Averting expenditures and desirable goods: Consumer demand for bottled water in the presence of fracking

A. Justin Kirkpatrick and T. Robert Fetter

May 20, 2017

Preliminary and Incomplete
Do Not Cite or Circulate

Abstract

Environmental conditions such as new industrial activity or drinking water quality violations may affect perceived water quality and cause individuals to invest in defensive expenditures. Several recent papers use household expenditures on bottled water to measure the welfare effect of changes in water quality, arguing that these averting expenditures are a lower bound on compensating variation. We offer a new perspective and argue that when consumers choose averting behaviors such as substituting bottled water for tap water, their willingness to pay incorporates other characteristics of the good. If consumers gain utility from these characteristics, defensive expenditures need not represent a lower bound on compensating variation for a change in water quality. Rather, the observed expenditures should be adjusted to account for consumers' increased utility due to other desirable characteristics. To provide an empirical context, we consider consumers' purchases of bottled water in response to the entry of hydraulic fracturing (fracking) activity in Pennsylvania and Ohio in the past decade. We develop a structural model of demand for bottled water and estimate it using fine-resolution supermarket scanner data, and compare the resulting estimates to averting expenditures from a reduced-form model. Our results suggest consumers increase purchases of bottled water in response to the entry of fracking, with an average increased expenditure of \$3.43 per household per quarter in the preferred model; however, only \$0.39 per household per quarter is attributable to the specific attribute of avoiding exposure to tap water potentially contaminated (or perceived to be contaminated) by fracking.

Keywords: Averting Expenditures, Structural vs Reduced-Form Models, Hydraulic Fracturing

JEL codes: C18, D62, Q33, Q51, Q53

Kirkpatrick (`justin.kirkpatrick@duke.edu`) and Fetter (`rob.fetter@duke.edu`) are at the Duke University Program on Environmental Policy, Durham, NC. We are grateful to Douglas H. Wrenn, Lori Bennear, Bryan Bollinger, and Chris Timmins for their helpful comments. All errors and omissions are our own. The empirical analysis herein is based on data from The Nielsen Company (US), LLC and marketing databases provided by the James M. Kilts Center for Marketing Data at the University of Chicago Booth School of Business.

1 Introduction

The rise of hydraulic fracturing ('fracking') in the United States has led to a dramatic increase in oil and gas extraction as well as a significant spatial shift in extraction activity. Many areas located over shale deposits have seen a host of benefits including jobs, lower energy prices, royalties, and tax revenues, as well as attendant increased burdens on public infrastructure including roads, water supplies, and wastewater treatment (Hausman and Kellogg, 2015). Furthermore, public fears of potential drinking water contamination have grown prominent.

Fracking involves extracting gas and/or oil from geologic formations below the caprock that forms the geologic floor for groundwater. This requires a well bore to pass through groundwater strata. When well bore casings are not properly sealed, it is possible for either components of drilling and fracking fluids or geologic methane to intrude into groundwater aquifers. Contamination risks also arise from other aspects of the production process, such as possible leaks in plastic liners on holding ponds, or spills from trucks or pipelines. As the magnitude and consequences of these risks are not well understood, households build perceptions of risks to water quality; these perceptions, in turn, influence household choices over consumption of tap water and alternative water sources. Previous research suggests that nearby drill rigs and fracturing activity cause decreasing property values (Muehlenbachs, Spiller, and Timmins, 2015; Gopalakrishnan and Klaiber, 2013), decreased property rents (Muehlenbachs, Spiller, Steck, et al., 2015), and increased purchases of bottled water (Wrenn, Klaiber, and Jaenicke, 2016).

One way to estimate welfare effects from changes in environmental quality, such as water quality, is to measure averting or defensive expenditures: that is, household expenditures on goods or services that ameliorate or substitute for the change in environmental quality. Intuitively, if averting expenditures reduce households' exposure to pollution—but households can reduce their exposure only so far—the averting expenditure represents a lower bound on consumers' willingness-to-pay for improved environmental quality. However, when the behavior or substitute has other characteristics that provide positive utility, the averting expenditure may not represent a lower bound after all (Bartik, 1988). When consumers choose averting behaviors, their willingness to pay includes other characteristics of the substitute good, besides perceived higher quality. For instance, consumers who perceive water quality risks from fracking and choose to substitute bottled water for tap water may also enjoy the taste of bottled water, or the portability of bottles. If these characteristics enter consumers' utility functions directly, then expenditures on bottled water do not represent a lower bound on the compensating variation for a change in water quality. Rather, the observed expenditures should be adjusted downward to account for the component of utility that is due to other desirable characteristics.

We employ models common in the industrial organization field to estimate household demand functions for bottled water, and include "avoiding potentially contaminated tap water" as a time- and location-varying attribute of water purchases. We begin with a traditional horizontal consumer demand model extended to allow for consumer heterogeneity over observable and unobservable characteristics, as well as estimating household-level fixed effects in preferences for the "outside good", tap water. We estimate our model using supermarket scanner data with fine resolution both spatially (zip code) and temporally (weekly). We are in the process of extending our model to allow for dynamic stockpiling and consumption.

We also evaluate the change in pure expenditure due to the entry of fracking in a reduced-form context. This allows us to compare the WTP implied by a pure-expenditure model to the WTP implied by the demand model. Preliminary results from our preferred models (at this time) suggest that consumers respond to the entry of fracking by increasing purchases of bottled water, with an average expenditure of \$3.43 per household per quarter, but only \$0.39 per household per quarter is attributable to the specific attribute of avoiding exposure to tap water potentially contaminated (or perceived to be contaminated) by fracking.[1]

Our paper makes several contributions. First, we extend the environmental valuation literature by applying a structural model of consumer demand to a private retail good inextricably linked to environmental quality. As information on consumer purchases becomes more readily available, it becomes more feasible for researchers to assess willingness to pay for non-market goods, like environmental quality, by measuring changes in consumption of market substitutes. Some previous researchers have used supermarket scanner data in reduced-form models to estimate willingness to pay for environmental quality (Graff Zivin, Neidell, and Schlenker, 2011; Wrenn, Klaiber, and Jaenicke, 2016), suggesting the resulting estimates represent a lower-bound willingness to pay. These analyses do not explicitly address the degree to which the supposed "lower bound" may in fact be overstated due to joint production (of utility) that arises from desirable product characteristics. We address this, and offer an empirical demonstration of an alternative approach that addresses the concern noted by Bartik (1988).

Second, we contribute to research on the benefits and costs of increased fracturing activity. In addition to many economic benefits, other authors have documented costs associated with air pollution emissions, increased trucking, habitat fragmentation, and noise and crime (Mason, Muehlenbachs, and S. Olmstead, 2015). Perceived and real impacts to water quality, including drinking water resources, have also been a central concern among policymakers and local residents. In theory, as Hausman and Kellogg (2015) note, any observed changes in home values capture the value of all local environmental disamenities to the marginal resident. However, changes in housing prices may also capture the effects of local booms, in addition to environmental disamenities.[2] In addition, when consumers have heterogeneous preferences, changes in housing prices may not accurately capture marginal valuations or welfare effects (Kuminoff and Pope, 2014; Hausman and Kellogg, 2015). Furthermore, policymakers and others may wish to understand how much of a composite impact is attributable to concerns about water quality, and studying this (and potentially other component parts) provides useful information about the overall magnitude of the value of environmental disamenities.

