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

A. Justin Kirkpatrick and T. Robert Fetter

November 21, 2023

Abstract

Where markets are incomplete, one common approach for measuring willingness to pay (WTP) is to use agents' defensive or averting expenditures. Researchers typically interpret such defensive expenditures or "coping costs" as the lower-bound WTP for the good of interest, such as higher-quality public goods. We clarify that averting expenditures overestimate the lower-bound WTP for improved public goods if actors get additional utility from the defensive good. We demonstrate the effect empirically in the context of bottled water purchases in response to hydraulic fracturing, and find the reduced-form estimate exceeds the true value by a factor of 200 to 400 percent.

KEYWORDS: Averting Expenditures, Coping Costs, Structural vs Reduced-Form Models, Hydraulic Fracturing

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

Kirkpatrick (`jkirk@msu.edu`) - Michigan State University Department of Economics; Fetter (`rob.fetter@duke.edu`) - Nicholas Institute for Environmental Policy Solutions, Duke University.
We are grateful for helpful comments from Lori Bennear, Bryan Bollinger, Joe Herriges, David Keiser, Chris Timmins, and Doug Wrenn, as well as seminar participants at Duke University, Michigan State University, and at the conferences of the American Economics Association and the Association of Environmental and Resource Economists. 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 naturally occurring 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 decreased 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 (WTP) includes other characteristics of the substitute good, besides perceived higher environmental 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. Yet this adjustment is not common in the literature; papers that use defensive expenditures to measure WTP for public goods mention the issue infrequently, and even more rarely account for it empirically.

We begin by evaluating the change in pure expenditure due to the entry of fracking in a reduced-form context. Results from our preferred models suggest that consumers respond to the entry of fracking by increasing purchases of bottled water, with an average expenditure of $56.89 per household per quarter, but only a fraction of this is attributable to the specific attribute of avoiding exposure to tap water potentially contaminated (or perceived to be contaminated) by fracking.[1]

We then 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). This allows us to compare the WTP implied by a pure-expenditure reduced form model to the WTP implied by the demand model.

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 that is linked to environmental quality. As information on consumer purchases becomes more readily available, it becomes more feasible for researchers to assess WTP for non-market goods, like environmental quality, by measuring changes in consumption of market substitutes. We build on efforts of previous researchers who have used supermarket scanner data in reduced-form models to estimate WTP for environmental quality, particularly those who have used bottled water expenditures to measure WTP for water quality (Graff Zivin, Neidell, and Schlenker, 2011; Wrenn, Klaiber, and Jaenicke, 2016). These researchers suggest the resulting estimates represent a lower-bound WTP, but do not address the degree to which the so-called "lower bound" may 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 fracking 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 $16.99 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.

[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.

Finally, we contribute to the literature on demand estimation by addressing a specific type of household-level heterogeneity with a novel specification that allows us to capture household-level taste for tap water. Consideration of this type of heterogeneity is important for consumer demand models where households may have varied tastes for the "outside good," the alternative to not purchasing any of the products in the model. The situation is particularly relevant to the health insurance purchase decision, where those opting to not purchase health insurance may do so due to persistent (but unobserved) tastes. We use simulation to show that failing to account for this heterogeneity leads to biased and potentially inconsistent estimates.

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, 2019). 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 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

[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

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 shale gas production decreases the likelihood of methane infiltration into groundwater by relieving pressure exerted by gas formations, 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 WTP 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 WTP 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 WTP.

Several recent empirical analyses have used logic along these lines to impute preferences for environmental quality, including WTP for improved air quality based on expenditures for face masks (Zhang and Mu, 2018), 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 some consumers prefer bottled to tap water for characteristics other than possibly higher quality, for instance, but the same may not be true for 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 WTP for improved environmental quality, 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,

[5]In addition, we hope that policymakers are interested in not just 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).

air conditioning produces cooler and drier air, thus averting some of the adverse consequences of climate change, but can also reduce some indoor air pollutants, and this latter benefit also enters the utility function directly. Therefore, in cases of joint production analysts ought to account for how households value "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. Suppose the consumer is willing to pay \$.99 for the bottle of water based on the taste, portability, and brand of the good, and 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, concerned about the safety of their own water supply, is willing to pay \$.02 to avoid consuming their tap water. Having a new WTP 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' WTP 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.

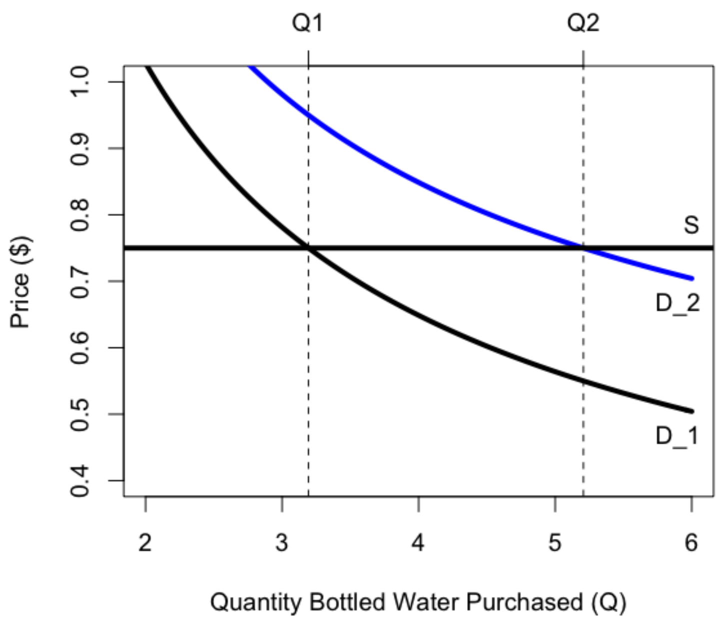

Figure 1: Simple Example of Structural Versus Reduced Form

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 a world with and without fracking.

3 Models

3.1 Reduced form model of averting expenditures

We begin with a reduced-form model to estimate households' change in bottled water consumption, as well as averting expenditures, in response to the arrival of fracking. We use a difference-in-differences approach to isolate the effect of fracking from other demand shifters that vary over space and time.

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 approach, we distinguish households' exposure to fracking as well as our measure of well-dependent exposure to fracking. As discussed in detail in Section 4 make ref, we cannot directly discern a household's exposure to fracking nor its well dependence, but we can calculate the unconditional probability that a household is exposed to fracking, is well-dependent, or (importantly) both. We include a rich set of fixed effects to capture time-varying and non-time-varying unobservables. Our reduced-form specifications are of the general form:

$$y_{it} = \beta_0 + \beta_1 Fracking_{it} + \iota_i + \omega_t + \gamma_{zt} + \epsilon_{it} \tag{1}$$

The dependent variable y_{it} is household expenditures on bottled water by week or total ounces of water purchased by week, and the i and t subscripts denote households and time, respectively. $Fracking_{it}$ is the unconditional probability of either exposure to fracking or a probability of well-dependent exposure to fracking, both of which vary by zip code. ι_i is a household fixed effect, ω_t is a weekly fixed effect, and γ_{zt} the interaction of zip code status (ever-fracked, adjacent to ever-fracked, and never-fracked + non-adjacent) and week, which controls flexibly for differential time-varying unobserved trends common to zip codes based on their fracking type. Our fracking exposure measures vary at the zipcode-week level, precluding zipcode-specific weekly fixed effects.

