



Munich Personal RePEc Archive

Mercury advisories: Information, education, and fish consumption

Shimshack, Jay P. and Ward, Michael B. and Beatty,
Timothy K.M.

Tufts University, The Australian National University, University of
British Columbia

2007

Online at <https://mpra.ub.uni-muenchen.de/25995/>
MPRA Paper No. 25995, posted 23 Oct 2010 13:43 UTC

Mercury advisories: Information, education, and fish consumption

Jay P. Shimshack*

Department of Economics 206 Tilton Hall
Tulane University
New Orleans, LA 70118, USA
jshimsha@tulane.edu

Michael B. Ward

Crawford School of Economics and Government
The Australian National University
Canberra, ACT 0200, Australia
michael.ward@anu.edu.au

Timothy K.M. Beatty

Department of Applied Economics
University of Minnesota
St Paul, MN 55108, USA
tbeatty@umn.edu

This is a pre-print. The final peer-reviewed version of this paper is published in
Journal of Environmental Economics and Management
Volume 53, Issue 2, March 2007, Pages 158-179
doi:10.1016/j.jeem.2006.10.002
©Elsevier

*Shimshack and Ward contributed equally to this paper.

Mercury Advisories: Information, Education, and Fish Consumption

JEL No. D12 - Consumer Economics: Empirical Analysis
I18 - Government Policy; Regulation; Public Health
Q53 - Water Pollution and Hazardous Waste

ABSTRACT

This paper examines responses to a national FDA advisory that urged at-risk individuals to limit store-bought fish consumption due to the dangers of methyl-mercury. We investigate consumer response using both parametric and nonparametric methods. Some targeted consumers significantly reduced canned fish purchases as a result of the advisory, suggesting that information-based policies can achieve the issuing agency's goals. Education and newspaper readership were important determinants of response, suggesting that information acquisition and assimilation are key factors for risk avoidance. While some groups reduced consumption as a result of the advisory, we do not find a response among the relatively large group of at-risk households which met neither the education nor readership criteria. The advisory also had unintended spillover effects; some consumers not considered at-risk reduced consumption in response to the advisory.

Keywords: mercury, information, advisory, environmental health, fish consumption, children's health, environmental risk, pollution

I. Introduction

Information provision is an integral part of many state and federal programs to mitigate environmental and public health dangers. Examples include the toxics release inventory, lead paint disclosures, drinking water quality notices, food nutritional labeling, and product safety warnings. This paper examines the determinants of consumer response to one such information policy, the 2001 FDA methyl-mercury fish advisory.

Government agencies maintain that mercury exposure from environmental pollution is a prominent public health risk. A 2001 Center for Disease Control (CDC) study found that one in ten American women of childbearing age has elevated levels of mercury in her blood. At current agency reference doses and margins of safety, the CDC findings suggest that every year at least 85,000 U.S. children are born at risk of neurological damage from mercury exposure. As the consumption of contaminated fish is the primary source of environmental exposure to mercury, it is a health risk that households could readily limit. Young children, nursing mothers, and pregnant women are the most susceptible to mercury toxicity.

Reducing mercury exposure among at-risk groups requires reduced fish consumption because mercury persists in the environment. Even completely eliminating emissions would not eliminate mercury risks in the near term.¹ In January 2001, the Food and Drug Administration (FDA) issued a commercial fish consumption advisory that warned of the health hazards from mercury and urged at-risk individuals to limit fish consumption. Changes in consumption patterns following this first major national mercury advisory are the focus of our study.

To what extent did the FDA advisory reduce exposure to at-risk groups? We address the question by examining household-level fish consumption from the U.S. Consumer Expenditure Survey (CEX). Specifically, we analyze how certain groups' consumption of canned fish products changed in response to the advisory. An advisory can only achieve the issuing agency's goals if

¹ Domestic emissions controls alone are unlikely to eliminate the risk, even in the long term, because many fish are imported. Further, mercury emissions from foreign sources may be deposited into U.S. waters.

consumers are aware of it and are willing and able to translate awareness into behavior. We therefore focus on proxies for access to information and ability to assimilate information, which are suggested by the literature as response predictors. For example, since news readership is a proxy for information acquisition, we investigate differential responses among readers and non-readers. Education serves as a proxy for both information acquisition and assimilation, so we investigate differential responses among educated and less educated consumers. We also investigate health choices since they may serve as a proxy for access to health information.

We address consumption response empirically with a regression analysis and non-parametric tests. Our simplest non-parametric analysis is a comparison of means before and after the advisory for various groups. We also use a difference-in-differences comparison of means approach to sweep out overall consumption shocks not directly attributable to the advisory, statistically mimicking a control group correction. A limitation of these standard mean tests is that they focus only on measures of central tendency. To provide a broader view of consumer response, we apply quantile treatment effect (QTE) analyses. We also extend the QTE tests to the difference-in-differences framework in order to provide robustness to common shocks.

We find three main results. First, *information-based policies can be effective*. We find that targeted consumers most likely to be aware of and to understand the advisory responded by significantly reducing fish consumption on average. Second, *information policies have pronounced distributional consequences*; certain groups are more exposed to methyl-mercury simply because advisories do not “reach” them. We find that a large group of at-risk consumers, including the least educated, did not respond to the advisory. Third, *information advisories produce significant unintended spillover effects*. We find that some consumers did reduce consumption because of the advisory, despite not being considered at-risk in the advisory.

This is the first economic study of advisory-induced consumer responses for store-bought fish, the primary source of mercury exposure to the public. The most closely related research measured responses of recreational anglers to localized safety advisories. See, for example,

Belton *et al.* (1986) and May and Berger (1996). Using assumptions based on such recreational demand studies, Jakus, McGuinness, and Krupnick (2002) developed health and welfare benefit estimates of a striped bass advisory to Chesapeake Bay anglers.

This study extends a broader literature on public advisories as a policy tool. Adler & Pittle (1984) have a pessimistic view of the efficacy of advisories in practice. It is debated whether even the surgeon general's warning for tobacco was in and of itself a "watershed event" (Fenn *et al.* (2001) and Sloan *et al.* (2002)). Our findings indicate that advisories can achieve the issuing agency's goals, but the short-run response is nuanced. Some sectors of the at-risk population strongly respond, while others respond minimally, if at all. Readership and education are the primary response predictors.

This research also makes a contribution to the product and food safety literature. Experimental work by Viscusi *et al.* (1986) shows that, given information about product hazards, subjects undertake precautionary behavior generally consistent with basic economic theory. Our research confirms these experimental findings in a revealed preference setting. In previous empirical work, Foster and Just (1989) (milk), Brown and Shrader (1990) (eggs), and Kinnucan *et al.* (1997) (meat) all show that adverse health information is correlated with reductions in overall consumption. These studies were based on aggregate data. Our data allow us to disentangle information-related response determinants at the household level.

The paper proceeds as follows. Section II examines the context for our analysis. Parts A and B review sources of mercury exposure, health consequences, and key policy milestones. Part C presents a brief conceptual framework for investigating household health and consumption decisions. Section III summarizes our consumer expenditure data. Section IV examines several methodological approaches, each with their own strength. Graphical analyses, non-parametric statistical tests, and a standard parametric analysis are included. Section V presents our results by answering a series of key questions. Finally, section VI concludes by interpreting our results for economics and policy.

II. Background

A. Sources and Consequences of Mercury Exposure

Levels of mercury circulating in the environment have increased considerably over the last century. Coal-fired electrical plants are currently the largest source of anthropogenic mercury. Mercury binds with sulfuric compounds in coal, and burning releases the mercury into the atmosphere. When atmospheric mercury is deposited into surface water, bacteria convert the mercury into organic methylmercury. It then enters a fish's bloodstream from water passing over gills and accumulates in the tissues. Methylmercury bio-accumulates up the food chain. Even in water where ambient mercury levels are extremely low, mercury concentrations may reach high levels in predatory species like tuna, mackerel, and shark.

For the general public, fish consumption is the primary source of exposure to mercury. Cooking and other forms of preparation do not mitigate exposure. Once consumed, mercury is a neurotoxin, which is absorbed into the bloodstream. In adults, abnormally high concentrations can contribute to brain damage, heart disease, blurred vision, slurred speech, and other neurological ailments. Such concentrations in adults are rare. However, the FDA maintains that even modest mercury concentrations pose a risk of significant harm to the developing neurological systems of fetuses, infants, and children. Consequences may include learning and attention disorders, or generally slow intellectual and behavioral development, as well as severe neurological illnesses such as cerebral palsy. Fetuses and nursing infants are at risk because mercury readily passes through the placenta, concentrates in umbilical tissues, and leaches into breast milk.

B. Mercury & Public Policy

Mercury has recently drawn considerable regulatory scrutiny. For example, the Clean Air Mercury Rule was touted as “the first ever national cap on mercury emissions.” Similarly, the EPA has established power plant mercury emissions standards as a top national priority.

However, even very strict standards cannot eliminate mercury exposure because mercury persists in the environment. Further, most large fish consumed domestically are caught abroad. For these reasons, demand-side consumer policy is, and will remain, an important tool for managing mercury exposure.

Major milestones in consumer policy are reported in Figure 1.² There was a period in which government agencies maintained that mercury consumption risks were minimal. Indeed, FDA scientists counseled in 1994 that “normal patterns of consumption” do not pose a health threat. This official stance persisted until mid-2000, when the FDA weighed the cumulative findings of an EPA report (1997) and a National Academy of Sciences (June 2000) study that asserted significant dangers from consuming contaminated fish. In August of 2000, the FDA announced it was considering a new methyl-mercury advisory and solicited comment.

