000   Calories from Fat 000

	% Daily Value*
Total Fat 00g	**00%**
Saturated Fat 0g	**00%**
Cholesterol 00mg	**00%**
Sodium 000mg	**00%**
Total Carbohydrate 00g	**00%**
Dietary Fiber 0g	**0%**
Sugars 00g	
Protein 00g	

Vitamin A 0%   ◆   Vitamin C 0%

Calcium 00%   ◆   Iron 0%

Percent Daily Values are based on a 2,000 calorie
diet. Your daily values may be higher or lower
depending on your calorie needs:

	Calories:	2,000	2,500
Total Fat	Less than	65g	80g
Sat Fat	Less than	20g	25g
Cholesterol	Less than	300mg	300mg
Sodium	Less than	2,400mg	2,400mg
Total Carbohydrate		300g	375g
Dietary Fiber		25g	30g

Calories per gram:
Fat 9 • Carbohydrate 4 • Protein 4

Figure 1-3: The Sample Nutrition Facts Panel Provided by the USDA (post-NLEA)

Figure 1-4: A Sample Pre-NLEA Nutrition Label
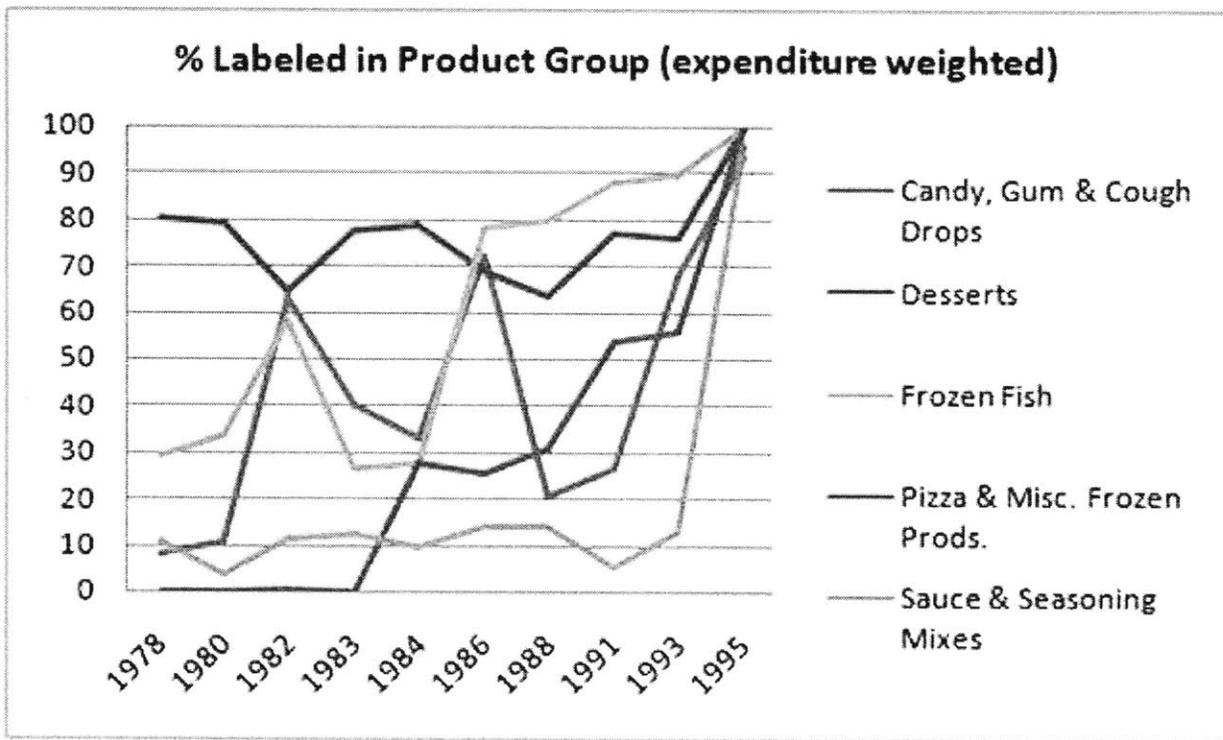
Figure 1-5: Proportion of Products Labeled by Year for a Sample of Product Groups
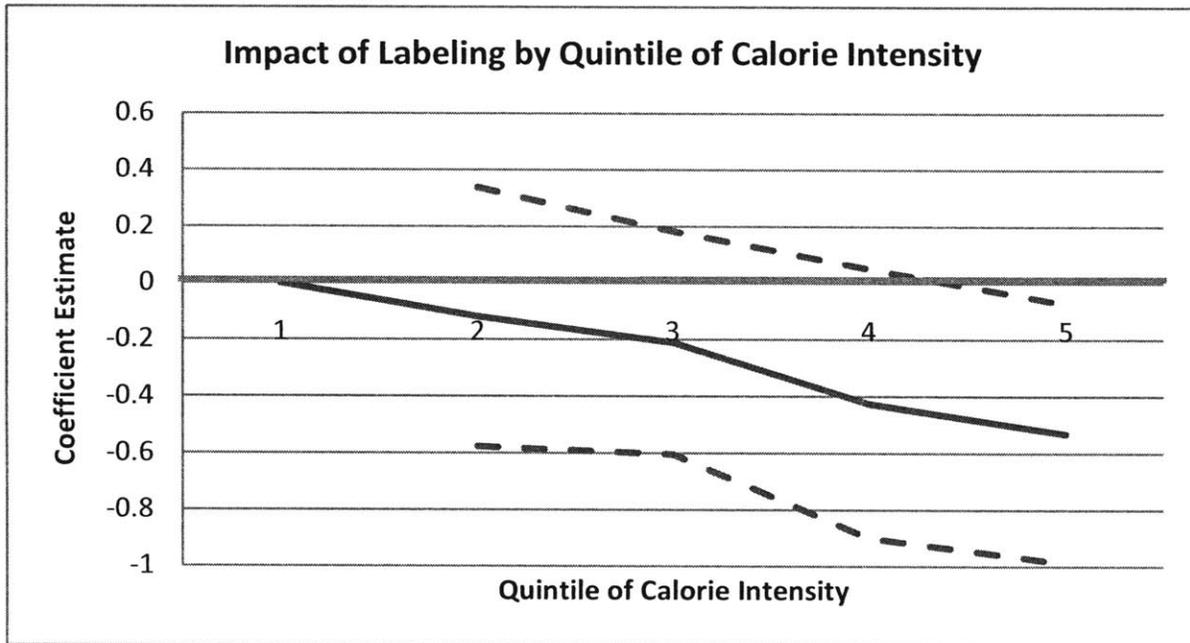
Figure 1-6: This graph shows the coefficients from estimating equation 1.1, but substituting quantiles of calorie intensity interacted with the proportion of products labeled for the independent variable. The dotted lines indicate 95% confidence intervals. The coefficient can be interpreted as the change 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 1.1: 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 1.2: 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 1.3: 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 1.4: 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 1.5: 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.

66

Table 1.6: 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 1.7: 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 1.8: 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 1.9: 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

68

Table 1.10: 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 1.5, 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 1.11: 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 1.6 and 1.7 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.

Chapter 2

Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program

2.1 Introduction

The Medicare Modernization Act of 2003, better known as the legislation that added the Part D prescription drug benefit to the Medicare program, represents the single most significant expansion of public insurance programs in the U.S. in the past 40 years. The most novel, and controversial, feature of this legislation was the use of multiple private insurance providers to deliver this new public insurance product. Unlike the traditional model of government mandated uniform insurance packages for all enrollees, under the Part D program dozens of private insurers were allowed to offer a wide range of products with varying prices and product features. Perhaps most well-known was the extent to which plans covered the donut hole, a broad uncovered range of expenditures in the minimum mandated plan.

This unprecedented privatization of the delivery of a public insurance product raises a host of important policy questions. Primary among these is the impact of allowing choice across so many private insurance options. The typical elder in our data (described below) faces a choice of over 40 stand-alone drugs plans, and our estimates suggest that the range of cost from the most to least expensive option facing an elder is comparable to the mean of those costs. Choice is clearly meaningful in this context. Yet, to date, we know almost nothing about how elders are making these crucial choices.

This paper investigates the choices of elders for the newly formed Part D program in 2006. We analyze data that provides information on the Part D plans chosen and prescription drug utilization for a large sample of elders in the U.S. These data were collected by Wolters Kluwer (WK), a switch agent that lies between pharmacies that fill prescriptions and the insurance companies and prescription benefit managers that pay for them. WK collects information on almost one-third of all third party prescription drug transactions, and we will use the universe of their data for those over age 65 during 2005-2006 to examine choice of Part D plan. We match to this data set a comprehensive set of information from the Centers for Medicare and Medicaid Services (CMS) on the Part D plans available to each person in our data set.

Specifically, for each elder whose claims appear in our sample, we model the financial implications of each of the plans in their choice set, based on both 2005 and 2006 drug utilization and several different models of expectations. We begin by presenting the basic facts on choice, documenting that the vast majority of elders are choosing plans that are not on the efficient portfolio of plan choice for that elder. We then turn to more rigorous multinomial models of individual choice to incorporate non-financial characteristics, preference heterogeneity and unobserved plan characteristics into our analysis.

Our findings are striking: along three dimensions, elders are making choices which are inconsistent with optimization under full information. First, elders place much more weight on plan premiums than they do on the expected out of pocket costs that they will incur under the plan. Second, consumers appear to value plan financial characteristics far beyond any impacts on their own financial expenses or risk. Third, consumers substantially under-value variance reducing aspects of alternative plans. The first two of these conclusions are robust

72

to a variety of specifications and econometric approaches; the third is more sensitive.

Our paper proceeds as follows. Part I provides background on the Part D program and reviews the growing literature on its impacts. Part II discusses our data sources, and Part III presents initial results on choice set variation and choice behavior. Part IV describes our choice framework, and Part V presents results and robustness checks. Part VI considers issues of misspecification and measurement error, and Part VII assesses robustness to heterogeneity concerns. Part VIII concludes with a discussion of the policy implications of our findings. The appendices referenced throughout are included in the online version of this paper.

2.2 Background

2.2.1 The Medicare Part D Program

Medicare, which provides universal health insurance coverage to those over age 65 and to those on the disability insurance program, was established in 1965. The original program covered most medical needs for the elderly and disabled, including hospital and doctor costs, but it excluded coverage for prescription drugs. This omission was not perceived as a major one in the early years of the Medicare program, but in the 1990s the advancement of prescription drug treatments for common illnesses among the elderly drew attention to this gap in Medicare coverage. Medicare recipients, for example, spent an average of $2,500 each on prescription drugs in 2003, more than twice what the average American spent on all health care in 1965.[1]

In 2003, the Bush administration and Congress reached agreement on a far-reaching prescription drug benefit package at a projected cost to the federal government $40 billion per year for its first ten years. The most noticeable innovation of the Part D plan is that this new Medicare benefit is not delivered by the government, but rather by private insurers under contract with the government. Beneficiaries can choose from three types of private insurance plans coverage of their drug expenditures. The first is stand-alone plans called

---

[1]Data for prescription drug spending comes from the Congressional Budget Office (2002). Data for average Americans health spending comes from the National Health Expenditures section of the Centers for Medicare and Medicaid Services National Health Accounts.

Medicare Prescription Drug Plans (PDP) (a plan that just offers prescription drug benefits). In 2006, there were 1429 total PDPs offered throughout the nation, with most states offering about forty PDPs. The majority of PDPs are offered by a dozen national or near national companies. The second alternative is Medicare Advantage (MA) plans, plans that provide all Medicare benefits, including prescription drugs, such as HMO, PPO, or Private FFS plans. There were 1314 total plans nationally in 2006. Finally, beneficiaries could retain their current employer/union plan, as long as coverage is creditable or at least as generous (i.e. actuarially equivalent) as the standard Part D plan, for which they would receive a subsidy from the government.

Under Part D, recipients are entitled to basic coverage of prescription drugs by a plan with a structure actuarially equivalent to the following: none of the first $250 in drug costs each year; 75% of costs for the next $2,250 of drug spending (up to $2,500 total); 0% of costs for the next $3,600 of drug spending (up to $5,100 total, the donut hole); and 95% of costs above $5,100 of drug spending. Over 90% of beneficiaries in 2006, however, are not enrolled in the standard benefit design, but rather are in plans with low or no deductibles, flat payments for covered drugs following a tiered system, or some form of coverage in the donut hole. The main requirement for plans is that they must have equal or greater actuarial value than the standard benefit.[2] The government also placed restrictions on the structure of the formularies that plans could use to determine which prescription medications they would ensure. Overall, Part D sponsors have great flexibility in terms of plan design. Many insurance companies sponsored multiple plans of differing levels of premiums and coverage generosity. Arranging the data into cells by plan sponsor and state, we find that only a quarter of the cells have only one plan, and 58% contain three plans. In those sponsor/state cells with multiple plans, most sponsors offer one standard plan and one or two enhanced plans.

Enrollment in Part D plans was voluntary for Medicare eligible citizens, although Medicare recipients not signing up by May 15, 2006 were subject to a financial penalty if they eventually joined the program (to mitigate adverse selection in the choice of joining the

---

[2]Cover Memo for Medicare Part D Benefit Parameters: Annual Adjustments for Standard Benefit in 2007 (CMS).

program). One group, however, was automatically enrolled: low income elders who had been receiving their prescription drug coverage through state Medicaid programs (the dual eligibles). These dual eligibles were enrolled in Part D plans by default if they did not choose one on their own. The Part D plans for dual eligibles could charge copayments of only $1 for generics/$3 for name brand drugs for those below the poverty line, and only $2 for generics/$5 for name brand drugs for those above the poverty line, with free coverage above the out of pocket threshold of $3600.[3]

Despite reluctance voiced before the legislation passed, there was enormous interest from insurers in participating in the Part D program. By November 2006, 3,032 plans were being offered to potential Part D enrollees. Every county in the nation had at least 27 plans available; the typical county had 48 plans, while some counties featured more than 70 choices, primarily due to high number of MA plans.[4]

Enrollment in the new Part D program was initially fraught with problems, but in the following months the federal government was able to iron out many of the problems that had arisen during the initial transition. As of June 2006, there were 10.4 million people enrolled in stand alone PDP plans, 5.5 million people enrolled in MA plans and about 6 million dual eligibles.[5]

Yet 73% of people over 65 felt that the Medicare prescription drug benefit was too complicated, while 91% of pharmacists and 92% of doctors expressed this concern. When asked if they agree with the statement Medicare should select a handful of plans that meet certain standards so seniors have an easier time choosing, 60% of seniors answered Yes.[6]

Despite these reservations, there were no signs of diminished plan choice in subsequent

---

[3]In addition, two other groups receive substantial subsidies those found eligible for Low Income Subsidy (LIS) or for Partial Subsidy by the SSA. To qualify for LIS, beneficiaries must have income less than 135resources less than $7,500/individual or $12,000 couple. This group received benefits comparable to the dual eligibles with incomes above 100% of poverty. To qualify for Partial Subsidy, beneficiaries must have income at 135%-150% of poverty and resources less than $11,500/individual or $23,000/couple. This group can enroll in plans with a $50 deductible, a 15% copayment up to the out of pocket threshold, and $2/$5 copayments above that point. In addition, premiums are fully paid by the government up to 135% of poverty, and then partially subsidized up to 150% of poverty.

[4]Details on number of plans in a median county obtained from Prescription Drug Plan Formulary and Pharmacy Network Files for 2006, provided by CMS.

[5]Enrollment data (rounded) taken from CMS, State Enrollment Data spreadsheet at http://www.cms.hhs.gov/PrescriptionDrugCovGenIn/02_EnrollmentData.asp#TopOfPage. Enrollment numbers also available at http://www.kff.org/medicare/upload/7453.pdf.

[6]Kaiser Family Foundation and Harvard School of Public Health (2006).

years. The number of PDPs increased by about 30% in 2007, from 1,429 to 1,875 and remained at this level in 2008.[7]

2.2.2 Issues of Elder Choice in Part D

Duggan et al. (2008) provide a detailed overview of many of the economic issues raised by the Medicare Part D Program. The use of this private delivery device, with such a multiplicity of choices, is a novel feature of the Part D legislation. Standard economic theory would suggest that this is a beneficial plan feature: allowing individuals to choose across a wide variety of plans that meet their needs, rather than constraining them to a limited set of choices being made by the government, can only increase welfare in the standard model in a partial equilibrium setting.

But there are reasons to believe that the standard model is insufficient, particularly for a population of elders. There is growing interest in behavioral economics in models where agents are better off with a more restricted choice set, as nicely reviewed in Iyengar and Kamenica (2006). Recent theoretical work shows that the traditional more is better principle may be reversed in some choice set contexts, for example when the presence or absence of options conveys information (Kamenica (2008); Kuksov and Villas-Boas (2010)) or when agents have preferences with regret (Irons and Hepburn (2007); Sarver (2008)). And a growing body of empirical work shows that individuals are less likely to participate in markets where they face more choice; decisions to purchase a good (Iyengar and Lepper (2000); Boatwright and Nunes (2001), take a loan (Bertrand, Karlin, Mullainathan, Shafir, and Zinman (2005)) or enroll in a 401(k) plan (Sethi-Iyengar, Huberman, and Jiang (2004)) are found to decrease when participation requires choosing from a larger set of alternatives.

Iyengar and Kamenica (2006) find that not only the decision to participate in a market, but also the nature of choice itself, is affected by the size of the option set. They investigate choice over asset allocation in both laboratory and real-world (pension plan choice) settings, and find that individuals opt for safer investments when faced with a larger range of risky choices. In particular, they find that the presence of more investment options in a 401(k)

---

[7]Data on 2008 plans taken from CMS 2008 PDP Landscape Source (v. 09.25.07) available at http://www.cms.hhs.gov/prescriptiondrugcovgenin/.

76

plan leads to more frequent choice of money market or bond options rather than equity investment. Iyengar and Lepper (2000) also find that satisfaction with choices made falls with the size of the choice set in several experimental settings.

Another recent literature has shown that the nature of how choices are presented can have important impacts on choice. For example, Hastings and Ashton (2008) examine financially illiterate individuals choosing across retirement funds in Mexicos privatized Social Security system. They find that presenting information on plan administrative fees in pesos, rather than in percentage terms, causes a significant shift in choice towards lower fee plans. In another study, Chetty, Looney, and Kroft (2007) find that consumers are much more sensitive to tax rates when the tax burden is included in posted prices rather than added at the register.

These issues may be paramount within the context of the elderly, given that the potential for cognitive failures rises at older ages. Salthouse (1996) shows clear evidence that the performance on a series of memory and analytic tasks declines sharply after age 60. Part of the reason for this may be the rise in incidence of dementia with age; starting at age 60, dementia rates roughly double every five years (Fratiglioni, De Ronchi, and Aguero (1999)). A recent study by Agarwal, Driscoll, Gabaix, and Laibson (2010) shows that in ten different contexts, ranging from credit card interest payments to mortgages to small business loans, the elderly pay higher fees and face higher interest rates than middle-aged consumers. These types of findings raise particular concern about choice in the Part D context.