[1]The \$0.39 figure is based on the monetary value of the change in indirect utility associated with consumption of bottled water per ounce, and assumes constant marginal utility of bottled water consumption. As we discuss in detail later, declining marginal utility is a more realistic assumption, and would result in a lower estimate. On the other hand, it may also be appropriate to multiply the per-ounce indirect utility by total household water consumption (e.g., 32 ounces per person per day) rather than observed bottled water consumption, and this would result in a higher estimate.

[2]Local boom effects may also affect consumers' expenditures on bottled water. However, we can control for this explicitly because we observe (time-varying) household-level income and expenditures.

Although not reflected in this draft, we expect future drafts will make additional contributions. One of these is to develop a novel measure of the salience of fracturing activity to households. We plan to combine detailed data on the spatial location of fracking sites and timing of visible development activity with LiDAR satellite data to assess the visibility of fracking activity over terrain contours (accounting for hills and valleys) and vegetative obstructions. This will allow us to develop a fine-grained measure of consumer exposure: the proportion of households within each geographic market that experience visible evidence of shale gas development over specific days or weeks. This metric offers substantially improved precision over measures used in previous analyses (Muehlenbachs, Spiller, and Timmins, 2015; Wrenn, Klaiber, and Jaenicke, 2016).

The remainder of the paper proceeds as follows: Section 2 provides background on the empirical setting and discusses related literature. In Section 3 we describe our model; in Section 4 we document the data we use. In Section 5 we provide results and discussion, and Section 6 concludes.

2 Background

2.1 Hydraulic Fracturing

The rise of hydraulic fracturing for shale gas in US energy production has been dramatic. Shale gas grew from 5% of total US dry gas supply in 2004 to 56% in 2015.[3] Thanks to the suite of technologies that has allowed production from formations that were previously judged uneconomic, natural gas has largely replaced coal in the production of electricity [4]. The largest contribution to shale gas has been from the Marcellus Shale, which underlies Pennsylvania, New York, Ohio, and West Virginia.

The dramatic growth in production has brought significant economic benefits as well as environmental concerns. Among the more prominent environmental concerns is the potential for contamination of surface water and groundwater. The fracturing process involves the high-pressure injection of millions of gallons of fluid down a wellbore, including chemicals that may be toxic or regulated (Stringfellow et al., 2014; Fetter, 2016). After the fracture has been completed, much of this water, as well as other produced water from the shale formation, may return to the surface, bearing contaminants from deep underground (sometimes including heavy metals or radionuclides (S. M. Olmstead et al., 2013)). Since the wellbore and production casing must extend through groundwater resources to reach the productive shale, concerns have been raised that improper casing or other errors in the production process could result in groundwater contamination.

Another source of possible water contamination arises from the disposal of flowback fluid. In the Marcellus region, especially in Pennsylvania, geologic features constrain the ability of operators to reinject the flowback fluid back underground. The flowback can sometimes be recycled into fracturing fluid for a subsequent fracture, but this is not always feasible, due to high concentrations of dissolved solids

[3] EIA Natural Gas Monthly data through December, STEO through May 2015 and Drilling Info; http://www.eia.gov/conference/2015/pdf/presentations/staub.pdf.)

[4] http://www.cnbc.com/2015/07/14/natural-gas-tops-coal-as-top-source-of-electric-power-generation-in-us.html

that may hinder its effectiveness (Blauch, 2010). Alternative disposal options include expensive truck transport across state lines to Ohio or West Virginia, where injection wells are more readily available. The other major disposal option is discharge to a wastewater treatment facility, but a number of studies have found that municipal facilities lack the technology to adequately remove contaminants frequently present in flowback water. As a result, flowback fluid disposal may threaten both surface water and groundwater resources.

Other research, as well as popular media such as the film *Gasland*, has addressed the possibility that methane could migrate through strata into groundwater resources. Although this possibility is contested in scientific literature, with several papers suggesting that producing shale gas resources relieves pressure exerted by gas formations and thus decreases the likelihood of methane infiltration into groundwater, it remains a matter of widespread public fear that may lead consumers to invest in defensive expenditures so as to avoid perceived risks to water quality.

2.2 Welfare Effects and Averting Behavior

There is a long history of economic literature on measuring averting expenditures on market goods to measure willingness to pay for improvements in non-market goods such as environmental quality. The basic idea recognizes that demand for 'defensive' products such as air filters is a function partly of the utility from consuming the outside option, such as unfiltered air. Intuitively, one way to estimate societal willingness to pay for a public good like high-quality ambient air would be to measure private expenditures on defensive technologies that allow individual households to avoid air pollution. However, to the extent that these defensive technologies do not allow households to avoid all of the ill effects, defensive or averting expenditures would represent a lower bound for willingness to pay.

Several recent empirical analyses have used logic along these lines to impute preferences for environmental quality, including willingness to pay for improved air quality based on expenditures for face masks (Mu and Zhang, 2015), for averting climate change (warmer temperatures) based on residential electricity consumption (Deschênes and Greenstone, 2011), and for higher quality drinking water based on expenditures for bottled water (Graff Zivin, Neidell, and Schlenker, 2011; Wrenn, Klaiber, and Jaenicke, 2016). Although the critique we highlight here—that defensive expenditures may offer additional desirable characteristics to consumers and therefore do not necessarily represent a lower bound—has been known in the literature since at least Bartik (1988), these analyses do not explicitly recognize the implications for the use of averting expenditures. Admittedly, this concern is likely to have a larger impact in some contexts than others: it is easy to believe bottled water has desirable characteristics other than possibly higher quality, for instance, but not so easy to believe the same is true of disposable face masks used by healthy individuals.[5]

Courant and Porter (1981) argued that averting expenditures may not, in fact, represent a lower bound for willingness to pay for improved environmental quality,

[5] In addition, we hope that policymakers are interested not just in the lower bound of WTP, but the true value, which would also include, for instance, health consequences that consumers do not manage to avoid by defensive expenditures (Graff Zivin, Neidell, and Schlenker, 2011; Reynolds, Mena, and Gerba, 2008).

depending on the consumer's utility function and the properties of the technology by which averting expenditures achieve their purpose. Bartik (1988) noted several reasons that using averting expenditures could be problematic. One of these is that the lower bound argument on averting expenditures does not hold when the expenditure enters the consumer's utility function directly (e.g., the case of joint production). For instance, air conditioning may reduce air pollutants, but also provides cooler air, which enters utility directly. Therefore, "information on a household's valuation of defensive measures for non-defensive reasons" is necessary in cases of joint production (Bartik, 1988; Dickie, 2003).