The FDA formally released the new mercury advisory on January 12, 2001.³ The advisory singled out infants, small children, pregnant or nursing mothers, and women who may become pregnant. It states in part, “.... the primary danger from methylmercury in fish is to the developing nervous system of the unborn child, it is prudent for nursing mothers and young children not to eat these fish as well.” The advisory named several large fish that these targeted consumers should avoid entirely. More generally, it stated that consumers should limit their consumption of all fish, including canned fish, to no more than 12 ounces per week (less than two average meals). This advisory was an unusual response by the FDA; while agency inspections,

² Figure 1 and the accompanying discussion emphasize consumption advisories for fish commercially caught and marketed. EPA and state advisories for methylmercury contamination in locally, recreationally caught fish have been periodically issued as well. Due to their relatively limited scope and scale, we consider these recreational advisories of secondary importance. The interested observer may wish to check the EPA’s ‘Local Fish Advisory Programs’ page at <http://www.epa.gov/waterscience/fish/states.htm> .

³ Our analysis compares pre-advisory to post-advisory consumption. While the advisory occurred on a specific date, we account for the possibility that dissemination and consumption responses occurred with some delay. In a sensitivity analysis, we also address the possibility that the advisory or its content was anticipated.

approvals, and sanctions are common, this type of broad and direct consumer campaign was, and remains, very rare.⁴

Figure 1. Consumer Policy Milestones

Time Period	Consumer Advisory Policy Event
Sept. 1994	FDA Releases ‘FDA Consumer’ ... “Eating commercially available fish should not be a problem.”
Dec. 1997	EPA Releases ‘Mercury Study Report to Congress’ ... “A snapshot of our current understanding of mercury.”
1998-2000	Interest groups and the EPA debate the appropriate reference dose for mercury exposure and policy decisions.
June 2000	National Academy of Sciences (NAS) Releases ‘Toxicological Effects of Methylmercury’ ... “60,000 U.S. children may be at risk.”
Aug-Dec 2000	FDA debates existence and language of new consumer advisory, soliciting comments from consumer advocates, public health professionals, environmental groups, and industry organizations. Focus groups conducted.
Jan 2001	FDA issues new consumption advisory stating that pregnant women, women of childbearing age, and young children should limit consumption of all fish, and should not eat fish known to contain high levels of mercury.
Jan-Mar 2001	Phase I of FDA Mercury Advisory Education Plan.
Jan-Dec 2002	Phase II of FDA Mercury Advisory Education Plan.

The FDA’s outreach program consisted of a two-phase information campaign. Over the course of three months following the advisory, the FDA communicated its message by releasing pre-prepared newsprint and television press releases. Similar media kits were sent to weekly print news sources, parenting magazines, and women’s health periodicals. Phase I of the information campaign also included letters to physicians and health organizations. Phase II was a methodologically similar, but less intense, “reminder” campaign conducted in 2002.

C. Household Health and Consumption Decisions

Household consumption is a function of perceived risk, along with price, demographics, and other such factors. We consider changes in perceived risk due to the advisory in a well-

⁴ FDA inspections can identify localized public health threats, and product- or location- specific consumption advisories are not infrequent. For example, the FDA recently publicized a number of branded almond recalls due to the possibility of *salmonella enteritidis* contamination. Advisories specifically advocating the reduction or elimination of certain foods are rare.

established framework for evaluating consumption decisions in the presence of health information. In seminal works, Michael Grossman (1972a, 1972b) applied the household production function approach of Lancaster (1966) to health and health capital, which serve as production “commodities”. Health information is viewed as an input in the household production function, and factors that impact a household’s perceived risk due to health information enter as demand shifters. Consequently, household consumption is a function $C (PR, P, D \dots)$ of perceived risk PR , prices P , demographics D , and other factors.

Characteristics impacting a household’s perceived risk are based upon the health and consumption literature. Grossman (1972a) hypothesizes that educated households may be likely to process information more efficiently. Further, one proxy for access to health information might be healthy choices made in other consumption areas like fruits, vegetables, and tobacco. Finally, as Ippolito and Mathios (1990) point out, if health information is concentrated in the news and print media, agents who regularly read news sources will face lower information acquisition costs and so are more likely to absorb the relevant risk warning. Consequently, a household’s perceived risk is a function $PR (E, R, HC | T, A)$ of education E , news readership R , and health choices HC , that is conditioned on target status T , and advisory status A . Substituting the perceived risk function into the household consumption function then implies that consumption is affected by these same factors.

III. Data

A. CEX Diary Surveys

Our research assesses the impact of the FDA advisory on consumption of canned fish. We analyze data from the Bureau of Labor Statistics' Consumer Expenditure Survey (CEX). This annual survey asks a cross-section of households to record all expenditures over a two-week period in daily diaries. We sum these data to reflect total household purchases of each item over the sample period.

Using the CEX diaries offers a number of advantages. First, CEX data are widely used for economic and statistical analyses. Second, the unit of observation is the household, allowing us to account for a diverse set of demographic and expenditure variables.⁵ Third, CEX households are geographically diverse, and weighting allows the dataset to approximate a nationally representative sample. Finally, purchase snapshots provide unbiased estimates of consumption.

B. Sample & Definitions

The most direct measure of fish consumption in the CEX is expenditure on canned fish. We choose canned fish because it is widely consumed, it was specified in the advisory language, and data are readily available. To translate expenditures into quantities, we divide by price. We use the BLS regional average price for canned tuna by month, since the CEX does not contain price information.⁶ We construct an adult-equivalence scaling factor for tuna consumption by regressing total in-home meat consumption on the number of adults, babies, young children, medium-aged children, and old children living in the household. Adults are normalized to one, and children are scaled accordingly.⁷ Since the mercury advisory may induce changes in the

⁵ Datasets tracking landings and exports are available, but these contain no household-level data. Further, these aggregate statistics reflect institutional as well as household consumption and do not account for possible warehousing.

⁶ Tuna has consistently comprised over 80% of canned fish consumption over the last decade. The ratio of canned tuna consumption to other canned fish has remained quite stable.

⁷ Children ages 1 to 5 consume approximately 24 percent of an adult's meat consumption, children ages 6 to 11 consume 39 percent of an adult's meat consumption, and children ages 12 to 18 consume 61 percent

decision to consume and the quantity conditional on consuming, our analysis considers three separate indicators: total consumption, a consumption decision indicator, and consumption conditional on non-zero expenditures.

The literature and our conceptual framework suggest that demographics, education, news readership, and health choices may be important determinants of information response. Since households with young or nursing children are directly targeted by the advisory, our analysis includes a dummy for the presence of children ages 0-5. Our analysis also includes a dummy for college graduates, a dummy for newspaper or magazine purchases, and an ad-hoc proxy index for health choices. We consider households ‘health conscious’ if their food expenditure share of fresh fruits of and vegetables is larger than 70 percent of demographically similar households, and have no tobacco expenditures.⁸

One of our objectives is to sweep out consumption shocks not directly attributable to the advisory. The ideal reference households will not have young or nursing children and will be less educated, non-readers, and less health conscious. The ideal reference group would also be as similar demographically to our target group as possible, save the presence of young or nursing children. Therefore, in order to construct the most comparable reference group, our sample only includes traditional households (we exclude multiple adult households headed by single people or households with three or more adults). Similarly, because the advisory targets pregnant women, women of child-bearing age, and young children, we restrict our analysis to households where the woman (if present) is no older than 45.⁹

of an adult’s meat consumption. Our method for constructing these factors follows USDA practices (Lino 2004).

⁸ This breakpoint is admittedly somewhat arbitrary. In a later sensitivity analysis, we confirm that our results are robust over a range of variable definitions.

⁹ Our data do not allow us to identify directly women who are or may become pregnant. Presented results omit childless married women, the control demographic most likely to be pregnant or become pregnant in the near future (that is not already identifiable as target). As a sensitivity analysis, we confirmed that this relatively small demographic group responded similarly to other targeted consumers.

Our sample covers the period 1999-2002; two years before and after the FDA advisory. Observations are approximately evenly distributed over the sample period: there are 5297 two week household expenditure snapshots in the two years prior to the advisory and 5240 expenditure observations in the two years after the advisory.¹⁰

C. Summary Statistics

Summary statistics and variable definitions are presented in Table 1. The table illustrates the stability of household demographic composition over time. All nine variables reflecting households' physical composition, news purchases, education, and health consciousness have similar means before and after the warning. Average changes are an order of magnitude smaller than their standard deviations. This suggests that variability in consumption behavior over time is unlikely to be attributable to variability in sample composition.

The statistics in Table 1 also show that average aggregate canned fish quantity was approximately 8.5 percent higher after the advisory than before. This is of course consistent with the price drop illustrated in the table. Shares, which incorporate both prices and quantities, remained relatively constant over time. The ensuing analysis does account for underlying price movements and other potential common shocks; we examine changes in expenditure patterns for sub-populations *relative to* relevant control sub-populations.

¹⁰ Our sample of interest excludes households with incomplete diaries (about 3 percent of the original data) and households that report no in-home food purchases for the diary period (about 4 percent of the original data). Further, we omit households with more than twelve members total and those with per-capita quantities more than four standard deviations above the mean for households with positive fish expenditures. These latter two categories represent approximately 0.2 percent of the data.