2.2.3   Previous Studies of Part D Choice

We are aware of only three previous studies of these issues in the context of Part D. The first is a set of studies by Dan McFadden and colleagues, as summarized in Heiss, McFadden, and Winter (2006). These researchers surveyed a set of elders about their plans for enrolling in Part D programs, and evaluate whether enrollment intentions in the plan were rational given the penalties for delay. They find that 71% of potential enrollees were making the appropriate decision (under various assumptions about discount rates, etc.), while 10% of enrollees did not intend to enroll when it would be in their interests to do so, and 19% intended enroll when it would be in their interest to delay. Thus, for most potential enrollees, the decision over whether to enroll seems to be made rationally.

Their findings are less sanguine, however, for choice of Part D plan. This survey offered individuals a choice of the standard plan described above versus alternatives that provide different levels of insurance coverage (e.g. catastrophic only, complete coverage, etc.), with corresponding actuarially fair premiums. They find that only about 36% of enrollees choose the cost-minimizing plan, and they do not place much value on the insurance aspects of more comprehensive plans. They conclude that consumers are likely to have difficulty choosing among plans to fine-tune their prescription drug coverage, and do not seem to be informed about or attuned to the insurance feature of Part D plans.

While this is an interesting set of findings, it provides only a preliminary look at the crucial issue of plan choice. These conclusions are based on data which do not contain precise detail about the prescription drugs used by individuals; assumptions about utilization are made using aggregate imputations from other sources. Moreover, this is based on hypothetical choices across a set of non-existing plans; individuals may become educated about the program when they are actually faced with plan choices. Thus, the failures of choice documented by Heiss et al. (2006) may not hold when we use data on actual individual utilization and choices.

A recent paper by Lucarelli, Prince, and Simon (2008) uses aggregate data on plan market shares to conduct a study of how plan features impact demand and to undertake a welfare analysis of choice restrictions. They estimate sizeable welfare losses from limiting the option set facing seniors. But they do so in a framework which assumes that seniors are choosing optimally so that by definition restricting the choice set can only be harmful. Without individualized data on plan choices, they are unable to evaluate the underlying efficacy of plan choice.

Most closely related to our work is a recent field experiment by Kling, Mullainathan, Shafir, Vermeulen, and Wrobel (2008). They examine how providing people with information about the relative costs of each of the available plans in 2007 computed using their 2006 claims impacts their choices. They find that individuals who receive this intervention are more likely to switch plans, and more likely to end up with lower predicted and realized costs. Using our richer dataset on patient claims, we are able to model the individualized risk characteristics of plans in addition to looking just at average costs. Our model is also

more general in terms of sample and implications. While they investigate the consequences of one particular intervention on a sample of patients at a single hospital, our model allows us to calculate the potential welfare gains from reforms which change the structure of the choice set, and to do so for a large fraction of Medicare Part D enrollees.

2.3 Data

Our primary data source is a longitudinal sample of prescription drug records from the Wolters Kluwer (WK) Company. They are the largest switch operator in the prescription drug market: they collect the electronic claims from pharmacies and pass them on to the Pharmacy Benefit Managers (PBMs) and insurance companies that will pay the claims. After adjudicating the claim, it is passed back through the switch to the pharmacy. WK performs this function for a large sample of pharmacies throughout the U.S. Once pharmacies are in their sample, there is a 93% chance that they remain enrolled, so this is effectively a longitudinal sample of pharmacies. On average the claims captured by the WK system represent almost 31% of all 3rd party prescription claims filled in the U.S.[8] The geographic distribution of these data is very closely representative of the geographic distribution of 3rd party claims as well; the correlation between the WK market share and the overall 3rd party market share across each of the states is 0.86. WK keeps a longitudinal file that tracks prescription drug use for more than 100 million persons in the U.S. They have made available to us for research purposes a longitudinal sample of prescription claims for any individuals age 65 and over in 2005. These data are crucial because they are the only available data (of which we are aware) that contain information both on specific drug utilization by elders and on plan choice. Information about specific drug utilization is key because plan costs vary tremendously based on drug utilization, as we discuss below.

We begin with a sample of 1.53 million elders who (a) have a Part D claim with coverage of any sort (e.g. past the deductible), (b) are not employer-insured, dual eligibles or eligible for low-income subsidies/partial subsidies, (c) have claims for only one region of the country, (d) have no claims with missing payment information, (e) are in the sample of consistently

---

[8]Figure based on data provided by WK for Q3:2006.

reporting pharmacies, and (f) have data for both 2005 and 2006. This data file has a rich set of information about every drug claim for individuals in the longitudinal sample, including information on: month in which the prescription was filled; county of location of the pharmacy; a de- personalized patient id which allows longitudinal patient linkages; patient age; NDC code for the drug; quantity measures (days supply, dosage, package size); patient and insurer payments; and insurer or PBM name. The total drug price is computed as the sum of patient and insurer payments; one or both of these values are blank for a small fraction of claims, and individuals with these claims are dropped. We constructed a crosswalk between the drug ID variable in the claims data and the NDC numbers listed on the formulary using a file from First Data Bank. To allow for the possibility of substitution if individuals were enrolled in a different plan, we constructed a coarser drug ID variable which is unique only up to drug name and assumed that individuals could substitute to the cheapest drug with the same name were they enrolled in an alternative plan.

WK has created a sample for us that links longitudinally all claims from elders that fill prescriptions at a pharmacy in their sample. Thus, there are three types of attrition from the sample. First, elders may die (in which case we still observe all of their claims). Second, pharmacies may enter or leave the sample. This can be addressed by using only pharmacies that are continuously in their sample.[9] Finally, individuals may switch pharmacies. If the switch is to a pharmacy within the WK sample, then the company does a detailed statistical match to ensure that the patient is captured and matched to other prescriptions (based on the de-identified form of data fields such as first name, last name, date of birth, year of birth, gender, health insurance id and zip code). If the switch is outside of the WK sample, then the individuals will be lost to this sample.

Unfortunately, there is no way to capture such transitions. So long as these transitions are not correlated with premiums, it will not bias our price elasticity estimate; so long as they are not correlated with plan cost sharing provisions, they will lead only to a general understatement of out-of-pocket spending that will lead us to overstate that coefficient. We can assess their importance by taking advantage of the fact that Wolters Kluwer provided us

---

[9]A store is flagged as continuously enrolled provided that the store does not miss more than 11 days (including weekends and holidays) of reporting in a month.

with a coverage level variable which indicates the proportion of pharmacies in each county which are covered by Wolters Kluwer. We have rerun our models on the 10% of counties where WK covers at least 40% of all third-party prescriptions, and our results are very similar to what we report below. This suggests that attrition is not significantly biasing our results.

2.3.1 The CMS Plans Database

We obtain information on availability of Part D plans and specific plan features directly from four files provided by CMS: the plan information file, the beneficiary cost file, the formulary file and the geographic locator file. The plan information file lists plan names and identifiers, and regions/counties in which plans are offered. The beneficiary cost file contains copays and coinsurance rates for different tiers of each plan. The formulary file contains a list of all the drugs that are included on the formulary for each plan. The geographic locator file allows us to identify all the Social Security Administration (SSA) counties that correspond to different PDP and MA regions.

The major strength of the CMS data is that it allows us to fully parameterize any elders plan choice set based on their location. We have used these data to build a cost calculator that mimics the calculator provided on Medicares web site. This calculator uses a given set of prescriptions for a given elder to compute their projected out of pocket spending in each plan available in their county.

2.3.2 Matching patients to their Part D Plan

One challenging aspect of the WK data is that we know each patients county and the name of the company that provides the Part D plan that is covering each prescription, but not specifically which Part D plan offered by that company is covering the prescription. For example, we know that an elder is covered by a Humana product, but not whether it is Humana Complete, Humana Enhanced, etc.

Fortunately, we can resolve this matching problem in most cases by using a combination of county code, company name, and copayment structure. For each claim and each of the plans within the same company offered in a particular county, we check if the copay that the

patient paid for this claim matches any of the prescribed copays of the plan. We assign a person to a plan if most of their claims match to the same unique Part D plan. We carry out this exercise for each month. To confirm that a person has been matched to a correct Part D plan, we look at all the months together and insist that a person be consistently matched to the same plan in each month from June 2006 on, since enrollment into Part D plans was open until May 15th 2006.

Of the approximately 1.53 million individuals in our sample, 50.5% were matched to Part D plans. The remainder were excluded either because they had a large number of non-Part D claims (implying that they have some other form of coverage), because they had too few claims to reliably match, or because their copays were inconsistent with the copays listed for Part D plans in their region.

Of the matched individuals, 57.1% were uniquely matched to a Part D plan, and 42.9were multiply matched (meaning that more than one Part D plan was consistent with their co-pays). While the unique matches are clear, excluding multiple matches leads us to misstate the proportion of enrollment in some plans. This problem is especially severe among Humana plans because Humana offers several plans which differ only in the deductible and donut hole coverage and thus cannot generally be distinguished on the basis of copays. While comprising 20% of all matches, Humana plans are only 10% of unique matches. To deal with this problem, we include both unique and multiple matches, with multiple matches randomly assigned to one of the plans to which they are matched with probability equal to the proportion of total national enrollment in that plan in 2006.

2.3.3   Construction of Out of Pocket Cost Variables

The total enrollee costs of Part D can be decomposed into premiums, which are known for certain at the time of plan choice, and the distribution of out of pocket costs given the information available at the time when plans are chosen. Our focus is on estimating the distribution of costs given all of the information potentially available to individuals at the time when they make their choice. There are three reasons that estimating this distribution is challenging: first, we only observe realized out of pocket costs for the plan in which an individual is enrolled; second, we observe only a single realization of out of pocket costs for

each individual (making it impossible to compute a variance measure); and third, we do not observe all of the information available to individuals at the time when they make their choice.

To handle the first difficulty, we assume that the set of 2006 claims is fixed and would remain constant had the individual in question chosen a different plan; that is, we assume no moral hazard. This assumption allows us to use the calculator to determine what each individuals realized costs would be for each plan in their choice set. Given typical estimates of the elasticity of prescription drug utilization in the range of 0.2 to 0.5, and considering that this would only impact our results to the extent that individuals have sufficient foresight to take into account future utilization effects in their plan choices, this is a fairly innocuous assumption, as shown in Appendix A available in the online version of this paper.

To handle the second difficulty, we sample realized costs from 200 individuals who are identical to the individual in question at the time when the plan choice is made. In practice, we define identical as individuals with the same decile of 2005 drug expenditures, 2005 days supply of branded drugs and 2005 days supply of generic drugs; after extensive searching, we found that this combination provided the best prediction of 2006 prescription drug spending based on 2005 characteristics. We therefore assign each individual to one of 1000 cells demarcated by the interacted deciles of these measures. We restrict our sample to individuals for whom there are at least 200 other individuals in their cell, and we use these 200 individuals in each cell to compute both our rational expectations measure of utilization in 2006 (described below) and our variance measure.

The third difficulty is that individuals may actually know more than can be predicted given 2005 costs at the time when they make their plan choices. Intuitively, we can attempt to determine whether individuals know more than can be predicted given 2005 costs by analyzing whether their choices are sensitive to the component of the variation in realized costs across plans which cannot be predicted given 2005 characteristics. We discuss a model of this type in

2.4　Final Sample Creation

Under Part D individuals could enroll not only in a stand alone PDP plan, but also in a more comprehensive MA plan; we distinguish between individuals matched to MA and those matched to PDP plans based on copay and exclude the former. We focus just on PDP plans (and therefore, just on individuals who chose PDP plans) because MA plans involve broader health care decisions which are beyond the scope of our data (e.g. regarding HMOs and fee-for-service plans). This exclusion is justified by the independence of irrelevant alternatives assumption that underlies our logit modeling, as discussed (and tested) further below. We also exclude individuals who have fewer than 500 observations in their state or fewer than 100 observations in their brand/state cell to increase the speed of estimation of the model by reducing the required number of brand/state fixed effects; this restriction has no effect on our final results.

Our final sample consists of 477,393 individuals. The typical patient in this sample is almost 75 years old, three-fifths are female, and they have an average of 34 claims per year. Their total prescription drug spending averages $1,711 per year. While some individuals were enrolled in Part D for the full year, others enrolled as late as May. The average total premiums paid after enrollment was $287 and the average OOP costs paid out over the same period was $666. This is the sample used in the efficient frontier analysis below. In our conditional logit models, we randomly subsample 20% of these individuals for computational reasons. We estimate the more computationally demanding random coefficients models on a randomly chosen subsample of 15,000 patients.

The distribution of enrollees across Part D plans is highly correlated in this final sample with the national facts on PDP enrollment provided by CMS. The correlation between the share by brand in our sample and the CMS sample is 0.98, and the correlation between the share of our sample in the top 10 plans is correlated with the CMS reported share in those plans at 0.89 (the correlation for the top 100 plans is 0.91).

2.5 Facts on Plan Choice

To motivate our regression framework, Figure 2-1 shows the basic facts on the relationship of plan choice to total plan costs. For each individual in the data, we estimate the total cost of enrolling in each PDP plan in their county, adding both premiums and expected out of pocket costs. We then estimate the difference in total costs between the plan chosen by that individual and the lowest cost plan in their county. For this exercise, we use a perfect foresight model of expectations, using actual 2006 expenditures to estimate the costs that individuals face in each plan.

As Figure 2-1 shows, only 12.2% of individuals choose the lowest cost plan in their state. Indeed, on average, individuals could save 30.9% of their total Part D spending by choosing the lowest cost plan rather than the plan they chose. If we redo these calculations using actual 2005 expenditures, or predicted 2006 expenditures based on 2005 expenditures rather than actual 2006 expenditures, we find even stronger deviations from the lowest cost plan.[10]

Of course, individuals are not simply choosing a fixed payment stream when choosing a Part D plan; individuals who are highly risk averse may explicitly be choosing plans with higher mean expenditure to protect themselves against variance in expenditure. Yet this does not seem to be the case. Even if we only include plan choices where the variance is non-increasing, over 70% of enrollees could have chosen a lower cost plan, and the typical enrollee could have saved 23.3% of their Part D expenditures without raising their variance.[11]

The explanation for these facts is shown in Figure 2-2, which shows the choice set for individuals in California. The X axis in this graph is the mean of total costs for each plan, and the Y axis is the average standard deviation in costs (where the standard deviation is computed using the 1000 cell method, and the average is taken across individuals). In this graph, there is a clear efficient frontier of plans which dominate others in terms of both

---

[10]It appears that some plans may have offered low premiums in 2006 in order to entice consumers to choose their plan in the first year of the Part D program before raising their premiums in subsequent years. This behavior should not impact our analysis except insofar as there are large switching costs because consumers have the option to switch plans after each year, but one might still wonder to what extent the above results are driven by such plans. To assess this issue, we repeated the above analysis using the 2007 premiums for all plans and found that the average potential cost savings fell slightly from 30.9% to 25%.

[11]The fact that this number is smaller than the 30.9% number is because we are searching for cost savings over a small set of plans, not because individuals are especially sensitive to risk, a point we document further below.

cost and variance. This graph masks considerable heterogeneity across individuals: different plans lie on the efficient frontier for different individuals, so the fact that a plan lies off the efficient frontier in this graph does not imply that it is suboptimal for each individual. Nonetheless, most of the plans are well off the efficient frontier, meaning individuals could have either lowered their mean costs or their variance by picking a different plan.

As we will document below, one reason for the large amount of choice off the efficient frontier is that individuals consider plan characteristics in making their choices but not how those plan characteristics matter for themselves. This is perhaps best illustrated by a simple examination of the decision to choose a plan with donut hole coverage. Figure 2-3 shows the probability of choosing donut hole coverage, and the financial implications of doing so, sorted by 2006 spending percentiles; the results are once again similar for other measures such as 2005 actual spending or 2006 predicted spending. The bottom line shows the percent of the population at each percentile choosing donut hole coverage; the top line shows the average savings for individuals in that quantile from switching from the lowest cost plan in their region which offers donut hole coverage to the lowest cost plan that does not.

The plans which offer donut hole coverage actually have slightly inferior coinsurances relative to the lowest cost non-donut hole plans in the initial coverage range, and so the cost of donut hole coverage is rising with expenditures until the point when individuals become likely to enter the donut hole.

The results here are striking: the percentage choosing donut hole coverage is virtually flat throughout the spending distribution at around 10%. Even if individuals are willing to pay extra in mean costs for the protection provided by donut hole coverage, it is hard to rationalize the fact that the same proportion of individuals in the 10th and 85th percentile of the spending distribution choose donut hole coverage.

2.6 Base Model of Part D Plan Choice

In this section, we extend the efficient frontier analysis presented above by considering several discrete choice models. These models serve three general purposes in our setting. First, they allow us to control for additional plan characteristics such as plan quality. Second, they allow

us to understand more precisely how preferences combine with choice set characteristics so we can forecast how individuals might choose in counterfactual choice environments. Third, they allow us to quantify the welfare consequences of choices. We begin by specifying a CARA utility model with a normally distributed cost distribution:

$$U(C) = -\exp(-\gamma(W - C)) \text{ where } C \sim N(\mu, \sigma^2) \tag{2.1}$$

In this case, indirect utility is given by:

$$u(\mu, \sigma^2) = EU(C) = -\alpha \exp(\gamma\mu + \frac{1}{2}\gamma^2\sigma^2) \tag{2.2}$$

where $\alpha = -\exp(\gamma W)$ is a constant. A first-order Taylor expansion about the point $(\mu', \sigma^{2'})$ yields:

$$u(\mu, \sigma^2) \approx u(\mu', \sigma^{2'}) - \alpha\gamma u(\mu', \sigma^{2'})(\mu - \mu') - \frac{1}{2}\alpha\gamma^2 u(\mu', \sigma^{2'})(\sigma^2 - \sigma^{2'}) \tag{2.3}$$

We can write total costs as $C = \pi + OOP$ and since π is known for any given plan, $Var(C) = Var(OOP) = \sigma^2$ and $\mu = E(C) = \pi + E(OOP) = \pi + \mu^2$. Adding an error term (and dropping constant terms) we can rewrite the approximate indirect utility as:

$$u = -\alpha\gamma u(\mu', \sigma^{2'})(\pi + \mu^*) - \frac{1}{2}\alpha\gamma^2 u(\mu^{*'}, \sigma^{2'})\sigma^2 + \epsilon \tag{2.4}$$

This maps into a conditional logit model of plan choice where the utility of individual i from choosing plan j is given by:

$$u_{ij} = \pi_j\beta_0 + \mu_{ij^*}\beta_1 + \sigma_{ij}^2\beta_2 + x_j\lambda + q_{b(j)}\delta + \epsilon_{ij} \tag{2.5}$$

with $\beta_0 = \beta_1 = -\alpha\gamma u(\mu^{*'}, \sigma^{2'})$ and $\beta_2 = -\frac{1}{2}\alpha\gamma^2 u(\mu^{*'}, \sigma^{2'})$. In this equation x represents any financial plan characteristics which impacts choice, $q_{b(j)}$ represents plan quality ratings and other non-financial aspects of plans (which vary only across brands), and ϵ_{ij} are i.i.d. type I extreme value random variables.