A simple example shows the joint production problem in averting expenditures applied to the present context. Let the market for bottled water consist of one product type and one consumer. The consumer is willing to pay \$.99 for the bottle of water based on the taste, portability, and brand of the good. The exogenous price of the good is \$1.00. Therefore, the consumer does not purchase the bottled water. However, once fracking appears in the consumer's vicinity, the consumer is willing to pay \$.02 to avoid consuming their tap water. Having a new willingness to pay of \$1.01 and a constant price of \$1.00, the consumer now purchases the bottled water. In an averting expenditures framework, a full \$1.00 is deemed an averting expense - it was not spent prior to fracking, but was spent after fracking appears. This does not account for the \$.99 of joint production. Although the compensating variation of the bottled water *as an alternative to consuming tap water possibly affected by fracking* is \$.02, the "lower bound" in an averting expenditures framework is \$1.00.

A graphical example is shown in Figure 1, where D_1 is the initial demand curve for bottled water, and D_2 is the demand after fracking arrives in the vicinity, increasing individuals' willingness to pay for bottled water. The reduced form averting expenditures estimate is $price \times (Q_2 - Q_1)$. However, the dollar equivalent is the area between D_1 and D_2 over Q_{max}. In this example, the reduced form estimate is not the lower bound - rather, the area between the curves is less than the area below S and between Q_2 and Q_1.

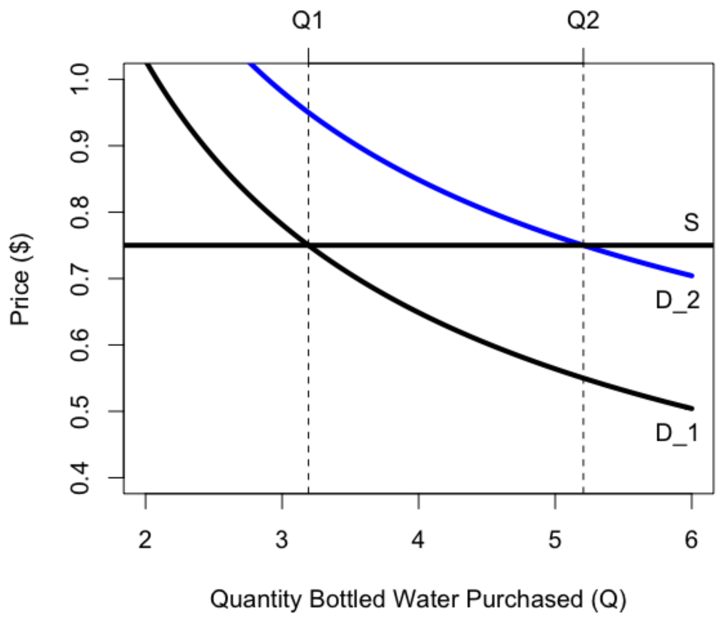

Figure 1: Simple Example of Structural Versus Reduced Form

Disentangling the change in utility resulting from the arrival of fracking requires estimating the parameters of the underlying indirect utility functions over heterogeneous consumers (Timmins and Schlenker, 2009), necessitating a structural model of consumer demand. Estimating on data in multiple markets with multiple consumers both pre- and post-fracking allows for parameterization of the individual indirect utility function with respect to fracking. With an indirect utility function including the marginal utility of income in hand, it is straightforward to calculate the dollar equivalent of utility changes between with- and without-fracking.

3 Model

3.1 Discrete choice structural model of consumer demand

We model consumer choices within a horizontal discrete choice framework common in the industrial organization and consumer demand literature. The use of a horizontal model allows for products that are differentiated along multiple dimensions and with varying characteristics. The discrete choice model used here is a type of random utility model (RUM) first established by Lancaster (1966) and McFadden et al. (1973). The advantage of a RUM model is that the specification is grounded in economic theory and is consistent with utility maximization while allowing for unobserved components Train (2009). Consumer tastes are represented by individual-level taste parameters, β_i. We employ a heterogeneous parameters logit, allowing β_i to vary over observable household characteristics with an eye towards extension of the model to a random parameters mixed logit Timmins and Murdock (2007).

Current consumer demand literature allows for heterogeneous unobserved consumer tastes for product characteristics via random parameters mixed logit (Walker and Ben-Akiva, 2002; Train, 2009). The advantage of a random parameter model is that unobserved consumer tastes for product characteristics may be correlated, which is approximately equivalent to correlations in the unobserved utility. This allows for more realistic substitution patterns (Train, 2009). However, in our case (and given computational constraints), we employ the heterogeneous parameters model and acknowledge that estimates may imply unrealistic substitution patterns. While this problem is significant in the classic example of luxury versus budget vehicles, it is not as severe in the context of substitution between bottled water products and brands.

The indirect utility for consumer i gained from consuming bottled water product j in market t is defined as $U(x_{jt}, p_{jt}, D_{it}; \theta)$, a function of observed product characteristics (x_{jt}), price (p_{jt}), individual consumer characteristics (D_{it}), and unknown parameters to be estimated (θ). Throughout this analysis, a market is defined at the store-week level and notated as t. This indirect utility is specified as:

$$u_{ijt} = \alpha_i(y_i - p_{jt}) + x_{jt}\beta_i + \epsilon_{ijt} \tag{1}$$

Where ϵ_{ijt} is an unobserved stochastic term distributed Type 1 Extreme Value. For any choice occasion, a high value of ϵ_{ijt} induces greater utility for consumer i for product j. We first normalize the outside good of "no purchase" to zero and note that the choice of product j depends only on relative utilities across products $j \in J$. Income enters the utility of each choice identically and, because only relative utilities matter, may be dropped from the specification[6]. Notating the choice outcome as $d_{ijt} = 1$ if product i is chosen and zero otherwise, the T1EV assumption on ϵ_{ijt} yields the following familiar logit probability:

$$Pr(d_{ijt} = 1) = Pr(u_{ijt} > u_{ikt} \quad \forall k \neq j) = \frac{exp(\alpha_i p_{jt} + x_{jt}\beta_i)}{1 + \sum_{k=1}^{J} exp(\alpha_i p_{kt} + x_{kt}\beta_i)} \tag{2}$$

Each choice occasion occurs within a market $t \in T$. We define a market at the store-week level and allow the choice set J to vary. In some stores (and in some weeks), bottled water offerings are sparse and may consist of only a few brand, size, and packaging choices. In others, choices are rich and varied. We take the consumer's choice of store to be exogenous and designate the choice set available in market t as J_t.

Consumer heterogeneity is captured by (α_i, β_i). Following Nevo (2000) but removing the random component of β_i, let D_i be a $[d \times 1]$ vector of demographic characteristics (e.g. number of children) and Π be a $[(k+1) \times d]$ matrix of coefficients which relate the $k+1$ taste characteristics to the d product characteristics. In a simple logit, β_i is a function of population mean parameters (α, β), observed characteristics D_i, and taste shifting parameters Π. We write tastes (α_i, β_i) as a heterogeneous parameters logit:

$$\begin{pmatrix} \alpha_i \\ \beta_i \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} + \Pi D_i \tag{3}$$

The presence of fracking is hypothesized to drive increased purchases of bottled water by decreasing the utility of the outside good (i.e., tap water). Because consumer choices are based only on *relative* utility, a reduction in the utility of the outside good is the equivalent of an *increase* in the utility of an inside good. That is, when fracking arrives in a consumer's vicinity, any choice of bottled water effectively has the attribute of "avoiding consuming potentially contaminated tap water". With the introduction of this attribute, the utility predicted by the product characteristics x_{jt} varies over individuals *even within the same market and with identical observable characteristics* D_{it} as the presence of fracking is measured in the consumer's zip code, but the market itself serves consumers from multiple zip codes. We assume the presence or absence of fracking around the store does not enter the consumer's utility.