Table 1. Summary Statistics ¹¹

Variable	Description	ENTIRE SAMPLE		PRE-ADVISORY		POST-ADVISORY	
		Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
PURCHASED?	Dummy; '1' if canned fish purchased in 2-week diary period, '0' otherwise	0.168	0.374	0.169	0.375	0.167	0.373
QUANTITY	Canned Fish Quantity (lbs.)	0.264	0.758	0.252	0.724	0.275	0.789
	Quantity Conditional on Purchase (lbs.)	1.57	1.17	1.49	1.12	1.64	1.21
SHARE	Canned Fish Expenditure Share	.004	0.017	.004	0.016	.004	0.017
	Share Conditional on Purchase	.026	0.033	.025	0.032	.027	0.034
PRICE	Real Regional Price (per lb.)	1.94	0.155	2.02	0.139	1.86	0.133
SUB PRICE	Index of Substitute Prices – Base Period Normalized to 1	1.09	0.057	1.04	0.030	1.14	0.034
FOOD	Real In-home Food Expenditures (\$100s)	1.15	0.967	1.15	0.974	1.14	0.961
AGE	Age of Respondent	38.8	13.4	38.6	13.4	38.9	13.3
CHILDREN	Dummy; HH with Young/Nursing Child?	0.303	0.458	0.306	0.461	0.300	0.458
READER	Dummy; Newspaper or Magazine Purchase?	0.242	0.428	0.249	0.432	0.235	0.424
EDUCATED	Dummy; Respondent College Graduate?	0.299	0.458	0.290	0.454	0.308	0.462
HEALTHY	Dummy; Particularly Healthy Household?	0.225	0.418	0.225	0.418	0.225	0.418
RCHILD	Reader/Children Interaction	0.078	0.268	0.081	0.273	0.074	0.262
ECHILD	Educated/Children Interaction	0.097	0.296	0.093	0.290	0.101	0.301
HCHILD	Healthy/Children Interaction	0.076	0.266	0.079	0.270	0.074	0.262
PERSONS	Number of Equivalent Adults	1.90	0.906	1.91	0.908	1.90	0.905

¹¹ Summary Statistics Weighted in Standard Manner. 'Persons' is not directly a variable in the model, but is used for demographic scaling.

IV. Empirical Methods

Our empirical analysis addresses the following questions: First, did the groups targeted by the FDA mercury advisory language reduce their consumption in response to the advisory? Second, what are the determinants of advisory response? Did news readership influence consumption choices? Did education levels influence consumption choices? Did health consciousness influence consumption choices?

We address these questions using a sequence of three fundamentally connected approaches: graphs, non-parametric statistical tests directly corresponding to the graphs, and regressions which extend the non-parametrics. First, we simply illustrate changes in the empirical distribution of pre- and post-advisory consumption by visual inspection of a graph. The graphs provide easily interpretable intuition about changes both in mean consumption and in the overall distribution of consumption. Of course, graphical analysis does not account for statistical noise. So second, we employ non-parametric statistical tests which formalize the intuition of the graphical analysis. Using a difference-in-differences approach to control for unobserved shocks, we test for changes in mean consumption. Similarly, we use a control approach to evaluate changes in the overall distribution of consumption, as measured by quantile treatment effects.¹² Third, we extend the non-parametric comparison of means approach with standard regression analysis. Each comparison of means can, in fact, be calculated as a simplified regression based on a limited set of covariates. Regression essentially runs the comparison of means simultaneously, accounting for potential correlations in the full set of covariates. In this sense, the non-parametric tests can be viewed as a simple statistical bridge between a visual inspection and a full regression analysis.

¹² See Heckman *et al.* (1997) and Bitler *et al.* (2005, 2006) for quantile treatment effects. Millimet *et al.* (2004) provides an example of a difference-in-differences distributional analysis.

A. Graphically Comparing Cumulative Distribution Functions

Our analysis of each question begins with a graphical presentation of fish expenditure shares. We compare post-advisory empirical cumulative distribution functions (cdfs) with pre-advisory cdfs. If, on average, households meaningfully altered their behavior after the advisory, the post-advisory cdf will differ from the pre-advisory cdf *ceteris paribus*.¹³ Given a fixed quantile q , a horizontal shift to the left indicates that q percent of households now consume at or below a smaller expenditure share. Thus, a broad shift to left indicates that consumers reduced their consumption. Alternatively, a shift to the right signifies increased consumption.

Changes in empirical cdfs can be difficult to identify visually in absolute cdf graphs. For this reason, we present graphs such as Figure 2, which plot the horizontal difference between pre- and post- advisory consumption cdfs. For these difference graphs, negative expenditure share differences indicate a leftward shift in the post-advisory empirical cdf and positive expenditure share differences indicate a rightward shift in the post-advisory cdf. In graphs such as Figure 2, negative values imply a decrease in consumption and positive values imply an increase in consumption.

Since the area to the left of any cdf, to the right of the vertical axis, and below probability 1 can be interpreted as a mean (here, mean fish expenditure share), the integrated area between the horizontal zero-axis and the cdf difference curve can be interpreted as the change in mean consumption between the pre- and post- advisory periods. The areas below the axis contribute towards a reduction in mean consumption after the advisory, and the areas above the axis contribute towards an increase. In Figure 2, areas both above and below the horizontal axis are quite small. It therefore appears that average fish expenditure shares over all demographic groups did not change significantly after the advisory.

¹³ To be precise, the weighted empirical cdfs will differ. Throughout our analyses, all graphical data include probability weights.

Figure 2. Difference Graph: Entire Sample




Figure 2 demonstrates that aggregate fish shares over all expenditure groups did not significantly change after the advisory. The figure presents the horizontal difference between the pre- and post-advisory empirical cdfs. At every quantile with positive fish consumption, the difference is quite close to 0.

B. Non-parametric Statistical Tests Mean Comparisons

Since the information in the graphs represents differences in means, we use simple statistical methods to formally test graphical insights. For example, we could test the mean reduction in the overall share of food expenditures allocated to fish. For subscript 0 indicating ‘pre-advisory’, subscript 1 indicating ‘post-advisory,’ and X indicating mean fish expenditure share, this test statistic would be $X_0 - X_1$, and its value corresponds to the net sum of the integrated areas in Figure 1.

Of course, changes in fish consumption over time may not be fully attributable to the mercury advisory. For example, canned fish prices, substitute prices, and advertising expenditures all changed over this time period. Therefore, we “sweep out” these common by computing a difference-in-differences of means (DDM). For example, we will examine consumption responses of demographic groups directly targeted by the advisory language, after netting out consumption

changes for demographic groups untargeted by the advisory. Formally, we examine the *inter*-group difference of the *intra*-group changes in mean consumption.¹⁴

This technique can best be understood by a simple example. Consider two basic sub-groups of our data: those targeted, families with young children, and those not targeted. We wish to isolate the impact of the mercury advisory on targeted families (Δ_{advisory}), while controlling for other potential factors such as advertising campaigns or changes in price that affect both targeted and non-targeted sub-groups equally (Δ_{common}). Using the non-targeted group as a control, it is then possible to statistically sweep out their impact. Again, let X_0 and X_1 be mean consumption by the target sub-group in the pre- and post- advisory periods respectively. Similarly, let Y_0 and Y_1 be consumption by the non-target sub-group. Then,

$$X_1 = X_0 + \Delta_{\text{common}} + \Delta_{\text{advisory}}$$

$$Y_1 = Y_0 + \Delta_{\text{common}}$$

$$\Delta_{\text{advisory}} = (X_1 - X_0) - (Y_1 - Y_0)$$

It is irrelevant whether the initial average consumption levels X_0 and Y_0 differ across the two groups, so long as the common shock Δ_{common} impacts both equally. If both the sub-group of interest and the reference group respond to some extent to the mercury advisory, the DDM approach identifies the *relative* difference in response to the advisory. In what follows, we make frequent use of this fact. For example, we measure the incremental impact of education on advisory response within the target sub-group.

We use comparison tests in three ways. First, we compare unconditional difference-in-differences means as discussed above. These statistics intuitively parallel the difference between the integrated area in the presented graph for the group of interest and the integrated area in the presented graph for the control group. Second, we apply these same comparisons to means

¹⁴ The graphical analog to this test, not presented in the interest of space but potentially useful for a reader's intuition, is a "difference-in-differences" graph. We do present graphs side-by-side so that the reader can

conditional on consuming canned fish. Third, we apply analogous comparisons to the number of consumers purchasing any canned fish. For the three cases, relative to a control group the corresponding null hypotheses are: (1) No change in mean consumption, (2) No change in mean consumption, conditional on purchase, and (3) No change in the percent of the group purchasing any canned fish.

Quantile Treatment Effects

The methods described in the previous section emphasized changes in mean consumption. While the mean is important, comparisons only account for shifts in the central tendency of a distribution. One particular feature of interest here is behavior at the right tail of the consumption distribution, as this measures consumption of those considered most at-risk. In order to capture responses across the entire distribution of consumption, we use the quantile treatment effects (QTE) approach. This method tests whether various quantiles of the consumption cdfs in pre- and post-advisory periods differ significantly.¹⁵

As a concrete example, the 90th quantile is the smallest expenditure share such that at least 90% of households have lower shares. Since consumption patterns may differ overall in the pre- and post-advisory periods, the 90th quantile expenditure share in the post-advisory period may differ from the 90th quantile share in the pre-advisory period. The QTE is the difference between these two expenditure shares; it measures the advisory “treatment”. Note that the QTE does not identify individual impacts because it does not measure changes by those *specific* individuals initially at a given quantile, an effect which would require perfect rank preservation (Heckman *et al.* (1997)). Rather a QTE reflects a shift in the overall consumption pattern across individuals, an impact of considerable importance from a public policy perspective.

While the standard QTE analysis for our dataset provides a robust non-parametric check for shifts in consumption patterns, it only captures absolute changes, rather than changes relative

mentally make this comparison.

to a reference group. Thus, it does not control for potential common shocks.¹⁶ Following Meyer, Viscusi, and Durbin (1995), we adapt the classical difference-in-differences approach to provide the necessary correction. Thus, the corrected QTE is simply the difference between the QTE for the treatment group and the QTE for the baseline group. The underlying assumption is that common shocks impact the quantiles for the treatment and baseline groups to the same extent.¹⁷ This pairing of two well-established methodologies provides a broad test for shifts in consumption patterns that is robust to the possibility of common shocks.