[6]

3.2 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

[6] In the specifications without this household fixed effect, we include a measure of each household's "taste" for tap water, using the household-level parameter that is output by the random-parameters model. do we still? Not at the moment. Maybe post-re-estimation of structural model

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 is that the specification is grounded in economic theory and is consistent with utility maximization, but also allows for unobserved components. The use of random terms rationalizes different utility-maximizing choices even when made by consumers with identical characteristics. The generalized extreme value (GEV) RUM (Walker and Ben-Akiva, 2002) incorporates complex correlations in the random component over products and consumer characteristics. For an in-depth description of discrete choice RUM models, see K. E. Train (2009).

The indirect utility for consumer i gained from consuming bottled water product j in market t is defined as $U(x_{jt}, p_{jt}, \tau_{it}, 1^F_{it}; \theta)$, a function of observed product characteristics ($\mathbf{x}_{jt}$, exclusive of x^{oz}_j, the size of bottled water j), price (p_{jt}), individual consumer characteristics (τ_{it}), consumer exposure to fracking, (1^F_{it} taking the value of 1 when exposed to fracking and zero otherwise), and unknown parameters to be estimated (θ). Throughout this analysis, a market is defined at the store-week level and notated as t.

Water consumption has two unique properties that set it apart from traditional consumer goods. First, humans have a biological requirement to consume water. Whatever water not consumed through purchases of bottled water (or other beverages) is largely obtained from the household's tap. Second, humans have an upper limit on the quantity of water that can rationally be consumed. Tap water is essentially zero cost on the margin, but near-infinite quantities of water are rarely consumed. These properties help to simplify the specification of utility for bottled water purchases. Let W_i be household i's total water budget - the biologically-required quantity of water that must be consumed, which is allowed to vary by household. Because a household may consume this water either from their tap/well, from bottled water, or some combination of each, we allow the consumer to obtain different utility from meeting W_i by consuming tap and bottled water. This indirect utility is specified as:

$$u_{ijt} = \alpha_i(y_i - p_{jt}) + \beta_i^{Tap}(W_i - x_j^{oz}) + \beta_i^{Bottled} x_j^{oz} + \beta_i^C \mathbf{x}_j + \epsilon_{ijt} \tag{2}$$

Where y_i is income, 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. If we allow the household's utility per ounce of tap water consumed to change in the presence of fracking (1^F_{it}), equation 2 becomes:

$$\begin{aligned} u_{ijt} &= \alpha_i(y_i - p_{jt}) + \left(\beta_i^{Tap} + \beta_i^{TF} 1^F_{it}\right)(W_i - x_j^{oz}) + \beta_i^{Bottled} x_j^{oz} + \beta_i^C \mathbf{x}_j + \epsilon_{ijt} \\ &= \alpha_i(y_i - p_{jt}) + \left(\beta_i^{Tap} + \beta_i^{TF} 1^F_{it}\right) W_i + \left(\beta_i^{Bottled} - \beta_i^{Tap} - \beta^{TF} 1^F_{it}\right) x_j^{oz} + \beta_i^C \mathbf{x}_j + \epsilon_{ijt} \end{aligned} \tag{3}$$

We set the outside good of "no purchase" to the utility of consuming entirely from the tap:

$$u_{i0t} = \alpha_i y_i + \left(\beta_i^{Tap} + \beta_i^{TF} 1^F_{it}\right) W_i + \epsilon_{i0t}$$

The choice of product j depends only on relative utilities across products $j \in J$. Income and W_i enter the utility of each choice identically and, because only relative utilities matter, can be normalized to the utility of the outside good $j = 0$:

$$\tilde{u}_{ijt} = -\alpha_i p_{it} + \left(\beta_i^{Bottled} - \beta_i^{Tap} - \beta^{TF} 1_{it}^F\right) x_j^{oz} + \beta_i^C \mathbf{x}_j + \tilde{\epsilon}_{ijt}$$

The unobserved household water budget, W_i, cancels out in $\tilde{u}_{ijt}$ and only the net difference in per-ounce preference between *tap* and *bottled* remains. Setting $\beta_i^{BT} = \beta_i^{Bottled} - \beta_i^{Tap}$:

$$\tilde{u}_{ijt} = -\alpha_i p_{it} + \beta_i^{BT} x_j^{oz} - \beta^{TF} 1_{it}^F x_j^{oz} + \beta_i^C \mathbf{x}_j + \tilde{\epsilon}_{ijt} \quad (4)$$

Justin: reinforce how this lets quantity of water purchased drive identification of utility

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:

$$Pr(d_{ijt} = 1) = Pr(u_{ijt} > u_{ikt} \quad \forall k \neq j) = \frac{exp(-\alpha_i p_{it} + \beta_i^{BT} x_j^{oz} - \beta^{TF} 1_{it}^F x_j^{oz} + \beta_i^C \mathbf{x}_j)}{1 + \sum_{k=1}^{J} exp(-\alpha_i p_{it} + \beta_i^{BT} x_j^{oz} - \beta^{TF} 1_{it}^F x_j^{oz} + \beta_i^C \mathbf{x}_j)} \quad (5)$$

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), 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 demographic 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 further extend heterogeneity into unobserved consumer characteristics by allowing (α_i, β_i) to include a stochastic component, v_i. The stochastic component allows for variation in taste beyond that predicted by observed demographics in D. Furthermore, this component also allows for correlations in the random, unobserved taste components. Letting Σ indicate the $[(k+1) \times (k+1)]$ variance-covariance matrix of consumer tastes, we write tastes (α_i, β_i) as a random-parameters mixed logit (K. E. Train, 2009; Nevo, 2000):

$$\begin{pmatrix} \alpha_i \\ \beta_i \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} + \Pi D_i + \Sigma v_i \quad (6)$$

Where v_i is *iid* and follows some distribution, and D_i is observed in our data.

The advantage of the random parameters specification is that it allows unobserved consumer tastes to be correlated. That is, a consumer who places greater preference for multi-unit bottled water relative to other consumers with identical demographic characteristics may also have greater preference for larger total quantities of water (again, beyond the preference common within the consumer's demographic cohort). In this example, Σ has at least one non-zero off-diagonal value which represents the correlation between these two tastes. Diagonal values in Σ represent the variance of each taste parameter, and off-diagonals represent the covariance.

It can be shown that this specification can approximate underlying correlations in the T1EV error terms, ϵ_{it}, over a choice set J_t (K. E. Train, 2009). Accounting for these correlations eliminates the assumption of independence of irrelevant alternatives that is implicit in a simple logit model. That is, by allowing for correlations in ϵ_{it}, the model no longer imposes the unrealistic cross-product substitution implied by the simple logit. Instead, if a product is eliminated from the choice set, a consumer is predicted to substitute into other products *with similar product characteristics.*

These substitution patterns are important in our context. Total ounces of water is a key product characteristic. A consumer whose product choice process is nested may first decide to purchase some quantity of water e.g. 128 ounces (1 gallon). With this in mind, the consumer finds the water aisle at the supermarket, finds the 128 ounce containers, and makes a choice within the 128 ounce offerings. Absent a preferred brand, the consumer will most likely remain within the 128 ounce nest, but will substitute with a similar brand. A nested logit specification can account for this choice, but at the expense of an interpretable coefficient on x_j^{oz}, since this characteristic defines the nest. A random parameters mixed logit with flexible correlations can produce identical cross-substitution patterns and provides coefficients that can be used to derive compensating variation.

