Heterogeneous Misperceptions of Energy Costs: Implications for Measurement and Policy Design

Sébastien Houde,* Erica Myers†

May 18, 2017

Preliminary and incomplete, please do not cite

Abstract A widely-used test of consumer misperception compares the demand response to potentially misperceived costs such as energy operating costs, sales tax, or shipping and handling, versus salient, correctly perceived purchase costs. The ratio of the responsiveness coefficients has been considered a sufficient statistic to determine the degree of consumer misperception. In this paper, we show that this statistic corresponds to a first-order approximation of the average degree of misperception, and this approximation can lead to an economically important bias. The direction of the bias depends on the sign of the correlation between the two responsiveness coefficients. Further, we show that even if measured correctly, the average amount of misperception in a market is not sufficient for optimal policy design, and can in fact be misleading. Using data for the U.S. refrigerator market, we quantify the bias of the first order approximation and demonstrate the importance of accounting for heterogeneity of misperception. We find substantial heterogeneity in perception of energy costs and show that this heterogeneity is not driven by income. In our context the first-order approach provides a downward bias of the average degree of misperception. We use the estimated distribution the two responsiveness coefficients to simulate different optimal policies and show that heterogeneity in the degree of misperception can significantly impact optimal policy design and estimates of welfare effects.

JEL Codes: Q41, Q50, L15, D12, D83.

Key words: demand estimation, energy efficiency gap, attention allocation, behavioral welfare economics.

*Department of Agricultural and Resource Economics, University of Maryland, shoude@umd.edu.

†Department of Agricultural and Consumer Economics at the University of Illinois at Urbana-Champaign, ecmyers@illinois.edu.

1. Introduction

There is a large literature estimating misperception of certain complex and potentially less salient aspects of product costs, such as shipping and handling (Hossain and Morgan 2006), sales tax (Chetty et al. 2009), highway tolls (Finkelstein 2009), and medical expenses (Loewenstein et al. 2013), to name a few. In the energy context, misperception of the operating costs of energy using durables has been debated for more than three decades (e.g., Hausman 1979; Dubin and McFadden 1984; Dubin 1992; Goldberg 1998; Allcott 2011; Grigolon et al. 2014; Busse et al. 2013; Houde 2016) and is still central to the debate on how to design policies to account for various negative externalities associated with energy use.

The general test used to quantify the degree of misperception consists of comparing the demand response to an assumed-to-be correctly perceived aspect of product cost (e.g. posted purchase price) against the potentially mis-perceived aspect of product cost. This paper makes two methodological contributions for the execution of these types of studies. First, we show that the methods that have been used to quantify misperception, both structural and reduced-form, implicitly rely on a first-order approximation of the distribution between the coefficients on the correctly perceived aspects of cost and the potentially mis-perceived aspects of cost. We demonstrate that this approximation can lead to an economically important estimation bias of average misperception and the direction of this bias depends on the sign of the correlation between the sensitivity to assumed-to-be correctly perceived and potentially mis-perceived costs. Our second contribution is to show that even if measured correctly, the average amount of misperception in a market is not sufficient for optimal policy design, and can be in fact very misleading. Recent papers by Allcott and Taubinsky (2014), Farhi and Gabaix (2016), and Taubinsky and Rees-Jones (2016) investigating optimal taxation in the presence of behavioral biases make a similar point. Building on the framework proposed in these papers, we derive new results in the context of environmental taxation that show how heterogeneity in misperception affects the economic efficiency of different types of policies if not properly accounted for.

Our empirical investigation focuses on the U.S. refrigerator market and proceeds in three steps. First, we replicate the standard approach to estimating the average degree of misperception of energy costs. Using a unique administrative data set with micro-level sales data from a large appliance retailer, we are able to recover more credible estimates of the coefficients on price and energy costs than has been possible in the past. Previous work has focused largely on cross-sectional variation in energy price (e.g., Hausman 1979; Houde 2016; Rapson 2014), or has used aggregated state-level annual data (Jacobsen 2015). Similar to recent papers in the context of cars and housing (e.g., Busse et al. 2013; Allcott and Wozny 2014; Grigolon et al. 2014; Sallee et al. 2016; **?**), our empirical strategy exploits fine-grained panel data, allowing us to control for product, region, and time specific unobservables. Using the first-order approximation, we find that consumers do respond to energy costs, but tend to undervalue them. The fact that we find a response to energy costs in this context and that the coefficient is credibly estimated and robust is in itself a contribution to the debate of misperception of energy costs. Previous studies for the appliance sector provided mixed evidence that consumers respond to energy costs.

In the second step of our empirical investigation we recover the full joint distribution of preferences for prices and energy costs. We find substantial heterogeneity in perception of energy costs and show that this heterogeneity is not driven by income. Misperception does appear to be important and to vary widely across consumers—some consumers appear to not pay attention to energy costs, a large fraction respond to energy costs, but tend to undervalue them, and another fraction overvalue energy costs. These patterns are strikingly similar across income groups and imply that the first-order approach provides a downward bias of the average degree of misperception in the appliance market. In the third step, we use the distribution to simulate different optimal policies and show that heterogeneity in the degree of misperception can significantly impact optimal policy design and the estimates of the welfare effects.

The paper is organized as follows. In the next section, we present the standard framework used to quantify misperception and show how this framework has commonly been used. In particular, we show why the statistic used to quantify misperception corresponds to an approximation and present a formal expression to evaluate the bias of this approximation. In Section 3, we discuss the data and choice environment for our empirical investigation. In Section 4, we develop an empirical strategy to quantity the average amount of misperception using the first-order approximation. In Section 5, we estimate the full distribution of heterogeneity that allows us to compute the bias of the first-order approximation. Section 6 follows and investigates how heterogeneity in misperception impacts the design of policies.