This gives us $\gamma = \frac{2\beta_2}{\beta_1}$ which allows us to map the ratio of the coefficients on the variance of costs and the coefficient on the mean of costs into the coefficient of absolute risk

87

aversion. This assumes that wealth is constant across all states of the world: the only risk facing individuals is uncertainty about the distribution of out of pocket costs. The same expression would hold if we added idiosyncratic risk that was uncorrelated with prescription drug expenditures, but it is not implausible that there would be correlated risks: in states of the world where prescription drug expenditures are higher, other medical expenditures are higher as well. Such correlated risks would tend to bias upwards our already low estimates of risk aversion.

We include in our model several financial plan characteristics beyond premiums, out of pocket costs, and the variance of out pocket costs. These are: the deductible of the plan; a dummy for whether the plan covers all donut hole expenditures; a dummy for whether the plan covers generic expenditures in the donut hole only; and a cost-sharing index. The cost sharing index is calculated for each plan as the average percentage of expenditures covered by the plan between the deductible and the donut hole. This variable differs from expected out of pocket costs in that it has the same value for everyone in the sample for each plan, and because it is not directly impacted by whether plans have deductibles or donut hole coverage. We also include two measures of plan quality: the share of the top 100 drugs used by elders that is included in the plans formulary and a quality index. The quality index is computed by CMS on a 1-5 scale by aggregating consumer ratings at the brand level collected along 17 dimensions which are categorized as Customer Service, Drug Pricing Information (availability / rate of price changes), and Using Your Plan to Get Your Prescription Filled.

Identification is a natural concern in this context. All of the plan characteristics included in our model may be endogenous due to unobserved demand factors, and they may be biased by correlation with unobserved plan characteristics. To address this concern, we observe and include in our model all of the publicly available information that might be used by individuals to make their choices. We also consider models where we control for a full set of brand dummies, as well as a full set of interactions of state dummies with brand dummies. When we include brand dummies, the coefficient on the quality index (which is measured at the brand level) is no longer separately identified although it can be recovered by a GLS regression of these dummies on the quality variable. When brand-state dummies are included, coefficients on plan characteristics such as the premium, deductible and donut

hole coverage are identified by the variation across plans offered by the same brands in a given state.[12]

Even with these fixed effects, it is possible that premiums are endogenous because they are set based on brand-state specific assessments of demand conditions. If premiums are higher in regions where insurers anticipate more demand for their particular plan (relative to other plans offered by the same insurer), our estimate of the coefficient on premiums will be biased towards zero since individuals will appear to be less averse to higher premiums. To the extent that these factors make high premiums appear less undesirable than they actually are, our conclusion that premiums are overweighted relative to out of pocket costs would be strengthened, as would our estimates of the welfare loss due to consumer mistakes.[13]

2.6.1 Restrictions on Preferences

The model laid out above suggests three natural restrictions on preferences which extend the efficient frontier concept to the discrete choice setting.

$$\text{Restriction } 1 : \beta_0 = \beta_1 \tag{2.6}$$

This restriction states that the coefficient on premiums should equal the coefficient on expected out of pocket costs. Controlling for the risk characteristics of plans, individuals should be willing to pay exactly one dollar in additional premiums for coverage which reduces expected out of pocket costs by one dollar. If this restriction fails to hold, individuals are not choosing on the efficient frontier: they could switch to alternative plans with comparable risk characteristics but lower total costs.

$$\text{Restriction } 2 : \lambda = 0 \tag{2.7}$$

---

[12]For instance, in many states Humana offers a Standard plan with lower premiums but limited coverage, an Enhanced plan with higher premiums but no deductible, and a Complete Plan which offers superior cost sharing and full donut hole coverage at much higher premiums.

[13]We did attempt estimating the models reported below using two instruments using the control function approach: these were the average premium for a given plan in all states where the plan is offered (designed to avert local demand shocks) and a marginal cost instrument constructed using the average covered expenditures for individuals enrolled in the plan. In both cases, the magnitude of the coefficient on premiums increased in the IV models. We are not confident that the exclusion restriction is satisfied for either of these instruments, so we continue to estimate the model without an instrument below.

This restriction states that financial plan characteristics other than premiums, expected out of pocket costs and the variance of out of pocket costs do not impact choices. Individuals should not care about deductibles, donut hole coverage or copays per se; they should only care about these factors to the extent that they impact the distribution of out of pocket costs. Once we control for this distribution, these factors should be redundant.

$$\text{Restriction } 3 : \beta_2 < 0 \tag{2.8}$$

This restriction states that individuals should be risk averse.

While these restrictions follow naturally from utility maximization with full information and standard preferences, the model from which they are derived makes several important functional form assumptions: we assume that the distribution of out of pocket costs can be summarized by its mean and variance, that indirect utility is a linear function of this mean and variance, and that the errors are i.i.d. type I extreme value. In Appendix A, we show that the restrictions assumed in the previous section still hold even when these functional forms assumptions are weakened.[14]

Of course, it is always possible to write down preferences that would violate the above restrictions, but these restrictions are generally compatible with commonly used expected utility functions given the observed cost distributions.

2.7 Results

In this section we present the results from the conditional logit model described above. At the outset, it is important to mention that the results we present here are not very sensitive to misspecification or measurement error. Appendix A-C (available in the online version) investigate these issues extensively via simulation. We take the distribution of realized costs observed in the data for each of the 1000 cells, assume it is the true distribution, and simulate individuals choices using known utility functions by assuming individuals maximize expected

---

[14]In particular, we simulate choices using the actual distribution of costs and several commonly used utility functions (CRRA, CARA) with varying levels of risk aversion. In some cases, the restrictions do not hold exactly, but the violations are much smaller in magnitude than we observe when we estimate the model using actual choices.

utility. We then investigate to what degree the above restrictions on choice are violated if we add various types of known misspecification or measurement error. For example, if the true utility function is CARA with risk aversion 1, does estimating the model in the logit framework with linearized indirect utility generates choice inconsistencies? What if we only observe a noisy measure of OOP costs which contains attrition and measurement error? The upshot of this analysis is that with very large risk aversion or substantial amounts of measurement error, we do sometimes observe statistically significant choice inconsistencies, but these are always smaller in magnitude than we report below and have inconsistent signs. In our discussion below, we make this comparison more explicit.

2.7.1 Base Results

Table 2.1 reports the results from several conditional logit models. Model (1) includes only the premium, realized out of pocket costs,the variance of out of pocket costs, and the quality variables. As noted in the discussion of the cost variables, expected out of pocket costs meaning the individuals expectation of out of pocket costs at the time of plan choice is not directly observed, so we use realized costs as a proxy for expected out of pocket costs. This proxy has noise (where noise includes the component of realized costs unknown to the individual at the time when the choice is made) and so its coefficient is biased downwards. We address this problem at length in the next section and show that it does not much impact our conclusions.

The cost variables premiums and out of pocket costs are measured in hundreds of dollars. Model (1) therefore shows that a $100 increase in premiums leads to a 32% reduction in the probability that a given plan is chosen, implying an average elasticity of -0.75.[15]

There are two ways to interpret the remaining coefficients. First, we can divide by the premium coefficient in order to compute the willingness to pay in dollars for a one unit increase in the characteristic. Second, the coefficient itself can be interpreted as the percentage increase in the probability that a plan is chosen from a one unit increase in

---

[15]The implied elasticity varies across plans based on premium level and market share. The 32% number given in the text is derived from the equation . Thus, for which holds for a large number of plans, we can interpret as the percentage change in associated with a one unit change in .

the characteristic provided that probability is small (as it is for most plans). When we compare models estimated using actual choices with simulations, we compare the implied dollar value of plan characteristics computed by dividing by the coefficient on premiums. This is because the scale of the coefficients is determined by the proportion of choices explained by the included variables, and in the simulations, this scaling factor is arbitrarily set by whatever standard deviation we assume for the structural error term. We can and do use the simulations to compare ratios of coefficients with the actual results, but nothing substantive can be inferred from the absolute magnitude of the coefficients in the simulation. Two points about the model (1) results are noteworthy. First, the coefficient on out of pocket costs is only about as large as the coefficient on premiums, violating Restriction 1. Second, the coefficient on the variance term is negative and significant, but extremely small, implying risk aversion substantially less than we obtained in the simulations with CRRA = 1.

Model (2) adds additional covariates to control for deductibles, donut hole coverage, average cost sharing, formulary coverage and plan quality. Many of these covariates enter the model with significant coefficients. When we add plan characteristics, the coefficient on premiums increases suggesting that it was initially biased downward due to omitted variable bias. The coefficient on the variance term drops even further once we add a control for the # of the most popular 100 drugs which are included in the plans formulary. One explanation for this is that that while individuals prefer plans which cover more drugs, they do not have sufficient foresight to choose plans which cover drugs which they (or at least people in their cell) might need in the future but are not already taking. Alternatively, it may be that there is substantial measurement error in the variance term, and that # of top 100 drugs is a proxy for the variance. Models (3) and (4) add brand dummies and brand-state dummies respectively. The coefficient on premiums actually shrinks once we include brand-dummies, but the effects of the premium remain large; a $100 increase in annual premiums leads to a 50% reduction in the probability that a plan is chosen, corresponding to an average elasticity of -1.17. The coefficient on out of pocket costs has similar magnitude across all of the models, which reflects the fact that it is identified based on individual variation. In columns (3) and (4) the coefficient on the premium is more than five times as large as the coefficient on out of pocket costs.

The coefficients on plan characteristics are also very large in all specifications. Controlling for the out of pocket cost consequences, model (4) which has the smallest plan characteristics - suggests that individuals are willing to pay over \$300 for full donut hole coverage, \$50 for generic donut hole coverage, about \$80 to go from a deductible of 250 to a deductible of 0, about \$80 to go from the plan with the least cost sharing (25%) to the plan with the most cost sharing (65%), and \$12 for each of the top 100 drugs which appear on the formulary. These numbers are not enormous, but they are an order of magnitude larger than the results in the simulations,[16] and have non-trivial consequences for the welfare evaluation of plan choice as we investigate in the welfare analysis section.

It is important to underscore the fact that these numbers are not the full hedonic value of those plan characteristics these are the willingness to pay above and beyond the implications of those plan characteristics for out of pocket costs.[17] Because individuals appear to be underweighting the individualized component of out of pocket costs, we can interpret these numbers as saying: individuals are willing to pay a price in premiums for desirable plan characteristics, but this price is insufficiently sensitive to their individual circumstance.

We noted above that in some of the Appendix simulations, plan characteristics had statistically significant coefficients even controlling for out of pocket costs, due to (imposed) measurement error or misspecification of the utility function. We might worry that the results using actual choices likewise reflect these factors rather than choice inconsistencies. The coefficients on plan characteristics estimated using actual choices imply larger dollar values for plan characteristics than do the simulations, however. For example, our estimate

---

[16]For comparison, in the simulations where true utility is CRRA with risk aversion of 10 (misspecification is increasing in risk aversion), we estimate - controlling for OOP costs - a \$9 value for full donut hole coverage, a -\$8 value for generic donut hole coverage, a \$33 value of moving from a \$250 deductible to a 0 deductible, a \$32 value of going from the plan with the least cost sharing to the most cost sharing, and -\$1 for each of the top 100 drugs which appear on the formulary (since these values are driven entirely by misspecification in the simulations there is no reason the signs should be sensible).

[17]To recover the hedonic value of plan characteristics, it would be necessary to add the values reported below to the values of plan characteristics implicit in our out of pocket cost measure. We can attempt to recover these values by regressing our OOP cost measure on plan characteristics controlling for individual fixed effects. This procedure will give biased results to the extent that plan characteristics not included in the regression impact out of pocket costs so we try not to lean heavily on this exercise when interpreting the above results. Nonetheless, we report the results of this regression here for reference: \$88 increase in OOP cost for a \$250 deductible, \$99 decrease in OOP costs for full donut hole coverage, \$12 decrease for generic donut hole coverage, \$810 decrease moving from cost sharing of 0% to cost sharing of 100% (or \$324 moving from the 25th percentile plan to the 75th percentile), and a \$1.3 decrease for each of the top 100 drugs covered.

of the implied value of donut hole coverage controlling for OOP costs is larger than the average OOP cost savings from donut hole coverage observed in our data, so even with infinite measurement error in OOP costs, the simulations could not match what we observe.

Thus, formal modeling of choice reveals a violation of all three of the preference restrictions we laid out above. The coefficient on premium is an order of magnitude larger than the coefficient on out of pocket expenditures; generalized plan characteristics enter the model highly significantly, even conditional on individual out of pocket risk; and individuals are not willing to pay more for plans with lower variance in expected spending.

One potential shortcoming of our last conclusion, however, is the reliability of our variance measure. We compute the variance by assessing the variability in out-of-pocket spending across a sample of similar individuals; we have tried several alternative specifications of risk preferences (such as including quantiles in the right-tail of the distribution of costs) and this does not appear to alter our results. We have also constructed alternative cells for a subset of the sample which also take into account information on the particular drugs each individual uses, and have found that our current method captures more than 90variance term across plans. It is not a foregone conclusion that the coefficient on our variance variable would be zero even if elders do not explicitly consider the cost of alternative plans under a range of hypothetical outcomes: if patients with a greater risk of getting sicker in the following year also chose plans with more coverage, we might expect this to show up in the variance variable even after we control for the average value of plan characteristics. The result appears to buttress our finding that individuals are insensitive to the individualized consequences of plan choice. Nevertheless, we place less weight on this last choice inconsistency because of concerns about the appropriate specification and measurement of risk in our setting some of which are discussed further in the robustness section.

2.7.2 Robustness

While our unique data set makes this analysis possible, the data do have a number of shortcomings. In this section we show the robustness of our basic conclusions to efforts to address these shortcomings.

First, we are able to match only 50.5% of our sample of 1.53 million individuals to Part D

plans, partly because we use very stringent criteria designed to minimize false matches. The cost of such an approach is that our matched sample may not be representative of the full sample of 1.53 million; in particular, the individuals in our sample have more claims because that makes it easier to match them to a plan. We therefore consider a less stringent matching strategy; in this less stringent match, we consider a copay as matched for a given claim if the appropriate copay appears anywhere on the claims formulary, even if the copay listed for drugs in that particular tier and days supply is not correct. We also accept matches if just 50% of overall claims are matched rather than requiring that this threshold be exceed in every month after the first Part D claim is observed. Using this strategy, we are able to match 1.28 million of our 1.53 million individuals, or 84%. The first column of Table 2.2 shows the robustness of our findings to this alternative measure. None of our main conclusions are altered. Our coefficients of interest are somewhat smaller than in the original sample, which may be because we of noise introduced by being less stringent in our willingness to accept matches, but the fundamental choice inconsistencies persist.

Second, a limitation of our approach is that we exclude individuals with low numbers of claims, since our matching algorithm requires enough claims to identify the plan copayment structure. This clearly leads to bias to our variance measure which is constructed by selecting matched individuals in the same cell; this is yet another reason why we have less confidence in the variance results than in our results for other forms of choice inconsistency. But there is no reason why this should lend a particular bias to our other results; if anything, we might think that failing to correctly specify the variance would make low OOP cost plans seem more desirable (since they also have lower variance), thus biasing upwards the coefficient on OOP costs. If this were the case, our results would be too conservative in reporting mistakes.

The exclusion of individuals with a small number of claims also raises selection issues, since these individuals may make better choices or have systematically different price elasticities. To address this problem, we have reestimated our model including additional individuals who we have excluded thus far. First, assuming some serial correlation in claims behavior, we can mimic the inclusion of low claims individuals in 2006 by including those with no claims in 2005 (we have excluded them thus far to allow for the creation of our variance measure and the rational expectations measure used below). We reran our base-case discrete choice

model using a random sample drawn from the larger sample including individuals with 0 claims in 2005. In our original sample, the variance variable was constructed by assigning each individual to one of one thousand cells based on 2005 claims. Because all of these individuals are identical in 2005, we assigned all of them to a single cell, and computed the variance by running 200 randomly chosen individuals in that cell through every plan. The results of this analysis are shown in column 2, and they do not differ much from our original results.

In column (3), we extend this analysis further by restricting the analysis to those who have zero claims in 2005 and fewer than 12 claims in 2006; this is the closest we can get to the zero claim sample in 2006 while still matching plans. The results are once again quite similar.

Third, we include in our analysis both unique matches and multiple matches, imputing the latter based on market shares. This adds a degree of noise to our estimation that could plausibly bias the results. We address this issue in two ways in Table 2.2. First, we reestimate the model including only those multiple matches where we can identify the plan with 95based on the relative enrollment given by CMS for the matched plans in the patients state (e.g. one of the plans has at least 19 times the enrollment of the other matched plans combined). Including unique matches, this was 75% of the original sample. Going further, in column 5, we reestimate the model only with those observations for which we can make a unique match. Neither of these changes alters any of our main conclusions.

Fourth, we do not know with certainty which individuals in our sample are dual eligibles. If we mistakenly include dual eligibles in our sample, we will measure them has having a lot of variation in OOP costs when they really have none, so we wrongly interpret them as insensitive to OOP costs. We are fairly confident, however, that we are excluding dual eligibles from our sample because they have such a limited range of possible copayments. For example, for 90% of the observations we use, more than one-quarter of their claims have copayments which are inconsistent with being a dual eligible (e.g. more than $5). To further ensure that problems identifying dual eligibles were not biasing our results, we have reestimated our model only on individuals where at least 50% of their claims are inconsistent with the copayment rates for duals. As we show in column 6 of Table 2.2, this has little

impact on our results.

Finally, in column 7, we show the impacts of decomposing further the aggregate quality index that we have used in our work thus far. We decompose the index into its three primary components. Doing so, we find that choice is positively associated with each of these quality components. The most important characteristic of quality to consumers appears to be the ease in filling prescriptions. Most importantly, decomposing the quality measure has no impact on our results.