Because the presence of fracking is identical over all products for any choice occasion (consumer i in market t), the presence of fracking must be interacted with a relevant characteristic of each product j. In all specifications, we interact the measure of fracking with the total quantity (ounces) of water characteristic present in x_{jt}. This yields a parameter that measures the increase in the marginal utility of an ounce of water per unit increase in the measure of fracking in a consumer's zip code. The ratio of this coefficient to α_i yields the marginal value of the attribute of bottled water associated with "avoiding consuming potentially contaminated tap water".

In many consumer demand applications, prices are endogenous to the system, especially in models which rely on aggregated market-share data (Berry, Levinsohn,

[6]This implies constant marginal utility of income, which is reasonable in light of the small budget share of bottled water expenditures we observe

and Pakes, 1995). In our application, we rely on individual-level observed purchases wherein each consumer is a price-taker. Prices are assumed to be exogenous in this context, though a possible extension of the model would include instrumenting for prices. Because prices vary at a fine spatial level (store) and temporal scale (week), instrumenting for endogeneity may prove difficult.

3.1.1 Model Specification

In our parsimonious specification we allow utility to vary over price, total ounces of water, the number of units in the package, and whether a good is a store brand, a national brand, a flavored brand, or an "other" brand. We specify observable consumer heterogeneity only with the interaction of the presence of >10 wells in the consumer's zip code and the use of private well water in that zip code. The consumer's traits are constant over all products in J. We therefore interact the well presence interaction with total ounces. The coefficient estimated, $\beta^{w,o}$ is the per-ounce of bottled water marginal utility of the presence of fracking wells. To calculate the dollar equivalent of the change in utility (compensating variation) associated with an increase of the number of wells by one percent, we take the ratio of $\beta^{w,o}$ to α_i. Since water is a biological requirement, consumers cannot substitute out of water altogether. Therefore, re-optimization is likely minimal, and compensating variation is close to the true welfare measure.

3.1.2 Dynamics

Bottled water is a highly storable good. This allows consumers to purchase larger quantities at once and then stockpile the good. Failure to model the dynamic process that consumers use to optimize over time leads to potentially significant bias (Hendel and Nevo, 2002; Hendel and Nevo, 2006). Further work in this area is necessary, but remains outside of the scope of this paper for the moment.

3.1.3 Estimation

Estimation of the model is straightforward and follows Train (2009), using Maximum Likelihood (ML) to find the parameters $\theta = \{\alpha, \beta, \Pi\}$ which maximize the likelihood of observing the data.

3.2 Reduced form model of averting expenditures

For purposes of comparison to the structural model estimates, we build a reduced-form model to estimate households' averting expenditures in response to the arrival of fracking in close proximity. We use a difference-in-differences approach to isolate the effect of fracking from other demand shifters that vary over space and time.[7]

[7]Wrenn, Klaiber, and Jaenicke (2016) also use a reduced-form approach to estimate averting expenditures in response to the same event in the same geographic area (i.e., fracking in Pennsylvania and Ohio). However, we appear to have a slightly different data set on household shopping trips and expenditures; thus, since our main purpose is to compare results from a reduced-form and a structural model, we provide reduced-form estimates from our own data set. This also allows us to take advantage of having a longer time series than Wrenn, Klaiber, and Jaenicke (2016) (household purchase data through 2014) and consumer expenditure data at a finer geographic resolution (zip code rather than county), which means we can add a richer set of fixed effects.

Muehlenbachs, Spiller, and Timmins (2015) and Wrenn, Klaiber, and Jaenicke (2016) identified more substantial impacts for households served by private well water rather than municipal water. Following their findings, we distinguish between geographic units (zip codes) based on the prevalence of municipal water service. To facilitate comparison with our structural model we employ a binary treatment measure equal to one if a zip code contains ten or more unconventional wells.[8] We include individual (time-varying) household characteristics, as well as fixed effects to capture time-varying and non-time-varying unobservables. Our reduced-form specification takes the form:

$$\begin{aligned} Y_{it} = \beta_0 + \beta_1 TenWells_{it} + \beta_2 TenWells_i t * WellWater_i + \\ \beta_3 WellWater_i + \beta_4 ShaleZip_i + \gamma X_{it} + \mu_t + s_t + \nu_i + \nu_i * \mu_t + \epsilon_{it}, \end{aligned} \tag{4}$$

where the i and t subscripts denote households and time, respectively. $TenWells$ is an indicator for whether household i is located in a zip code with at least ten unconventional wells spud as of quarter t; $WellWater_i$ is a binary variable that takes the value one if household i is located in a zip code in which the majority of households use well water; and $ShaleZip_i$ is a binary variable equal to one if household i is located in a zip code that eventually has at least ten unconventional wells spud (i.e., is eventually 'treated' by the entry of fracking activity). X_{it} represents time-varying and non-time-varying household characteristics, such as income, age, and the number of persons per household. We add fixed effects for year (μ_t), season of the year (α_t), county (ν_i), and county-year ($\nu_i * \mu_t$). Note that the t subscript represents different units of time in the same equation: quarter in Y_{it} and $TenWells_{it}$, season in s_t, and year in μ_t.[9]

4 Data

4.1 Nielsen data

Data on household purchases and consumer demographics were drawn from the Nielsen consumer HomeScan panel dataset, which is provided by the Kilts Center for Marketing Data at the University of Chicago Booth School of Business. This dataset records all food and beverage purchases for a panel of nearly 60,000 households across the US. We focus our study on the years 2005-2014 and in the states of Pennsylvania and Ohio, representing about 13,000 households, each of which is represented in the sample for about four years on average. This subset contains purchases that occurred before the entry of substantial fracking operations (2007 in Pennsylvania and 2011 in Ohio), and continues well into the maturation of fracking.

HomeScan data contains detailed information on all purchases made by panelists. We use all trips by panelists, including those trips where no purchase of bottled water occurred. For trips in which one or more bottled water products were purchased, we

[8] We tried an alternative specification in which a zip code is "treated" when it receives its first unconventional well, as well as specifications where treatment is a linear or log function of the number of wells (following Muehlenbachs, Spiller, and Timmins (2015), who represent treatment as a linear function of the number of wellbores or well pads), and found comparable results.

[9] As a robustness check, we also run a model with an additional interaction of fixed effects $WellWater_i * ShaleZip_i$. This model produces essentially identical results.

identify the chosen purchase by UPC. For trips where more than one bottled water product was purchased, we use only the largest water product (in total ounces). This is necessary as horizontal models of consumer demand require only a single, discrete purchase. Trips in which a panelist purchased two or more water products will appear in our data as only a single purchase, potentially biasing estimates of consumer demand. However, in all cases, the bias will be downward in coefficients relating to total ounces of bottled water, including the estimate of interest.