We base our statistical inference for QTE's on exact distribution-free quantile confidence intervals. From a probabilistic view, the question of whether a particular observation falls above or below some point s_i is a Bernoulli trial. Conditional on s_i being the q^{th} quantile, the probability that K_i shares out of T total trials are less than s_i is given by the binomial formula

$\binom{T}{K_i} q^{K_i} (1-q)^{T-K_i}$, where for example $q = 0.9$ for the 90th quantile. The confidence that the true

quantile lies below a value s_i is then $\sum_{k=1}^{K_i} \binom{T}{k} q^k (1-q)^{T-k}$. See Hogg and Craig (1978) p. 304 for

a formal treatment. Confidence intervals for the final QTE's are then constructed from these individual quantile confidence intervals. Recall that each difference-in-differences QTE involves observations from four groups. The probability that the QTE lies below a value t is evaluated by convolution over the quantile probabilities for each of those groups; see e.g. Poe *et al.* (1994).¹⁸

That is, we evaluate each combination of quantiles over the grid of observed shares. We then

¹⁵ Note that we present tests for a leftward shift in the cdf, which corresponds in a reduction in expenditure shares. More commonly, studies test for rightward shifts. See for example Maasoumi and Millimet (2005).

¹⁶ If treatments are assigned randomly to a control and treatment group, the standard analysis suffices. The mercury advisory does not satisfy this condition as treatments are not randomly assigned and there are systematic initial differences between groups.

¹⁷ In the difference-in-differences in means, $\bar{X}_1 - (\bar{Y}_0 - \bar{Y}_1)$ is sometimes interpreted as the predicted mean based only upon the common shock. One difficulty with this interpretation is that it is logically possible to produce negative mean consumption. The same issue arises here. In this broader context, it is possible that, after the QTE correction, the implied cdf may not be a proper cumulative distribution function. Thus, one must be careful not to over-interpret the underlying structure.

check whether the implied QTE, the difference-in-differences at that combination of points, is less than t . If so, we increment the probability that the QTE is less than t by the probability of that combination.¹⁹ In all tables of results, the presented significance values for the QTE's correspond to a one-sided test.

C. Regression Methods

We generalize the previous non-parametric DDM method with a standard regression analysis, consistent with the bulk of the mainstream demand literature. Moreover, to ensure the robustness of our results to indicators of consumption, we run a regression with quantity purchased as the dependent variable to supplement the previous analyses based on expenditure shares.²⁰

The main advantage of regression over the previous mean tests is that it controls for correlations between explanatory variables. In addition, traditional regression allows us to experiment with continuous definitions of the explanatory variables, rather than the dichotomous definitions necessary for graphs and the DDM's which formalize them. Note, however, that the fundamental identification approach of sweeping out common shocks in regression is still implicitly a difference-in-differences.

The choice of explanatory variables is motivated by the conceptual framework in Section IIC. We include total food expenditure, region, household demographics, and factors that

¹⁸ The method of Krinsky and Robb (1986) is a Monte Carlo approach to convolution.

¹⁹ Since we are evaluating over a grid of observed points, we use the change in the quantile cdf at each gridpoint as a probability approximation. In other words, we act as if the data generating process were discrete, with support over the observed points.

²⁰ We perform complementary analyses with two different consumption indicators. Regressions use absolute quantities as the dependent variable. Non-parametric analyses use expenditure shares as the dependent variable. Since demand for canned fish is inelastic and price was lower after the advisory, one would expect expenditure shares to fall while quantities rise. While this may be a source of concern, the primary non-parametric difference-in-differences tests and the regression differences are both designed exactly to sweep out this type of common shock. Further, results across consumption indicators are similar.

influence perceived consumption risk due to health information.²¹ In all cases, since households vary in size, we demographically scale household composition covariates multiplicatively by the adult-equivalent measure discussed in Section III (Pollack and Wales, 1981).

Regressions are parameterized following the differences-in-differences literature.²² We include both the targeted dummy alone and interacted with the post-advisory dummy. We also allow the targeted treatment effect to differ by factors hypothesized to influence perceived risk. For example, we include education, education interacted with the target dummy, education interacted with the post-advisory dummy, and education interacted with both the post-advisory dummy and the target dummy. Our ultimate reference group is then non-targeted, non-reading, less educated, less healthy consumers. Time dummies for the pre- and post-advisory periods capture price and non-price shocks such as advertising for this reference group.

As with many household expenditure datasets, we observe a large number of zero purchases. Here, zeros may arise in two ways. One possibility is infrequency of purchase, since a diary survey represents only a snapshot of a given household's canned fish expenditure. A second possibility is abstention from the good entirely. To capture the dichotomous purchase choice, we begin the analysis with a standard probit regression. Of course, conditional on purchasing canned fish, we are also interested in the impact of the FDA warning on the quantity purchased. Therefore, we run a second stage continuous regression. We allow the same covariates to influence both the discrete purchase and the continuous quantity decision, but we do not impose cross-equation restrictions on the covariates of interest.

The error term in this conditional demand equation is potentially correlated with the error term in the probit equation. In this case, our model is exactly that suggested by Blundell and

²¹ Recall that we assume the impact of prices and non-price shocks such as advertising is the same in the reference group and the group of interest. Thus, prices and non-price shocks will be swept out of the difference-in-differences regressions.

²² Bertrand, Duflo, and Mullainathan (2004) demonstrated that serial correlation may cause problems with inference from parametric difference-in-differences estimation. However, as a repeated cross-section, we

Meghir (1987) for the case of a good with non-negative desired demand.²³ Mathematically, this is equivalent to Heckman’s (1976) selectivity model. See Deaton and Irish (1984) for a discussion, and Fry and Pashardes (1994) for an application. To summarize, our empirical model can be represented by:

$$\begin{aligned} Z &= X\beta + \varepsilon_C \\ D &= X\Gamma + \varepsilon_D \\ Q &= \begin{cases} Z & \text{if } D > 0 \\ 0 & \text{otherwise,} \end{cases} \end{aligned}$$

for observed quantity Q , binary purchase decision D , continuous quantity choice Z , and explanatory variables X .

We estimate two primary specifications. The first specification examines the entire sample, whereas the second specification highlights pre- versus post-advisory changes by removing the year immediately prior to and immediately after the advisory. We employ this latter specification to allow for potential lags in consumer responses.

V. Empirical Results

A. Advisory Response

Regression results are presented in Table 2. Graphical results are presented in Figures 3-6, and statistical test results are summarized in accompanying tables 3-6. We motivate our discussion of the results by asking a series of policy relevant questions. Sensitivity analysis follows in sub-section B.

have no individual based serial correlation. Further, one of their suggested remedies is to collapse time series information into pre-event and post-event time periods. Our analysis is structured this way.

²³ Another possible source of zero expenditures is the standard Tobit-style censoring where observation error may drive consumption to zero. We believe that this is not a major concern in our analysis. However, in the sensitivity section, we confirm that results are robust to this assumption.

Table 2. Difference-in-Differences Regression Results

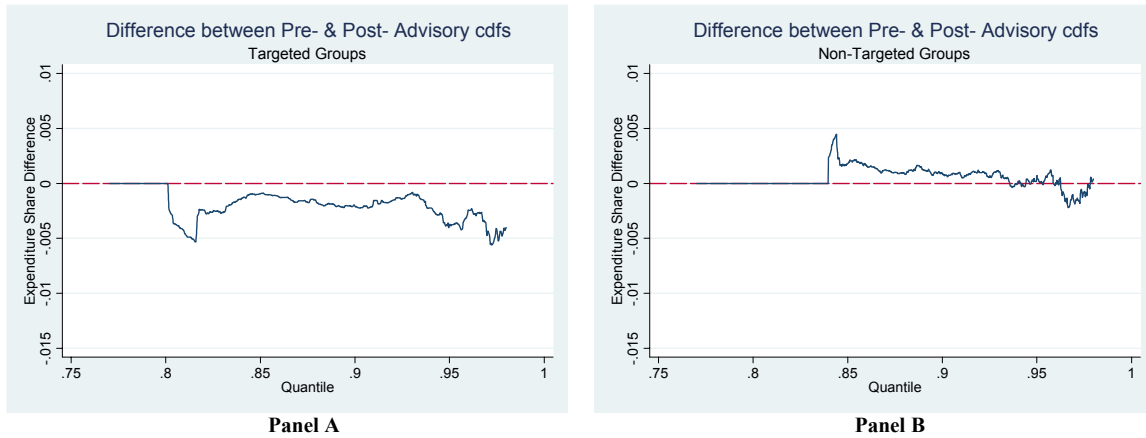
Variable	Binary Consumption Decision		Quantity Conditional on Consuming	
	Entire Sample	Cored Sample	Entire Sample	Cored Sample
FOOD EXPENDITURES	0.390** (0.020)	0.431** (0.028)	0.660** (0.040)	0.738** (0.063)
AGE OF RESPONDENT	0.001 (0.001)	0.001 (0.001)	0.002 (0.002)	0.001 (0.002)
CHILD (TARGETED)	0.065* (0.033)	-0.000 (0.046)	0.099* (0.059)	0.057 (0.081)
CHILD*POST	-0.041 (0.045)	0.017 (0.068)	-0.043 (0.080)	-0.007 (0.108)
READER	0.060* (0.032)	0.047 (0.044)	0.084 (0.055)	0.099 (0.077)
READER*POST	-0.101** (0.044)	-0.116* (0.064)	-0.151* (0.082)	-0.199* (0.111)
EDUCATED	-0.029 (0.035)	-0.039 (0.046)	-0.020 (0.061)	-0.023 (0.080)
EDUCATED*POST	0.054 (0.048)	0.055 (0.064)	0.104 (0.083)	0.036 (0.111)
HEALTHY	0.010 (0.039)	-0.007 (0.055)	0.003 (0.070)	-0.030 (0.100)
HEALTHY*POST	-0.043 (0.054)	-0.053 (0.076)	-0.016 (0.096)	-0.032 (0.136)
READER*CHILD	-0.079* (0.046)	-0.056 (0.064)	-0.132* (0.078)	-0.153 (0.109)
READER*CHILD*POST	0.041 (0.069)	0.033 (0.095)	0.069 (0.117)	0.108 (0.164)
EDUCATED*CHILD	0.034 (0.048)	0.089 (0.067)	0.052 (0.082)	0.138 (0.114)
EDUCATED*CHILD*POST	-0.111* (0.069)	-0.210** (0.096)	-0.229** (0.117)	-0.378** (0.167)
HEALTHY*CHILD	0.015 (0.052)	-0.023 (0.077)	0.068 (0.092)	0.026 (0.141)
HEALTHY*CHILD*POST	0.039 (0.075)	0.159 (0.108)	-0.037 (0.132)	0.139 (0.200)
PRE-ADVISORY DUMMY	-0.437** (0.065)	-0.399** (0.087)	-0.889** (0.117)	-0.884** (0.150)
POST-ADVISORY DUMMY	-0.397** (0.063)	-0.394** (0.087)	-0.803** (0.113)	-0.798** (0.148)
TIME INVARIANT CONS.	-0.633** (0.058)	-0.615** (0.093)	-	-