The presence of fracking is hypothesized to drive increased purchases of bottled water by decreasing the utility of the outside good. 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, 1_{it}^F, any choice of bottled water effectively has the attribute of "avoiding consuming potentially contaminated tap water".[7] With the introduction of this attribute, the utility predicted by the product characteristics x_{jt}^{oz} varies over individuals *even within the same market and with identical observable characteristics* τ_{it}, as the presence of fracking is measured in the consumer's zip code, but the market serves consumers from multiple zip codes. We assume the presence or absence of fracking around the store does not enter the consumer's utility.

The coefficient of interest is β_i^{TF}, which represents the change in taste for the household's tap water when it is potentially affected by local fracking. If the arrival of fracking reduces the household's utility from consuming tap water, then $\beta^{TF} < 0$. The ratio of this coefficient to α_i yields the change in the marginal value of consuming an ounce of tap water. This measure establishes a water-specific compensating variation for fracking.

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

3.2.1 Household-level Heterogeneity for Tap Water

Each household can be assumed to have different, unobserved quality of tap water, either due to unobserved variation in the quality of public supply, or due to variation

[7]Note that for the purposes of the random parameters logit, we define "fracking" as the unconditional probability of a household in the zip code being (i) well-water dependent and (ii) within 500m of a fracked well.

in well and filter quality for households on well water. If this unobserved taste is truly random, it will be captured in the taste for larger sizes of bottled water, β_i^{BT}, with greater preference for larger quantities of bottled water when unobserved household tap water is disliked. However, in a classic omitted variables problem, if a household's unobserved taste for its tap water is correlated with the presence of fracking, the coefficient on the interaction of fracking and size will be biased. This could be of concern for a variety of reasons. If households located in areas that were formerly coal-producing regions are more likely to have been exposed to water quality issues over the last 30 years, they may have a negative preference on their tap water. If fracking is more likely to occur in areas with greater fossil fuel resources (e.g., due to existing infrastructure, legal frameworks, experienced workforces, or natural resource endowments), this would induce a correlation between the treatment (the arrival of fracking) and negative tastes for tap water, biasing the parameter of interest.

We address this with a household fixed effect for the "outside good." The concept mirrors the canonical product-specific constants (or "alternative-specific constants") which account for unobserved product-level utility common across all households (S. T. Berry, 1994). Rather than controlling for common unobserved characteristics, the aim is to control for household-specific characteristics in taste for the outside good. Our study represents a novel application of fixed effects at the level of individual household, without adding restrictive assumptions on these fixed effects. Goolsbee and Petrin (2004) allow alternative-specific constants to vary over regional markets, while Petrin and K. Train (2010) use a similar alternative-specific model but allow the variance of the shocks for the "inside" goods to vary by household characteristics. Lutzeyer, Phaneuf, and Taylor (n.d.) employ a product-specific constant that summarizes a group of products, and allow it to vary over latent class membership of each household. Dube et al. (2002) discuss a household-product specific constant, but only in passing, and do not estimate a model with this flexibility. The method we employ here also has useful applications in other problems. For instance, heterogeneity of preferences for the outside good would be integral to fields such as health insurance policy choice, where each customer has potentially very different preferences and tastes for being uninsured.

Rather than parameterize the product-specific constants, we instead include individual-level product-specific constants. Rewriting (4) and dropping the constant utility component related to income:

$$\tilde{u}_{ijt} = \delta_{ij} - \alpha_i p_{ijt} + \beta_i^{BT} x_j^{oz} - \beta^{FT} 1_{it}^F x_j^{oz} + \beta_i^C \mathbf{x}_j + \tilde{\epsilon}_{ijt} \tag{7}$$

Estimating (7) would require $[(J-1) \times H]$ parameters in addition to α and β. This would be computationally intensive. Noting that our goal is to account for heterogeneity in household preference for the "outside" good, tap water, and noting that only differences in δ_j matter, we normalize to δ_{i0}.

$$\tilde{u}_{ijt} = (\delta_{ij} - \delta_{i0}) - \alpha_i p_{ijt} + \beta_i^{BT} x_j^{oz} - \beta^{FT} 1_{it}^F x_j^{oz} + \beta_i^C \mathbf{x}_j + \tilde{\epsilon}_{ijt} \tag{8}$$

Product characteristics of the inside goods are well-defined by the characteristic space. Thus, we set $\delta_{ij} = 0 \quad \forall j \in J \neq 0$, allow the β_i's to capture preferences for characteristics of the "inside" goods, and note that setting $\delta_{ij} = 0 \quad \forall j \neq 0$ is identical to a model with a constant value of δ_{i0} relative to all δ_{ij}. It is straightforward to see that choice probabilities that are identified only up to scale (such as in the logit form) are not constant over values of δ_{i0}. As δ_{i0} increases, the value of the inside goods decrease relative to the "outside" good, and the choice probabilities decrease for the inside goods proportionally.[8]

Assuming a T1EV distribution for ϵ_{ijt} yields the familiar choice probabilities, indexed by household i.

$$Pr(d_{ijt}=1) = \frac{exp\left((\delta_{ij}-\delta_{i0}) - \alpha_i p_{ijt} + \beta_i^{BT} x_j^{oz} - \beta^{FT} 1_{it}^F x_j^{oz} + \beta_i^C \mathbf{x}_j\right)}{\sum_{k=0}^{J} exp\left((\delta_{ik}-\delta_{i0}) - \alpha_i p_{ikt} + \beta_i^{BT} x_j^{oz} - \beta^{FT} 1_{it}^F x_k^{oz} + \beta_i^C \mathbf{x}_k\right)} \quad (9)$$

We propose a practical means of using estimated results to confirm that household-level heterogeneity in tastes for tap water is captured - since our expectation is that the term δ_{i0} represents unobserved tastes over unobserved tap water quality, we would expect to see values correlated in areas with the same water supplier (for households on municipal water) or similar geology (for households using well water). We observe each household's zip code and will estimate δ_{i0}. Therefore, we can map spatial relationships between δ_{i0} and examine their coincidence with municipal water supply boundaries or geological maps. A statistical test of spatial clustering also provides insight into the characteristics captured by the values of δ_{i0}.

3.2.2 Identification of Household Heterogeneity

We rely on repeated observations of household purchases over time to identify household-specific, persistent heterogeneity in taste for tap water by leveraging the share of non-purchases by household i over the observed time horizon. Identification in this context is intuitive - if we observe any two observably identical households where one has a low observed share of non-purchases relative to the other, it must be the case that the household with a high share of non-purchases has higher preference for their tap water relative to bottled water. Identification follows from this notion.

3.2.3 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 some "other" brand. We specify observable consumer heterogeneity only with the interaction of the presence of fracking (i.e., $\geq$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 presence of fracking with total ounces of water. The coefficient estimated, $\beta^{w,o}$, is the marginal (dis)utility of the presence of fracking, per ounce of bottled water. To calculate the dollar equivalent of the change in utility (compensating variation) associated with a change in fracking from "not present" to "present", we

[8] The utilities given by the first two parentheses which differ only in the difference between the 2^{nd} and 3^{rd} entries relative to the 1^{st}, do not result in the same choice probabilities. The second two parentheses, however, do result in the same choice probabilities. This emphasizes the role of utility of the outside good (1^{st}) relative to the utility of the inside goods.

$$\begin{pmatrix}0\\5\\7\end{pmatrix} \neq \begin{pmatrix}0\\6\\8\end{pmatrix} = \begin{pmatrix}-1\\5\\7\end{pmatrix}$$

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.