2. Framework

2.1. Set-Up

Consumer i chooses among J different energy-intensive durables. The price (capital cost) of product j is p_j and the future operating energy cost over the entire expected lifetime of the product is E_j. For ease of exposition and without loss of generality, we abstract away from uncertainty and consumer-specific heterogeneity in the product lifetime, future energy prices, utilization, depreciation, and discount factor. We simply assume that E_j is the exact measure of expected energy operating costs discounted with a normal rate of return. Consumer i values product j as follows:

$$U_{ij} = \eta_i P_j + \theta_i E_j + \gamma_j + \epsilon_{ij} \tag{1}$$

where γ_j is the vertical quality of product j and captures all the attributes of the products and ϵ_{ij} captures idiosyncratic preferences. A large literature has estimated variants of this model with the goal of identifying the preference parameters for price and energy cost: η and θ, respectively. The common approach used to quantify the degree of misperception is to test whether consumers trade off one dollar of energy cost the same way they trade off one dollar of purchase price. Formally, the test of misperception is whether $\theta/\eta \neq 1$. An alternative, but equivalent test that has been widely used in the literature solves for the implicit discount rate such that the ratio $\theta/\eta = 1$. Using this latter approach, an implicit discount rate that is markedly above the normal rate of return is seen as a sign that consumers undervalue energy costs.

The literature has focused on a wide range of issues associated with the estimation of η and θ. The seminal paper that started the debate about misperceptions of energy costs is Hausman (1979), which focused on addressing the endogeneity of utilization and purchase decisions. To address this problem, he developed an estimator that jointly modeled the utilization and purchase decisions and found an implicit discount rate of about 20%. Following Hausman (1979), a large number of studies have performed a similar exercise and found a wide range of discount rates, but that all tend to exceed normal rate of return (Train 1985). One shortcoming of Hausman (1979) and other early studies on the topic is that most estimators did not control for unobserved product attributes correlated with E_{ij}. This is in contrast with most recent studies (Busse et al. 2013; Allcott and Wozny 2014; Grigolon et al. 2014; Sallee et al. 2016; **?**), which have paid close attention to unobservables that might bias the estimate of θ. In these studies, rich panel data are exploited where the same technology is sold across different regions and time periods subject to credible exogeneous variation in energy prices or technology features. This allows to use empirical strategies that control for time trends, region specific unobservables, product-specific time invariant characteristics, and selection issue due to the utilization decision. One important point to note is that the recent literature has focused less on the estimation of the parameter η. In fact, in several papers (Busse et al., Lin et al.), the coefficient on price is treated as a free parameter and calibrated using price elasticities found in the literature.

The point we raise is that even if η and θ are credibly estimated, using these two parameters alone to measure the *average degree* of misperception can lead to a biased quantification of misperception.

2.2. **Bias in Measuring Average Misperception**

If the preference parameters for prices and energy costs η and θ vary in the population, the statistic used to measure misperception: $m = \theta/\eta$, is the ratio of two random variables. A closed form solution for the distribution of m exists only for a few specific distributions (e.g., two lognormals), but moments of that distribution can be easily approximated given any distributions of η and θ, as we show below.

Proposition 1. *The first-order Taylor approximation of $E[m] = E[\theta/\eta] \approx E[\theta]/E[\eta]$.*
The second-order Taylor approximation of $E[m]$ is:

$$\frac{E[\theta]}{E[\eta]} - \frac{cov(\eta, \theta)}{E[\eta^2]} + \frac{Var(\eta)E[\theta]}{E[\eta^3]}$$

The approach commonly used to quantify misperceptions of energy costs is the first-order Taylor approximation shown in Proposition 1. The size of the bias of this approximation is therefore of the order:

$$-\frac{Cov[\eta,\theta]}{E[\eta^2]}+\frac{Var[\eta]\cdot E[\theta]}{E[\eta^3]}+\mathcal{O}(n^{-1}) \tag{2}$$

The size and the sign of the correlation between the parameters η and θ play an important role in determining the bias of the first-order approximation. If these parameters are positively correlated, the first-order approximation will be biased upward. From Equation 2, it is also clear that heterogeneity in the perception of price plays a particularly important role in quantifying $E[m]$ as the first three moments of the distribution of η enter the expression. Heterogeneous responses to both prices and energy costs and how these responses are correlated are thus crucial to infer accurately measure misperception. To illustrate, consider this simple example where η and θ are strongly positively correlated because some consumers do not pay attention to prices or energy costs. Suppose that there are only two consumer types who make a purchase decision according to Equation 1, where $\beta_i = \{\eta_i, \theta_i\}$, $i = 1, 2$. A fraction σ of consumers of type 1 only pay attention to quality so that $\beta_1 = \{0, 0\}$ and a fraction $1-\sigma$ pay equal attention to both prices and energy costs, so that $\beta_2 = \{\eta, \eta\}$. In this market, the first-order approximation implies:

$$E[m] \approx \frac{E[\theta]}{E[\eta]} = \frac{(1-\sigma)\eta}{(1-\sigma)\eta} = 1.$$

Clearly, this leads to a misleading conclusion even for the case where the share of inattentive consumers is very large, i.e., $\sigma \approx 1$, we will still conclude that the degree of inattention is low in this market.

3. Data and Environment

Our empirical investigation focuses on the U.S. refrigerator market, which offers several advantages. First and foremost, refrigerators are one of the few appliance categories that consume a large amount of energy, but for which the utilization decision is not likely to be a main driver of the overall energy operating costs over the lifetime of the appliance. Although refrigerator energy costs could be subject to idiosyncratic variation across households, the characteristics of a refrigerator such as it size, door design, presence of ice maker are the main determinants of its energy costs. Endogeneity of the utilization and purchase decisions are therefore not a main concern, which simplifies the estimation. Second, the U.S. refrigerator market is subject to rich variation in refrigerator prices, energy costs, rebates for energy efficient appliances, and choice sets that allow us to identify the preference parameters and infer the degree of misprception to energy costs. Third, refrigerators is an important market in the U.S. and elsewhere, which is expected to grow particularly fast in developing countries in the upcoming decades. Contributing to the design of policies that improve the energy efficiency of refrigerators is thus important to reduce the negative externalities associated with household energy use.

The main data source used for the estimation consists of transaction level data from a large U.S. appliance retailer. The sample includes all transactions where a refrigerator was purchased during the period 2007-2012. We observe each transaction, which contains information about the price paid by the consumer, the zip code of the store where the purchase was made, the manufacturer model number of the model purchased, and a transaction identifier that tracks consumers making multiple purchases. For a large subset of transactions, the identifier is matched with household demographics collected by a data aggregator (Table 1). Detailed attribute information for each manufacturer model number is also available and include: manufacturers' reported energy use, dimensions (width, height, depth), whether a product is certified Energy Star, the presence of ice maker, color, brand, door design, and several other features pertaining to design and technology options.