2.8   Misspecification and Measurement Error

In the previous section, we presented results from a conditional logit model of plan choice and identified three apparent irregularities in choices. Our interpretation is that these results reflect consumer errors plan characteristics are more salient than are their implications for the distribution of out of pocket costs, and individuals are unable to compute the individualized risk characteristics of the alternative plans. In this section we consider two related alternative explanation: we have misspecified out of pocket costs because we have failed to appropriately model the information available to individuals at the time when they make their plan choice; and that our findings are driven by measurement error in out-of-pocket costs. Modeling Private Information

Thus far we have measured out of pocket costs using the realized cost measure constructed from 2006 claims. An alternative measure that we consider in this section we label our rational expectations measure. Recall that to create our variance measure we classified all individuals into 1000 cells defined by deciles of 2005 total spending, generic prescriptions and branded prescriptions, and ran the 2006 claims of 200 persons in each cell through the cost calculator for that plan. This procedure generates a distribution of costs for each patient and plan. Our rational expectations measure is defined as the mean of this distribution. Under the strong assumptions discussed above (CARA utility and a normal distribution of costs), the mean and the variance would completely summarize the impact of the cost distribution on utility; our simulations in the Appendix A show that they summarize this distribution well anyway even if these assumptions are relaxed.

It is useful to compare this rational expectations measure to the perfect foresight/realized costs measure we have been using. The latter measure is too broad in the sense that it includes information not available to individuals at the time when they choose (provided that is, that they do not know exactly what their drug needs and drug prices will be for the coming year). The former measure is too narrow in the sense that individuals may have private information at the time they choose beyond what can be inferred from their 2005 costs. If a patient learns they have cancer just prior to choosing their 2006 plan, they would correctly forecast that their drug needs would likely exceed the average of those with similar 2005 spending. We address these concerns by developing a model with which we can identify the information available to consumers at the time when they choose. The intuition behind this model is that we can determine if individuals know more than we can predict given just their 2005 spending by evaluating whether their plan choices are responsive to the component of 2006 spending which is not known in 2005.

The formal derivation of this model is presented in Appendix B. In summary, we augment our model in two ways. The first is to include a normally distributed term which captures the degree of private information: the difference between actual out of pocket spending in 2006 and what we would have predicted for 2006 based on our rational expectations model. If we were estimating a linear model, this would be comparable to estimating our model by instrumental variables, where we instrument the perfect foresight level of costs with our rational expectations cost measure, which is independent of private information. In our non-linear setting, the comparable correction is to include this noise term, which essentially amounts to estimating a random coefficients model with one extra parameter to identify the degree of private information.

Second, the measured variance from the 1000 cell exercise overstates the true variance in costs because some of this variation represents variation in realized costs which is unpredictable based on 2005 costs but is known to the individual at the time when they choose. We develop a correction for the variance based on the estimated degree of private information.

Table 2.3 reports the results from estimating this model. For computational reasons, we estimate this model on a much smaller sample by randomly selecting 15,000 patients from our earlier sample. Column (1) reports our earlier results, column (2) reports the earlier

specification on the new sample, and column (3) the results from adding the term for private information and correcting the variance. The model is estimated using the Laplace approximation developed in Jerry A. Hausman and Matthew C. Harding (2007) with bootstrapped standard errors, including controls for the various plan characteristics.

The results in Table 2.3 suggest that there is substantial private information: individual choices take into account about 60% of the variation in out of pocket costs which cannot be predicted given their cell. We also continue to find that the coefficient on realized costs is well below that on premiums, and that financial plan characteristics such as the donut hole and deductible continue to enter highly significantly in this model. Therefore, two of the major choice inconsistencies persist even when we model private information.

This model implies, however, that individuals know much of what their costs will be to each plan in their choice set in the coming year so there is little insurance motive. Under this interpretation, the variance in out of pocket costs is small for all plans because there is little uncertainty. This means that any measured response to the variance term would imply high levels of risk aversion, and that the standard errors in our estimates of risk aversion are much larger than we concluded in the model ignoring private information. The risk index in these models (obtained by dividing the variance coefficient by the premium coefficient and multiplying by 200) is comparable to what we obtained in our Appendix B simulations for CRRA = 3 with wealth = 17000.[18]

The bottom line from our models of private information is that our conclusions about the gap between the premium and out of pocket expenditure coefficients, and the powerful role for general plan financial characteristics in driving choice, are robust to a wide variety of specifications of out of pocket spending risk. Our conclusion about the low degree of estimated risk aversion, however, is more sensitive to the precise specification of the model.[19]

---

[18]As we highlight in Abaluck and Gruber (2008), the results reported in Table 2.2 are also consistent with an alternative model of information where individuals are not using all available information, but rather are paying attention only to a part of their prescription drug expenditures. For that portion to which they are attentive, individuals are rationally weighting premiums and out of pocket costs in the same way in making their decision. Yet individuals do not respond to variation in out of pocket costs beyond that portion. Under this interpretation, we find that the degree of private information is smaller and the coefficient on the risk measure once again becomes very small.

[19]To further explore robustness here, we also consider an alternative measure of predicted out of pocket costs: predicting those costs based only on use of regular drugs. A regular drug is defined as any drug for which the individual consumed at least 90 days supply in 2005. In our new measure, we construct the OOP

2.8.1 More General Measurement Error

The private information model can also be interpreted as correcting for a specific form of measurement error in our model: that arising from idiosyncratic variation across individuals in their knowledge about expected out of pocket costs at the time that they choose their Part D plan. Our private information model is the non-linear equivalent to a linear model that addresses measurement error in 2006 realized costs by instrumenting them by predicted costs based on 2005 characteristics. The fact that our conclusions are robust to controlling for private information is therefore equivalent to saying that instrumenting for idiosyncratic measurement error across individuals does not change our conclusions. However, idiosyncratic measurement error is only one of several types of measurement error in our out of pocket cost coefficient. In this section we consider robustness to alternative forms of measurement error. To model the impact of measurement error, we draw on the simulation model developed in Appendix A, which captures the no inconsistency baseline. Without measurement error, this model illustrates that choices under a variety of specifications of risk would not demonstrate the inconsistencies we see in our data. We can augment that analysis by adding measurement error in out of pocket costs to this simulation model, and using predicted costs as our measure of OOP costs in the simulation. We consider three alternative specifications of measurement error, and present the detailed results of our analysis in Appendix C.

The first is purely idiosyncratic individual-specific measurement error in forming expectations of out-of-pocket costs. Consistent with the discussion above, even with very large error of this form we find that our simulated out of pocket and premium coefficients are similar, and the coefficients on plan-specific plan charcteristics are very small; that is, our predicted out- of-pocket cost is effectively acting as an instrument for measurement error in this case.

We then consider a form of multiplicative error designed to mimic what might be observed if there were attrition due to patients having claims at pharmacies not included in our data. Our simulations then show that even sizeable attrition bias causes only a small upwards

---

cost variable assuming that drug use in 2006 will consist only of regular drugs in 2005. Our results are very similar using this alternative measure.

bias to the out-of-pocket coefficient, and causes only very modest coefficients on the plan characteristics, an order of magnitude smaller than what we observe in our logit models. The final simulation we consider is one in which there is a systematic plan-specific error, perhaps because of errors in assigning individuals to the correct plans.

We consider a multiplicative specification to capture the fact that the impact of such errors on OOP costs would likely be proportional to the number of claims an individual possessed. In most cases, this once again leads to an upwards-biased out-of-pocket cost coefficient, although if the error becomes large enough the bias becomes slightly downward (but much less than in our regressions). In this case we do estimate some sizeable coefficients on plan characteristics, with the coefficient on full donut hole coverage rising to 2/3 of what we observe in our regressions, but the plan characteristic coefficients are not consistently signed; we estimate a large positive coefficient on the deductible, for example, and a large negative coefficient on generic donut hole coverage.

Thus, our simulations do not provide any evidence to suggest that our consistent pattern of a small out-of-pocket cost coefficient and large plan characteristics coefficients are due to measurement error. Rather, they appear to correspond to true choice inconsistencies.

2.9 Heterogeneity

The independence of irrelevant alternatives assumption that underlies the conditional logit model places strong restrictions on how elasticities vary across plans and will lead to inconsistent estimates if preferences are heterogeneous across the population. To address this concern, we assess the robustness of our model to heterogeneity driven by both observed and unobserved factors. We first note that our model already allows for a substantial amount of individual variation: we estimated the coefficients on individualized out of pocket cost parameters. Nonetheless, it may still be the case that preferences vary in ways not included in our model. In terms of observed heterogeneity, we have reestimated our model for a number of separate samples: by gender; by age; and by tercile of the 2005 prescription drug expenditure distribution. In every case, we find that our results are very similar across all samples. In particular, each of these samples illustrates the three choice inconsistencies documented

thus far: the premium coefficient is many multiples of the out of pocket cost coefficient; financial plan characteristics enter significantly; and the estimated degree of risk aversion is very low.

We therefore turn to considering unobserved heterogeneity. We use the Laplace approximation developed by Hausman-Harding (2007) to estimate a model with normally distributed random coefficients on all included characteristics. Our goal here is primarily a robustness check: does accounting for heterogeneity change any of our qualitative conclusions? Table 2.4 shows the results of this analysis. As before, column 1 is the original model on a small sample. Column 2 adds random coefficients on premium, perfect foresight OOP, variance and quality, while column 3 adds random coefficients on all variables. Again, we see that the choice inconsistencies are present even after accounting for unobserved heterogeneity. Further, the magnitude of the coefficients estimated in the model without heterogeneity (which correspond to the mean of the random coefficients in this model) is not much affected. We do estimate significant heterogeneity in the coefficients on premium, quality, the deductible and the generic donut hole term; allowing for this heterogeneity turns out not to have a significant impact on the welfare results we report below.

We can also interpret the results in Table 2.4 as a test of the IIA assumption. To the extent that any of the coefficients are significant, this suggests that the IIA assumption does not hold exactly. Nonetheless, the fact that the magnitude of the coefficients does not change substantially once we allow for random coefficients suggests that this assumption is not altering our conclusions.

2.10   Conclusions and Implications

The new delivery mechanism for a public insurance benefit introduced by the Medicare Part D program is a radical departure from the traditional public insurance model and an exciting opportunity to understand the role of choice in the delivery of public insurance. Using a unique data set we have provided the first evidence on the efficacy of the choices made by individuals under Part D. While individual choices are consistent with maximizing behavior such as preferring plans with lower premiums, lower out of pocket exposure and

higher quality, they are inconsistent with the standard model in three important respects: individuals underweight out of pocket spending relative to premiums; they overweight plan characteristics beyond their own circumstances; and they do not fully appreciate the risk-reducing aspects of plans for themselves.

Our conclusions do imply that the distribution of health insurance plan coverage would be quite different if there were no choice inconsistencies. We estimate that the share of our sample with some coverage in the donut hole gap would fall by 40% if these inconsistencies were corrected. This would have led to a major shift away from Humana, the insurer that offered the most generous donut hole coverage, towards other insurers.

Yet we also note that our are results are not inconsistent with those of Heiss, McFadden and Winter (2009), who use survey data of individuals to document significant adverse selection in plan choice. The fact that we estimate a non-zero coefficient on out of pocket costs in our logit models is consistent with some adverse selection: controlling for premiums and plan characteristics, individuals prefer plans which offer better coverage for the drugs they plan to take. Our estimated degree of adverse selection is lower than that estimated by Heiss et al., however. They estimate that among those with more than 3 drug claims in a year, the odds of choosing some gap coverage is 10.5% higher than for those with 3 or fewer claims; that difference is only 2.7% in our data. They also estimate that among those with more than $2250 in prescription drug spending in a year, the odds of choosing some gap coverage is 8.8% higher than for those with less spending; that difference is 7.4% in our data. This difference in results between our analyses is partially driven by the fact that they have individuals with zero claims in their data, while we do not in ours; although, as discussed in the robustness section, our conclusions do not appear sensitive to that exclusion.

One means of assessing the implications of these findings is to consider the partial equilibrium welfare gains that would occur were individuals making fully informed and rational decisions about plan choice (ignoring, for now, any supply side considerations or computation costs; these are discussed later). If individuals were fully informed, their choices would be given by the model estimated above but satisfying three additional restrictions: the coefficient on premiums is equal to that on expected out of pocket costs; financial plan characteristics other than premiums are excluded from the utility function once we control

for the individuals expected out of pocket costs; and individuals exhibit risk aversion in their plan choice. We assume that the coefficient on premiums represents the marginal utility of a dollar if individuals were fully informed (this in turn determines the dollar value of quality variables and risk characteristics).We define a normative utility function to include premiums and out of pocket costs (equally weighted), variance and quality, and value the latter characteristics in terms of dollars of premiums. We then ask: if individuals had chosen the plan which maximizes this normative utility function rather than the plan which they did in fact choose, by how much would utility be improved when assessed according to the normative utility function?[20] The answer in this model is about 27% of total costs this is comparable to the 30.9% we found when we looked only at cost savings. The small difference is due to the fact that the lowest cost plans also have slightly lower quality ratings on average.[21] We can interpret the 27the potential partial equilibrium utility gains. If there were some intervention that would make individuals fully informed and fully rational, this is the amount by which their utility could be improved (in partial equilibrium). This large effect suggests that policy makers consider reforms that realize some of these gains. Some possibilities include directly providing individualized information about costs (as in Kling 2008) or appointing surrogates such as doctors or pharmacists to play some role in plan choice.

A more difficult question is whether these findings justify actual restrictions in the choice set facing seniors. As discussed in Abaluck and Gruber (2008), if policy makers are able to restrict choice to the plans on the efficient frontier, there are sizeable welfare gains for seniors (in partial equilibrium). It is unclear, however, whether policy makers would be able to effectively

A full modeling of policy implications must also consider the general equilibrium implications. For example, restricting the size of the choice set may lower competitive pressure on the supply side. Of course, there are possible reforms which would preserve the competitive nature of the bidding process while reducing the number of plans ultimately offered to con-

---

[20]Appendix D of Abaluck and Gruber (2008) contains the formula used to make this calculation as well as a derivation.

[21]The 27alternatively impose a coefficient which corresponds to a coefficient of absolute risk aversion of .0003 (roughly CARA = 3 with wealth of 17,000). In that case, the number rises to 27.6lowest cost plans offer comparable risk protection to the plans which are actually chosen. distinguish such plans; simply restricting choice to a random subset of plans does not raise welfare.

sumers, such as first stage bidding across plans to offer one of a limited set of plan structures. On the demand side, this analysis assumes that the estimated choice process is fixed. We assume that individuals choose according to the same positive utility function regardless of the size of the choice set any utility increases from smaller choice sets arise because there is less scope for error. If individuals are in fact better able to evaluate alternatives in a smaller choice set, then our analysis would understate the potential gains. Moreover, surveys indicate that elders spend an average of 3 hours selecting their Part D plan (Kling et. al. 2008), so the dollar value of the hours saved by dramatically simplifying the choice process may be non-trivial as well. Our models do not distinguish between the case of boundedly rational consumers choosing plans they trust as a heuristic given the time-costs of fully evaluating choices, and the case where consumers simply err in underweighting out of pocket costs due to a lack of cognitive ability. While this distinction is important for evaluating the potential efficacy of providing consumers with additional information, it is less relevant to considering the welfare impact of altering the choice set: in either case, our estimates imply that consumers would be better off if there were less scope for choosing the wrong plan.

Table 2.1: Conditional Logit Results

	(1)	(2)	(3)	(4)
Premium	-.4330**	-.7663**	-.4990**	-.5218**
(hundreds)	-0.0029	-0.0038	-0.0061	-0.0069
OOP Costs (realized)	-.2127**	-.1172**	-.0961**	-.0967**
(hundreds)	-0.152	-0.0015	-0.0015	-0.0016
Variance	-.0189**	-0.0004	-0.0006	-0.0005
(times 10^6)	-0.0027	-0.0007	-0.0007	-0.0007
Deductible	x	-.2899**	-.1628**	-.1674**
(hundreds)		-0.0049	-0.0067	-0.0072
Donut Hole	x	3.023**	1.762**	1.865**
		-0.0181	-0.0277	-0.0303
Generic Coverage	x	.4203**	.3004**	.2700**
		-0.014	-0.0175	-0.0177
Cost Sharing	x	3.282**	1.189**	1.057**
		-0.0538	-0.0741	-0.0778
# of top 100 on Form	x	.0937**	.0587**	.0644**
		-0.0007	-0.0017	-0.0018
Avg. Quality	.4091**	.7398**	x	x
	-0.0032	-0.0039		
Brand Dummies	NO	NO	YES	NO
Brand-State Dummies	NO	NO	NO	YES
Risk Index	10	0	0	0
# of patients	95742	95742	95742	95742
# of plans	702	702	702	702
# of states	47	47	47	47
# of brands	36	36	36	36

Table shows conditional logit results from estimating the model given in equation (6) by maximum likelihood. Each column shows coefficients from a single regression. The coefficients reported are the parameters of the utility function, not marginal effects. Standard errors are in parentheses. * indicates significance at the 5% level and ** indicates significance at the 1% level. The first column includes only premium, realized out of pocket cost and the variance measure. The second column adds controls for the indicated plan characteristics, the third column adds brand fixed effects and the fourth column adds brand-state fixed effects. Premiums, out of pocket cost and deductibles are in hundreds of dollars and the variance term is in millions. The cost sharing variable is computed as the average value of covered expenditures divided by total drug expenditures for individuals in the choice set. The average quality variable is a normalized version of the "average rating" index provided by CMS. The risk index is twice the coefficient on the variance divided by the coefficient on premiums scaled by 100. In the model in the text, this value equals one million times the coefficient of absolute risk aversion.