To control for brand effects, we categorize bottled water purchases into four categories: store brand, national brand (including Dasani, Nestle, Glaceau Smart Water, and Aquafina), luxury brands (Evian and Fiji), and flavored waters (including Propel, Glaceau Vitamin Water, Sobe Life Water, and other flavored brands). This allows us to parsimoniously control for national brands, and to develop store-level prices for store brands. The underlying assumption is that consumers may view national brands in one light (due to advertising campaigns or familiarity with a brand), but may view store brands as the "same", even in different stores. That is, a bottle of Safeway Select water has the same "brand" utility as a bottle of Kroeger brand water, but Dasani brand water differs. Our specification captures this relationship. We categorize all other non-national, non-store brands (e.g. Ozarka, Deer Park, Arrowhead) as "other".

Demographic information is drawn from HomeScan data on panelists. We use demographic household information on household size, race, head-of-household education, whether or not a household has kids under the age of 18, whether or not the household lives in a single-family home, the age of the oldest head-of-household, and household income. For household income, we take the median of the reported income "bin" to generate a continuous measure of income.

Estimating consumer preferences on observed choices requires knowledge of the consumer's choice set. We assume a consumer chooses a market independent from their demand for bottled water, and generate the consumer's choice set from the Nielsen Retail Scanner dataset. This dataset contains sales data for participating supermarket and similar retailers reported at the end of every week. Therefore, it contains all products offered which had non-zero sales for a given week, and is assumed to be an accurate representation of a *market* (store-week) choice set. These data are linked to HomeScan purchases by a unique store code and week-end. If a consumer reported a purchase from a store-week in which the good purchased was not present in the scanner data (possibly due to discrepancies in the reporting week), then the panelist's chosen good is added to the scanner data with a price derived from the panelist's reported purchase price. This good is also included in the choice set of all other panelists purchasing in that (store-week) market. All products are defined by total ounces, number of containers, whether the product is a single bottle (e.g. "jug" of water), and brand category - for instance, "96-12-F-Other" is an offering with a total of 96 ounces of water over 12 bottles in a multi-unit package of "other" brand (e.g. Ozarka, Deer Park, Arrowhead, etc.). In cases where multiple brands from the same category are offered, a market sales-weighted price is generated. In this case, an observed purchase of an "other" brand offering is assumed to be made at the sales-weighted price, regardless of consumer-reported price.

Many trips and bottled water purchases in the data are made at stores which do not participate in the scanner data collection program. When no bottled water is purchased, these trips do not have associated choice sets, and therefore provide no information on a consumer's choice of products. These trips are dropped from the data in the non-dynamic model. For trips with bottled water purchases that occurred at stores which do participate, but which did not report for a given week, the choice set for that observed purchase is simply the observed purchase plus the outside good. Under the assumption that non-participation in the scanner data for a given market (store-week) is not systematically correlated with consumer choices, the use of a limited choice set does not bias the results (Train, 2009).

The data yields complete demographics on 13,383 panelists over 10 years of participation for a total of 50,678 panel-years. A total of 1,982,548 trips and choice sets are observed.

4.2 Wells and municipal water boundaries

We obtain information on unconventional wells from state regulatory agencies in Pennsylvania (Department of Environmental Protection SPUD report) and Ohio (Department of Natural Resources). Both states provide information on unconventional wells including location (latitude and longitude) and spud date. Although recent concern among media and the public has focused on hydraulic fracturing, the drilling rig is generally the most visible element of onsite infrastructure (outside of the immediate vicinity of the well pad), so—like Wrenn, Klaiber, and Jaenicke (2016) and Muehlenbachs, Spiller, and Timmins (2015)—we use the spud date (the start of drilling operations) as the relevant date rather than the date of the fracturing operation. Fracking usually occurs within a few weeks after drilling commences, so the two operations would generally occur within the same quarter (and the same year) in any case. In some cases, we do not observe the actual spud date for wells in Ohio and we instead use the date the drilling permit was issued. Where we do observe both spud date and permit date, we find that drilling typically occurs within two to three months of receiving the drilling permit. We observe a total of about 11,000 unconventional wells spud by the end of 2014: 1,959 in Ohio and 8,815 in Pennsylvania.

We obtain municipal water boundaries from the Department of Environmental Protection (Public Water Supplier Service Areas) in Pennsylvania. We overlay the municipal water boundaries with zip codes and classify each zip code by the proportion of its area that overlaps public water system areas. Zip codes are classified as being on public water if at least half the area is located within a public water system service area, and as being served by well water otherwise. We were unable to obtain a shapefile of public water system boundaries in Ohio, so we classified zip codes as being served by well water if they are "shale zip codes" (i.e., if they eventually contain at least ten unconventional wells).

4.3 Summary statistics

Table 1 provides information on the number of zip codes by state and eventual treatment status. As noted previously, we lack data on public water system boundaries in Ohio, and thus assign wellwater status in that state based on whether a zip code eventually receives an unconventional well; thus, we assume there are no "shale" zip codes in Ohio that are on public water. Other than the absence in that cross-tabulation, there are reasonable numbers of households in each of the subcategories, although in 2005 a fairly small number of households in Pennsylvania shale zip codes on public water (most shale areas are served by private well water). In terms of average household expenditures on bottled water, there is no clear relationship between bottled water expenditure and either public water service or being a shale zip code. Over time from 2005-2014, average household expenditures on bottled water increase in all types of areas, with some unevenness to the trend, especially in shale zip codes in Ohio that are served by well water.

Table 1: Zip Codes, Households, and Bottled Water Expenditures

	Public Water				Well Water			
	Non-Shale Zips		Shale Zips		Non-Shale Zips		Shale Zips	
	OH	PA	OH	PA	OH	PA	OH	PA
Zip codes	728	381	0	27	54	507	30	79
HH panelists								
2005	1538	687	0	15	68	465	28	32
2006	2116	965	0	42	95	666	34	65
2007	2995	1415	0	79	144	1017	50	131
2008	3122	1513	0	82	154	1137	55	145
2009	3045	1541	0	81	151	1183	52	146
2010	3077	1598	0	82	145	1222	50	150
2011	2812	1474	0	82	129	1120	43	139
2012	2809	1497	0	84	138	1154	38	146
2013	2714	1467	0	75	130	1109	39	135
2014	2405	1307	0	70	119	977	33	120
Avg. BW exp per HH (all years):	0.603	0.652	n/a	0.651	0.488	0.616	0.744	0.600
Avg. BW exp per HH:								
2005	0.301	0.302	n/a	0.0793	0.0171	0.341	0.000	0.117
2006	0.555	0.396	n/a	0.548	0.765	0.568	0.747	0.292
2007	0.747	0.725	n/a	0.357	0.45	0.567	0.657	0.277
2008	0.647	0.739	n/a	0.855	0.583	0.735	1.049	0.624
2009	0.505	0.591	n/a	0.329	0.438	0.668	0.591	0.328
2010	0.532	0.585	n/a	0.629	0.204	0.509	0.605	0.919
2011	0.715	0.761	n/a	0.473	0.394	0.59	0.466	0.758
2012	0.706	0.755	n/a	0.806	0.696	0.742	1.355	0.673
2013	0.590	0.644	n/a	0.808	0.549	0.622	1.022	0.646
2014	0.581	0.767	n/a	1.178	0.663	0.639	0.779	0.830

Note: For this draft, we lack information on public water system boundaries for Ohio, and instead assume that a zip code uses well water if and only if shale gas development eventually arrives.