We match the transaction data with local energy prices and rebate information. Energy prices are constructed from the form 861 of the Energy Information Administration (EIA), which contains revenue and quantity of kWh consumed by residential consumers. Together, these variables provide a measure of average electricity price for each electric utility operating in the U.S. The EIA also provides information about which utility is operating in each county, which allows us to compute average electricity prices at the county level. One important feature of the U.S. electricity market is that prices vary widely across regions due to difference in generation technologies across unit and market structure (Figure 2). Variation in prices over time tend, however, to be modest, with the exception of a few local markets. This contrasts with gas prices that vary widely over time, but much less across regions. For our estimation, we exploit both the persistent differences across regions and the existing within county variation over time.

Rebates for energy efficient appliances were offered during the sample period by both state governments and electric utilities. The State Energy Efficiency Appliance Rebate Program (SEEARP) was funded as part of the stimulus package of the American Recovery Act. This program led to generous rebates for Energy Star certified products during the year 2010 and 2011. Several electric utilities also offer rebates for Energy Star certified-refrigerators. Both of the rebate programs vary across time and regions. Houde and Aldy (2017) find that consumers responded to SEEARP, but a large fraction of consumers were inframarginal to the program. Houde and Aldy (2017) and Datta and Filippini (2016) also find that rebates offered by electric utilities have a modest impact on demand due to a low take-up rate.

One particular feature of the U.S. appliance market is that appliance retailers, such as ours, have a national price policy and retail prices are subject to large and frequent variations. Houde (2016) shows that the price of each refrigerator model at this same retailer is subject to weekly variation that can exceed 20% and that variation is model-specific and not perfectly correlated across brands. Therefore, even after controlling for week-of-sample fixed effects interacted with brand dummies, large variation in price remains (Figure 3). This model-specific variation is highly idiosyncratic and is generated by the retailer's dynamic pricing algorithm. We exploit this variation to identify the coefficient on price.

4. **Empirics Part I: Homogeneous Model**

Our empirical strategy is based on a simple discrete choice model, where utility of consumer i for product j in region r and time t is given by:

$$U_{ijrt} = \eta p_{jrt} + \theta E_{jrt} + \epsilon_{ijrt} \tag{3}$$

where P_{jrt} is the purchase price and E_{jrt} is the annual operating cost (\$/year), and ϵ_{ijrt} represents the unobservable portion of utility. Our goal is to first estimate η and θ without considering heterogeneity in the distribution of preferences.

To obtain our estimating equation, we first assume that ϵ_{ijrt} is extreme value distributed, which gives rise to the multinomial logit:

$$P_{jrt} = \frac{e^{U_{jrt}}}{\sum_{k=0}^{J} e^{U_{krt}}} \tag{4}$$

The number of units sold for a particular product can be obtained by multiplying P_{jrt} by the size of the market in region r and time t: M_{rt}: $M_{rt} \cdot P_{jrt} = q_{jrt}$. The local market size, M_{rt} is defined as all potential buyers of a refrigerator in a given region and time period. The outside option in this context ($j = 0$) refers to the decision to: 1) not purchase a refrigerator in this particular week, 2) purchase a refrigerator at another retailer, or 3) to purchase a refrigerator model at the retailer, which has very low market share, and has been excluded from our final sample. We normalized the utility from the outside option in each market ($r \times t$) to a constant: $U_{0rt} = \alpha_{rt}$. Using the Berry (1994) transformation, we obtain a log-linear expression with the log of the quantity of a product sold as the dependent variable:

$$\begin{aligned}
&log(M_{rt} \cdot P_{jrt}) - log(M_{rt} \cdot P_{0rt}) = U_{jrt} - U_{0rt} \\
&log(q_{jrt}) - log(q_{0rt}) = U_{jrt} - U_{0rt} \\
&log(q_{jrt}) - log(q_{0rt}) = \eta P_{jrt} + \theta E_{jrt} + \epsilon_{ijrt} - \alpha_{rt} \\
&log(q_{jrt}) = \eta P_{jrt} + \theta E_{jrt} + \epsilon_{ijrt} - \alpha_{rt} + log(q_{0rt})
\end{aligned}$$

Our estimation equation is a version of this log-linear expression, where the unobservable component of utility is made up of several sets of fixed-effects as well as an idiosyncratic error term as follows.

$$log(q_{jrt}) = \eta P_{jrt} + \theta E_{jrt} + \gamma_j + \xi_{rt} + Week \times Brand_{jt} + EStar \times State_{jrt} + \varepsilon_{ijrt}. \tag{5}$$

where

$$\epsilon_{ijrt} = \gamma_j + Week \times Brand_{jt} + EStar \times State_{jrt} + \varepsilon_{ijrt}$$

and

$$\xi_{rt} = log(q_{0rt}) + \alpha_{rt}$$

The term ξ_{rt} represents region by time fixed effects. By including these fixed effects in our estimation, we can flexibly control for trends in preferences for the outside option, without having to make explicit assumptions about local market size or the level of utility from the outside option. In our preferred specification, we control for these trends using zip-code by six-month fixed effects. We will also present specifications with zip code month-of-sample fixed effects.

Model-specific fixed effects, γ_j represent time and region invariant preferences for a particular model. Week by brand specific fixed effects, $Week \times Brand_{jt}$, control for trends in the attractiveness of particular brands, which may effect market share. This set of fixed effects is notably important to rule out the effect of promotions and marketing campaigns that might be correlated with price changes.

We also include Energy Star by region specific fixed effects, $EStar \times State_{jrt}$. We do this primarily to ensure that the cross-sectional variation in local electricity prices can be treated as exogeneous. One concern is that preferences for energy efficient products are correlated with local electricity prices. This correlation may arise for a variety of reasons. In regions with high electricity prices, electric utilities, retailers, and local governments may market energy efficient products more, for instance. In the appliance market, the Energy Star certification plays a particularly important role in helping various stakeholders to promote energy efficiency and there is substantial variation across regions in how the program is publicized. Figure 1 shows a map compiled by the U.S. EPA that classifies each designated market area (DMA) has a high, low, or medium (coded as "other" by the EPA) level area of publicity for Energy Star. Regions with the highest electricity prices, such as New England and California, also tend to have a large number of high publicity DMAs. On the other hand, regions with low electricity prices tend to have low level of publicity for Energy Star. To account for the possible effect of Energy Star, and region-specific preferences for energy efficiency more generally, we add Energy Star (ES) dummies interacted with region fixed effects. In our analysis, we will consider both ES-state and ES-county-year fixed effects.