Table 2.2: Robustness Checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Premium	-.4099**	-.5717**	-.5099**	-.4935**	-.5517**	-.5220**	-.7422**
(hundreds)	-0.007	-0.0076	-0.0074	-0.0072	-0.0078	-0.007	-0.0043
OOP Costs (realized)	-.0508**	-.0808**	-.0846**	-.0977**	-.0986**	-.0944**	-.1100**
(hundreds)	-0.0016	-0.0018	-0.0018	-0.0016	-0.0016	-0.0015	-0.0015
Variance	-0.0009	-.0067**	-0.0011	-0.0003	0.0002	-0.0007	-0.0004
(times 10^6)	-0.001	-0.0029	-0.0014	-0.0007	-0.0012	-0.0008	-0.0008
Deductible	-.0456**	-.2680**	-.2179**	.4378**	.4876**	-.1723**	-.2730**
(hundreds)	-0.0072	-0.007	-0.0068	-0.0095	-0.01	-0.0072	-0.0051
Donut Hole	1.224**	2.023**	1.907**	2.872**	3.956**	1.837**	2.935**
	-0.0302	-0.0299	-0.0296	-0.0346	-0.038	-0.0304	-0.0204
Generic Coverage	-0.0241	.1168**	.1345**	.7570**	.8362**	.2597**	.4206**
	-0.0158	-0.0173	-0.0296	-0.0233	-0.0265	-0.0178	-0.0141
Full Cost Sharing	.9973**	2.129**	1.319**	1.357**	-0.0986	1.076**	3.633**
	-0.1313	-0.1316	-0.1295	-0.0838	-0.0897	-0.0781	-0.0555
# of top 100 on Form	.0737**	.0664**	.0599**	.1173**	.1126**	.0654**	.0952**
	-0.0013	-0.0015	-0.0014	-0.0023	-0.0022	-0.0018	-0.0007
Customer Service	x	x	x	x	x	x	.1644**
							-0.0066
Prescription Filling	x	x	x	x	x	x	.5666**
							-0.0037
Pricing Avail./Changes	x	x	x	x	x	x	.4142**
							-0.0054
Missing Quality Dummies	NO	NO	NO	NO	NO	NO	YES
Brand Dummies	NO	NO	NO	NO	NO	NO	NO
Brand-State Dummies	YES	YES	YES	YES	YES	YES	NO
Risk Index	0	0	0	0	0	0	0
# of patients	99999	99981	99981	99988	99999	99994	95742
# of plans	888	861	861	702	702	702	702
# of states	49	49	49	47	47	47	47
# of brands	48	42	42	36	36	36	36

Table shows conditional logit results from estimating the model given in equation (6) by maximum likelihood on different subsamples to check robustness. Each column shows coefficients from a single regression. The coefficients reported are the parameters of the utility function, not marginal effects. Standard errors are in parentheses. * indicates significance at the 5% level and ** indicates significance at the 1% level. Premiums, out of pocket cost and deductibles are in hundreds of dollars and the variance term is in millions. The cost sharing variable is computed as the average value of covered expenditures divided by total drug expenditures for individuals in the choice set. The first column estimates the model on a random sample of 100,000 patients selected from the sample of individuals who were matched to a PDP plan in the "Lax match" discussed in the text. The second column estimates the model on the original sample used in Table 1 plus individuals with zero claims in 2005. The third column uses the same sample as the second column, but dropping individuals with zero claims in 2005 and more than 12 claims in 2006. The fourth column includes only unique matches and multiple matches which could be assigned with 95% certainty. The fifth column includes only unique matches. The sixth column uses the original sample but restricting to individuals for which more than 50% of their claims were non-dual. The seventh column disaggregates the quality variable. Note that the seventh column also does not include brand or brand-state dummies since these are collinear with the quality variables. The risk index is twice the coefficient on the variance divided by the coefficient on premiums scaled by 100. In the model in the text, this value equals one million times the coefficient of absolute risk aversion.

Table 2.3: Results with Private Information

	Final Sample	Restricted Sample	Private Info
% Private Information			.5818**
			-0.0618
Premium	-.7383**	-.7156**	-.7489**
(hundreds)	-0.0038	-0.0094	-0.0132
OOP Costs (realized)	-.1169**	-.1040**	-.1687**
(hundreds)	-0.0016	-0.0039	-0.0094
Variance	-0.0026	-.1103*	-.8574**
(times 10^6)	-0.0014	-0.0517	-0.2947
Deductible	-.2677**	-.3079**	-.2767**
(hundreds)	-0.0014	-0.1257	-0.0137
Donut Hole	2.823**	2.805**	2.870**
	-0.0181	-0.049	-0.0478
Generic Coverage	.3066**	.4743**	.4784**
	-0.0143	-0.0341	-0.0347
Full Cost Sharing	2.990**	3.391**	2.829**
	-0.0546	-0.1417	-0.1743
# of top 100 on Form	.0939**	.0995**	.1005**
	-0.0007	-0.0019	-0.0021
Avg. Quality	.7167**	.7418**	.7512**
	-0.0039	-0.0098	-0.0095
Brand Dummies	NO	NO	NO
Brand-State Dummies	NO	NO	NO
Risk Index	1	31	229
# of patients	95742	15001	15001
# of plans	702	702	702
# of states	47	47	47
# of brands	36	36	36

Table compares conditional logit results with results from estimating the random coefficients model given in equations (11) and (14) using the Laplace approximation to the likelihood function developed by Hausman and Harding (2007) with bootstrapped standard errors. Each column shows coefficients from a single regression. The coefficients reported are the parameters of the utility function, not marginal effects. Standard errors are in parentheses. * indicates significance at the 5% level and ** indicates significance at the 1% level. The first column is identical to the second column of Table 1. The second column estimates the same model on a random subsample of 15,000 and the third column estimates the random coefficients model on this same subsample. Variable definitions are identical to Table 1. The "Percent Private Information" field corresponds to the variable τ_{frac} in the model in the text.

Table 2.4: Random Coefficients Results

	(1)	(2)	(3)
Premium	-.7156**	-.7677**	-.7354**
(hundreds)	-0.0094	-0.0111	-0.0841
Std. Deviation of Premium	x	.2940**	.2659**
		-0.056	-0.033
OOP Costs (realized)	-.1040**	-.1091**	-.1193**
(hundreds)	-0.0039	-0.0052	-0.0171
Std. Deviation of OOP Costs	x	0.0001	0.0226
		-0.0003	-0.0506
Variance	-.1103*	-0.1486	-0.0866
(times 10^6)	-0.0517	-0.0926	-0.1237
Std. Deviation of Variance	x	.2035**	0.2603
		-0.0685	-0.1769
Deductible	-.3079**	-.2524**	-0.2354
(hundreds)	-0.1257	-0.0107	-0.1903
Std. Deviation of Deductible	x	x	.2922*
			-0.1452
Donut Hole	2.805**	2.775**	2.523**
	-0.049	-0.063	-0.3538
Std. Deviation of Donut Hole	x	x	0.8078
			-0.5997
Generic Coverage	.4743**	.4845**	-0.0037
	-0.0341	-0.0427	-0.232
Std. Deviation of Generic Coverage	x		1.106*
			-0.5108
Full Cost Sharing	3.391**	2.083**	1.829
	-0.1417	-0.1852	-1.304
Std. Deviation of Cost Share	x		0.129
			-0.4828
# of top 100 on Form	.0995**	.0992**	.1405**
	-0.0019	-0.0023	-0.0078
Standard Deviation of top 100	x		0.1071
			-0.079
Avg. Quality	.7418**	.7729**	.7622**

	-0.0098	-0.009	-0.0358
Standard Deviation of Quality	x	.3753**	0.3503
		-0.0146	-0.2182
Brand Dummies	NO	NO	NO
Brand-State Dummies	NO	NO	NO
Risk Index	124	124	124
# of patients	15001	15001	15001
# of plans	702	702	702
# of states	47	47	47
# of brands	36	36	36

Table shows results from estimating the random coefficients model discussed in the heterogeneity section, estimated using the Laplace approximation developed in Hausman-Harding (2007) with bootstrapped standard errors. This model is identical to the model in equation (6), adding the normally distributed noise term (which is a function of) and the variance adjustment. Each column shows coefficients from a single regression. The coefficients reported are the parameters of the utility function, not marginal effects. Standard errors are in parentheses. * indicates significance at the 5% level and ** indicates significance at the 10% level. In columns (2) and (3), each set of four rows reports the mean and standard deviation of random coefficient (and their standard errors). The first column estimates the conditional logit model on the subsample of 15000 (and is identical to the second column of Table 2). The second column adds random coefficients on financial characteristics and quality and the third column adds random coefficients on all variables. Variable definitions are otherwise identical to Table 1.
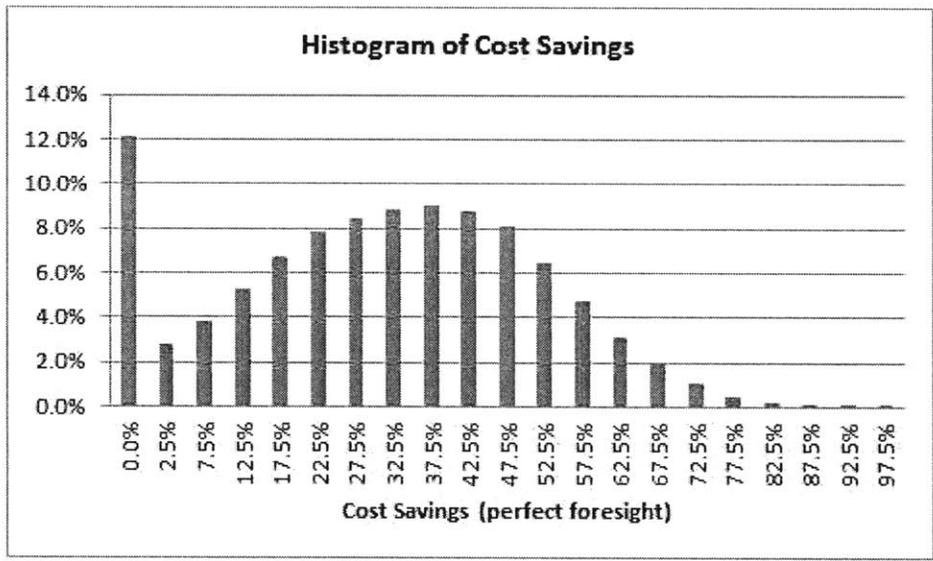
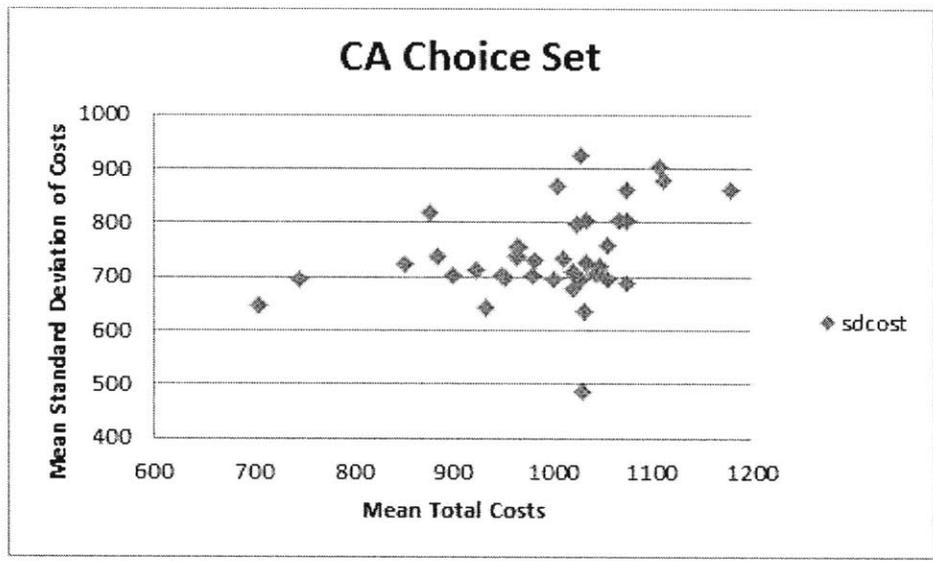Figure 2-1: Histogram of cost savings from switching to lowest cost plan

Figure 2-2: Average mean and standard deviation for each PDP plan in CA

114
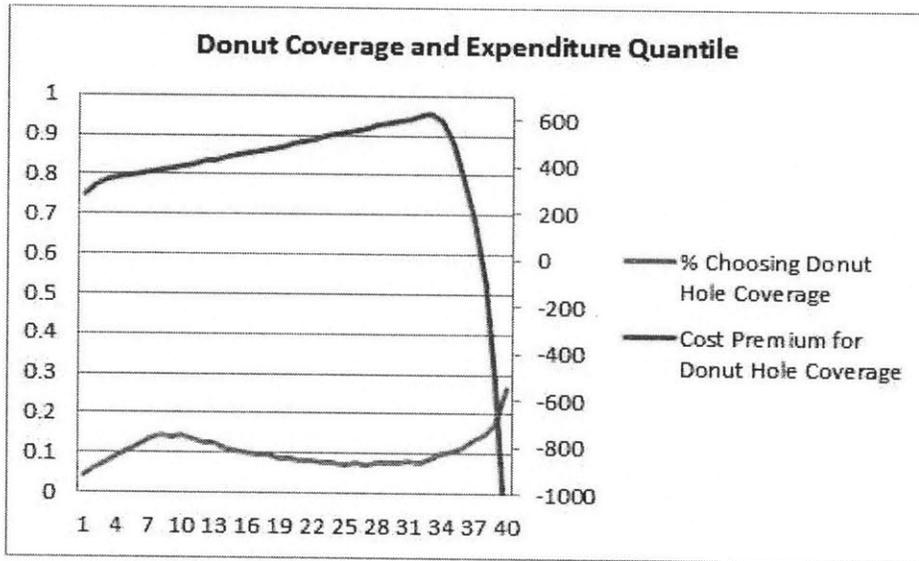
Figure 2-3: Percent Choosing Donut Hole Coverage and Added Cost by Expenditure Quantile

Chapter 3

The General Equilibrium Impact of Information on Prices, Quality and Welfare

3.1 Introduction

When does information provision lead to welfare gains? More generally, is it social welfare maximizing in a general equilibrium setting for consumers to choose as they would if they were fully informed? Lipsey and Lancaster (1956) suggests otherwise; if producers price above marginal cost or firms do not internalize the full benefits of investment in quality, then the theory of the second best says that it will in general not be social welfare maximizing for consumers to choose optimally either. These paper delves more deeply into the relationship between the partial and general equilibrium welfare gains from information.

The motivation is as follows. Suppose we observe that consumers receive information about a given product characteristic and that ex post, consumers' choices are more sensitive to that product characteristic. It is common practice to take ex post preferences as the appropriate normative benchmark (e.g. Chernewa, Gowrisankaran, and Scanlon (2008)).

We can use those preferences to compute the partial equilibrium consumers surplus gain from information, and if we observe costs as well, we can compute partial equilibrium total surplus. In many cases however, the general equilibrium response of prices and quality to a change in information is difficult to estimate. Among other reasons, this might occur if the time frame is too short, the scale of the information intervention too small, the data too coarse, or if the intervention itself is only simulated (as with the second chapter of this thesis, Abaluck and Gruber (2009)). Because of this, it is helpful to understand theoretically the relationship between the partial and general equilibrium social welfare impact of information.

The main contribution of this paper is to derive a set of sufficient conditions on demand under which the general equilibrium response of quality to information provision must be welfare improving relative to the world in which quality is held fixed. As long as one has estimated partial equilibrium demand as a function of product characteristics, one can check whether these conditions are satisfied. If these conditions are met then the welfare calculation taking into account changes in prices but not changes in quality will be a lower bound on the general equilibrium welfare gain from the provision of information. The primary case I consider are models in which quality is undersupplied if consumers value quality appropriately (this occurs because competitive firms internalize only a fraction of the surplus from each unit of quality). In such models, the sufficient conditions I identify imply that it is optimal in general equilibrium for consumers to overweight quality relative to its true impact on their utility and that if consumers currently either underweight or correctly appraise product quality, any intervention which causes them to weight quality more in their product choices will be welfare improving.

The results apply in more general settings than just information provision; comparing the weight attached to a given product characteristic before and after an information intervention and taking ex post preferences as correct is one way of determining the value of a given characteristic to consumers, but it is not the only way. The other chapters in this thesis consider alternatives in which the value of particular characteristics (nutrient content in the case of foods and out of pocket costs in the case of insurance plans) can be computed ex ante given the relationship between the characteristic in question and alternative characteristics whose value we know from other settings. The welfare results in this paper apply equally well

118

in those cases. In any setting where we can distinguish between the positive weight attached to product characteristics and the normative weight appropriate for computing consumer surplus, we can use these results to evaluate any policy that would alter the positive weight attached to those characteristics. The results here relate the partial equilibrium welfare calculation - in which it is always optimal for the positive and normative weights to coincide - with the general equilibrium welfare calculation, where generically this will no longer be true.

In Section 3.2, I briefly consider some related literature. Section 3.3 discusses how I will model perceptions of product quality as a function of information and how this relates to earlier approaches. Section 3.4 presents the main theorems along with some discussion of the assumptions and why they are necessary, Section 3.5 presents an explicit model which allows a closed form solution which is used to study the relationship between allocative and productive inefficiency, and Section 3.6 concludes.

3.2 Literature Review

This paper is most closely related to a small number of earlier studies which consider the general equilibrium impacts of information. Dranove, Kessler, McClellan, and Satterthwaite (2002) considers the general equilibrium impact of health plan report cards; while patients respond to these report cards, hospitals are incentivized to improve their score by selectively treating healthier patients. They find that overall, more resources are utilized for reduced health benefits. The main disanalogy between their environment and the models considered here is the disconnect between reported quality (that observed by consumers) and actual quality (which matters for social welfare). In the Dranove et al. (2002) world, firms can manipulate reported quality without improving actual quality. In the models considered in this paper, this disconnect is not present. For this reason, the paper most closely related to this study is Jin and Leslie (2003), which analyzes the impact of mandating restaurant hygiene reporting on consumers' choice of restaurants and the level of product quality. They observe that quality increases following the provision of information. This paper sets out general conditions under which we would expect this to occur and attempts to draw some

119

conclusions about the relative importance for welfare of the allocational impact of information provision (i.e. helping people choose better) and the productive impact (i.e. firms altering their quality choices).

This paper also relates to a broader literature in behavioral industrial organization, reviewed by Ellison (2006). Whether the analysis here is fully behavioral is partly a semantic distinction. The welfare impact of information in general equilibrium is a topic of neoclassical concern, but the model applies equally well to settings where consumers underweight product characteristics for more behavioral reasons (such as hyperbolic discounting or framing effects). An important question underlying the analysis presented here is why - if consumers are biased - firms would not seek to correct this bias. Ellison (2005) and Gabaix and Laibson (2006) develop models in which firms do not necessarily want to de-bias consumers in equilibrium, but these features are not explicitly incorporated into the models considered here. I take as given the empirical fact that mandated information provision impacts demand and ask - whatever the reason why firms chose not to provide this information ex ante - what can be said about welfare given the direct impact of information on demand, and given the firms response along price and quality margins.

3.3 When Does Imperfect Information Impact Demand?

In this section, I consider the question of when consumer errors impact demand, and when they tend to offset when another, leading the aggregate demand of firms to be unchanged. This will motivate the specification used in subsequent sections. By errors, I mean any deviation between positive and normative utility. As discussed in the introduction, these deviations could be identified by comparing demand before and after an information intervention, or via some other method.