Figures in parentheses are robust standard errors. ** - Significant at $\alpha = .05$ * - Significant at $\alpha = .10$. Both regression specifications use a maximum likelihood procedure mathematically identical to Heckman's Selectivity model. Specifications also include 4-1 regional dummies and 4-1 race dummies, but we omit these control results to conserve space.

Did the target group respond to the FDA advisory?

Taken as a whole, it appears that households with young or nursing children responded to the advisory relative to non-target households, but the statistical evidence is modest. Panel A of Figure 3 shows a general reduction for the target group at most expenditure share levels. Recall that the integrated area between the difference curve and the horizontal axis is equal to the difference in means. Here, the net integrated area is negative; the sample mean clearly fell after the advisory for this group. In contrast, Panel B shows a slight increase for the non-target control group. Comparing Panel B to Panel A, we see that mean expenditure share for the target group fell relative to the non-target group.²⁴

Figure 3. Difference Graphs: Target and Non-Target Groups



²⁴ All difference graphs are presented up until the 98th quantile. Beyond this quantile, statistical noise predominates.

Table 3. Non-Parametric Tests Summary: Target and Non-Target Groups²⁵

	%Δ Mean Share	%Δ Mean Share, Given Positive Share	%Δ Fraction of Consumers with Positive Share	Q80?	Q85?	Q90?	Q95?
Target, Net of Non-Target	-21.8 (0.08)	-13.0 (0.17)	-9.7 (0.13)	No (.24)	No (.12)	Yes (.05)	Yes (.10)

P-values are in parentheses. Quantile Treatment Effects evaluated at $\alpha=0.10$.

The difference-in-difference of means (DDM) statistic is the normalized numerical value of the difference in means for the target group, net of mean changes for the non-target group (here, the reference group). The net drop in the overall mean expenditure share allocated to fish for the target group was 21.8 percent. The corresponding DDM statistic is statistically different from zero at the 8 percent significance level.

While the overall mean fell, the disaggregated components of this mean did not fall significantly when considered individually. Neither the percent drop in net share conditional on purchase nor the percent drop in the net proportion of households purchasing is significant at conventional levels. Similarly, the quantile treatment effect results are modest. We find evidence of significant quantile treatment effects above the 90th percentile. At Q90 and Q95, more than 90 percent of the test statistic's confidence distribution lies below 0. Thus, after accounting for common shocks to non-target groups, target consumers' total expenditure on fish fell significantly across this range of the distribution.

The parametric regression results in Table 2 are also suggestive, but frequently insignificant. The coefficient on the child*post-advisory interaction represents the impact of the advisory for targeted, non-reading, less educated, less healthy consumers, relative to the pre-advisory baseline for this group. The aggregate impact of the advisory on targeted consumers is the joint impact of the coefficients on child*post, education*child*post, reader*child*post, and

²⁵ To enhance the economic interpretation, Tables 3-6 report percent changes for mean expenditure share, mean expenditure share conditional on consuming, and proportion of group purchasing any fish. The actual test statistics, however, are based upon absolute differences rather than percents.

healthy*child*post. Only in the full sample continuous quantity regression are these four coefficients jointly significant at the 10 percent level.

Did health-conscious consumers respond to the advisory?

No. We find no evidence that health-conscious households, as a group, responded to the advisory. Recall that we define healthy households by a function of good diet and tobacco abstinence. Panel A of Figure 4 illustrates the change in expenditure patterns for healthy households in the non-target group, while Panel B represents less healthy households in the non-target group. In Panel A, we do not observe a consistent pattern. In Panel B, the net integrated area is relatively small. On balance, there is little change in mean expenditure behavior among non-targeted households. Panels C and D of Figure 4 represent expenditure changes by healthy and less healthy target groups, respectively. While both panels show a reduction in mean expenditure by targeted groups, a clear differential response among healthy and less healthy target consumers is not apparent.

Figure 4. Difference Graphs: Healthy and Less Healthy Groups




Table 4. Non-Parametric Tests Summary: Healthy and Less Healthy Groups

	% Δ Mean Share	% Δ Mean Share, Given Positive Share	% Δ Fraction of Consumers with Positive Share	Q80?	Q85?	Q90?	Q95?
Non-Target: Healthy, Net of Less Healthy	-13.2 (0.27)	-6.4 (0.36)	-7.2 (0.30)	No (.39)	No (.32)	No (.27)	No (.11)
Target: Healthy, Net of Less Healthy	-16.8 (0.22)	-14.6 (0.20)	-1.8 (0.44)	No (.34)	No (.16)	No (.44)	No (.28)
Healthy: Target, Net of Non-Target	-18.8 (0.26)	-14.6 (0.29)	-6.5 (0.36)	No (.33)	No (.20)	No (.25)	No (.50)
Less Healthy: Target, Net of Non-Target	-21.7 (0.12)	-11.5 (0.24)	-10.5 (0.15)	No (.33)	No (.13)	No (.11)	No (.12)

P-values are in parentheses. Quantile Treatment Effects evaluated at $\alpha=0.10$.

Non-parametric statistical tests confirm the visual insights. Among non-targeted consumers, the DDM statistic for the percent change in the overall mean expenditure share for the

healthy subgroup (after accounting for mean changes for the less healthy subgroup) is not significant at conventional levels (p-value of 0.27). The corresponding overall DDM statistic is also not statistically significant for the target group (p-value of 0.22). All other non-parametric DDM statistics are similar. Neither the healthy group's net mean share conditional on purchase nor the healthy group's net proportion of consumers purchasing any fish changes significantly, regardless of whether the particular households are targeted or non-targeted. These findings are not restricted to the mean; we find no evidence of significant quantile treatment effects for all tested percentiles.

Regression results are consistent with these findings. In Table 2, the row labeled 'Healthy*Post' indicates no significant impact of the advisory among health conscious non-targeted consumers. Similarly, we find no significant incremental advisory response from target group membership for health-conscious households, as summarized in the interaction row labeled 'Healthy*Child*Post.'

Did readers respond to the advisory?

Yes. Households purchasing newspapers or magazines reduced fish expenditure shares after the advisory. Panel A of Figure 5 indicates that shares fell after the advisory among readers in the non-target group. In contrast, Panel B shows that expenditure shares rose among non-targeted non-readers. Thus, non-target readers' share fell considerably after netting out changes to non-target non-readers. Panels C and D of Figure 5 represent changes in share by reading and non-reading target consumers, respectively. While both panels indicate a reduction in post-advisory expenditure shares, it appears that targeted readers responded somewhat more than targeted non-readers. Collectively, the figures suggest a fall in post-advisory shares for readers and a differential response among readers and non-readers.