We aggregate our transaction data to the product-zip code-week level and our energy cost data for a particular product vary by county-month (the level of variation of our average energy price data). For the estimation, we compute E_{jrt} as the annual operating energy cost of product j, which is the annual kWh consumption reported by the manufacturer multiplied by the county average electricity price in a particular month. Not all products sell every week in every zip code, so there are a large number of zero sales in our dependent variable.

For our preferred specification, we estimate Equation 5 with a Poisson regression model. The Poisson estimation has several advantages in our context. First, since we are estimating a log-linear relationship, it is better to model the conditional expectation of the dependent variable directly rather than applying a transformation to the dependent variable. Simply taking the log of sales counts (y) is not possible, since the value is zero for a non-trivial fraction of our data. We could apply other transformations that are defined for all $y \geq 0$, such as $log(1+y)$, but it is not obvious how to recover $E(y|\mathbf{x})$ from a linear model for $E(log(y+1)|\mathbf{x})$. Second, like OLS, which is consistent and aymptotically normal even if the normality assumption is violated, Poisson has the nice property that quasi-maximum likelihood estimation recovers consistent, asympotically normal coefficient estimates even if the Poisson distribution does not hold and standard errors can be adjusted for violations of the Poisson variance assumption that the variance is equal to the mean (Wooldridge 2010). Third, we are able to efficiently estimate models with high-dimensional fixed effects using the algorithm proposed by (**?**). Unlike other count models, e.g., negative binomial, the Poisson model is very computationally efficient with this algorithm. One down-side of directly modeling $E(y|\mathbf{x})$ is that the identification of the parameters is not as transparent as with OLS, since it is being driven in part by the non-linear functional form assumptions. However, as we show in Appendix XX our results are qualitatively consistent with an OLS estimate of both a linear probability model and a linear model with $log(1+y)$ as the dependent variable.

Measurement Error. Since we constructed our county-level electricity price as an average price for the utilities that serve a particular county, there will be some measurement error in the calculated operating costs for individual purchases. Another potential source of measurement error is that we use the average county price of the county in which the appliance was purchased, which may differ in some cases from the county in which the consumer lives. This type of classic measurement error could attenuate our estimates of the coefficient on operating cost, and bias our estimates of the "perception of energy costs" parameter towards zero. In order to address this possibility, we estimate our model limiting our sample to those counties that are served by only one utility, so that our average energy price is not subject to this type of measurement error for everyone living in that county. If the measurement error were a significant biasing factor in our estimates, we would expect the coefficient estimates on operating cost to increase in magnitude using this sub-sample.

Another potential source of measurement error results from our calculation of average electricity price for a utility as residential revenue divided by residential sales. In states where there is retail competition, incumbent utilities earn revenue for providing distribution services, even when they are not directly selling electricity to a customer. In these cases, our estimation of average utility price may be systematically biased upwards, since not all of the revenue is directly tied to sales. If the effect of this type of non-classical measurement error were large, it would bias our estimates of the "perception of energy costs" upwards, away from under-valuation. In order to address this possibility, we estimate a model limiting our sample only to states without retail competition. [1]

Interpretation of Model Parameters. The coefficient on annual energy cost, θ, reflects how future energy operating costs affect the probability of purchase. Assuming that consumers form time-invariant expectations about the annual operating electricity expenditure and do not account for the effect of depreciation, the lifetime energy operating cost (LC_j) for the durable j is given by:

$$LC_{jrt} = \sum_{t=1}^{L} \rho^t E_{jrt} = \rho \cdot \frac{1-\rho^L}{1-\rho} \cdot E_{jrt},$$

where L is the lifetime of the durable, $\rho = 1/(1+r)$ is the discount factor, and E_{jrt} is the product of the electricity price paid by a household in region r at time t and the manufacturer's expected annual electricity consumption for durable j.[2] In the estimation model specified by Equation 5, the coefficient on electricity cost is a reduced form parameter that relates to the discount factor and marginal utility of income as follows:

[1]The following states (and the District of Columbia) have retail competition and are omitted for this part of the analysis: Texas, Illinois, Ohio, Pennsylvania, Maryland, Washington DC, Delaware, New Jersey, New York, Rhode Island, Connecticut, Massachusetts, New Hampshire, and Maine.

[2]E_{jrt} maps one-to-one to $Elec_{jrt}$ based on the region of residence for the household.

$$\theta = \eta \cdot \rho \cdot \frac{1-\rho^L}{1-\rho}. \tag{6}$$

The estimates of η and θ can then be used to infer a value of an implicit discount rate r. Alternatively, we can take a stand on the what is the appropriate normal rate of return for the consumers in the sample, use it to compute ρ, and test if the ratio $\theta/\eta \cdot \rho \cdot \frac{1-\rho^L}{1-\rho}$ is equal to one.

Results: Main Specification. Table 2 displays the results from our main estimation. The first column estimation includes zip code by six-month fixed effects and Energy Star by State FE along with product fixed effects and brand by week fixed effects. The estimation in column 2 is the same as column 1 except for zip code specific trends at the zip code by month level rather than zip code by 6 months, and the estimation in column 3 is the same as column 1 except for Energy Star by county by year fixed effects rather than Energy Star by state fixed effects. The results are consistent across all three of these specifications. We find that a $1 increase in purchase price corresponds to a .2% change in purchase probability and a $1 increase in annual operating cost corresponds to a 1.6% change in purchase probability. Using Equation 6, we transform these estimates into an applied discount rate reported in the last row of Table 2. The results imply a 8.5% to 9.5% discount rate, which is consistent with recent work on automobile purchases (Busse et al. 2013; **?**; Sallee et al. 2016; Grigolon et al. 2014) and housing purchases (**?**).