Suppose that given perceived quality $\tilde{v}$, consumer i's utility from purchasing a product from firm j is given by:

$$U_{ij} = \theta \tilde{v}_{ij} - p_j \tag{3.1}$$

The primary case I consider will be a setting in which information improves consumers ability to differentiate between products. Prior to receiving information consumers may distinguish

imperfectly between products. In the first Chapter of this thesis, I present evidence that consumer's imperfectly distinguish between the nutrient content of products within product groups. We can write their perceived nutrient content as $\tilde{v}_j = (1 - \alpha)E(v_j) + \alpha v_j$ where v_j is the actual nutrient of product j and $E(v_j)$ is an expectation within product groups. Note that these beliefs might be unbiased within product groups, but they will be systematically biased for particular products. In this case, within product groups, consumers will respond to a difference in quality of Δv as if the quality difference were only $\alpha \Delta v$. In the above utility framework, substituting this expression in for $\tilde{v}_j$ gives:

$$U_{ij} = \alpha \theta v_j + (1 - \alpha)E(v_j) - p_j \tag{3.2}$$

An alternative model of information commonly studied in the literature allows for only idiosyncratic errors across products. Prior to an information intervention, $\tilde{v}_{ij} = v_j + e_{ij}$ where the e_{ij} are independently and identically distributed across consumers and firms. In that case, utility becomes:

$$U_{ij} = \theta v_j - p_j + \theta e_{ij} \tag{3.3}$$

which is the random utility framework of Luce (1959) and McFadden (1980). Gabaix, Laibson, and Li (2005) characterize the behavior of mark-ups in this model as the number of firms go to infinity. It is clear in this model that the variance of the error term will impact overall demand for a product for even a small amount of noise: as this variance rises, a smaller fraction of observed choices will be explained by quality or prices, so elasticities of demand with respect to quality and prices will fall.

The impact of noisy quality perceptions on demand in this model is dependent on the assumption that the errors are independently distributed. Consider for example a model of horizontal differentiation, where a unit mass of consumers are uniformly distributed on a line of length 1 with firms at both ends. Utility is given by:

$$U_{ij} = \theta v_j - p_j - dt_{ij} + \theta e_{ij} \tag{3.4}$$

Where t_{ij} is the distance of consumer i from firm j. In this model, demand is given by:

$D_1(\Delta \tilde{v}_i) = \frac{\theta \Delta \tilde{v}_i - \Delta p + d}{2d}$. Because demand is linear in $\tilde{v}_i$, a small amount of mean 0 noise will no longer have any impact on demand. Formally, provided range of the noise is small enough that some consumers always go to the nearest firm regardless of their realizations of e_{ij}, noise will have no impact on the demand for each product (and thus, no impact on pricing or quality decisions). Noise will still matter for welfare because it will create allocational inefficiencies, but that is all it will do.

The upshot of this discussion is that there are two reasons to study systematic rather than purely idiosyncratic misperceptions. Firstly, there is empirical evidence that such misperceptions are relevant. Secondly, the impact of such misperceptions will be less dependent on difficult to verify assumptions about the functional form and correlational structure of the error term. As the next section makes clear, it will be possible to draw fairly general conclusions about the welfare impact of systematic misperceptions.

3.4 The General Case: Endogeneous Prices and Quality

Suppose there are two firms $j = 1, 2$, and a continuum of consumers of mass 1. Consumer i's utility from firm j is given by:

$$U_{ij} = \alpha f(v_j) - p_j + \epsilon_{ij} \qquad (3.5)$$

where f is strictly increasing in v_j. I make no assumption about ϵ_{ij} except that it is independent of v_j and p_j. Thus, among other cases, this model includes logit models or models with horizontal differentiation. As discussed in the previous section, I assume that welfare is appropriately evaluated at $\alpha = 1$, either because this is how consumers choose ex post with full information, or because this is the standard implied by some alternative criterion. I assume that firms maximize profits given by $(p_i - c_i)D_i$ where p_i, c_i and D_i are all potentially functions of quality. Assume without loss of generality that $v_1 > v_2$. The first result I will prove is as follows:

Proposition 1. *Suppose that quality is fixed but prices can adjust endogenously in response*

122

to changes in demand. Provided mark-ups are larger for the high quality firm and $\epsilon_{-\frac{\partial D_2}{\partial p_2},\alpha} = \frac{\partial \ln\left(-\frac{\partial D_2}{\partial p_2}\right)}{\partial \ln \alpha} \leq 0$, social welfare is maximized at $\alpha^* > 1$ and increasing in α for all $\alpha < \alpha^*$.

The intuition for this result is straightforward: because mark-ups are higher for the higher quality firm, prices are too high for the high quality firm from the standpoint of allocative efficiency (that is, relative to the prices the social planner would choose). If $\alpha = 1$, too few consumers will choose the high quality firm because of the high price. $\alpha > 1$ can offset this distortion so that the competitive allocation matches the social planner's allocation. The condition: $\epsilon_{-\frac{\partial D_2}{\partial p_2},\alpha} < 0$ guarantees that α does not increase price elasticities so much that the endogenous response of prices to changes in quality offsets the direct impact of α causing demand for the high quality firm as α increases. This condition holds under a wide-variety of commonly used demand functions - for example, it holds if demand is a linear function of prices and qualities as in the models of uniformly distributed horizontal and vertical differentiation considered below. The proof of this claim is in Appendix B.1.

The next proposition gives the conditions under which endogenizing quality only strengthens this result. That is, consider the decomposition: $\frac{dSW}{d\alpha} = \frac{\partial SW}{\partial \alpha} + \sum_j [\frac{\partial SW}{\partial p_j}\frac{\partial p_j}{\partial \alpha} + \frac{\partial SW}{\partial v_j}\frac{\partial v_j}{\partial \alpha}]$. Proposition 1 sets the third term of this decomposition equal to 0, and shows that under this condition, $\frac{dSW}{d\alpha} > 0$ for all $\alpha < \alpha^*$ where $\alpha^* > 1$. The next proposition explicitly considers $\frac{\partial SW}{\partial v_j}\frac{\partial v_j}{\partial \alpha}$. I consider in particular settings where $\frac{\partial SW}{\partial v_j} \geq 0$ for $\alpha = 1$ so that quality is undersupplied if consumers value quality appropriately. I will show that:

Proposition 2. *Suppose that mark-ups are higher for the high quality firm and weakly increasing in quality $\left(\frac{\partial(p_j - c_j)}{\partial v_j} > 0\right)$ and $\epsilon_{-\frac{\partial D_2}{\partial p_2},\alpha} = \frac{\partial \ln\left(-\frac{\partial D_2}{\partial p_2}\right)}{\partial \ln \alpha} < 0$ and that demand satisfies the following additional conditions:*

1. $-\epsilon_{D_j,\alpha} = \frac{\partial \ln D_j}{\partial \ln \alpha} \leq \alpha$

2. $\epsilon_{\frac{\partial D_j}{\partial v_j},\alpha} = \frac{\partial \ln \frac{\partial D_j}{\partial v_j}}{\partial \ln \alpha} > \alpha$

Then we have: $\frac{\partial v_1}{\partial \alpha} > 0$ and $\frac{\partial v_2}{\partial \alpha} > 0$, which implies that social welfare with endogeneous prices and quality is increasing in α for all $\alpha < \alpha^*$ with $\alpha^* > 1$.

The high quality firm (firm 1) always wants to produce higher quality when α increases. For the low quality firm, there are offsetting effects. There are competitive effects wherein

mark-ups and demand increase more with quality when α is larger (since raising quality lures more consumers at the margin), but these are off-set by the fact that demand and mark-ups are lower overall so the return each unit increase in demand or mark-ups is lower. With $\pi(v_1, v_2, \alpha) = (p_2 - c_2)D_2$, we can write:

$$\frac{\partial \pi_2}{\partial v_2} = \frac{\partial(p_2 - c_2)}{\partial v_2} D_2 + (p_2 - c_2)\frac{\partial D_2}{\partial v_2} \tag{3.6}$$

The first new condition implies that the increase in the marginal impact of quality on mark-ups when α increases outweighs the decrease in demand (so the impact of α on the first term is positive). The condition says that demand does not decrease "too much." The second new condition implies that the increase in the marginal impact of quality on demand outweighs the decrease in mark-ups (so the impact of α on the second term is positive). In particular, it says that the increase in demand when quality increases is "large enough." The proof of this claim is in Appendix B.2.

3.5 An Explicit Example

In this section, I develop a model of differentiated product competition in which we can decompose the welfare impact of consumer misweighting into the partial equilibrium welfare loss from misallocation, and the general equilibrium impact of misweighting on prices and product quality. This will serve to illustrate the general theorem in the previous section. Because the model can be solved in closed form, it will also be possible to say more about the relative magnitude of the allocative and productive inefficiencies and how these relate to exogenous determinants of competition such as the degree of product differentiation.

The model explored in this section has two related unusual properties. First, each firm's choice of quality does not depend on the quality choices of competitive firms (that is, $\frac{\partial^2 \pi_i}{\partial v_i \partial v_{-i}} = 0$, so quality choices are neither strategic complements nor strategic substitutes). Because of this property, quality choices will not depend on either the number of firms or the degree of differentiation between firms; they will depend only on own-costs and demand. Second, the qualities chosen by optimizing firms match those chosen by the social planner. These

properties dramatically simplify the analysis.

3.5.1 Set-up

A mass 1 of consumers are located on a circle with two firms located at opposite ends. If consumers purchase the good from firm j, they obtain utility:

$$U_{ij} = \alpha\theta v_{ij} - dt_{ij} - p_{ij} \tag{3.7}$$

where v_j is quality, p_j is price, and t_j is the distance of consumer i to firm j. For ease of exposition, I will focus on the case with two firms located at opposite ends of the circle and with constant θ. The main results also hold in the n-firm case, and the case where θ is uniformly distributed in $[a, b]$ so there are both horizontal and vertical differentiation. These extensions do not change the main results. Profits for firm j are given by $(p_j - c_j v_j^2)D_j$. Suppose that firm j chooses (v_j, p_j). I will focus on the case where both firms have positive demand in equilibrium.

We can think of the $c_j v_j^2$ functional form as capturing any product for which it costs more to make each unit well: it could apply directly to foods (where the ingredients are more costly for better foods), cars, electronics or housing among many other products. We can also think of this functional form as a kind of reduced form model of an insurance market with selection where quality is the proportion of expenditures which are insured. In such a model, insurer costs for each consumer would be $v_j D_j$ where D_j gives the average expenditure of insured consumers. If $D_j(v_j) = cv_j$ due to adverse selection, we will obtain the functional form considered here. Depending on the product involved, we can think of the horizontal differentiation as arising due to literal location decisions (e.g. geographic networks of pharmacies with insurance plans), or due to consumer preference for some brands over others.

3.5.2 Competitive Demand, Prices and Qualities

Let t denote the distance of a consumer from firm 1. Consumers of type t will purchase from firm 1 rather than firm 2 if: $\alpha\theta v_1 - dt - p_1 \geq \alpha\theta v_2 - d(\frac{1}{2} - t) - p_2$ which yields

125

$t^* = \frac{\alpha\theta\Delta v - \Delta p + \frac{d}{2}}{2d}$ where $\Delta v = v_1 - v_2$ and $\Delta p = p_1 - p_2$. Demand for firm 1 is then given by: $D_1 = 2t^*$ and demand for firm 2 is given by $D_2 = 1 - 2t^*$. Firms choose p to maximize $\pi(p_i, p_{-i}) = (p_i - c_i v_i^2)D_i$ which yields mark-ups:

$$p_i - c_i v_i^2 = \frac{d}{2} + \frac{\alpha\theta\Delta_i v - \Delta_i c}{3} \tag{3.8}$$

where $\Delta_i v = v_i - v_{-i}$ and $\Delta_i c = c_i v_i^2 - c_{-i} v_{-i}^2$. Markups are increasing in the degree of horizontal differentiation between firms, increasing in the quality differential between firms and decreasing in the cost differential (i.e. increasing in rival's costs and decreasing in own-costs).

This implies $\Delta_i p = \frac{\Delta c + 2\alpha\theta\Delta v}{3}$, and substituting into the demand functions gives: $D_i = \frac{1}{d}\left(\frac{d}{2} + \frac{\alpha\theta\Delta_i v - \Delta c}{3}\right)$ and profits are given by:

$$\pi_i = \frac{1}{d}\left(\frac{d}{2} + \frac{\alpha\theta\Delta_i v - \Delta_i c}{3}\right)^2 \tag{3.9}$$

Differentiating gives $v_i = \frac{\alpha\theta}{2c_i}$. We can use these to reexpress mark-ups and demand as a function of costs. Mark-ups as a function of costs are equal to:

$$p_i - c_i v_i^2 = \frac{d}{2} + \frac{(\alpha\theta)^2}{12}\left(\frac{c_{-i} - c_i}{c_1 c_2}\right) \tag{3.10}$$

And demand is given by: $D_i = \frac{1}{2} + \frac{(\alpha\theta)^2}{12d}\left(\frac{c_2 - c_1}{c_1 c_2}\right)$.

In Appendix B.3.1, I show that the social planner chooses $v_i = \frac{\theta}{2c_i}$. Thus, when consumers value quality appropriately ($\alpha = 1$), the competitive qualities match the social planner's qualities. This is because of two off-setting effects. Consider for simplicity the world where $c_1 = c_2$, so firms are symmetric and we can ignore transportation costs. Relative to the social planner, competitive firms do not internalize the full surplus from each unit of quality for each existing consumer. The social planner has net benefit $\frac{\theta - 2cv}{2}$ from raising quality for each firm, while competitive firms gain only $\frac{\partial(p-c)}{\partial v}D = \frac{\theta - 2cv}{6}$. However, unlike the social planner, plans benefit from stealing demand from rival firms (on the margin, the consumers who switch are indifferent so social welfare is unchanged). This gives benefit: $(p - c)\frac{\partial D}{\partial v} = \frac{\theta - 2cv}{3}$.

126

So on net, firms have exactly the same net benefit to quality as the social planner. This is a "knife-edge" result coming from the fact that the increased quality incentive from business stealing exactly offsets the decreased quality incentive from not internalizing the full surplus for existing consumers.

3.5.3 Welfare Cost of Mistakes

I first show that the general theorem developed in Section 3.4 applies to this model. First note that mark-ups are larger for the high quality firm: $\Delta p - \Delta c = \frac{2(\alpha\theta\Delta v - \Delta c)}{3} = \frac{\alpha^2\theta^2\left(\frac{1}{c_1}-\frac{1}{c_2}\right)-\frac{\alpha^2\theta^2}{4}\left(\frac{1}{c_1}-\frac{1}{c_2}\right)}{3} = \frac{\alpha^2\theta^2}{4}\left(\frac{1}{c_1}-\frac{1}{c_2}\right) > 0$. Further, $\epsilon_{-\frac{\partial D_2}{\partial p_2},\alpha} = 0$ so Proposition 1 holds and $\alpha^* > 1$ is optimal ignoring quality changes. Further, one can check that $-\epsilon_{D_j,\alpha} \leq \alpha$ and $\epsilon_{\frac{\partial D_j}{\partial v_j},\alpha} > \alpha$ so Proposition 2 holds as well.

The other papers in this thesis consider the cost to consumer surplus of mistakes with prices and quality held fixed. In this section, I will discuss how considering producer surplus, endogenizing prices, and endogenizing quality impact this calculation.

The consumer surplus for fixed p and v is given by: $2(t(1) - t(\alpha))(\theta\Delta v - \Delta p - E[dt - d(1/2 - t)])$. $E[dt - d(1/2 - t)] = 2dE[t] - \frac{d}{2}$. $E[t] = \frac{t(1)+t(\alpha)}{2} = \frac{(1+\alpha)\theta\Delta v - 2\Delta p + d}{4d}$. So $E[dt - d(1/2 - t)] = \frac{(1+\alpha)\theta\Delta v - 2\Delta p + d}{2} - \frac{d}{2} = \frac{(1+\alpha)\theta\Delta v - 2\Delta p}{2}$.

$$CS(\alpha = 1) - CS(\alpha) = \frac{[(1 - \alpha)\theta\Delta v]^2}{2d} \tag{3.11}$$

For α different from 1, consumers purchase too little quality and the welfare loss is proportional to the value of quality. The welfare loss is also larger if firms are less horizontally differentiated, since consumers will more readily choose based on quality. The social welfare difference for fixed prices and qualities is given by:

$$SW(\alpha = 1) - SW(\alpha) = \frac{[(1 - \alpha)\theta\Delta v]^2}{2d} + \frac{(1 - \alpha)\theta\Delta v}{d}(\Delta p - \Delta c) \tag{3.12}$$

Provided mark-ups are increasing in quality so that $\Delta p > \Delta c$, the sign of producer surplus from informing consumers depends on whether consumers overweight or underweight quality characteristics. If consumers underweight quality - as was the case in the empirical work

127

in this thesis - then informing consumers will generate gains for producers as well. The intuition is that holding quality and prices fixed, giving more weight to quality will lead more consumers to choose the higher quality product. Because mark-ups are higher for the higher quality product, producers gain more as well. If however consumers currently overweight quality, the sign of this expression need not be positive. Informing consumers (meaning setting $\alpha = 1$) may reduce welfare if consumers currently overweight quality, because this overweighting offsets the distortion from higher mark-ups for the high quality firm.

3.5.4 Productive vs. Allocative Efficiency

To endogenize prices and quality, I consider social welfare in the competitive equilibrium relative to a social planner's benchmark. This makes explicit how the competitive equilibrium deviates from the optimal quality.

Let SW(allocation,qualities) denote social welfare as a function of the allocation and the quality choices. The allocation determines which firm consumers purchase from, while the qualities enter utility conditional on purchasing from a given firm. The allocation itself is a function of the allocation rule and the qualities at which the allocation rule is applied (e.g. the social planner's allocation changes as the qualities change). The qualities to which the allocation rule is applied need not be the same as the qualities at which the social welfare function is evaluated. We could compute social welfare at the allocation determined by the social planner's allocation rule applied to the social planner's qualities - this determines who purchases from which firms - while still evaluating the social welfare function given this allocation rule at the competitive qualities. Thus, I write $SW(a(b), c)$ where each of $a, b, c \in$ {social planner, competitive} to denote social welfare with allocation $a(b)$ and qualities c. I denote the social planner's allocation by "sp" and the competitive equilibrium allocation by "ce". I will decompose the welfare difference between the competitive equilibrium and the social planner's choices of allocation and quality into the productive inefficiencies and the allocative inefficiencies. That is, I write:

$$SW(sp(sp), sp) - SW(ce(ce), ce) = [SW(sp(sp), sp) - SW(sp(ce), ce)] +$$
$$[SW(sp(ce), ce) - SW(ce(ce), ce)] \qquad (3.13)$$

128

The first line gives the productive inefficiency. Given the social planner's allocation rule - what is the welfare loss if the competitive qualities are not equal to the social planner's qualities? The second line gives the allocative inefficiencies - given the competitive qualities - what is the welfare loss from the fact that the good is misallocated relative to what the social planner would choose?