Table 2 provides a summary of demographic information for the household-years in the Nielsen HomeScan dataset. Because the observations in this table are at the household-year level, households are generally represented more than once, and the table includes demographic characteristics that may vary over time (e.g., income). Shale areas tend to have lower income, slightly larger households, and slightly lower levels of education than non-shale areas, on average. Shale areas also tend to have a higher proportion of married adults, as well as lower proportions of racial minority panelists.

Table 2: Summary Information for Household-Years in Sample

	Non-shale Zips		Shale Zips	
	OH	PA	OH	PA
Number of households	27,906	23,514	422	1,901
Income	55,135	57,587	43,868	52,672
Household size	2.30	2.40	2.52	2.50
Age	52.32	51.41	51.86	51.43
Some high school	0.03	0.03	0.04	0.02
High school graduate	0.31	0.35	0.50	0.39
Some college	0.29	0.25	0.19	0.22
College graduate	0.27	0.27	0.23	0.28
Post-Graduate education	0.10	0.10	0.04	0.09
Children	0.20	0.23	0.24	0.22
Married	0.62	0.64	0.75	0.71
White	0.89	0.89	0.97	0.97
Black	0.08	0.08	0.00	0.02
Asian	0.01	0.01	0.00	0.00
Other	0.02	0.02	0.02	0.01

Notes: Data reflect Nielsen household panel-years. Age and education variables reflect female head of household where applicable, and male head of household if there is no female head of household.

5 Results

5.1 Discrete choice structural model

Results from the heterogeneous parameters logit are primarily as expected. All parameters are significant in Model 1, owing largely to the sample size of 1,982,548 choice occasions. For all models, the coefficient on price is negative. The coefficient on total water is negative, but of very small magnitude. This is likely a result of the inclusion of *multi*, the quantity of bottles in the product. For a consumer, the total quantity of water may not be of as much importance as the number of bottles. For instance, a 16-pack of 20 ounce bottles may not be as desirable as a 20-pack of 15 ounce bottles, despite having more total water as the latter provides 20 perceived "servings" of water.

The per-ounce measure of each category follow reasonable patterns. The baseline (omitted) category is "store brand". Flavored water (e.g. Propel) is preferred over all other categories, followed by the "luxury" category (e.g. Evian), the "national

Table 3: Heterogeneous Parameters Logit

	Model 1
price (α)	-4.22405^{***}
	(0.00115)
total water	-0.00089^{***}
	(0.00002)
multipack qty	0.36383^{***}
	(0.00019)
Flavored brand	0.02122^{***}
	(0.00005)
Luxury brand	0.01394^{***}
	(0.00025)
National brand	0.00924^{***}
	(0.00001)
Other brand	0.00322^{***}
	(0.00001)
Hh size x total water	0.00034^{***}
	(0.00000)
College x total water	-0.00125^{***}
	(0.00001)
Black x total water	0.00083^{***}
	(0.00002)
Asian x total water	0.00026^{***}
	(0.00004)
Fracking x private wellwater x total water	0.00231^{***}
	(0.00003)
Log Lik	11924436.70255
AIC	-8.58820
N	1982548

brand" category (e.g. "Dasani"), and finally the "other" category containing regional and local brands. Household size positively effects utility when interacted with total water (larger households prefer larger quantities of water). Black households have greater preference for bottle water, as do Asian households, but to lesser degree.

The parameter of interest is the interaction of the presence of wells with total water. Here, "wells" is a binary measure of the presence of greater than 10 wells in a given zip code *for those households that do not have municipal water service*. The effect is small but significant—when fracking is "present" in a consumer's zip code, consumers have higher utility per ounce of bottled water. Because utility in a logit model is relative, this is equivalent to a disamenity for well-water users, per ounce of tap water consumed.

A simple compensating variation measure can be calculated from the results. To find the dollar equivalent of the utility lost per ounce of water consumed as a result of the appearance of fracking, we take the ratio $\frac{\beta^{w,o}}{\beta^{p}}$. For Model 1, the compensating variation measure is -\$0.00055. Because we do not observe in our data increases in the consumption of water-based products such as soda or (not-from-concentrate) water, we consider this measure to be biased downwards.

5.2 Reduced Form Model Results

Table 4 provides a summary of the reduced-form model results. Each observation is a household-quarter. Individual results for the effect of "treatment" (the entry of ten unconventional wells in a zip code) and well water status are presented. However, we focus here on the effect of treatment and well water (i.e., the sum of the two coefficients), which is provided at the bottom of the table. In all specifications, we observe a positive and significant treatment effect, indicating that the "treatment" of unconventional wells increased households' expenditures on bottled water. Our point estimate for the treatment effect within zip codes on private well water ranges from \$3.23 per household per quarter (in the simplest model) to \$4.12 per household per quarter (in the model with county fixed effects, but not county-by-year fixed effects).

These estimates are consistent with those of Wrenn, Klaiber, and Jaenicke (2016), who also use a reduced-form model with Nielsen Homescan data to estimate household averting expenditures on bottled water that arise from the entry of shale gas. That paper finds annual averting expenditures ranging from \$7.85 to \$18.36 per household, depending on the exact specification, arising from the entry of unconventional wells. In comparison, our per-quarter estimate equates to \$12.92 to \$16.48 per year per household. We used a subset of our data, modified to match the sample used by Wrenn, Klaiber, and Jaenicke (2016), and were able to replicate their results approximately. However, we found some differences in both our sample (i.e., different summary statistics) and our regression results (when we replicated their methods, to the best of our ability, with our data set). In addition, we were able to obtain a data set with finer spatial resolution—zip code rather than county—which allows us to add a richer set of fixed effects.

Our reduced-form results confirm the findings of Wrenn, Klaiber, and Jaenicke (2016), and we do not consider our reduced-form results to be a contradiction to the earlier paper nor a substantial contribution to the literature. Rather, our intent in this paper is to offer a new perspective on the use of averting expenditures to measure welfare effects, and to compare the results of structural demand and reduced-form models in this context.

5.3 Discussion

In an averting expenditures framework, the reduced form model seeks to calculate the lower bound on the perceived disamenity associated with tap water consumption in the presence of fracking. The preferred model (column (4) of Table 4) estimates a per-household disamenity value of \$3.43 per quarter. The core critique in Section 2.2 notes that the lower bound argument does not hold when the good enters the consumer's utility function directly. The structural model in Section 5.1 yields a dollar-denominated increase in utility from consuming bottled water when fracking is present above and beyond the other desirable traits of bottled water—the characteristics of bottled water that would enter the consumer's utility directly - an amount equal to \$0.00055 per ounce.