Results: Robustness to Measurement Error in Electricity Price. Table 3 displays the results from our estimates limiting our sample in ways that will minimize measurement error. Column 1 displays the results from the estimation of 5 using just those counties served by a single utility and column 2 displays the results of the same estimation using just those states with no retail competition. The results are consistent with the results using all counties in all states. The results from the sub-sample with counties served by just one utility also suggest that a $1 increase in annual operating cost results in a 1.6% change in purchase probability. This suggests that attenuation bias from classical measurement error in defining local energy price is not significantly driving down our estimates of θ. The results from the sub-sample with states with no retail competition suggest that a $1 increase in annual operating cost results in a 1.7% change in purchase probability. Since the point estimate of the coefficient on energy operating cost is slightly higher than in the main specification rather than substantially lower, it does not seem like a systematic upward bias in energy price from retail competition states is driving our moderate discount rate estimates.

5. Empirics Part II: Recovering Heterogeneity

To assess the bias due to the first-order approximation, we need to characterize the joint distribution of the parameters η and θ. In the literature pertaining to energy efficiency durables, it is common to account for preference heterogeneity using random coefficient models, especially the random coefficient logit with normally distributed coefficients. Proposition 1, however, suggests that the normal distribution is not well suited to evaluate the bias as its symmetry imposes a skewness that is close to zero. The coefficient of skewness of the distribution of the parameter η matters in evaluating the bias and setting it to zero may lead to overvalue or undervalue the size of the bias.[3] In sum, to accurately characterize the bias it is important to recover the joint distribution of η and θ without imposing distributional assumptions that determine the high order moments.

Our strategy to estimate the joint distribution of η and θ consists of using the estimator proposed by Fox, Kim, Ryan, and Bajari (2010) (FKRB, therafter), which recovers a fully non-parametric distribution of heterogeneity. Using this non-parametric distribution, the quantity of interest $E[\frac{\theta}{\eta}]$ can be exactly computed and compared to its first-order approximation.

The FKRB estimator relies on the intuition that a continuous distribution of random parameters can be approximated by a discrete distribution defined over the discretization of the support of the continuous distribution. This suggests a simple estimator that consists of first discretizing the support of the random parameters into a large number of grid points, say K, evaluating the choice model at each grid point k, and then integrating over the discrete distribution by simply weighting and summing the choice model at each grid point. The estimator returns the weights, which correspond to the discrete probability density function of the random parameters. To illustrate, consider our setting where we are interested in evaluating the joint distribution of $F(\eta, \theta)$. Discretize the support of η and θ into K grid points, and define $\beta^k = \{\eta^k, \theta^k\}$ as one grid point. We

[3]The expression $E[\eta^3]$ in Equation 2 can be expressed as a function of the coefficient of skewness: $E[\eta^3] = Skew[\eta] \cdot Var[\eta]^{3/2} + 3E[\eta] \cdot Var[\eta] - E[\eta]^3$. Given that the coefficient of skewness can be positive or negative, setting $Skew[\eta]$ to zero has an ambiguous effect on the estimate of $E[\eta^3]$ and thus the overall expression for the bias.

can compute the choice model for each β^k using a parametric model (e.g., the conditional logit), where the probability of choosing product j given β^k is noted $P_j(\beta^k) = P_j^k$. The observed choice probability, P_j, is then equal to:

$$P_j = \int P_j(\eta, \theta) dF(\eta, \theta) \approx \sum_k^K \alpha^k P_j^k \tag{7}$$

where $\sum_k^K \alpha^k = 1$ because the weights are a discrete probability density function. By choosing a parametric form for the choice model, each P_j^k can be first be computed for each grid point and treated as data in the estimation. The estimator is thus semi-parametric and the estimation can simply proceed by running a linear regression with P_j as the dependent variable, P_j^k as regressors, and α^k as coefficient to be estimated. To ensure that the weights α^k sum to one, constrained linear least squares must be used with the constraint: $\sum_k^K \alpha^k = 1$. The estimator is appealing for a number of reasons. First, it is a computationally simple way to recover a fully non-parametric distribution of heterogeneity. Second, constrained linear least square is guaranteed to provide a global optimum if a solution exists, a non-trivial advantage over other estimators of random coefficients models, such as maximum likelihood, the EM algorithm, GMM, or bayesian methods, which are all prone to local optima and convergence issues.

The main weakness of the FKRB estimator is that it suffers from the curse of dimensionality. The number of grid points increases exponentially with the number of random coefficients. For instance, a model with three random coefficients discretized with hundred grid points in each dimension must be evaluated for a combination of 100^3 grid points. If the number of preference parameters is large, it becomes rapidly intractable to model each parameter with a random coefficient. FKRB propose two solutions. Some parameters can be estimated in a model without heterogeneity in a first step and then treated as data when estimating the weights α_k. For instance, a simple conditional logit can be estimated in a first step, and then the FKRB can be implemented where only a subset of the parameters are treated as heterogeneous. This approach is particularly appealing with models with a large number of fixed effects. It should, however, be noted that this two-step estimation may not produce consistent estimates and standard errors need to be adjusted in the second step. The second approach proposed by FKRB is estimate the model in one step with non-linear constrained least-squares. The feasibility of this estimator is context specific. Consider the case where the number of grid points is 1,000 and it takes 1 second to evaluate the choice model for each grid point, computing the choice probability will then take 1,000 seconds, which will result in a very long computing time for a non-linear optimization.[4] The non-linear optimization is also not guaranteed to converge to a global optimum.

For our application, the number of preference parameters is very large due to the number of product fixed effects (> 500). Moreover, the choice probabilities take a relatively long time to compute given the large number of observations used for the estimation. We thus favor the two-step estimator discussed above. We implement the estimator as follows.

From the entire sample of transactions, we take a large random subsample ($N = 66,000$). As with our homogeneous estimation, the subsample is restricted to transactions made by households owning their housing unit[5] with the goal of focusing on transactions made by consumers who are likely to pay the energy operating costs of their appliances.

For each transaction, we inter a zip code-trimester-specific choice set, i.e., all models offered in a given zip code[6] and during a trimester are considered to be in the consideration set of the consumer.[7] The parametric choice model is the conditional logit with alternative-specific utility given by

$$U_{ijrt} = \eta P_{jrt} + \theta E_{jrt} + \tau ES_{jt} + \phi Rebate_{jrt} + \gamma_j + Demo_i \times Att_t + \epsilon_{ijrt} \tag{8}$$

[4]The computing time can be greatly reduced by using parallel processing to compute the choice model over all the grid points.

[5]The data do not explicitly identify transactions that are made by households. We infer this information using a transaction identifier that tracks multiple purchases of customers. We classify customers that purchase more than one refrigerator during the sample period as non-households. This criterion is a conservative way to rule out contractors and other entities that buy a large number of appliances in bulk.