Figure 5. Difference Graphs: Reading and Non-Reading Groups

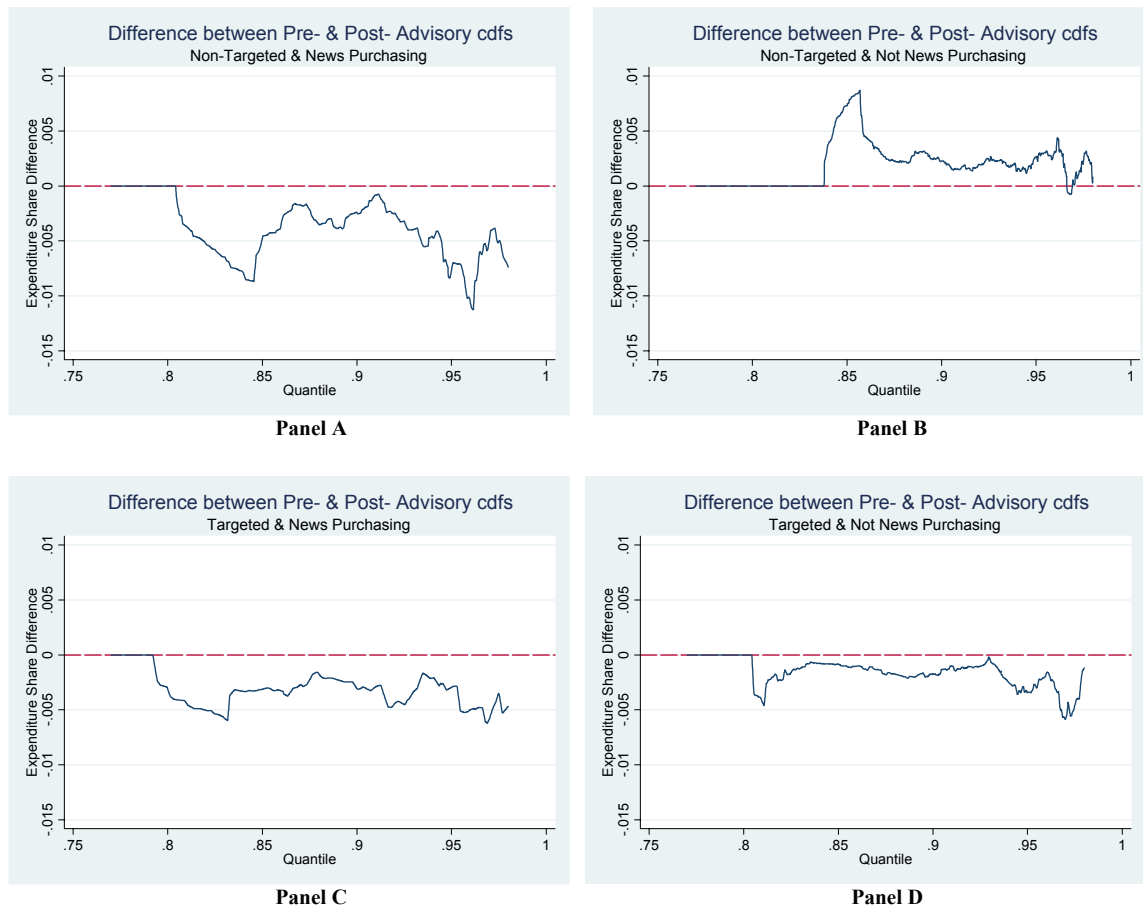


Table 5. Non-Parametric Tests Summary: Reading and Non-Reading Groups

	% Δ Mean Share	% Δ Mean Share, Given Positive Share	% Δ Fraction of Consumers with Positive Share	Q80?	Q85?	Q90?	Q95?
Non-Target: Readers, Net of Non-Readers	-28.6 (0.06)	+5.9 (0.64)	-30.7 (0.01)	No (.28)	Yes (.03)	Yes (.06)	Yes (.07)
Target: Readers, Net of Non-Readers	-19.1 (0.19)	-5.6 (0.36)	-15.9 (0.15)	No (.40)	No (.17)	No (.32)	No (.41)
Reading: Target, Net of Non-Target	-7.6 (0.39)	-25.2 (0.13)	+0.9 (0.51)	No (.45)	No (.57)	No (.46)	No (.67)
Non-Reading: Target, Net of Non-Target	-25.4 (0.08)	-8.3 (0.30)	-12.9 (0.09)	No (.29)	Yes (.05)	Yes (.03)	Yes (.06)

P-values are in parentheses. Quantile Treatment Effects evaluated at $\alpha=0.10$.

The non-parametric tests confirm the visual evidence. Among non-targeted consumers,

the drop in the overall mean expenditure share allocated to fish for the reading group (net of mean changes for the non-reading group) was 28.6 percent. The corresponding DDM statistic is statistically different from zero at the 6 percent level. Most of the change is attributable to changes in the number of consumers; the net proportion of non-target readers consuming any fish fell more than 30 percent (p-value of 0.01), while the net share conditional on purchase remained relatively constant. Results are not restricted to the means of the expenditure share distributions; we find evidence of significant quantile treatment effects for percentiles above Q85. Results for the smaller target subgroup are less pronounced; all DDM statistics are not statistically different from zero.

It may initially seem curious that the marginal impact of reading is more significant among the non-targeted than among the targeted. However, the baseline comparison for targeted consumers should already be lower than the baseline comparison for non-targeted consumers. In other words, some of the targeted respond despite being non-readers, while few of the non-targeted do so. Further, among readers, we cannot reject the hypothesis that target and non-target groups respond in the same way.

Among non-readers, the drop in the overall mean expenditure share for the target group (net of mean changes for the non-target group) was a statistically significant 25.4 percent. Results are not restricted to the mean, as quantile treatment effects above Q85 are significant at 10 percent levels. Share changes among target groups, relative to non-targeted groups, are not significant for readers. This latter result may initially seem puzzling. However, the result is consistent with the pronounced spillover effects detected above. Non-targeted readers significantly reduce their consumption, despite not being considered at-risk by the advisory. Consequently, the baseline comparison for the impact of target membership among readers is already lower than the baseline comparison for the impact of target membership among non-readers.

Regression coefficients support this analysis. Looking at Table 2, we see that the ‘Reader*Post’ row indicates a consistently significant impact of the advisory among non-targeted readers. However, the ‘Reader*Child*Post’ row reflects no incremental impact of being in the targeted group on readers. Therefore, we find that readers, as a group, reduced expenditure after the advisory relative to non-readers. However, there is no detected difference among readers across the target and non-target groups.²⁶

Did educated households respond to the advisory?

Yes. Educated households responded strongly, but only if they are in the target group. First, consider Panels A and B of Figure 6. Neither educated nor less educated non-target households seem to substantially change expenditure shares, since the net integrated areas are quite small in each graph. In contrast, Panel C shows a sharp drop in shares among educated households with young or nursing children. Panel D shows little change for less-educated households with young or nursing children. Comparing Panel C to D suggests a very strong impact of education for the response of the target group relative to the non-target group.

²⁶ Both parametric and non-parametric results may be biased if readership and fish consumption are joint consumption decisions in an economically meaningful way. For example, suppose households purchase newspapers and magazines primarily to find out about food safety. Our working assumption is that such joint decisions represent, at most, a small portion of readership and should not bias results. Further, our data suggests this concern is perhaps not practically important here; for both target and non-target consumers, news readership remains relatively constant across periods.

Figure 6. Difference Graphs: Educated and Less Educated Groups

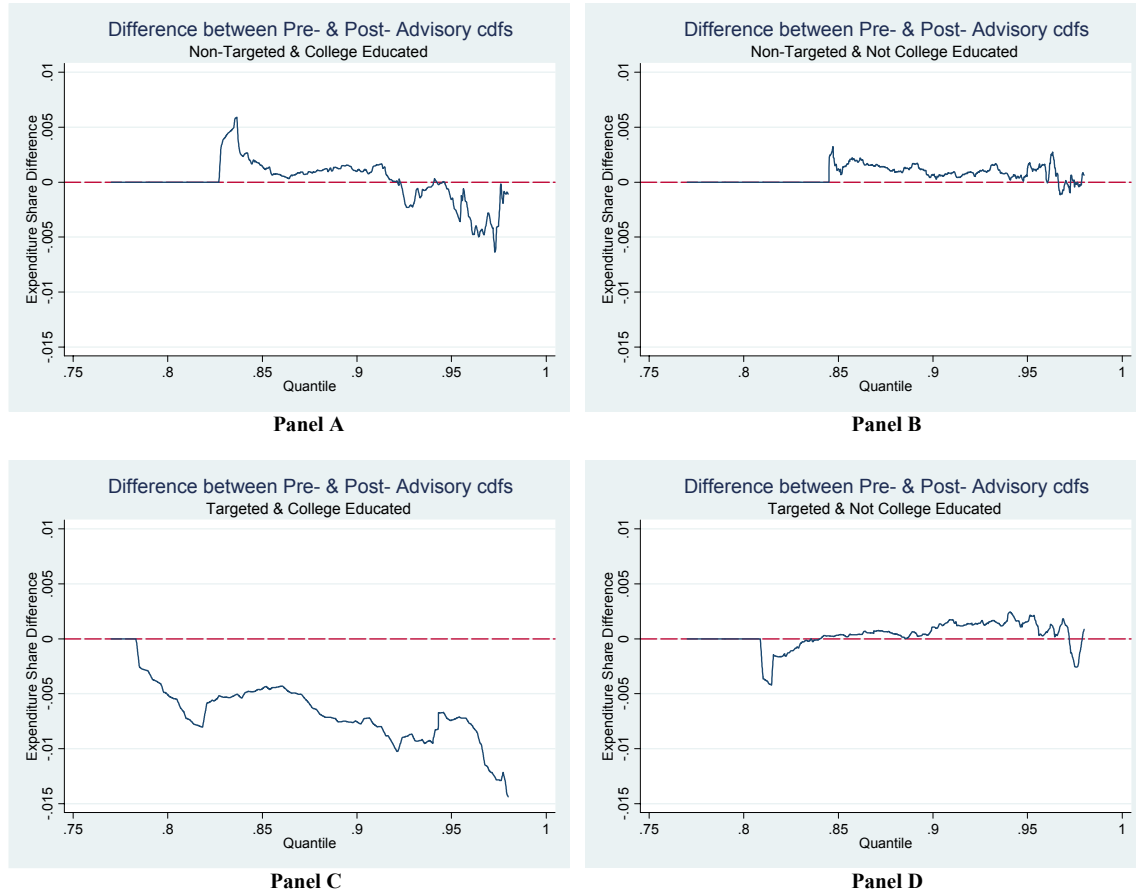


Table 6. Non-Parametric Tests Summary: Educated and Less Educated Groups

	%Δ Mean Share	%Δ Mean Share, Given Positive Share	%Δ Fraction of Consumers with Positive Share	Q80?	Q85?	Q90?	Q95?
Non-Target: Educated, Net of Less Educated	-1.5 (0.49)	-6.2 (0.35)	+4.2 (0.64)	No (.47)	No (.48)	No (.51)	No (.32)
Target: Educated, Net of Less Educated	-50.2 (0.01)	-43.2 (0.01)	-13.5 (0.17)	No (.27)	Yes (.02)	Yes (.01)	Yes (.01)
Educated: Target, Net of Non-Target	-51.2 (0.02)	-34.3 (0.06)	-20.4 (0.07)	No (.21)	Yes (.03)	Yes (.01)	Yes (.10)
Less Educated: Target, Net of Non-Target	-3.4 (0.43)	-0.5 (0.49)	-4.3 (0.34)	No (.32)	No (.36)	No (.48)	No (.52)

P-values are in parentheses. Quantile Treatment Effects evaluated at $\alpha=0.10$.