3.2.4 Estimation

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

Estimation of the individual alternative-specific constants requires a two-step process common in the consumer demand literature - given any set of parameter estimates, $\hat{\theta}$, (S. T. Berry, 1994) calculates a vector of product-specific constants, δ that generate the observed aggregate market shares. The process iterates between finding the parameters θ^r conditional on δ^r using Maximum Likelihood Estimation, and updating δ^{r+1} using a contraction mapping algorithm. It is shown in S. T. Berry (1994) that the following update process converges to the true parameters, θ:

$$\delta^{r+1} = \delta^r + log(S) - log\left(\hat{s}(\theta^r(\delta^r), \delta^r)\right)$$

In our context, however, we are not concerned with product-level unobserved heterogeneity. Our set of products numbers greater than 1,100, precluding product-specific constants. Furthermore, for bottled water, unlike vehicles (S. Berry, Levinsohn, and Pakes, 1995) and cereal (Nevo, 2001), it is much easier to capture heterogeneity in products explicitly in the characteristic space. Instead, we solve our household-level heterogeneity in unobserved tastes for the "outside good" by contracting out at the household level. Because our data has repeat purchases for each household, for any given value of θ, we calculate a single value of utility for tap water for each household which equates the observed share of purchases *of the outside good* (which, in our data, is "no purchase" of bottled water) to the predicted share of purchases for that household. Thus, δ is an H-dimensional vector, where H is the number of households in our data. The iterative update process follows (S. T. Berry, 1994):

$$\delta_h^{r+1} = \delta_h^r + log(S_{h0}) - log\left(\hat{s}_{h0}(\theta^r(\delta^r), \delta^r)\right)$$

Each update of the household fixed effect, δ_h, is a function of the prior value of δ_h and the share of "outside option" purchases for that household predicted by the estimate of θ, which is itself a function of δ, the household fixed effects for *all* households.

Using simulated data, we tested the application of this model and two-stage estimator. A description and detailed results are included in Appendix A. Using simulated data with known correlations between a household's taste for tap water and the likelihood of fracking occurring within the household's zip code, it is shown that failure to account for time-constant household heterogeneity leads to biased estimates of the parameter of interest, and that the described model and two-step estimation correctly estimates all parameters of interest up to scale.

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 2006-2014 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.

To facilitate estimation, we first select only households which (1) are located within a PA or OH zip code that is "fracked" between 2007 and 2014, (2) remain in the panel for at least 100 weeks, (3) average between one and eight shopping trips with expenditures greater than \$15.00 at a Nielsen-reporting store every four weeks, and (4) have at least one shopping trip where bottled water is purchased and at least one shopping trip where bottled water is not purchased. The final criteria (4) omits households for whom the household utility of the outside good would be positive or negative infinity and would thus not contribute to the identification of taste parameters. We omit all trips of less than \$15 to avoid counting a brief trip to the market as a full shopping trip. We also omit households that move across zip codes during the time period of our analysis, because including them would complicate the interpretation of the variable measuring the arrival of fracking activity. In the data, 397 households meet criteria (1), while the remaining criteria yield a total of 199 households. We randomly draw a sample of 198 households from within PA and OH, but outside of the "fracked" zip codes. This yields a sample of 397 household panelists.

HomeScan data contains detailed information on all purchases made by panelists. We use all trips of greater than \$15 by all sampled panelists, including those trips where no purchase of bottled water occurred. For trips in which one or more bottled water products were purchased, we 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. 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 (K. E. Train, 2009).

For all households in PA and OH, the data yields complete demographics on 13,383 panelists over 9 years of participation for a total of 50,678 panel-years. A total of 1,982,548 trips and choice sets are observed. For the sample used in estimating the structural model, the data yields 397 panelists over 9 years with a total of 84,046 trips.

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.[9] Fracking usually occurs within a few weeks after drilling commences, so the two operations would generally occur within the same quarter in any case. 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.

Municipal water service boundaries (Public Water Supplier Service Areas) are available for Pennsylvania from the Department of Environmental Protection. In principle, we could overlay the municipal water boundaries with zip code boundaries and calculate the ratio of the overlapping area to the total zip code area to approximate the proportion of households served by municipal water. However, this method is problematic for two reasons. First, there is no comparable information for municipal water service areas in Ohio.[10] Second, the method implicitly assumes a uniform distribution of households over both the zip code and the municipal water service area. Since fracking locations may be negatively correlated with housing density even within zip code, this assumption is unlikely to hold.

To better measure exposure to fracking and well-water dependence, we use the Microsoft US Building Footprint database[11]. This publicly-available dataset uses semantic segmentation and polygonization to identify and map the footprints of over 10.3 million structures in PA and OH. We filter this data to include only structures between 900 square feet and 3,500 square feet to avoid counting garages (<900 sq.ft.) and commercial buildings (>3,500 sq.ft.), and use the centroid of each footprint polygon as a potential household. In PA, we then overlay the Public Water Supplier Service Areas to designate each structure as well-dependent or municipally-supplied. We omit Public Water Suppliers that serve very small populations from a common well (e.g. mobile home parks) as these shared wells still render the customer well-dependent. In Ohio, we collect data on domestic water well locations from the Ohio Division of Water Resources. We then designate each structure in Ohio as being well-dependent if there are 2 or more wells within 500m. The fraction of structures in each zip code that are well-dependent forms our measure of unconditional household well dependence for that zip code.

Nielsen scanner data gives only the household's zip code of residence which is insufficient to attach the household's exact location, the household's dependence

[9]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.

[10]This was confirmed by several phone and email conversations with Ohio state officials, as well as extensive searching online. We are grateful to David Keiser for suggesting, as an alternative, the use of domestic water well locations.

[11]https://github.com/microsoft/USBuildingFootprints

on well-water, and the household's exposure to fracking. To assess the *probability* of being well-dependent and exposed to fracking, we generate an unconditional probability of well-dependence and fracking exposure over structures. For each structure in the data and for each week during our study period, we calculate the density of water wells in the area, designating a structure to be "well-dependent" if the density of wells in the area exceed 2 per 2km^2. Once a structure is exposed, it remains exposed in this measure, even if additional wells are drilled nearby. We divide by the total number of structures in the zip code, restricting the measure to the interval $[0, 1]$. An unconditional probability of 1 is observed only if (i) every structure in a zip code is well-dependent, and (ii) every structure has one or more fracked wells within 500m. The unconditional probability reflects the probability that a randomly-selected household in that zip code is both well-dependent *and* exposed to fracking. Appendix X provides ... We sensitivity-test measures of 250m, 1,000m, and 2,000m as well in an appendix.

4.3 Summary statistics

Sum. Stats needs to be updated once we nail down the of households - we have a more refined measure of exposure now, so there are more exposed Hh's

Table 1 provides descriptive statistics about bottled water expenditures and household characteristics among the 358 households in our main analysis. The mean bottled water expenditure per quarter is about \$170, equating to about \$2 per day per household. In expectation, a household in our subsample is on well water with probability 0.07. Average income is about \$60,000 per year, average household size is about 2.6 persons, and about three-quarters of the heads of households are married, while just over one-quarter have children in the home. The sample is largely of Caucasian descent, more so than the average population in Pennsylvania and Ohio.