[6]##% of zip codes have only one store. Our choice set are thus mostly store specific.

[7]We do not observe floor inventory. Therefore, a model is deemed to be offered if we observe at least one sale of that model at a given location and time period.

Relative to the homogeneous model, the model has a parsimonious set of controls: the energy star certification, energy star rebates, product fixed effects, and demographic information interacted with a subset of attributes. But we show that the model does well in replicating the estimates of Section 4 and is not subject to a large bias. To ease computations, only the parameters η and θ are random coefficients. We first estimate all coefficients with a conditional logit. In the second step, we fix all the regressors, except η and θ. The grid points for η and θ are determined by scaling the estimates of the conditional logit with various scaling factors ranging from -1 to 3. This means that we allow the maximum of the support of each parameter to be three times as large the estimated mean, and the minimum to take a positive value, although η and θ should be negative to be consistent with consumer optimization.

5.1. **Results**

Figure 4 shows the estimated probability density of the ratio θ/η. The distribution has three modes and shows that a large share of consumers respond to energy costs, but undervalue them. There is some probability mass close to zero, which suggests that some consumers do not pay attention to energy costs. It also appears that some consumers overvalue energy costs. This finding might appear surprising, but is consistent with other evidence for the U.S. appliance market. For instance, Houde (2016) shows that a fraction of consumers value Energy Star certified products well beyond the energy savings associated with the certification. Newell and Siikamäki (2014) also found that a coarse certification such as Energy Star can lead to overvaluation of energy costs.

The red and black horizontal lines in 4 show the bias in evaluating $E[\theta/\eta]$. According to the first-order approximation, the value of $E[\theta]/E[\eta] = \sum_k \alpha_k\theta_k/\sum_k \alpha_k\eta_k$ is 0.66 (red line), but if we evaluate $E[\theta/\eta]$ exactly by computing $\sum_k \alpha_k\theta_k/\eta_k$, we find a value of 0.75 (black line). Therefore, heterogeneity in misperception of energy costs is such that the first-order approximation leads to a downward bias in evaluating $E[\theta/\eta]$.

Another important take-away from the above results is that the quantity $E[\theta/\eta]$, even when computed accurately, is also misleading in diagnosing the degree of misperception of energy costs. For instance, we find that there is no probability mass at $E[\theta/\eta] = 0.75$. In other words, the distribution is highly skewed and there are no consumers that behave as the average consumer. If we were to design a policy solely based on the average degree of misperception and overlook the substantial heterogeneity in how consumers value energy costs, we would set a policy that is too stringent for some and too lenient for others.

There are many factors that drive the heterogeneity patterns found in Figure 4. For instance, households with different income levels might have different levels of credit constraints, which would lead to a different response to energy costs. We explore this hypothesis, by recovering the non-parametric distribution of η and θ controlling for income. We do so by estimating the FKRB estimator for six different income groups. Figure 5 shows the results, where each panel corresponds to the probability density for a specific income group. There are three important results. First, all six distributions follow a similar pattern to the one we found using the whole sample. This suggests that the existence of the three consumer types that we identify is not driven by access to credit or other factors related to income. Second, the first-order approximation is downward biased in all cases. Third, the ratio m is increasing with income. Note that the ratio is computed using a discount rate of 3% for all income groups. If we were to apply higher discount rates for lower income groups the average value of m would be constant across groups. Therefore, credit constraints do appear to play a role in the present context. However, even if we were to adjust the discount rates to reflect access to credit, substantial heterogeneity in the distribution of m would persist, which we take as evidence that misperceptions of energy costs are important in this market.

6. Implications for Policy Design

In this section, we assess how heterogeneity in misperception of energy costs impacts the design of policies used to address negative externalities associated with energy use. To illustrate the role that plays heterogeneity, we first focus on the simple case where the planner uses a single policy instrument, a Pigouvian tax, to address negative externalities and misperceptions. We show how heterogeneous misperceptions affect the level of the optimal Pigouvian tax and the measurement of its welfare effect.

Our measure of welfare is a direct application of Leggett (2002)'s formula to measure welfare in a discrete choice framework in the presence of imperfect information. This framework has been further developed by Allcott (2011), **?**, and Houde (2016) to measure welfare in the presence of consumers' biases. The first step in applying this framework is to take a stand on whether the utility function specified to model consumers' decisions coincides with the utility function that captures what consumers experience after making a purchase decision. If decision and experience utility do not coincide, because we believe that consumers misperceive some components of product costs for instance, the change in consumer surplus for a given policy change can be expressed as:

$$\Delta CS_{ktr} = \frac{1}{\eta_k} \cdot \left[ln \sum_j^J exp(\tilde{U}_{kjtr}) + \sum_j^J \tilde{P}_{kjtr}(\tilde{U}^E_{kjtr} - \tilde{U}_{kjtr}) \right] - \frac{1}{\eta_k} \cdot \left[ln \sum_j^J exp(U_{kjtr}) + \sum_j^J P_{kjtr}(U^E_{kjtr} - U_{kjtr}) \right]. \tag{9}$$

where the terms with a tilde are evaluated after the policy change, U^E_{kjtr} denotes experienced utility and U_{kjtr} corresponds to decision utility for consumer of type k. The expression in 9 corresponds to the standard measure of welfare for the multinomial logit (**?**) plus the term $\sum_j^J P_{kjtr}(U^E_{kjtr} - U_{kjtr})$, which we refers as the Leggett (2002)'s correction. The correction term arises because of the discrepancy between what consumers perceive they will experience and what they actually experience, and simply represents the expected (private) cost that consumers incur because of their misperceptions.

For our application, we will make the following two assumptions:

(1) If $m_k = \theta_k/\eta_k \neq 1$, consumer of type k misperceived energy costs and his decision utility differs from experience utility. To measure experience utility for type k, we set $\theta_k = \eta_k$.
(2) If $\eta_k > 0$, consumer of type k misperceived the product price and decision utility differs from experience utility. To measure experience utility for type k, we set $\eta_k = \bar{\eta}_{-k}$, where $\bar{\eta}_{-k}$ refers to the average value for the coefficient η for all types other than k.