For a moment, suppose that the marginal utility of bottled water consumption is constant at \$0.00055 per ounce. To make a meaningful comparison with the reduced-form estimate from Table 4, we wish to multiply this per-ounce utility by some quantity of water consumed. One approach would be to multiply the per-ounce indirect utility by total household water consumption (e.g., 32 ounces per person per day) rather than observed bottled water consumption. Assuming a per-person consumption of 32 ounces of water per day and a per-ounce value of \$0.00055, the compensating variation or disamenity-of-fracking value would be \$4.14 per household per quarter. Obviously, the calculated value is sensitive to the per-person consumption estimate; if we assume instead a per-person consumption of 64 ounces per day, we would calculate a compensating variation of \$8.18 per household per quarter.

We could also calculate a similar measure using the same measure of compensating variation but apply it only to observed purchases by households located in "fracked" zip codes. For every purchase made by a household in a "fracked" area, we observe in our data their purchases, including the total water. The product of total water purchased and our per-ounce compensating variation (\$0.00055) represents the component of observed expenditures that can be attributed to the presence of fracking. The measure is taken using only purchases made by households in a fracked zip code who rely on well-water. The per-year per-household average is \$0.96. The per-quarter average is \$0.39. The averages differ only over the observed purchases that are used to calculate them, with the preferred estimate being the per-quarter average. This figure represents the component of per-household utility derived from bottled water's use as a substitute away from a household's tap water.

Still need to incorporate the following points:

- **fully characterize nature of lower bound:** Assuming away joint production, the averting expenditure on bwater is a lower bound on WTP for environmental quality for three reasons:
  - A. We don't observe averting expenditures on other technologies, eg house filters.
  - B. It doesn't include damages, eg getting sick from drinking contaminated water.

Table 4: Reduced Form Results

	(1)	(2)	(3)	(4)
TenWells	1.144	1.178	0.881	0.954
	(1.500)	(1.390)	(1.269)	(1.406)
TenWells*WellWater	2.087	2.143	3.241*	2.480
	(1.915)	(1.802)	(1.678)	(1.913)
WellWater	-0.000725	-0.114	-0.0311	-0.0276
	(0.296)	(0.294)	(0.455)	(0.459)
ShaleZip	-0.697	-0.726	-0.446	-0.280
	(0.646)	(0.631)	(0.837)	(0.885)
Spring		1.623***	1.619***	1.609***
		(0.0879)	(0.0871)	(0.0877)
Summer		2.195***	2.196***	2.195***
		(0.0834)	(0.0829)	(0.0830)
Autumn		0.761***	0.765***	0.761***
		(0.0885)	(0.0886)	(0.0871)
Income		0.0321***	0.0304***	0.0306***
		(0.00402)	(0.00412)	(0.00416)
HH size 2		3.040***	2.929***	2.969***
		(0.379)	(0.378)	(0.382)
HH size 3		4.900***	4.764***	4.789***
		(0.555)	(0.551)	(0.560)
HH size 4		4.950***	4.903***	4.963***
		(0.593)	(0.593)	(0.599)
HH size 5		5.021***	4.912***	4.916***
		(0.736)	(0.733)	(0.739)
HH size 6		5.146***	5.002***	5.140***
		(0.886)	(0.891)	(0.904)
Age		-0.0559***	-0.0601***	-0.0610***
		(0.0135)	(0.0134)	(0.0135)
HS Grad		-1.194	-1.285	-1.348
		(0.952)	(0.965)	(0.975)
Some College		-1.427	-1.528	-1.639*
		(0.962)	(0.967)	(0.979)
College Grad		-1.929**	-2.039**	-2.119**
		(0.947)	(0.957)	(0.970)
PostGrad		-3.161***	-3.136***	-3.203***
		(0.992)	(1.000)	(1.010)
Children		-0.481	-0.560	-0.555
		(0.450)	(0.455)	(0.461)
Married		-0.767**	-0.628	-0.674*
		(0.382)	(0.384)	(0.386)
Black		1.969***	1.764***	1.802***
		(0.431)	(0.448)	(0.451)
Asian		-2.399***	-2.703***	-2.644***
		(0.844)	(0.858)	(0.867)
Other		0.283	0.0329	0.0544
		(0.637)	(0.648)	(0.654)
Constant	5.625***	5.263***	3.774**	5.571***
	(0.249)	(1.351)	(1.499)	(1.394)
Year FE	Yes	Yes	Yes	Yes
County FE			Yes	Yes
County-Year FE				Yes
Observations	207,118	207,118	207,118	207,118
R-squared	0.002	0.025	0.035	0.048
TenWells +	3.231***	3.321***	4.122***	3.434**
WellWater*TenWells	(1.109)	(1.074)	(1.042)	(1.374)

Robust standard errors, clustered at zip code, in parentheses
*** p<0.01, ** p<0.05, * p<0.1

  - C. We don't observe true WTP, only the expenditure, which (without joint production) is a lower bound on WTP. (But with joint production, the pure averting expenditure is no longer a lower bound.)
  - Let's assume away A and B (no other averting expenditures, and no damages). Now, if we can cleanly identify the effect of fracking—distinct from the effect of portability or taste—when we estimate a demand system with a structural model, we **do** observe true WTP. Furthermore, this is true even if there is joint production. (In that sense, a structural demand approach dominates a reduced-form approach, regardless of whether there is joint production. The structural model provides a superior estimate of WTP—it gets past both the joint production issue, and the issue that arises from expenditure being a lower bound for WTP

- in this context, explain the meaning if the lower bound essentially equal to the structural estimate of the true value.

6 Conclusions

[Reframe conclusion to match the foregoing discussion]

Hydraulic fracturing brings with it a variety of economic impacts that may confound economic assessments of its impacts. In particular, while fracking activity stresses public infrastructure, brings an influx of potentially temporary workers, and draws significant amounts of water which is then returned as potentially contaminated process water, it also brings with lease and royalty payments which may enter the local economy. Hedonic studies have been used to assess the "total basket" of amenities, considering "presence of fracking" to be the change in amenities, and examining the change in home sale values. Because the purchase of a home includes the basket of local amenities (parks, schools, roads, water quality, etc.), changes in home values associated with the presence of fracking will reflect the overall economic impact of fracking (Muehlenbachs, Spiller, and Timmins, 2015). This paper examines a component of that basket—the quality of drinking water—in more detail.

The per-household, per-quarter disamenity estimated in our structural model yields a value of $4.14, based on an assumption of 32 ounces of water consumed per day, an amount greater than the reduced form lower-bound estimate of $3.24. The results from the comparison of these models does not definitively validate the critique owing to Bartik (1988), but is not necessarily inconsistent with this critique.