Table 1: Descriptive statistics

Variable	N	Mean	Std. Dev.	Min	Max
Bottled water expenditure	358	169.6	226.4	0	1,335
Well water	358	0.0698	0.0952	0	0.728
Income	358	59.94	34.32	5	225
Household size	358	2.645	1.216	1	6
Married	358	0.763	0.426	0	1
Children	358	0.265	0.442	0	1
Black	358	0.0475	0.213	0	1
Asian	358	0.00559	0.0746	0	1

Bottled water expenditure is in dollars per household per quarter.
Well water is measured as domestic water wells per residential housing structure.

As noted in Section 4.1, our main results are for those households with sufficient observations such that a household fixed effect can be estimated. Table 2 compares the expenditures and household characteristics for the subsample of 358 households to the larger sample of 6,032 households. Compared to the full sample, the subsample has higher expenditures on bottled water (the mean value is about 22% higher), and is also more likely to be on well water. The households in the subsample have generally comparable demographic characteristics, although with a higher proportion of married heads of households, and a lower proportion of African-American heads of households.

Table 2: Comparing subsample to full sample

Variable	Full sample mean (SD)	Subsample mean (SD)	Difference in means
Bottled water expenditures	139.07	169.558	−30.488***
	(228.524)	(226.392)	
Well Water	0.049	0.070	−0.021***
	(0.083)	(0.095)	
Income	60.496	59.937	0.558
	(34.597)	(34.321)	
Household size	2.551	2.645	−0.095
	(1.261)	(1.216)	
Married	0.685	0.763	−0.078***
	(0.465)	(0.426)	
Children	0.279	0.265	0.013
	(0.448)	(0.442)	
Black	0.079	0.047	0.031**
	(0.269)	(0.213)	
Asian	0.010	0.006	0.005
	(0.101)	(0.075)	

$N = 358$ households in subsample, 6,032 households in full sample.
*** p<0.01, ** p<0.05, * p<0.1

5 Results

5.1 Reduced Form Model Results

Tables 3 and 4 provide a summary of the reduced form results using ounces of water purchased as the dependent variable. Columns (1)-(3) assume a regular panel (filling in weeks where no shopping trip was recorded with 0 purchases), while Columns (4)-(6) use actual trips. Column (6) features the richest set of fixed effects. While using only household fixed effects (Columns (1) and (4)) shows a significant increase in ounces purchased, including richer fixed effects reduces the magnitude of the point estimate and its significance. While all coefficients are positive, consistent with an increase in ounces of water purchased, the effect is not statistically significant.

A similar result is shown in Tables 5 and 6, which show results using dollars spent per week/trip rather than ounces purchased. Here, purchases are significant at the 10% level even when controlling for zip status x week fixed effects (Column (6)). Table 5 shows results using all zip codes, while 6 shows results using only fracked and adjacent zip codes.

Tables 7 and 8 show results for regressions using the well-dependent fracking exposure measure. Only results for actual trips are included. Magnitudes are slightly larger, suggesting well-dependence might play a part in determining household purchases of bottled water, but results and differences between the well-dependent measure and the overall measure are not significant.

Table 9 shows results including both fracking measures. This specification captures differential effects for fracking, which may include localized income effects,

Table 3: Ounces of water purchased; structure-specific well density

	Dependent variable:					
	Ounces of water purchased					
	(1)	(2)	(3)	(4)	(5)	(6)
fracCumExposedUnique	153.785***	97.604**	60.994	301.184***	183.960*	132.640
	(42.380)	(42.983)	(46.039)	(113.948)	(111.307)	(108.640)
Hh FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	No	Yes	Yes	No	Yes	Yes
ZipStatus x week FE	No	No	Yes	No	No	Yes
Panel	Weekly	Weekly	Weekly	Actual trips	Actual trips	Actual trips
Zips used	All	All	All	All	All	All
Observations	2,414,730	2,414,730	2,414,730	1,127,716	1,127,716	1,127,716
R^2	0.145	0.147	0.147	0.236	0.239	0.239
Adjusted R^2	0.141	0.142	0.142	0.227	0.230	0.230

Note: *p<0.1; **p<0.05; ***p<0.01
Errors clustered at zip level

Table 4: Ounces of water purchased; structure-specific well density

	Dependent variable:					
	Ounces of water purchased					
	(1)	(2)	(3)	(4)	(5)	(6)
fracCumExposedUnique	153.785***	65.066	60.994	301.184***	148.865	132.640
	(42.410)	(44.219)	(46.104)	(114.120)	(110.635)	(108.997)
Hh FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	No	Yes	Yes	No	Yes	Yes
ZipStatus x week FE	No	No	Yes	No	No	Yes
Panel	Weekly	Weekly	Weekly	Actual trips	Actual trips	Actual trips
Zips used	Fracked+Adj	Fracked+Adj	Fracked+Adj	Fracked+Adj	Fracked+Adj	Fracked+Adj
Observations	474,229	474,229	474,229	198,337	198,337	198,337
R^2	0.147	0.151	0.151	0.251	0.256	0.258
Adjusted R^2	0.143	0.145	0.145	0.241	0.245	0.244

Note: *p<0.1; **p<0.05; ***p<0.01
Errors clustered at zip

Table 5: Water expenditures; structure-specific well density

	Dependent variable:					
	Dollars spent per week/trip					
	(1)	(2)	(3)	(4)	(5)	(6)
fracCumExposedUnique	1.573***	1.272**	0.752	3.300**	2.716**	1.970*
	(0.503)	(0.496)	(0.478)	(1.307)	(1.271)	(1.171)
Hh FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	No	Yes	Yes	No	Yes	Yes
ZipStatus x week FE	No	No	Yes	No	No	Yes
Panel	Weekly	Weekly	Weekly	Actual trips	Actual trips	Actual trips
Zips used	All	All	All	All	All	All
Observations	2,414,730	2,414,730	2,414,730	1,127,716	1,127,716	1,127,716
R^2	0.140	0.141	0.142	0.226	0.228	0.229
Adjusted R^2	0.136	0.137	0.137	0.217	0.219	0.219

Note: *p<0.1; **p<0.05; ***p<0.01
Errors clustered at zip level

Table 6: Water expenditures; structure-specific well density

	Dependent variable:					
	Dollars spent per week/trip					
	(1)	(2)	(3)	(4)	(5)	(6)
fracCumExposedUnique	1.573***	0.867*	0.752	3.300**	2.253*	1.970*
	(0.504)	(0.483)	(0.479)	(1.309)	(1.233)	(1.175)
Hh FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	No	Yes	Yes	No	Yes	Yes
ZipStatus x week FE	No	No	Yes	No	No	Yes
Panel	Weekly	Weekly	Weekly	Actual trips	Actual trips	Actual trips
Zips used	Fracked+Adj	Fracked+Adj	Fracked+Adj	Fracked+Adj	Fracked+Adj	Fracked+Adj
Observations	474,229	474,229	474,229	198,337	198,337	198,337
R^2	0.145	0.148	0.149	0.247	0.251	0.253
Adjusted R^2	0.141	0.143	0.142	0.237	0.240	0.240

Note: *p<0.1; **p<0.05; ***p<0.01
Errors clustered at zip level

and fracking incident on well-dependence, which captures the effect of fracking unique to those households dependent on wells for drinking water. Results, though also not significant, may suggest that well-dependence plays a role in household responses to fracking.