Statistical tests once again support the graphical analyses. Among non-targeted consumers, the DDM statistic for the percent change in the overall mean expenditure share for the educated subgroup (net of mean changes for the less educated subgroup) is not significant at conventional levels (p-value of 0.49). Further, none of the non-parametric tests indicate any differential response among educated non-target consumers and less educated non-target consumers. In contrast, among target consumers, the drop in overall mean expenditure share allocated to fish in the educated group (net of changes for the less educated group) was more than 50 percent (p-value of 0.01). The fall is attributable to changes in both the proportion of at-risk educated consumers that purchase at all and shares conditional on consuming. Results are robust across the distributions; we find evidence of strongly significant quantile treatment effects above Q80.

Regression results tell a similar story. In Table 2, the interaction row labeled ‘Educated*Child*Post’ presents evidence about the impact of education on response patterns of targeted households, *beyond* any impact on educated non-target households. For both specifications, we find a statistically significant effect for both the number of consumers and the mean quantity. Educated households with young or nursing children strongly reduced consumption after the advisory, relative to the control group. Contrast these results to the row labeled ‘Educated*Post,’ which indicates the advisory response among non-targeted educated households was insignificant.

B. Sensitivity Analysis Single Differences

The primary non-parametric and parametric results in the previous section emphasize differential consumption changes between two groups. For example, we find a differential response between target and non-target consumers, readers and non-readers (whether targeted or not), and educated and less educated consumers (if targeted). The motivation for differential responses is robustness to common shocks.

However, sharper results can be obtained by assuming the unobserved common shocks are small on average. An additional benefit of this assumption is that estimated response magnitudes are absolute and more readily interpretable. In practice, imposing this restriction amounts to performing single difference non-parametric tests, without reference to a control group.

Single difference statistical test (SDM) results are presented in Table 7.²⁷ The most important feature is that the single difference results closely resemble the difference-in-differences results, so unobserved shocks are likely small on average. For responding groups, post-advisory consumption fell, both absolutely and relative to control groups. Non-target readers' mean expenditure share fell by approximately one-fourth (DDM 28.6%, SDM 13.7%), target readers' mean share fell on the order of one-fifth (DDM 19.1%, SDM 27.7%), and the target educated groups' mean share fell by approximately one-half (DDM 50.2%, SDM 43.6%).

The similarity between single difference and difference-in-differences results also suggests that there is no meaningful advisory response among those groups with statistically undetected consumption changes. Of course, this depends upon the power of the tests. For those least likely to be knowledgeable about the advisory (the non-reading, less educated group examined in the last rows of Table 7), we find an insignificant response. Tests of power at the 90

²⁷ We also ran the corresponding single difference parametric regressions. These regressions included observable potential time variant shocks like prices and substitute prices, but omitted time-varying

percent confidence level reveal that this group's overall mean expenditure share decrease is less than 13 percent, their mean share decrease conditional on fish purchase is less than 12 percent, and their mean fall in the proportion of consumers purchasing any fish is less than 8 percent.²⁸ In other words, changes for this uninformed group are relatively small, if not precisely zero.

constants. Results were consistent with both the difference-in-differences regression results and the non-parametric single difference results.

²⁸ Difference-in-differences results and analyses, not reported in the interest of space, are similar.

Table 7. Non-Parametric Tests Summary: Single Differences

	%Δ Mean Share	%Δ Mean Share, Given Positive Share	%Δ Fraction of Consumers with Positive Share
<u>Target and Non-Target</u>			
Target Group	-11.3 (0.14)	-4.1 (0.31)	-7.5 (0.15)
Non-Target Group	+8.8 (0.82)	+5.9 (0.78)	+2.8 (0.68)
<u>Healthy and Less Healthy</u>			
Non-Target Group: Healthy	-2.6 (0.46)	+0.3 (0.52)	-2.9 (0.40)
Non-Target Group: Less Healthy	+12.4 (0.87)	+7.6 (0.81)	+4.4 (0.75)
Target Group: Healthy	-21.7 (0.14)	-14.1 (0.19)	-8.9 (0.26)
Target Group: Less Healthy	-6.5 (0.31)	+0.6 (0.53)	-7.0 (0.21)
<u>Reading and Non-Reading</u>			
Non-Target Group: Readers	-13.7 (0.18)	+9.4 (0.77)	-21.1 (0.02)
Non-Target Group: Non-Readers	+16.4 (0.91)	+2.8 (0.63)	+13.2 (0.96)
Target Group: Readers	-27.7 (0.05)	-10.7 (0.17)	-19.0 (0.08)
Target Group: Non-Readers	-7.2 (0.29)	-4.0 (0.34)	-3.3 (0.34)
<u>Educated and Less Educated</u>			
Non-Target Group: Educated	+6.7 (0.65)	+1.1 (0.54)	+5.6 (0.70)
Non-Target Group: Less Educated	+9.6 (0.80)	+8.0 (0.82)	+1.5 (0.58)
Target Group: Educated	-43.6 (0.01)	-32.7 (0.01)	-16.2 (0.09)
Non-Target Group: Less Educated	+8.4 (0.72)	+11.8 (0.86)	-3.1 (0.36)
<u>Non-Reading and Less Educated</u>			
Non-Target Group: Less Educated Non-Readers	+16.0 (0.87)	+3.0 (0.62)	+12.6 (0.92)
Target Group: Less Educated Non-Readers	+14.2 (0.78)	+10.1 (0.79)	+3.7 (0.63)

p-values in parentheses

Sensitivity to Other Assumptions

The results of the preceding sections are consistent across both single-difference and difference-in-differences graphical, non-parametric, and regression analyses. Below, we provide evidence these results are robust to choices of proxy-variable definitions, model structure, error specification, and the precise nature of the ‘event’.

When the threshold for “educated” is defined as a college degree, we found a strong differential response compared to less educated target consumers. Increasing the threshold to some graduate education amplifies this difference. However, upon decreasing the threshold to high school graduation, the difference with the less educated group is no longer statistically significant.

Our definition of ‘health conscious’ is *ad hoc*. However, results are not sensitive to the construction of the proxy variable. A wide variety of plausible indices and thresholds were considered without finding any differential response between healthy and unhealthy consumers.

Our regression model assumes a mean-zero error, implying that the sample average is a consistent estimate of true market demand. If zero-censoring of the dependent variable due to observation error is a concern, a Tobit correction would be in order. Therefore, we tested a supplementary Cragg (1971) correlated double-hurdle model to address this concern. The results for this specification were quite similar to those reported.

Another possible concern is the sharpness of our study’s ‘event’ (the January 2001 advisory). For example, perhaps consumers were broadly aware of concerns about mercury exposure prior to the announcement, since a number of states had issued advisories for recreational fish before the FDA action. One might also be concerned that the possibility of a FDA advisory was widely publicized long before its actual release. However, experiments indicate that these concerns are unsupported. For example, we do not find differential responses between those in the eight states that issued their own commercial advisories and those in other

states. Further, target groups' expenditures remained unchanged or increased between each of pre-advisory years (1997/1998, 1998/1999, 1999/2000). We also detected no systematically differential response among educated and uneducated target consumers (the most divergent groups in our analysis) prior to the advisory itself.

There are other sound reasons to believe the event is properly defined. First, the FDA issued the advisory within months of initially considering action. Second, FDA focus groups conducted in October 2000 (two months before the advisory) indicated, "None of the [focus] groups showed much interest or concern about mercury as a hazard in fish before seeing the information pieces....There was little or no awareness in any group of a hazard due to low level mercury exposure from fish consumption that was not due to a specific [localized] pollution problem." (FDA 2000) Finally, if consumers had already reacted to the mercury hazard, it would be difficult to reconcile the observed differential responses after the advisory between educated and less educated consumers.

VI. Discussion & Conclusion

We find that some targeted consumers significantly reduced canned fish purchases as a result of the FDA mercury advisory of January 2001. In particular, college educated consumers in the target group responded strongly. Among households with young and nursing children, mean canned fish expenditure share fell by approximately 50 percent after the advisory, accounting for common shocks. In marked contrast, we detected no statistically significant response among those with less education.

We found that newspaper or magazine purchases were associated with a significant reduction in post-advisory consumption. Among households that purchased newspapers or magazines in the diary period, mean fish expenditure fell by 19 percent after the advisory, accounting for common shocks. However, we found no differential response among targeted readers and non-targeted readers.

Access to information and the ability to assimilate information were important limiting factors in the advisory response. We view newspaper readership as a reasonable proxy for exposure to information about the dangers of mercury in fish, and readers responded. We also view college education as a reasonable proxy for the ability to assimilate information appropriately, and educated individuals responded only if targeted by the advisory.

Can the observed changes in consumption be attributed to the FDA policy? The responses are consistent with increased information about mercury hazards. Further, FDA focus groups found no public awareness of the relevant risks two months prior to the advisory. Although we do not know whether individual responding consumers were aware of the advisory *per se*, there is no doubt that the advisory resulted in much greater general public awareness of mercury risk. In this sense, an advisory can achieve the issuing agency's goals through promoting awareness, even if indirectly.