When $\alpha = 1$, there are no productive inefficiencies, but there are still allocative inefficiencies because mark-ups are increasing in quality. The social planner would allocate consumers to firm 1 if $\theta v - dt - c_1 v_1^2 \geq \theta v - d(1/2 - t) - c_2 v_2^2$. Consumers choose firm 1 if $\theta v - dt - p_1 \geq \theta v - d(1/2 - t) - p_2$. Thus, unless $\Delta c = c_1 v_1^2 - c_2 v_2^2 = \Delta p = p_1 - p_2$, there will be allocative inefficiencies. In this model, $\Delta p > \Delta c$ and so too few consumers will choose firm 1 given fixed levels of qualities. The allocative inefficiencies are given by:

$$
\begin{aligned}
SW(sp(ce), ce) - SW(ce(ce), ce) &= \frac{1}{18d} \left((3 - \alpha)\theta \Delta v - 2\Delta c \right)^2 \\
&= \frac{\theta^4 (c_2 - c_1)^2}{72d(c_1 c_2)^2} \left(\alpha(3 - 2\alpha) \right)^2
\end{aligned}
\tag{3.14}
$$

When firms have symmetric costs, they produce the same quality and charge the same prices, so consumers always travel to the closest firm and there is no allocative inefficiency. As the degree of horizontal differentiation increases, the allocative inefficiency decreases. The price difference between the two firms is unaffected by the degree of differentiation, but as firms become more differentiated horizontally, demand for each firm moves closer to $1/2$ regardless of the cost differences. Thus, allocative inefficiency is actually more problematic for more competitive firms in this sense, because consumers more readily purchase the wrong product. Note that allocative inefficiencies are minimized at $\alpha = 3/2 > 1$ as Proposition 1 requires.

Next, let us consider productive inefficiencies. The welfare impact of these inefficiencies is given by:

$$
\begin{aligned}
&SW(sp(sp), sp) - SW(sp(ce), ce) = \\
&\frac{\theta^2 (1 - \alpha)^2}{32dc_1^2 c_2^2} \left((c_1 - c_2)^2 \theta^2 (1 - \alpha)^2 + 4dc_1 c_2 (c_1 + c_2) \right)
\end{aligned}
\tag{3.15}
$$

129

Note that at $\alpha = 1$, there are no productive inefficiencies because the competitive qualities match the social planner's qualities. Also, note that unlike the allocative inefficiencies, the productive inefficiencies from $\alpha! = 1$ do not go to zero as firms become sufficiently differentiated. Thus, while for more competitive firms allocative inefficiencies may predominate, as firms become more differentiated (and thus less competitive), the productive inefficiencies will be first order.

3.6 Conclusion

This paper examines the relationship between the partial equilibrium welfare gains from information and the general equilibrium welfare impact once one endogenizes prices and quality. There are two main theoretical results: first, holding fixed quality, it is generally optimal from a social welfare standpoint for consumers to overweight quality in order to offset the adverse allocational impact of high prices. A corollary of this result is that information provision tends to improve social welfare if consumers currently underweight quality. Second, I derive conditions under which firms provide higher quality when consumers are more informed, meaning that the partial equilibrium welfare impact of information holding quality fixed will be a lower bound on the general equilibrium welfare gain.

There are a number of immediate next steps: first, the "general" theorem is proven only in the case of 2 firms, so it should be extended to a model with n firms. Second, the cost structure is also somewhat restrictive: I allow marginal costs to vary flexibly with quality but I do not currently allow for fixed costs. The theorem should be extended to these cases. Third, the general equilibrium analysis incorporates prices and quality choices but does not yet incorporate entry. A more complete analysis would study how entry modifies the conclusions presented here. The model could also be specialized to particular industries which raise new general equilibrium considerations - for example, in health care markets, the model should be extended to explicitly incorporate adverse selection and moral hazard.

Appendix A

What Would We Eat: Appendices

A.1 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 (1.7) and (1.8), we can compute the predicted consumption of each food as follows:

$$E(N_{ijt}) = \Phi(\hat{Y}_{ijt}/\sigma_j)(\hat{Y}_{ijt} + \sigma_j \frac{\phi(\hat{Y}_{ijt}/\sigma_j)}{\Phi(\hat{Y}_{ijt}/\sigma_j)}) \qquad (A.1)$$

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

131

total caloric intake. This is given by:

$$E(C_{it}) = \sum_j x_{jt} E(N_{ijt}) \tag{A.2}$$

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.[1]

Appendix Table A.1 presents the results for the four models reported in Table 1.5 which 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.

---

[1]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.

A.2 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.[2] 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).[3]

---

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

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

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 A.4.

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 (.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.

A.3   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:

$$
\begin{aligned}
Y_{ijt} \quad \approx \quad & Y_{ij0} + \frac{\partial Y_{ijt}}{\partial p_{jt}}(z_0)dp_{jt} + \frac{\partial Y_{ijt}}{\partial \mu_{it}}(z_0)d\mu_{it} \\
& + \sum_n \frac{\partial Y_{ijt}}{\partial E_{ijt}(x_{nj})}(z_0)dE_{it}(x_{nj}) + \frac{\partial Y_{ijt}}{\partial \gamma_{ijt}}(z_0)(\rho_j t + d\epsilon_{ijt}) \quad\quad \text{(A.3)}
\end{aligned}
$$

134

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 (A.3), note that for any parameter z_{ijt} in θ_{ijt} we have:

$$\frac{\partial Y_{ijt}}{\partial z_{ijt}}(z^*) \;=\; \frac{\partial \theta_{ijt}}{\partial z_{ijt}}(z^*) \cdot -\frac{\eta_{ijt}(Y^*_{ijt} + K)}{\theta^*_{ijt}} \tag{A.4}$$

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

$$\frac{\partial Y_{ijt}}{\partial q_{ijt}}(z^*) \;=\; \frac{\partial \gamma_{ijt}}{\partial q_{ijt}}(z^*) \cdot -\frac{\eta_{ijt}(Y^*_{ijt} + K)}{\gamma^*_{ijt}} \tag{A.5}$$

This allows us to rewrite equation (A.3) as:

$$Y_{ijt} \approx Y_{ij0} + \eta_{ij0}\frac{K + Y_{ij0}}{\theta_{ij0}}\left[-\phi dp_{jt} + \sum_n \alpha_n(X_{in})dE_{ijt}(x_{nj}) + d\mu_{it} + \frac{t\rho_j + du_{ijt}}{(K + Y_{ij0})^{1/\eta_j}}\right] \tag{A.6}$$

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 (A.6) as:

$$\begin{aligned}
Y_{ijt} &\approx w_{ij0}\left[-\phi(p_{jt} - p_{j0}) + \sum_n \alpha_n(X_{in})(E_{ijt}(x_{nj}) - E_{ij0}(x_{nj})) + \frac{t\rho_j}{E(K + Y_{ij0})^{1/\eta_j}}\right] + e_{ijt} \\
&= w_{ij0}\left[-\phi p_{jt} + \sum_n \alpha_n(X_{in})E_{ijt}(x_{nj}) + t\hat{\rho}_j + \hat{\xi}_j + d\mu_{it}\right] + e_{ijt} \tag{A.7}
\end{aligned}$$

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

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:

$$\hat{\eta}_{ijt} \equiv -\frac{\partial E(N_{ijt})}{\partial p_{jt}} \frac{p_{jt}}{E(N_{ijt})} \tag{A.8}$$

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 (A.8) simplifies to:

$$\hat{\eta}_{ijt} = \phi w_{ijt} \frac{p_{jt}}{E(N_{ijt}|N_{ijt} > 0)} \tag{A.9}$$

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 (A.7), gives:

$$Y_{ijt} \approx \hat{\eta}_{ij0} \frac{E(N_{ij0}|N_{ij0} > 0)}{\phi p_{j0}} \left[-\phi p_{jt} + \sum_n \alpha_n(X_{in})E_{ijt}(x_{nj}) + t\hat{\rho}_j + \hat{\xi}_j + d\mu_{it}\right] + e_{ijt} \tag{A.10}$$

The full model is thus given by:

$$N_{ijt} = \max\{0, Y_{ijt}\} \tag{A.11}$$

where Y_{ijt} is given by equation (A.10) where

$$e_{ijt} = Y_{ij0} + w_{ij0}q_{ijt} \text{ and}$$

$$q_{ijt} = \sum_n \alpha_n(X_{in}) \left[(L_i - E_t(L_i))(E_j^L(x_{nj}) - E_j^U(x_{nj})) \right] +$$

$$\frac{t\rho_j}{(K + Y_{ij0})^{1/\eta_j}} - \frac{t\rho_j}{E_j(K + Y_{ij0})^{1/\eta_j}} + \frac{du_{ijt}}{(K + Y_{ij0})^{1/\eta_j}} \qquad (A.12)$$

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.

A.4 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 A.2. All specifications include the same control variables as specification (1) in Table 1.5. 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 1.5, 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 A.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 A.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

138

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.

A.5   Details of Behavioral Welfare Calculation

Provided $X_{in} \notin \{\bar{X}_{in}, \underline{X}_{in}\}$, we can compute the solution to equation 1.15 by implicitly differentiating the objective function in equation 1.15 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 1.14. The benchmark parameters are then

characterized by the system of linear equations:

$$(Sx)' q = 0 \qquad (A.13)$$

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 A.13 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:

$$w(d) = \sum_n \alpha_n \max(X_n^d - \bar{X}_n, \underline{X}_n - X_n^d, 0) \qquad (A.14)$$

where δ_n is the negative of the absolute value of the coefficient from the Martin et al. (2009)

140

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

141

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 1.8.2). 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:

$$(\tilde{h}_1, ..., \tilde{h}_J) \equiv \operatorname*{argmin}_{(\tilde{h}_1, ..., \tilde{h}_J)} \sum_n \left[Cov(h_j(A') - \tilde{h}_j(A')), \tilde{x}_{nj}) \right]^2 + \sum_j E(h_j(A') - \tilde{h}_j(A'))$$

$$s.t. \quad \sum_j N_j' \tilde{h}_j(A') \geq \sum_j N_j^* \tilde{h}_j(A^*) \tag{A.15}$$

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 A.15 over the set of diets $(N_1^*, ..., N_J^*)$ which are

142

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:

$$(\tilde{h}_1, ..., \tilde{h}_J) \equiv \underset{(\tilde{h}_1, ..., \tilde{h}_J, N_1^*, ..., N_J^*)}{\text{argmin}} \sum_n \left[Cov(h_j(A') - \tilde{h}_j(A')), \tilde{x}_{nj}) \right]^2 + \sum_j E(h_j(A') - \tilde{h}_j(A')) \ s.t.$$

$$\sum_j N_j' \tilde{h}_j(A') \geq \sum_j N_j^* \tilde{h}_j(A^*), \sum_j N_j^* h_j(A^*) \geq 0, \sum_j N_j^* = \bar{N} \quad (\text{A.16})$$

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 1.8 (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 A.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 1.5.

Appendix Table A.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 1.5. 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.

Appendix B

General Equilibrium Impact of Information: Appendices

B.1 Proof of Proposition 1

The proof is as follows: consumers of type ϵ will choose firm 1 if $\alpha f(v_1) - p_1 + \epsilon_{i1} > \alpha f(v_2) - p_2 + \epsilon_{i2}$, or equivalently, $\alpha \Delta v - \Delta p + \Delta \epsilon_i > 0$ where $\Delta v = f(v_1) - f(v_2)$, $\Delta p = p_1 - p_2$ and $\Delta \epsilon_i = \epsilon_{i1} - \epsilon_{i2}$. Let $\Delta \epsilon^* = \Delta p - \alpha \Delta v$. For $\Delta \epsilon_i < \Delta \epsilon^*$, consumers will choose firm 1, and for $\Delta \epsilon_i > \Delta \epsilon^*$, consumers will choose firm 2. Social welfare is thus given by:

$$\int_{-\infty}^{\infty} \left[\int_{-\infty}^{\epsilon_{i2}+\Delta \epsilon^*} [f(v_1) - c_1 + \epsilon_{i1}]f(\epsilon_{i1}, \epsilon_{i2})d\epsilon_{i1} + \int_{\epsilon_{i2}+\Delta \epsilon^*}^{\infty} [f(v_2) - c_2 + \epsilon_{i2}]f(\epsilon_{i1}, \epsilon_{i2})d\epsilon_{i1} \right] d\epsilon_{i2} \quad (B.1)$$

α enters this equation implicitly via $\Delta \epsilon^*$. Next, I show that we must have $\frac{\partial \Delta \epsilon^*}{\partial \alpha} = \Delta v - \frac{\partial \Delta p}{\partial \alpha} > 0$. Suppose instead that $\frac{\partial \Delta \epsilon^*}{\partial \alpha} \leq 0$. Then $\frac{\partial D_2}{\partial \alpha} \geq 0$ and $\frac{\partial D_1}{\partial \alpha} \leq 0$. Differentiating profit functions with respect to own prices gives the first order condition: $(p_j - c_j)\frac{\partial D_j}{\partial p_j} + D_j = 0$. Implicitly differentiating this with respect to α gives $\frac{\partial p_j}{\partial \alpha}\frac{\partial D_j}{p_j} + (p_j - c_j)\frac{\partial^2 D_j}{\partial \alpha \partial p_j} + \frac{\partial D_j}{\partial \alpha} = 0$ which we can rearrange to give:

$$\frac{\partial p_j}{\partial \alpha} = \frac{\frac{\partial D_j}{\partial \alpha}}{-\frac{\partial D_j}{\partial p_j}} + (p_j - c_j)\frac{\partial \ln\left(\frac{\partial D_j}{\partial p_j}\right)}{\partial \alpha} \quad (B.2)$$

Let F denote the distribution function for $\Delta \epsilon$. Then $D_1 = F(\alpha \Delta v - \Delta p)$ and $D_2 =$

145

$1 - F(\alpha \Delta v - \Delta p)$, which implies $\frac{\partial D_1}{\partial p_1} = \frac{\partial D_2}{\partial p_2} = -f(\alpha \Delta v - \Delta p)$. This implies that $\frac{\partial \ln\left(\frac{\partial D_i}{\partial p_i}\right)}{\partial \alpha}$ is equal for both firms. By assumption, this term is negative, so mark-ups larger for firm 1 imply that the second term of equation B.2 is more negative for firm 1 than firm 2. But since the first term is negative for firm 1 and positive for firm 2, this implies $\frac{\partial p_1}{\partial \alpha} < \frac{\partial p_2}{\partial \alpha}$ and $\frac{\partial \Delta p}{\partial \alpha} < 0$. Since $\Delta v > 0$, this implies: $\Delta v - \frac{\partial \Delta p}{\partial \alpha} > 0$ which is a contradiction. So we must have: $\frac{\partial \Delta \epsilon^*}{\partial \alpha} > 0$ as desired. Differentiating equation B.3 with respect to α thus gives:

$$\int_{-\infty}^{\infty} \left([f(v_1) - c_1 + \epsilon_{i2} + \Delta \epsilon^*]f(\epsilon_{i2} + \Delta \epsilon^*, \epsilon_{i2}) - [f(v_2) - c_2 + \epsilon_{i2}]f(\epsilon_{i2} + \Delta \epsilon^*, \epsilon_{i2})\right) d\epsilon_{i2} =$$

$$\int_{-\infty}^{\infty} \left([(1 - \alpha)\Delta v + (\Delta p - \Delta c)]f(\epsilon_{i2} + \Delta \epsilon^*, \epsilon_{i2})\right) d\epsilon_{i2} \qquad (B.3)$$

Because mark-ups are higher for the higher quality firm, $\Delta p - \Delta c > 0$, so this expression can only be zero if $(1 - \alpha)\Delta v < 0$ or $\alpha = \alpha^* > 1$ as desired. Finally, note that for $\alpha < \alpha^*$, $\frac{\partial SW}{\partial \alpha} > 0$ as desired.

B.2 Proof of Proposition 2

The proof is as follows: profits for each firm are given by $\pi_j = (p_j - c_j)D_j$. Differentiating with respect to v_j gives the first-order condition: $\frac{\partial \pi_j}{\partial v_j} = \frac{\partial (p_j - c_j)}{\partial v_j}D_j + (p_j - c_j)\frac{\partial D_j}{\partial v_j}$. Differentiating once more gives:

$$\frac{\partial^2 \pi_j}{\partial v_j \partial \alpha_j} = \frac{\partial (p_j - c_j)}{\partial v_j}\frac{\partial D_j}{\partial \alpha} + \frac{\partial^2 (p_j - c_j)}{\partial \alpha \partial v_j}D_j + \frac{\partial (p_j - c_j)}{\partial \alpha}\frac{\partial D_j}{\partial v_j} + (p_j - c_j)\frac{\partial^2 D_j}{\partial v_j \partial \alpha} \qquad (B.4)$$

By the monotone comparative statics theorem in Milgrom and Roberts (1990), if this expression is positive than we will have $\frac{\partial v_j}{\partial \alpha} > 0$.

The following argument suffices to show that this expression is positive for both firms (although in fact, fewer conditions are required for Firm 1). I start by considering the latter two terms. Condition 2 of the theorem is equivalent to the claim that: $\frac{\partial^2 D_j}{\partial v_j \partial \alpha} > \frac{\partial D_j}{\partial v_j} > 0$. Provided we have that: $(p_j - c_j) > |\frac{\partial (p_j - c_j)}{\partial \alpha}|$, the sum of the last two terms will be positive. This condition is equivalent to $-\epsilon_{M_j, \alpha} < \alpha$. Taking logs of the first-order condition of the profit function with respect to p_j gives $\ln(p_j - c_j) = \ln(D_j) - \ln(-\frac{\partial D_j}{\partial p_j})$. Differentiating with

respect to α and multiplying both sides by negative α gives: $-\epsilon_{M_j,\alpha} = -\epsilon_{D_j,\alpha} + \epsilon_{\frac{\partial D_j}{\partial p_j},\alpha}$ where $M_j = p_j - c_j$. The first term on the right-hand side is less than α and the second term is less than 0 by assumption, so we have the desired result and the third and fourth terms of equation B.4 sum to a positive number.