Table 7: Ounces of water purchased; structure-specific well density with water well dependence

	Dependent variable:					
	Ounces of water purchased					
	(1)	(2)	(3)	(4)	(5)	(6)
w.fracCumExposedUnique	160.192***	72.430	68.152	308.345**	157.284	139.308
	(45.020)	(46.678)	(48.470)	(124.659)	(119.978)	(118.265)
Hh FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	No	Yes	No	No	Yes	No
ZipStatus x week FE	No	No	Yes	No	No	Yes
Panel	Weekly	Weekly	Weekly	Actual trips	Actual trips	Actual trips
Zips used	Fracked+Adj	Fracked+Adj	Fracked+Adj	Fracked+Adj	Fracked+Adj	Fracked+Adj
Observations	474,229	474,229	474,229	198,337	198,337	198,337
R^2	0.147	0.151	0.151	0.251	0.256	0.258
Adjusted R^2	0.143	0.145	0.145	0.241	0.245	0.244
Note:	*p<0.1; **p<0.05; ***p<0.01 Errors clustered at zip level					

Table 8: Water expenditures; structure-specific well density with water well dependence

	Dependent variable:					
	Dollars spent per week/trip					
	(1)	(2)	(3)	(4)	(5)	(6)
w.fracCumExposedUnique	1.638***	0.933*	0.814	3.429**	2.376*	2.070
	(0.546)	(0.522)	(0.516)	(1.454)	(1.366)	(1.306)
Hh FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	No	Yes	No	No	Yes	No
ZipStatus x week FE	No	No	Yes	No	No	Yes
Panel	Weekly	Weekly	Weekly	Actual trips	Actual trips	Actual trips
Zips used	Fracked+Adj	Fracked+Adj	Fracked+Adj	Fracked+Adj	Fracked+Adj	Fracked+Adj
Observations	474,229	474,229	474,229	198,337	198,337	198,337
R^2	0.145	0.148	0.149	0.247	0.251	0.253
Adjusted R^2	0.141	0.143	0.142	0.237	0.240	0.240
Note:	*p<0.1; **p<0.05; ***p<0.01 Errors clustered at zip level					

calculate results in dollars per household per year to replace this old text: In all specifications, with regard to the interaction of fracking and well water (i.e., the sum of the two coefficients, shown at the bottom of the table), we observe a positive and significant treatment effect, indicating that the "treatment" of fracking increased households' expenditures on bottled water. Our preferred estimate for averting expenditures attributable to the arrival of fracking is 6, column (6), which reports the increase in water expenditures per trip when using only households in fracking and fracking adjacent zip codes. While statistically indistinguishable from zero at the 5% level, the point estimate suggests a quarterly increase (assuming the sample average of 3 reported trips per month) of $17.73.

Table 9: Water purchases; both fracking exposure measures

	Dependent variable:	
	Total ounces purchased	Total dollars spent
	(1)	(2)
fracCumExposedUnique	66.582	0.971
	(302.532)	(2.853)
w.fracCumExposedUnique	72.122	1.091
	(333.344)	(3.327)
Hh FE	Yes	Yes
Week FE	Yes	Yes
ZipStatus x week FE	Yes	Yes
Panel	Actual trips	Actual trips
Zips used	Fracked+Adj	Fracked+Adj
Observations	198,337	198,337
R^2	0.258	0.253
Adjusted R^2	0.244	0.240
Note:		*p<0.1; **p<0.05; ***p<0.01
		Errors clustered at zip level

Qualitatively, these estimates are consistent with those of Wrenn, Klaiber, and Jaenicke (2016) (WKJ), 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 in Pennsylvania and Ohio. 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. Do we need to just drop this?:::: Our analysis produces a similar estimate, despite obtaining household expenditure data at a finer geographic resolution (zip code rather than county), which allows us to assign "treatment" (by fracking) at a substantially finer resolution as well as condition on a richer set of fixed effects. We also use a more precise definition of the arrival of fracking: whereas WKJ consider that fracking arrives in 2007 across all Pennsylvania counties in which there is any fracking activity by 2010 (and uses Ohio, where fracking had not arrived as of 2010, as a control), we measure shale gas well activity by quarter, and allow for the possibility that fracking came to different zip codes at different times. WKJ also do not distinguish between households on municipal water and well water in most of their specifications, except in one series where they omit metropolitan area counties and estimate the change in expenditures on the rest (that is, assuming that all households in non-metro counties are on well water). Finally, we use a longer time series from the Nielsen HomeScan data: WKJ use data only from 2005-2010, whereas our panel extends through 2014. In addition to the foregoing differences of methods and data, which we believe generally represent improvements, we also introduce a constraint: our subsample of Nielsen households has higher bottled water expenditures than the average and is more likely to be serviced by well water (see Table [redacted]

Regardless of these differences of methods and data between our reduced-form estimates and those of WKJ, our primary 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 the averting expenditures context. To that end, in the following section we reflect on that comparison within the context of our own data and methods.

A paragraph, possibly pulled from below, summarizing the averting expenditure framework.

Point estimates from results in Tables 3 - 9 consistently show an increase in both expenditures and ounces of water purchased with varying degrees of significance. Bottled water is unique in that some purchases may be of larger size but be lower total cost. For instance, a 128 ounce jug of water frequently costs more than a 20 ounce bottle of water. We estimate a linear probability model for "jug" purchases using our fracking measures (Table **??** and note that, while the results are not statistically significant, point estimate of the probability of purchasing a 128 ounce or larger container of bottled water increases when fracking is introduced. Including both the fracking and the well-dependent fracking measure allows for a differential effect on those households who rely on well water. While the difference is not significant, the point estimates would indicate a stronger effect for well-dependent households. RBecause consumers can feasibly increase ounces of bottled water consumed while decreasing expenditures (or vice versa), it is necessary to estimate the underlying utility of purchase, rather than rely solely on the reduced form estimates. Justin: a little kludgy, but a good start on justifying the structural when the RF shows only increase in expenditures.

Table 10: Linear Prob. Model - isJug

	Dependent variable:		
	isJug (purchased 1+ gallon size)		
	(1)	(2)	(3)
fracCumExposedUnique	0.267		0.036
	(0.234)		(0.599)
w.fracCumExposedUnique		0.289	0.252
		(0.258)	(0.683)
Hh FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
ZipStatus x week FE	Yes	Yes	Yes
Panel	Actual trips	Actual trips	Actual trips
Zips used	Fr. + Adj.	Fr. + Adj.	Fr. + Adj.
Observations	198,337	198,337	198,337
R^2	0.315	0.315	0.315
Adjusted R^2	0.302	0.302	0.302
Note:			*p<0.1; **p<0.05; ***p<0.01 Errors clustered at zip level

5.2 Discrete choice structural model

Results from the random parameters logit, shown in Table 11, are primarily as expected. All parameters are significant in Model 1, owing largely to the sample size of 84,046 choice occasions. For all models, the coefficient on price is negative. The coefficient on total water is positive, while the coefficient on *Jug*, defined as any bottled water greater than 72 ounces in a single container, is negative.