Targeted consumers likely to be aware of and understand the advisory tended to reduce fish consumption. So, mercury advisories and education programs can achieve policy goals like reducing the contaminant exposure of nursing and young children. However, those at-risk consumers least likely to be knowledgeable about the advisory do not significantly reduce consumption. Notably, this large group of non-college educated, non-readers is also likely to be poorly equipped to withstand negative health shocks. At a minimum, these latter results suggest that a broader and more targeted educational outreach program may be necessary to reach many vulnerable members of society. Possible enhanced outreach methods include health-advertising campaigns (on public transportation, for example), in-store advisory signs, and mandatory product labeling. Mathios (2000) showed that labeling induces important consumption responses, and Teisl *et al.* (2002) showed that point of consumption labeling is particularly effective for canned fish.²⁹

²⁹ Teisl *et al.* (2002) examined the impact of “dolphin-safe” eco-labeling on tuna consumption.

More broadly, we find that well-informed consumers actively respond to environmental risk warnings. Prominent advisories can be an effective and low-cost means of achieving policy goals. However, particular attention must be paid to less educated and less informed consumers. On another cautionary note, our results also indicate that informed individuals may respond more broadly than intended, as non-targeted readers reduced fish consumption after the mercury advisory. While this may be a rational or even optimal response, it is not consistent with the stated intent of the advisory. Therefore, advisories and outreach programs should be carefully crafted with such spillovers in mind.

This paper suggests promising avenues for future research that are beyond the present scope. First, what is the social value of information based policies? Second, given this value, how do advisories compare to other policy options for managing mercury exposure? Advisories may significantly reduce exposure. However, a full welfare analysis requires evaluating the accuracy of the advisory's underlying scientific information. This paper identifies how consumers responded to a prominent advisory, but does not purport to assess the validity of that advisory's content. Further, there are a number of subtle theoretical issues in developing an accurate empirical estimate of the benefits of demand-side mitigation. These welfare questions form a critical element of an ongoing research agenda in this area.

Acknowledgements

We wish to thank numerous seminar participants for helpful comments and suggestions (the NBER, the National Academies, UCSD, UCB, UCSB, UBC, UW, WSU, AERE, and AAEEA). Special thanks to Chris Costello, Josh Graff Zivin, Jeff LaFrance, Carol McAusland, Gib Metcalf, Steven Yamarik, and two anonymous referees.

References

- Adler, Richard, and R. David Pittle, "Cajolery or Command: Are Education Campaigns an Adequate Substitute for Regulation?," *Yale Journal of Regulation* 1:2 (1984), 159-194.
- Belton, Thomas, R. Roundy, and Neil Weinstein, "Urban Fishermen: Managing the Risks of Toxic Exposure," *Environment* 28 (1986), 19-37.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Difference-in-Differences Estimates?" *Quarterly Journal of Economics* (2004), 249-275.
- Bitler, Marianne, Jonah Gelbach, and Hilary Hoynes, "Distributional Impacts of the Self-Sufficiency Project," *NBER Working Paper 11626* (2005).
- Bitler, Marianne, Jonah Gelbach, and Hilary Hoynes, "What Mean Impacts Miss: Distributional Effects of Welfare Reform Experiments," Forthcoming, *American Economic Review* (2006).
- Blundell, Richard, and Costas Meghir, "Bivariate Alternatives to the Tobit Model," *Journal of Econometrics* 34 (1987), 179-200.
- Brown, Deborah J., and Lee F. Schrader, "Cholesterol Information and Shell Egg Consumption," *American Journal of Agricultural Economics* 72:3 (Aug. 1990), 548-555.
- Cragg, John G., "Some Statistical Models for Limited Dependent Variables with Application to the Demand for Durable Goods," *Econometrica* 39:5 (Sept. 1971), 829-844.
- Deaton, Angus, and Margaret Irish, "Statistical Models for Zero Expenditures in Household Budgets," *Journal of Public Economics* 23 (1984), 59-80.
- Fenn, Aju J., Frances Antonovitz, and John R. Schroeter, "Cigarettes and addiction information: new evidence in support of the rational addiction model," *Economics Letters* 72 (July 2001), 39-45.
- Foster, William, and Richard E. Just, "Measuring Welfare Effects of Product Contamination with Consumer Uncertainty," *Journal of Environmental Economics and Management* 17 (1989), 266-283.
- Fry, Vanessa, and Panos Pashardes, "Abstention and Aggregation in Consumer Demand: Zero Tobacco Expenditures," *Oxford Economics Papers* 46:3 (July 1994), 502-518.
- Grossman, Michael, "On the Concept of Health Capital and the Demand for Health," *Journal of Political Economy* 80 (1972), 223-255.
- Grossman, Michael. *The Demand for Health: A Theoretical and Empirical Investigation*. New York: Columbia University Press. 1972.
- Heckman, James, "The Common Structure of Statistical Models of Truncation, Sample Selection, and Limited Dependent Variables and a Simple Estimator for Such Models," *Annals of Economic and Social Measurement* 5 (1976), 475-492.

- Heckman, James J., Jeffrey Smith, and Nancy Clements, "Making the Most Out of Programme Evaluations and Social Experiments: Accounting for Heterogeneity in Programme Impacts," *Review of Economic Studies*, 64:4, (Oct. 1997), 487-535.
- Hogg, Robert V. and Allen T. Craig, *Introduction to Mathematical Statistics*, 4th ed. Macmillan, New York, 1978.
- Ippolito, Pauline and Alan Mathios, "Information, Advertising, and Health Choices: A Study of the Cereal Market," *RAND Journal of Economics* 21 (1990), 459-480.
- Jakus, Paul, Meghan McGuinness, and Alan Krupnick, "The Benefits and Costs of Fish Consumption Advisories for Mercury," *RFF Discussion Paper 02-55* (Oct. 2002).
- Kinnucan, Henry W., Hui Xiao, Chung-Jen Hsia, and John D. Jackson, "The Effects of Health Information and Generic Advertising on U.S. Meat Demand," *American Journal of Agricultural Economics* 79:1 (Feb. 1997), 13-23.
- Krinsky, Itzhak and A. Leslie Robb, "On Approximating the Statistical Properties of Elasticities," *The Review of Economics and Statistics*, 68:4 (Nov., 1986), 715-719.
- Lancaster, Kelvin, "A New Approach to Consumer Theory," *Journal of Political Economy* 74 (April 1966): 132-157.
- Lino, Mark. 2006. "Expenditures on Children and Families, 2005," US Department of Agriculture, Center for Nutrition Policy and Promotion. Miscellaneous Publication No. 1528-2005.
- Maasoumi, Esfandiar and Daniel L. Millimet "Robust inference concerning recent trends in US environmental quality," *Journal of Applied Econometrics* 20(1), 55 – 77.
- Mathios, Alan D, "The Impact of Mandatory Disclosure Laws on Product Choices: An Analysis of the Salad Dressing Market," *Journal of Law and Economics* 43(Oct. 2000), 651-677.
- May, H., and Joanna Burger, "Fishing in a Polluted Estuary: Fishing Behavior, Fish Consumption, and Potential Risk," *Risk Analysis* 16 (1996), 459-71.
- Meyer, Bruce D., W. Kip Viscusi, and David L. Durbin, "Workers' Compensation and Injury Duration: Evidence from a Natural Experiment," *American Economic Review*, 85:3 (Jun., 1995), 322-340.
- Millimet, Daniel, Dipanwita Sarkar, and Esfandiar Maasoumi, "A Distributional Analysis of the Return to Marriage," *mimeo* Southern Methodist University, 2004.
- Poe, Gregory L., Eric K. Severance-Lossin, and Michael P. Welsh, "Measuring the Difference (X - Y) of Simulated Distributions: A Convolutions Approach" *American Journal of Agricultural Economics*, 76:4 (Nov., 1994), 904-915.
- Pollack, Robert and Terrence J. Wales, "Demographic Variables in Demand Analysis," *Econometrica* 49:6 (Nov. 1981), 1533-1551.

Sloan, Frank A., V. Kerry Smith, and Donald H. Taylor, "Information, addiction, and 'bad choices': lessons from a century of cigarettes," *Economics Letters* 77:2 (Oct. 2002), 147-155.

Teisl, Mario F., Brian Roe and Robert Hicks. "Can Eco-Labels Tune a Market? Evidence from Dolphin-Safe Labeling of Canned Tuna," *Journal of Environmental Economics and Management*, 43 (May 2002): 339 – 59.

U.S. Department of Labor. Bureau of Labor Statistics. Consumer Expenditure Survey, Diary Survey. (Washington, DC: BLS, 1999-2002). Inter-University Consortium for Political and Social Research (Ann Arbor, MI: ICPSR, 1999-2002) [distributor].

U.S. Environmental Protection Agency, "Mercury Study Report to Congress," (Washington, DC: EPA-452/R-97-003, Dec 1997).

U.S. Food and Drug Administration , "Report Submitted to Interagency Working Group on Mercury Contamination: Findings from the Focus Group Testing of Mercury-in-Fish Messages" (Washington, DC: FDA, Oct 23,2000).

U.S. National Research Council: Committee of the Toxicological Effects of Methylmercury, *Toxicological Effects of MethylMercury* (Washington, DC: National Academy Press, June 2000).

Viscusi, W. Kip, Wesley A. Magat, and Joel Huber, "Informational Regulation of Consumer Health Risks: An Empirical Evaluation of Hazard Warnings," *The Rand Journal of Economics* 17 (Autumn 1986), 351-365.