Next, consider the first two terms. Condition 1 of the theorem implies that $-\frac{\partial D_j}{\partial \alpha}\frac{\alpha}{D_j} < \alpha$ or $-\frac{\partial D_j}{\partial \alpha} < D_j$. The second term of equation B.4 will thus be larger than the first provided we have: $\frac{\partial^2 (p_j - c_j)}{\partial \alpha \partial v_j} > \frac{\partial (p_j - c_j)}{\partial v_j}$. We can equivalently write this condition as $\frac{\partial \ln\left(\frac{\partial (p_j - c_j)}{\partial v_j}\right)}{\partial \alpha} > 1$. Start with the first-order condition of the profit function with respect to prices: $(p_j - c_j)\frac{\partial D_j}{\partial p_j} + D_j = 0$. Differentiating this with respect to v_j gives: $\frac{\partial (p_j - c_j)}{\partial v_j}\frac{\partial D_j}{\partial p_j} + (p_j - c_j)\frac{\partial^2 D_j}{\partial p_j \partial v_j} + \frac{\partial D_j}{\partial v_j}$ which implies:

$$\frac{\partial (p_j - c_j)}{\partial v_j} = \frac{(p_j - c_j)\frac{\partial^2 D_j}{\partial p_j \partial v_j} + \frac{\partial D_j}{\partial v_j}}{-\frac{\partial D_j}{\partial p_j}} \tag{B.5}$$

Taking logs of both sides gives:

$$\ln\frac{\partial (p_j - c_j)}{\partial v_j} = \ln\left((p_j - c_j)\frac{\partial^2 D_j}{\partial p_j \partial v_j} + \frac{\partial D_j}{\partial v_j}\right) - \ln\left(-\frac{\partial D_j}{\partial p_j}\right) \tag{B.6}$$

The next step is to differentiate both sides of equation B.6. By assumption $\frac{\partial \ln\left(-\frac{\partial D_j}{\partial p_j}\right)}{\partial \alpha} < 0$ so the second term is positive. The derivative with respect to α of the first term on the right-hand side is givey by: $\frac{(p_j - c_j)\frac{\partial^3 D_j}{\partial p_j \partial v_j \partial \alpha} + \frac{\partial^2 D_j}{\partial v_j \partial \alpha}}{(p_j - c_j)\frac{\partial^2 D_j}{\partial p_j \partial v_j} + \frac{\partial D_j}{\partial v_j}} > \frac{\partial \ln\frac{\partial D_j}{\partial v_j}}{\partial \alpha} > 1$ where the second inequality follows by assumption and the first follows provided we have: $(p_j - c_j)\frac{\partial^3 D_j}{\partial p_j \partial v_j \partial \alpha} > 0$ and $(p_j - c_j)\frac{\partial^2 D_j}{\partial p_j \partial v_j} < 0$.

Start with: $(p_j - c_j)\frac{\partial^2 D_j}{\partial p_j \partial v_j} < 0$. Differentiating the first order condition with respect to p_j for firm j with respect to v_j and rearranging implies that this expression has the same sign as:

$$\frac{(p_j - c_j)}{\partial v_j}\frac{\partial D_j}{\partial p_j} + \frac{\partial D_j}{\partial v_j} \tag{B.7}$$

From the first order condition of profits with respect to v_j, we have: $\frac{\partial (p_j - c_j)}{\partial v_j} = \frac{-(p_j - c_j)\frac{\partial D_j}{\partial v_j}}{D_j}$, which plugged back into equation B.7 gives: $\frac{\partial D_j}{\partial v_j}\left[\frac{(p_j - c_j)\frac{\partial D_j}{\partial p_j}}{D_j} - 1\right] = \frac{\frac{\partial D_j}{\partial v_j}}{D_2}[\frac{\partial \pi_j}{\partial p_j} - 2D_j] < 0$.

147

Next, consider: $(p_j - c_j)\frac{\partial^3 D_j}{\partial p_j \partial v_j \partial \alpha} > 0$. Rearranging the definition of $\epsilon_{-\frac{\partial D_j}{\partial p_j},\alpha}$ and using the fact derived above that $-\epsilon_{M_j,\alpha} = -\epsilon_{D_j,\alpha} + \epsilon_{\frac{\partial D_j}{\partial p_j},\alpha}$ yields the equation:

$$\frac{\partial^2 D_j}{\partial \alpha \partial p_2} = \frac{\frac{\partial D_j}{\partial p_j}}{\alpha}\left(\epsilon_{D_j,\alpha} - \epsilon_{M_j,\alpha}\right) \tag{B.8}$$

Denote the term in parentheses by q (and note that this is equal to $\epsilon_{\frac{\partial D_j}{\partial p_j},\alpha}$ which is less than 0 by assumption). Differentiating both sides with respect to v_2 and denoting the derivative of the elasticities in parentheses by q gives:

$$\frac{\partial^3 D_j}{\partial p_j \partial v_j \partial \alpha} = \frac{1}{\alpha}\frac{\partial^2 D_j}{\partial p_j \partial v_j}q + \frac{\frac{\partial D_j}{\partial p_j}}{\alpha}\frac{\partial q}{\partial v_j} \tag{B.9}$$

I showed above that $\frac{\partial^2 D_j}{\partial p_j \partial v_j} < 0$ so the first term in this expression is positive. Since $\frac{\partial D_j}{\partial p_j} < 0$, we just need to show: $\frac{\partial q}{\partial v_j} < 0$. We can get this as follows: $\frac{\partial \epsilon_{D_j,\alpha}}{\partial v_j} = \frac{\alpha}{D_j}\frac{\partial D_j}{\partial v_j}\left[1 - \frac{\frac{\partial D_j}{\partial \alpha}}{D_j}\right]$. $\frac{\partial \epsilon_{M_j,\alpha}}{\partial v_j} = \frac{\alpha}{p_j - c_j}\frac{\partial(p_j - c_j)}{\partial v_j}\left[-1 + \frac{\frac{\partial(p_j - c_j)}{\partial \alpha}}{p_j - c_j}\right]$. Comparing these two quantities, we get the inequality:

$$\alpha\frac{\partial \ln D_j}{\partial v_j}\left[1 - \frac{\partial \ln D_j}{\partial \alpha}\right] > \alpha\frac{\partial \ln(p_j - c_j)}{\partial v_j}\left[1 - \frac{\partial(p_j - c_j)}{\partial \alpha}\right] \tag{B.10}$$

This equation holds if: $\frac{\partial \ln D_j}{\partial v_j} > \frac{\partial \ln(p_j - c_j)}{\partial v_j}$ and $-\frac{\partial \ln D_j}{\partial \alpha} > -\frac{\partial(p_j - c_j)}{\partial \alpha}$. Taking logs of the first-order condition with respect to own-price and differentiating with respect to v_2 and α respectively yields the desired results given the assumptions that: $\epsilon_{-\frac{\partial D_j}{\partial p_j},\alpha} < 0$ and the above result that $\frac{\partial^2 D_2}{\partial v_2 \partial p_2} < 0$.

B.3   Specific Example Calculations

B.3.1   The Social Planner's Problem

Let t denote the distance of a consumer from firm 1. It is efficient for a consumer of type t to purchase from firm 1 rather than firm 2 if: $\theta v_1 - dt - c_1 v_1^2 \geq \theta v_2 - d(\frac{1}{2} - t) - c_2 v_2^2$ which yields $t^* = \frac{\theta \Delta v - \Delta c + \frac{d}{2}}{2d}$ where $\Delta c = c_1 v_1^2 - c_2 v_2^2$. Ignoring transportation costs, if everyone

purchased the good from the nearest firm, social welfare would be given by: $\frac{\theta(v_1+v_2)-\Delta c}{2}$.
This expression misstates social welfare since consumers located at distances between $1/2$ and t^*, purchase from firm 1 rather than firm 2. Thus, we need to add: $2(t^* - 1/2) \cdot (\theta\Delta v - \Delta c) = \frac{1}{d}(\theta\Delta v - \Delta c - \frac{d}{2})(\theta\Delta v - \Delta c)$. Finally, we need to add in transportation costs, $dt^* + \frac{d(1/2-t^*)}{2} = \frac{d-2dt^*}{4} = \frac{\frac{d}{2}+\Delta c-\theta\Delta v}{4}$. Combining these gives the expression:

$$SW = \frac{\theta(v_1+v_2)-\Delta c}{2} + \frac{1}{d}\left[(\theta\Delta v - \Delta c)^2 - \frac{d}{2}(\theta\Delta v - \Delta c)\right] - \frac{\theta\Delta v - \Delta c = \frac{d}{2}}{4} \qquad \text{(B.11)}$$

Differentiating this expression gives the optimal quantities: $v_1 = \frac{\theta}{2c_1}$ and $v_2 = \frac{\theta}{2c_2}$.

B.3.2   Calculation of Allocative and Productive Inefficiencies

More generally, demand for firm 1 in the competitive case is given by: $D^{ve} = \frac{1}{2} + \frac{\alpha\theta\Delta v-\Delta c}{3d}$ while the social planner would choose $D^{sp} = \frac{1}{2} + \frac{\theta\Delta v-\Delta c}{d}$. The welfare loss is given by $(D^{sp} - D^{ce})(\theta\Delta v - \Delta c - E[dt - d(1/2 - t)])$. The average t for a switcher between the competitive and social planner's allocation is: $\frac{t^{sp}+t^{ce}}{2} = \frac{D^{sp}+D^{ce}}{4}$. So $E[dt - d(1/2 - t)] = 2dE[t] - \frac{d}{2} = d\frac{D^{sp}+D^{ce}-1}{2}$. $D^{sp} - D^{ce} = \frac{(3-\alpha)\theta\Delta v-2\Delta c}{3d}$. $D^{sp} + D^{ce} - 1 = \frac{(3+\alpha)\theta\Delta v-4\Delta c}{3d}$. $d\frac{D^{sp}+D^{ce}-1}{2} = \frac{(3+\alpha)\theta\Delta v-4\Delta c}{6}$, so $(\theta\Delta v - \Delta c - E[dt - d(1/2 - t)]) = \frac{(3-\alpha)\theta\Delta v-2\Delta c}{6}$.

$$\begin{aligned} SW(sp(ce),ce) - SW(ce(ce),ce) &= \frac{1}{18d}\left((3-\alpha)\theta\Delta v - 2\Delta c\right)^2 \\ &= \frac{\theta^4(c_2-c_1)^2}{72d(c_1c_2)^2}\left(\alpha(3-2\alpha)\right)^2 \qquad \text{(B.12)} \end{aligned}$$

When firms have symmetric costs, they produce the same quality and charge the same prices, so consumers always travel to the closest firm and there is no allocative inefficiency. As the degree of horizontal differentiation increases, the allocative inefficiency decreases. The price difference between the two firms is unaffected by the degree of differentiation, but as firms become more differentiated horizontally, demand for each firm moves closer to $1/2$ regardless of the cost differences. Thus, allocative inefficiency is actually more problematic for more competitive firms in this sense, because consumers more readily purchase the wrong product.

More generally, demand for firm 1 in the competitive case is given by: $D^{ce} = \frac{1}{2} +$

$\frac{\alpha\theta\Delta v - \Delta c}{3d}$ while the social planner would choose $D^{sp} = \frac{1}{2} + \frac{\theta\Delta v - \Delta c}{d}$. The welfare loss is given by $(D^{sp} - D^{ce})(\theta\Delta v - \Delta c - E[dt - d(1/2 - t)])$. The average t for a switcher between the competitive and social planner's allocation is: $\frac{t^{sp} + t^{ce}}{2} = \frac{D^{sp} + D^{ce}}{4}$. So $E[dt - d(1/2 - t)] = 2dE[t] - \frac{d}{2} = d\frac{D^{sp} + D^{ce} - 1}{2}$. $D^{sp} - D^{ce} = \frac{(3-\alpha)\theta\Delta v - 2\Delta c}{3d}$. $D^{sp} + D^{ce} - 1 = \frac{(3+\alpha)\theta\Delta v - 4\Delta c}{3d}$. $d\frac{D^{sp} + D^{ce} - 1}{2} = \frac{(3+\alpha)\theta\Delta v - 4\Delta c}{6}$, so $(\theta\Delta v - \Delta c - E[dt - d(1/2 - t)]) = \frac{(3-\alpha)\theta\Delta v - 2\Delta c}{6}$.

Next, let us consider productive inefficiencies. For consumers between 0 and $t^{sp(ce)}$ and between $t^{sp(sp)}$ and 1, there is no change in allocation between the $(sp(sp), sp)$ world and the $(sp(ce), ce)$ world. The only change in utility comes from the change in quality. The impact of this change on utility is given by: $2t^{sp(ce)}(\theta(v_1^{sp} - v_1^{ce})) + (1 - 2t^{sp(sp)})(\theta(v_2^{sp} - v_2^{ce})$. Consumers between $t^{sp(ce)}$ and $t^{sp(sp)}$ move from firm 2 to firm 1 and experience a change in quality. Their change in utility is thus given by: $2(t^{sp(sp)} - t^{sp(ce)})(\theta(v_1^{sp} - v_2^{ce}) - E[dt - d(1/2 - t)])$. The average t for a switcher is: $\frac{t^{sp(sp)} + t^{sp(ce)}}{2}$. So $E[dt - d(1/2 - t)] = 2dE[t] - \frac{d}{2} = d(t^{sp(sp)} + t^{sp(ce)} - \frac{1}{2}) = \frac{\theta\Delta v^{sp} - \Delta c^{sp}}{2} + \frac{\theta\Delta v^{ce} - \Delta c^{ce}}{2} = \frac{\theta^2}{4}(1 + \alpha(1 - \frac{\alpha}{2}))\left(\frac{c_2 - c_1}{c_1 c_2}\right)$. So in total, we have:

$$\begin{aligned}
\Delta SW &= 2t^{sp(ce)}(\theta(v_1^{sp} - v_1^{ce})) + (1 - 2t^{sp(sp)})(\theta(v_2^{sp} - v_2^{ce}) + \\
&\quad 2(t^{sp(sp)} - t^{sp(ce)})\left(\theta(v_1^{sp} - v_2^{ce}) - \frac{\theta\Delta v^{sp} - \Delta c^{sp} + \theta\Delta v^{ce} - \Delta c^{ce}}{2}\right) \quad \text{(B.13)}
\end{aligned}$$

$2t^{sp(ce)} = \frac{\theta\Delta v^{ce} - \Delta c^{ce} + \frac{d}{2}}{d} = \frac{1}{2} + \frac{\alpha}{2d}(1 - \frac{\alpha}{2})\left(\frac{c_2 - c_1}{c_1 c_2}\right)$. $\theta(v_i^{sp} - v_i^{ce}) = \frac{\theta^2(1-\alpha)}{2c_i}$. $1 - 2t^{sp(sp)} = \frac{\Delta c^{sp} - \theta\Delta v^{sp} + \frac{d}{2}}{d} = \frac{1}{2} - \frac{\theta^2}{4d}\left(\frac{c_2 - c_1}{c_1 c_2}\right)$.

$$\begin{aligned}
SW(sp(sp), sp) &- SW(sp(ce), ce) = \\
&\frac{\theta^2(1 - \alpha)^2}{32dc_1^2 c_2^2}\left((c_1 - c_2)^2\theta^2(1 - \alpha)^2 + 4dc_1 c_2(c_1 + c_2)\right) \quad \text{(B.14)}
\end{aligned}$$

Bibliography

Abaluck, J. and J. Gruber (2009). Choice inconsistencies among the elderly: evidence from plan choice in the Medicare Part D program.

Agarwal, S., J. Driscoll, X. Gabaix, and D. Laibson (2010). The age of reason: financial decisions over the life cycle and implications for regulation. *Brookings Papers on Economic Activity 2009*(2), 51–117.

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.*

Bertrand, M., D. Karlin, S. Mullainathan, E. Shafir, and J. Zinman (2005). What's psychology worth? a field experiment in the consumer credit market.

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.

Boatwright, P. and J. Nunes (2001). Reducing assortment: An attribute-based approach. *Journal of Marketing 65*(3), 50–63.

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.

Chernewa, M., G. Gowrisankaran, and D. P. Scanlon (2008). Learning and the value of information: Evidence from health plan report cards. *Journal of Econometrics 144*, 156–174.

Chetty, R., A. Looney, and K. Kroft (2007). Salience and taxation: Theory and evidence.

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.

Dranove, D., D. Kessler, M. McClellan, and M. Satterthwaite (2002). Is more information better? The effects of 'report cards' on health care providers.

Duggan, M., P. Healy, and F. Morton (2008). Providing prescription drug coverage to the elderly: America's experiment with medicare part d. *The Journal of Economic Perspectives 22*(4), 69–92.

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.

Ellison, G. (2005). A Model of Add-On Pricing*. *Quarterly Journal of Economics 120*(2), 585–637.

Ellison, G. (2006). Bounded rationality in industrial organization. In *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, Volume 2, pp. 142–74. Cambridge, UK: Cambridge University Press.

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).

Fratiglioni, L., D. De Ronchi, and T. Aguero (1999). Worldwide prevalence and incidence of dementia. *Drugs &# 38; Aging 15*(5), 365–375.

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.

Gabaix, X. and D. Laibson (2006). Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets*. *Quarterly Journal of Economics 121*(2), 505–540.

Gabaix, X., D. Laibson, and H. Li (2005). Extreme value theory and the effects of competition on profits. *Massachusetts Institute of Technology, Working Paper*.

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*.

Hastings, J. and L. Ashton (2008). Financial literacy, information, and demand elasticity: Survey and experimental evidence from mexico.

Heiss, F., D. McFadden, and J. Winter (2006). Who failed to enroll in medicare part d, and why? early results. *Health Affairs 25*(5), w344.

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.

Irons, B. and C. Hepburn (2007). Regret theory and the tyranny of choice*. *Economic Record 83*(261), 191–203.

Iyengar, S. and E. Kamenica (2006). Choice overload and simplicity seeking. *University of Chicago Graduate School of Business Working Paper*.

Iyengar, S. and M. Lepper (2000). When choice is demotivating: Can one desire too much of a good thing?. *Journal of personality and social psychology 79*(6), 995.

Jin, G. and P. Leslie (2003). The Effect Of Information On Product Quality: Evidence From Restaurant Hygiene Grade Cards*. *Quarterly Journal of Economics 118*(2), 409–451.

Kamenica, E. (2008). Contextual inference in markets: On the informational content of product lines. *The American Economic Review 98*(5), 2127–2149.

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*.

Kling, J., S. Mullainathan, E. Shafir, L. Vermeulen, and M. Wrobel (2008). Misperception in choosing medicare drug plans. *Unpublished manuscript*.

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.

Kuksov, D. and J. Villas-Boas (2010). When more alternatives lead to less choice. *Marketing Science 29*(3), 507–524.

Lipsey, R. and K. Lancaster (1956). The general theory of second best. *The Review of Economic Studies* *24*(1), 11–32.

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

Lucarelli, C., J. Prince, and K. Simon (2008). Measuring welfare and the effects of regulation in a government-created market: The case of medicare part d plans.

Luce, R. (1959). Individual choice behavior.

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.

McFadden, D. (1980). Econometric models for probabilistic choice among products. *The Journal of Business* *53*(3), 13–29.

Milgrom, P. and J. Roberts (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica* *58*(6), 1255–1277.

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*.

Salthouse, T. (1996). The processing-speed theory of adult age differences in cognition. *Psychological review* *103*(3), 403.

Sarver, T. (2008). Anticipating regret: Why fewer options may be better. *Econometrica* *76*(2), 263–305.

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.

Sethi-Iyengar, S., G. Huberman, and W. Jiang (2004). How much choice is too much? contributions to 401 (k) retirement plans. *Pension design and structure: New lessons from behavioral finance*, 83–95.

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.