Table 11: Bottled Water Expenditures: Random Parameters Logit

	Model 1	Model 2
Price (α)	-2.603^{***}	-3.090^{***}
	(0.027)	(0.029)
Total Water	0.009^{***}	0.011^{***}
	(0.0003)	(0.0003)
Jug	-1.417^{***}	-1.632^{***}
	(0.005)	(0.005)
Flavored brand	0.020^{***}	0.020^{***}
	(0.0001)	(0.0001)
Luxury brand	-0.069^{***}	-0.069^{***}
	(0.003)	(0.002)
National brand	0.006^{***}	0.006^{***}
	(0.00004)	(0.00004)
Other brand	0.001^{***}	0.001^{***}
	(0.00003)	(0.00003)
HH size x Total water	0.0005^{***}	0.0005^{***}
	(0.00005)	(0.0001)
Fracking x Total water	0.005^{***}	0.004^{***}
	(0.0001)	(0.0001)
Var(Price)	0.871^{***}	1.333^{***}
Var(Total water)	0.0001	0.0001
Cov(Price, Total water)	-0.006	-0.008
Household FE		✓
Log Likelihood	$-57,315.15$	$-55,091.21$
Number of households	397	397
Number of trips	84,046	84,046

*** p<0.01, ** p<0.05, * p<0.1

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 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).

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 fracking 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 arrival of fracking, we take the ratio $\frac{\beta^{w,o}}{\beta^{p}}$. For Model 1 (without household fixed effects) the compensating variation measure is -\$0.0016, and for Model 2 (with household fixed effects) it is -\$0.0023. Because we do not observe in our data increases in the consumption of water-based products such as soda or (not-from-concentrate) juice, we consider this measure to be biased downwards.

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 **??**) estimates a per-household disamenity value of \$56.89 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 5.2 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 \$0.00165 per ounce in the model without household fixed effects, and \$0.0023 per ounce in the model with household fixed effects.

Suppose that the marginal utility of bottled water consumption is constant at \$0.00235 per ounce. To make a meaningful comparison with the reduced-form estimate from Table **??**, 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 2.6 persons per household on average, the size of the compensating variation or disamenity-of-fracking value would be \$17.19 per household per quarter. (The comparable calculation using the value calculated without household fixed effects is \$12.14 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 that is twice as high.

The compensating variation owing to the arrival of fracking that we estimate from the random parameters model, \$17.19 per household per quarter, amounts to just 30% of the increase in expenditures we estimate in the reduced-form context (\$56.89). In the models without household fixed effects the difference is even starker: the compensating variation of \$12.14 per household per quarter associated with fracking amounts to just 8% of the increase in expenditures we estimate using the reduced-form model (\$146.70). These results suggest that other desirable attributes associated with bottled water, such as taste and portability, account for a relatively substantial portion of observed increased expenditures, at least in the case of households that experienced a rise in fracking activity in these states. This result has potentially broader implications for use of reduced-form models to measure averting expenditures on defensive goods, especially in contexts where joint production could be substantial.

5.3.1 Unobserved household heterogeneity

The household fixed effects display a spatially autocorrelated pattern consistent with geographic patterns in tap water quality. Figure 5.3.1 shows the mean magnitude of

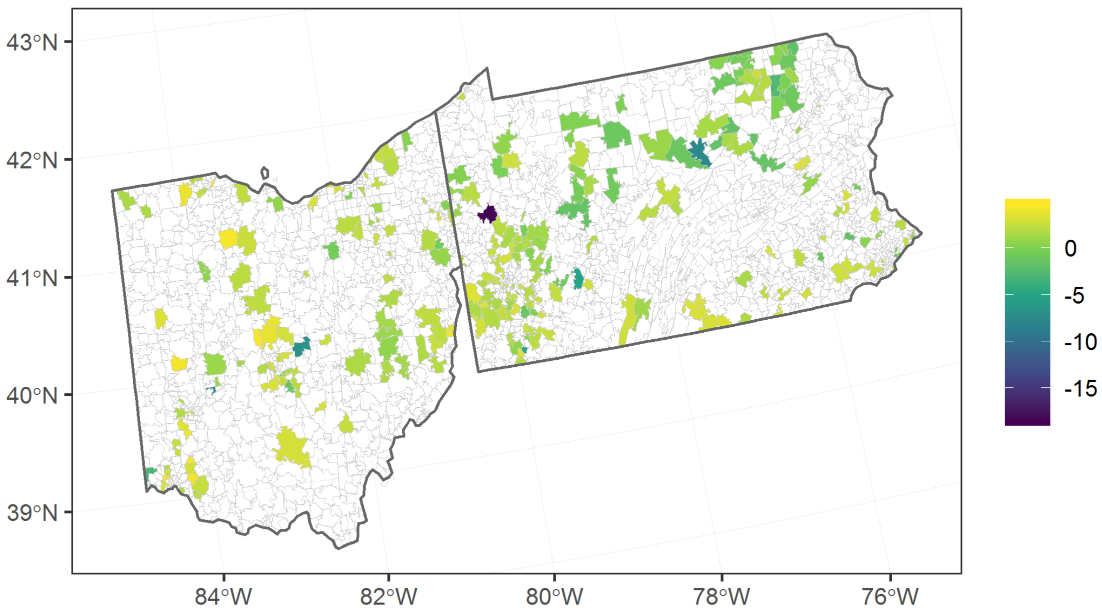

Figure 2: Household-specific fixed effects on outside option by zip code. Value reflects household preferences for own tap water that do not vary over time. In zip codes with multiple households, the mean is displayed. Clustering by value is visible, indicating spatial autocorrelation in household preference for outside option, possibly due to local water quality. Moran's I, a distance-weighted correlation statistic, is positive, indicating spatial autocorrelation, and is significant at the .01 level.

the fixed effects for each zip code in the sample. The northern portion of Pennsylvania generally has greater negative values relative to the southern portion, indicating a lower preference for tap water in that area. The Marcellus Shale underlies the north and western portion of the state, as well as the eastern portion of Ohio. These areas appear to have a negative trend in the household fixed effects. This is consistent with a persistent household perception of lower water quality in areas that have been historically associated with conventional gas, oil, and coal extraction. Omitting this time-invariant effect would introduce an omitted variables bias if historic gas and coal fields are more likely to be the site of modern fracking operations. In this situation, the omitted variable, time-invariant disamenity of own tap water, would confound identification of the effect of fracking. To test for spatial autocorrelation (e.g. clustering), we calculate a Moran's I (Moran, 1950). The test statistic is positive, indicating spatial autocorrelation, and significant at the .01 level.

5.3.2 Limitations of the structural method

To the extent that economists use averting expenditures to measure WTP for goods and services that are difficult to exchange in markets, it is useful to consider other ways (besides joint production) in which reduced-form measures of averting expenditures provide only partial insight on such WTP. Assuming away joint production (as many studies of averting expenditures implicitly do), averting expenditures on bottled water represent a lower bound on WTP for three reasons. First, such studies typically focus on one type of defensive product. For instance, our study (like most that analyze WTP for improved water quality) focuses on bottled water purchases, and does not include households' expenditures on other technologies such as home water filters or custom delivery of large water containers. Second, the averting expenditures method our analysis does not include damages that households incur from consumption of the outside good. Finally, the reduced-form framework does not allow observation of WTP in excess of the pure expenditure. In Figure 1, this is the area under the curve D_2 and above the price level (from Q_1 to Q_2). However, the structural demand estimate does allow us to estimate this component. Thus, the structural approach dominates a reduced-form approach, regardless of whether there is joint production. The structural model provides a superior estimate of WTP by incorporating concerns about joint production and issues that arise from expenditure being a lower bound for WTP.