The "total basket" of benefits and costs (including environmental costs) owing to fracking are important to understand for policymakers and homeowners in shale gas regions. Our results provide additional information on the environmental disamenity associated with fracking, and are consistent with hedonic estimates. Further consideration of dynamics in household bottled water purchases and a richer specification for household demand for bottled water will be necessary to refine the estimates, and to further elucidate the relationship between the structural and reduced form "averting expenditures" measures.

References

Bartik, Timothy J. (1988). "Evaluating the benefits of non-marginal reductions in pollution using information on defensive expenditures". In: *Journal of Environmental Economics and Management* 15.1, pp. 111–127. DOI: 10.1016/0095-0696(88)90032-0. URL: http://dx.doi.org/10.1016/0095-0696(88)90032-0.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995). "Automobile prices in market equilibrium". In: *Econometrica: Journal of the Econometric Society*, pp. 841–890.

Blauch, Mark (2010). *Developing Effective and Environmentally Suitable Fracturing Fluids Using Hydraulic Fracturing Flowback Waters*. Tech. rep. SPE 131784. Society of Petroleum Engineers.

Courant, Paul and Richard C. Porter (1981). "Averting expenditure and the cost of pollution". In: *Journal of Environmental Economics and Management* 8.4, pp. 321–329. URL: http://EconPapers.repec.org/RePEc:eee:jeeman:v:8:y:1981:i:4:p:321-329.

Deschênes, Olivier and Michael Greenstone (2011). "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US". In: *American Economic Journal: Applied Economics* 3.4, pp. 152–185. DOI: 10.1257/app.3.4.152. URL: http://dx.doi.org/10.1257/app.3.4.152.

Dickie, Mark (2003). "Defensive Behavior and Damage Cost Methods". In: *A Primer on Nonmarket Valuation*. Ed. by Patricia A. Champ, Kevin J. Boyle, and Thomas C. Brown. Norwell, MA: Kluwer Academic Publishers. Chap. 11, pp. 395–444.

Fetter, T. Robert (2016). "Fracking, Toxics, and Disclosure". Unpublished manuscript. URL: https://sites.duke.edu/trfetter/research.

Gopalakrishnan, Sathya and H. Allen Klaiber (2013). "Is the Shale Energy Boom a Bust for Nearby Residents? Evidence from Housing Values in Pennsylvania". In: *American Journal of Agricultural Economics* 96.1, pp. 43–66. DOI: 10.1093/ajae/aat065. URL: http://dx.doi.org/10.1093/ajae/aat065.

Graff Zivin, Joshua, Matthew Neidell, and Wolfram Schlenker (2011). "Water Quality Violations and Avoidance Behavior: Evidence from Bottled Water Consumption". In: *American Economic Review* 101.3, pp. 448–453. DOI: 10.1257/aer.101.3.448. URL: http://dx.doi.org/10.1257/aer.101.3.448.

Hausman, Catherine and Ryan Kellogg (2015). "Welfare and Distributional Implications of Shale Gas". In: DOI: 10.3386/w21115. URL: http://dx.doi.org/10.3386/w21115.

Hendel, Igal and Aviv Nevo (2002). *Sales and consumer inventory*. Tech. rep. National Bureau of Economic Research.

— (2006). "Measuring the implications of sales and consumer inventory behavior". In: *Econometrica* 74.6, pp. 1637–1673.

Kuminoff, Nicolai V. and Jaren C. Pope (2014). "DO "CAPITALIZATION EFFECTS FOR PUBLIC GOODS REVEAL THE PUBLIC'S WILLINGNESS TO PAY?" In: *International Economic Review* 55.4, pp. 1227–1250. DOI: 10.1111/iere.12088. URL: http://dx.doi.org/10.1111/iere.12088.

Lancaster, Kelvin J (1966). "A new approach to consumer theory". In: *The journal of political economy*, pp. 132–157.

Mason, Charles, Lucija Muehlenbachs, and Sheila Olmstead (2015). *The Economics of Shale Gas Development*. Tech. rep. Resources for the Future. URL: http://www.rff.org/files/sharepoint/WorkImages/Download/RFF-DP-14-42.pdf.

McFadden, Daniel et al. (1973). "Conditional logit analysis of qualitative choice behavior". In:

Mu, Quan and Junjie Zhang (2015). "Air Pollution and Defensive Expenditures: Evidence from Particulate-Filtering Facemasks". In: *SSRN Electronic Journal*. DOI: 10.2139/ssrn.2518032. URL: http://dx.doi.org/10.2139/ssrn.2518032.

Muehlenbachs, Lucija, Elisheba Spiller, Andrew Steck, et al. (2015). "The Impact of the Fracking Boom on Rents in Pennsylvania". URL: http://public.econ.duke.edu/~timmins/fracking_rents.pdf.

Muehlenbachs, Lucija, Elisheba Spiller, and Christopher Timmins (2015). "The Housing Market Impacts of Shale Gas Development". In: *American Economic Review* 105.12, pp. 3633–3659. DOI: 10.1257/aer.20140079. URL: http://dx.doi.org/10.1257/aer.20140079.

Nevo, Aviv (2000). "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand". In: *Journal of Economics & Management Strategy* 9.4, pp. 513–548.

Olmstead, S. M. et al. (2013). "Shale gas development impacts on surface water quality in Pennsylvania". In: *Proceedings of the National Academy of Sciences* 110.13, pp. 4962–4967. DOI: 10.1073/pnas.1213871110. URL: http://dx.doi.org/10.1073/pnas.1213871110.

Reynolds, Kelly A., Kristina D. Mena, and Charles P. Gerba (2008). "Risk of Waterborne Illness Via Drinking Water in the United States". In: *Reviews of Environmental Contamination and Toxicology*. Springer Nature, pp. 117–158. DOI: 10.1007/978-0-387-71724-1_4. URL: http://dx.doi.org/10.1007/978-0-387-71724-1_4.

Stringfellow, William T. et al. (2014). "Physical, chemical, and biological characteristics of compounds used in hydraulic fracturing". In: *Journal of Hazardous Materials* 275, pp. 37–54. DOI: 10.1016/j.jhazmat.2014.04.040. URL: http://dx.doi.org/10.1016/j.jhazmat.2014.04.040.

Timmins, Christopher and Jennifer Murdock (2007). "A revealed preference approach to the measurement of congestion in travel cost models". In: *Journal of Environmental Economics and management* 53.2, pp. 230–249.

Timmins, Christopher and Wolfram Schlenker (2009). "Reduced-form versus structural modeling in environmental and resource economics". In: *Annu. Rev. Resour. Econ.* 1.1, pp. 351–380.

Train, Kenneth E (2009). *Discrete choice methods with simulation*. Cambridge University Press.

Walker, Joan and Moshe Ben-Akiva (2002). "Generalized random utility model". In: *Mathematical Social Sciences* 43.3, pp. 303–343.

Wrenn, Douglas H., H. Allen Klaiber, and Edward C. Jaenicke (2016). "Unconventional Shale Gas Development, Risk Perceptions, and Averting Behavior: Evidence from Bottled Water Purchases". In: *Journal of the Association of Environmental and Resource Economists* 3.4, pp. 779–817. DOI: 10.1086/688487. URL: http://dx.doi.org/10.1086/688487.