The first assumption simply says that if consumers do not perceive a one dollar change in future energy operating costs (discounted with a normal rate of return) the same way they perceive a dollar change in product price, they are prone to a bias. Note that by setting $\theta_k = \eta_k$, we let consumers to be heterogeneous with respect to their response to price (i.e. marginal utility of income). The second assumption aims to deal with consumer types for which η is greater than zero. A positive coefficient on price is a manifestation of various consumers' biases with respect to the product price and notably high search costs. By setting $\eta_k = \bar{\eta}_{-k}$, we make the implicit assumption for all types $j \neq k$ for which $\eta_j \leq 0$, that the coefficient on price reflects their true marginal utility of income. Furthermore, we also implicitly assume that the distribution of the true marginal utility of income for consumers $\eta_k > 0$ follows the same distribution as the remaining of the population.[8]

Using the welfare measure in 9, Houde and Aldy (2017) derive an expression for the optimal Pigouvian tax in the presence of misperception of energy costs. Suppose that there are K consumer types and that the misperception parameter m_k applies to the energy costs inclusive of the tax. They show that if each product j consumes e_j amount of energy over its lifetime, consumers pay p_e for each unit of energy consumed, and e_j is associated with a constant marginal damage cost, ϕ, the optimal Pigouvian tax, τ^*, is:

$$\tau^* = \frac{\phi}{1-\mathcal{A}} + p_e \frac{\mathcal{A}}{1-\mathcal{A}} \tag{10}$$

with

$$\mathcal{A} = \sum_k \alpha_k (1-m_k) \sum_j \frac{\partial P_j^k}{\partial \tau} e_j$$

The expression 10 is similar to the results of Farhi and Gabaix (2016) except that the price of energy also enters the expression of the optimal Pigouvian tax. Allcott and Taubinsky (2014) also show that the optimal Pigouvian tax should be adjusted to account for biases, and it should be an upward adjustment if consumers undervalue the return to energy efficient investments.[9] The intuition in having the price of energy in Equation 10 is that there are two market failures at play, negative externalities and misperceptions, but only one policy instrument is used. In the presence of consumers' biases, the Pigouvian tax is second-best and multiple policy instruments are required

[8]**?** propose an alternative approach to deal with consumers' biases pertaining to product price. They first devise a procedure to identify choices that are considered a mistake in their sample. They tend exclude these choices from the sample, and estimate a choice model on what they consider a sample free of consumers' mistakes. They treat the coefficient on price from this estimation as the true marginal utility of income.

[9]It can be shown that Proposition 1 of Allcott and Taubinsky (2014) yields a similar expression to Equation 10, except that their expression for the optimal tax is also function of the degree of misperception in future utilization of the durable, in addition of the price of energy, and misperception m_k.

to achieve the first best outcome (?). Our goal here is to show how heterogeneity in misperception impacts the level of the optimal tax. We leave the question of the first-best policy design for future research. We compute the optimal Pigouvian tax under three scenarios: considering the full distribution of heterogeneity in m_k, considering the average value of m computed with the first-order approximation, and considering the exact value of $E[m]$. For each of those scenarios, we provide two measures of welfare: one that considers heterogeneity and the other evaluated for the average consumer (homogeneous model).

6.1. Results

For all scenarios, we fix the level of the externality at 0.02 $/kWh, this provides the benchmark for the level of the optimal Pigouvian tax without misperception. If we consider the full distribution of heterogeneity in misperception, we found that the optimal tax is approximately equal to zero. This result is driven by the fact that we found a small share of consumers that overvalue energy costs, which requires a downward adjustment to the tax. If we set the tax using the average value of m computed using the first-order approximation, the optimal tax is $\tau =$0.087 $/kWh, more than four times the externality cost. If we use the exact measure of $E[m]$, the optimal tax is $\tau =$0.063 $/kWh. Accounting for heterogeneity has thus a very dramatic effect on the level of the optimal tax.

Looking at the welfare estimates (Table 4), the level of the tax has a large impact of the various components of welfare, especially consumer surplus and government revenues. Accounting for heterogeneity in measuring welfare has however a modest effect on the estimates. Heterogeneity is thus very important in getting the stringency of the policy instrument, but much less for measuring welfare for a given policy.

7. Conclusion

The standard test of consumer misperception used widely in the literature compares the responsiveness of demand for changes in potentially misperceived aspects of cost against salient, correctly perceived aspects of cost. Consumers should be indifferent between an additional dollar of purchase price and an additional dollar of the potentially misperceived cost such as shipping and handling or, the present discounted dollar of energy expenditure, since total lifetime cost should be the relevant metric. The ratio of the responsiveness coefficients has been used as a sufficient statistic for the degree of consumer misperception.

This paper makes two methodological contributions for quantifying inattention. First, we show that ignoring heterogeneity in the distribution of the two responsiveness coefficients can lead to biased estimates of average misperception. The ratio of the two responsiveness coefficients is a first order approximation of the relationship between the two distributions. Taking into account higher order terms in a Taylor series approximation shows that the covariance between the two variables is an important parameter, which determines the direction of the bias. Second, we show that even if the average amount of misperception is measured correctly, heterogeneity in the degree of misperception can have significant implications for both optimal policy design and estimates of welfare effects.

We also contribute to the literature on the "energy efficiency gap", the observation that consumers seem to systematically undervalue operating costs relative to purchase price in energy using durables. In recent years, this issue has received attention because governments around the world have become interested in designing successful policy instruments for reducing greenhouse gas (GHG) emissions. The effectiveness of price-based instruments such as taxes or cap-and-trade programs depends crucially on whether consumers are responsive to fuel prices in markets for energy-using durables. Previous work in the appliance sector has shown mixed evidence as to whether consumers respond to local energy costs at all. Using a unique administrative data set with micro-level sales data from a large appliance retailer, we are able to recover more credible estimates of the coefficients on price and energy costs than has been possible in the past. Our empirical strategy exploits fine-grained panel data, allowing us to control for product, region, and time specific unobservables.

We find that the first-order approximation provides a downward bias of the average degree of misperception. Our results show that consumers are responsive to changes in local energy costs, even if they do undervalue them somewhat relative to purchase price. Our findings also show substantial heterogeneity in the degree of misperception, which significantly impact optimal policy design and estimates of welfare effects.

References

Allcott, H., “Consumers’ Perceptions and Misperceptions of Energy Costs,” *American Economic Review: Papers and Proceedings*, 2011, *101* (3), 98–104.