6 Conclusions

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 perceived quality of drinking water—in more detail.

The per-household, per-quarter disamenity estimated in our structural model yields a value of \$17.19, based on an assumption of 32 ounces of water consumed per day, an amount equal to just 20% of the disamenity estimated by our reduced-form approach. This suggests that joint production may represent a substantial component of averting expenditures, and provides an illustration of how structural demand models represent a considerable improvement in the use of averting expenditures to measure WTP for higher environmental quality.

References

Bartik, Timothy J. (Mar. 1988). "Evaluating the benefits of non-marginal reductions in pollution using information on defensive expenditures". In: *Journal of Envi-*

ronmental 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.

Berry, Steven T (1994). "Estimating discrete-choice models of product differentiation". In: *The RAND Journal of Economics*, pp. 242–262.

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.

Dube, Jean-Pierre et al. (2002). "Structural applications of the discrete choice model". In: *Marketing Letters* 13.3, pp. 207–220.

Fetter, T. Robert (2019). "Fracking, Toxics, and Disclosure". Unpublished manuscript. URL: https://sites.duke.edu/trfetter/research.

Goolsbee, Austan and Amil Petrin (2004). "The consumer gains from direct broadcast satellites and the competition with cable TV". In: *Econometrica* 72.2, pp. 351–381.

Gopalakrishnan, Sathya and H. Allen Klaiber (Sept. 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 (May 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 (Apr. 2015). "Welfare and Distributional Implications of Shale Gas". In: DOI: 10.3386/w21115. URL: http://dx.doi.org/10.3386/w21115.

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.

Lutzeyer, Sanja, Daniel J Phaneuf, and Laura O Taylor (n.d.). "THE AMENITY COSTS OF OFFSHORE WIND FARMS: EVIDENCE FROM A CHOICE EXPERIMENT". In: ().

Mason, Charles, Lucija Muehlenbachs, and Sheila Olmstead (Feb. 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:

Moran, P A P (1950). "Notes on Continuous Stochastic Phenomena". In: *Biometrika* 37.1/2, pp. 17–23. ISSN: 00063444. DOI: 10.2307/2332142. URL: http://www.jstor.org/stable/2332142.

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 (Dec. 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.

— (2001). "Measuring market power in the ready-to-eat cereal industry". In: *Econometrica* 69.2, pp. 307–342.

Olmstead, S. M. et al. (Mar. 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.

Petrin, Amil and Kenneth Train (2010). "A Control Function Approach to Endogeneity in Consumer Choice Models". In: *Journal of Marketing Research* 47.1, pp. 3–13. ISSN: 00222437. URL: http://www.jstor.org/stable/20618950.

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. (June 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 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 (Dec. 2016). "Unconventional Shale Gas Development, Risk Perceptions, and Averting Behavior: Evidence from Bottled Water Purchases". In: *Journal of the Association of En-*

vironmental and Resource Economists 3.4, pp. 779–817. DOI: 10.1086/688487. URL: http://dx.doi.org/10.1086/688487.
Zhang, Junjie and Quan Mu (2018). "Air pollution and defensive expenditures: Evidence from particulate-filtering facemasks". In: *Journal of Environmental Economics and Management* 92, pp. 517–536. ISSN: 0095-0696.

A Household Heterogeneity Simulation

A search of the literature yielded no examples of a household-specific fixed effect being implemented in a discrete choice RUM. Simulations were performed to observe the convergence of the household-level fixed effects. Results show (1) convergence to the true parameters up to scale when household-level fixed effects are included, and (2) biased estimates when naively estimated using a traditional multinomial logit.

Data was simulated using the following parameters:

Variable	Symbol	Value
Households	H	400
Choice occasions	T	40
Choice occasions with fracking	T^{frack}	10
Number of products	J	5
Number of observed characteristics	K	5
Variance of logistic preference shock (ϵ_j)	σ_ϵ	1
Tapwater taste ($\delta_0 \sim N(\mu_\delta, \sigma_\delta, \iota_\delta)$) parameters		
Mean	μ_δ	2
Variance	σ_δ	3
Skew	ι_δ	2

Data was drawn from these parameters and $400 \times 40 = 1,600$ choice occasions were simulated. $K-1$ observable product characteristics remained constant throughout the simulated dataset, but K_{price} was drawn from a random uniform distribution centered on a vector of mean prices based on the attributes (e.g. larger, luxury products had a higher mean price) with a range of plus or minus \$0.50. Price was not correlated with any unobserved household or product characteristics.

Figure 3 shows the ratio of the parameter estimates to the true parameters at each update of δ. The red line is the first parameter estimate from the Maximum Likelihood stage, and is the MLE estimate of the parameters without household-specific tastes for tap water. Each progressively-darker gray line is an iteration of the two stages, and the blue line is the final ratio of the parameter estimates to the true parameters. Because discrete choice models only identify to scale, a constant scaling is expected. A "perfect" estimate, then, results in a flat, horizontal line. As expected, the coefficient on the interaction between product size and the presence of fracking, $\beta_{sizeFracking}$ is biased towards zero before the first iteration of the two-step process.

The ratio of the parameter of interest to the parameter on price is the estimated willingness to pay. Fig 4 shows the ratio of this estimated ratio (from $\hat{\beta}$) to the true ratio. The estimated ratio converges to a constant value near unity within 30 iterations.

Figure 3: Convergence in Simulated Data with Household-level Tap Water Taste Heterogeneity

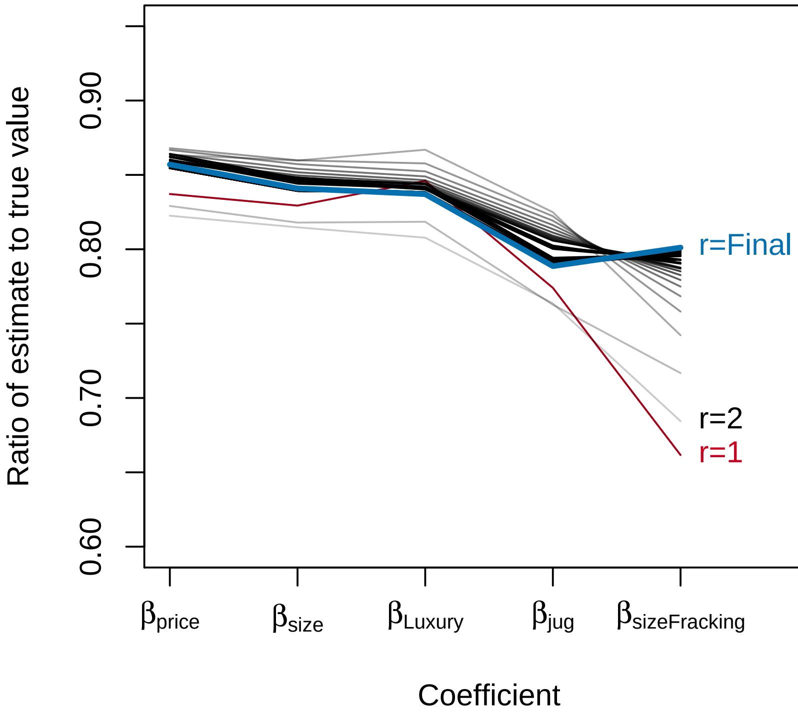

Figure 4: Convergence of Willingness to Pay Estimate in Simulated Data with Household-level Tap Water taste Heterogeneity