——— **and N. Wozny**, “Gasoline Prices, Fuel Economy, and the Energy Paradox,” *The Review of Economics and Statistics*, 2014, *96* (5), 779–795.

Busse, M., C. Knittel, and F. Zettelmeyer, “Are Consumers Myopic? Evidence from New and Used Car Purchases,” *American Economic Review*, 2013, *103* (1), 220–256.

Chetty, R., A. Looney, and K. Kroft, “Salience and Taxation: Theory and Evidence,” *American Economic Review*, 2009, *99* (4), 1145–1177.

Datta, S. and M. Filippini, “Analysing the Impact of ENERGY STAR Rebate Policies in the U.S.,” *Energy Efficiency*, 2016, *9*, 677–698.

Dubin, J., “Market Barriers to Conservation: Are Implicit Discount Rates too High?,” *Proceedings of the POWER Conference on Energy Conservation, University wide Energy Research Group, University of California, Berkeley*, 1992, *June 26, 1992.*

——— **and M. McFadden**, “An Econometric Analysis of Residential Electric Appliance Holdings and Consumption,” *Econometrica*, 1984, *52* (2), 345–362.

Farhi, E. and X. Gabaix, “Rare Disasters and Exchange Rates,” *Quarterly Journal of Economics*, 2016, *131* (1), 1–52.

Finkelstein, A., “E-ztax: Tax Salience and Tax Rates,” *Quarterly Journal of Economics*, 2009, *124*, 969–1010.

Goldberg, P., “The Effects of the Corporate Average Fuel Economy Standards in the U.S.,” *Journal of Industrial Economics*, 1998, *46*, 1–33.

Grigolon, L., M. Reynaert, and F. Verboven, “Consumer Valuation of Fuel Costs and the Effectiveness of Tax Policy: Evidence from the European Car Market,” *Working Paper: University of Leuven*, 2014.

H., Mullainathan S. Allcott and D. Taubinsky, “Energy Policy with Externalities and Internalities,” *Journal of Public Economics*, 2014, *112* (April), 72–88.

Hausman, J., “Individual Discount Rates and the Purchase and Utilization of Energy Using Durables,” *Bell Journal of Economics*, 1979, *10* (1), 220–225.

Hossain, T. and J. Morgan, "Plus Shipping and Handling: Revenue (Non)Equivalence in Field Experiments on eBay," *Advances in Economic Analysis and Policy*, 2006, *6* (2).

Houde, S., "Consumers' Response to Quality Disclosure and Certification: An Application to Energy Labels," *Working Paper: University of Maryland*, 2016.

——— **and J. Aldy**, "Consumers' Response to Sate Energy Efficient Appliance Rebate Programs," *forthcoming: American Economic Journal: Economic Policy*, 2017.

Jacobsen, G., "Do Energy Prices Influence Investment in Efficiency? Evidence from Energy Star Appliances," *Journal of Environmental Economics and Management*, 2015, *74* (C), 94–106.

Leggett, Christopher G, "Environmental valuation with imperfect information the case of the random utility model," *Environmental and Resource Economics*, 2002, *23* (3), 343–355.

Loewenstein, George, Joelle Y Friedman, Barbara McGill, Sarah Ahmad, Suzanne Linck, Stacey Sinkula, John Beshears, James J Choi, Jonathan Kolstad, David Laibson et al., "ConsumersâĂŹ misunderstanding of health insurance," *Journal of Health Economics*, 2013, *32* (5), 850–862.

Rapson, D., "Durable Goods and Long-run Electricity Demand: Evidence from Air Conditioner Purchase Behavior," *Journal of Environmental Economics and Management*, 2014, *68* (1), 141–160.

Sallee, J.M., S. West, and W. Fan, "Do Consumers Recognize the Value of Fuel Economy? Evidence from Used Car Prices and Gasoline Price Fluctuations," *Journal of Public Economics*, 2016, *135*, 61–73.

Taubinsky, D. and A. Rees-Jones, "Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment," *NBER Working Paper*, 2016, (Number 22545).

Train, K., "Discount rates in consumers' energy-related decisions: A review of the literature," *Energy*, 1985, *10* (12), 1243–1253.

Wooldridge, J., *Econometric Analysis of Cross Section and Panel Data, Second Edition*, MIT Press, 2010.

8. Tables and Figures

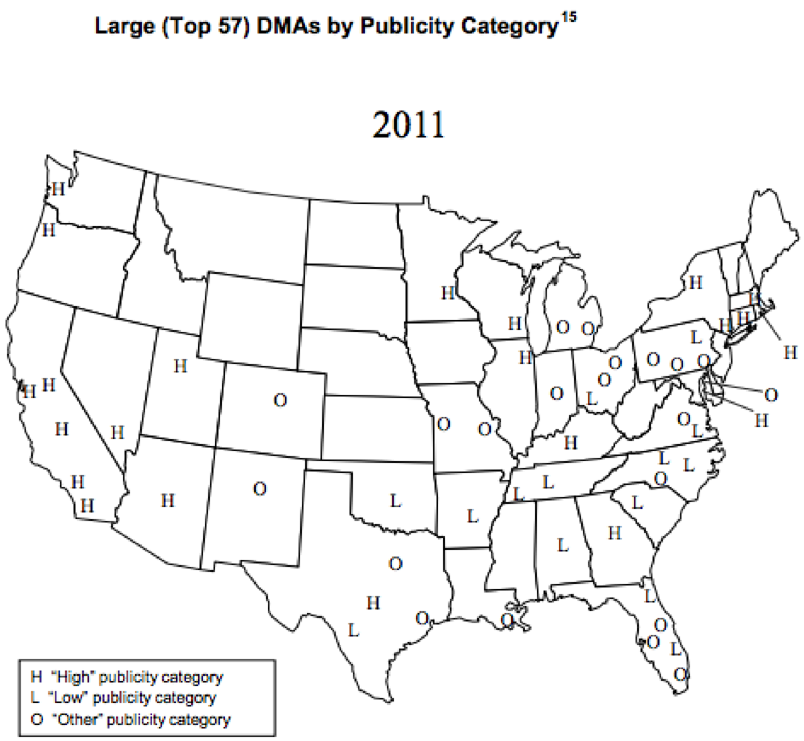

FIGURE 1. Source: EPA

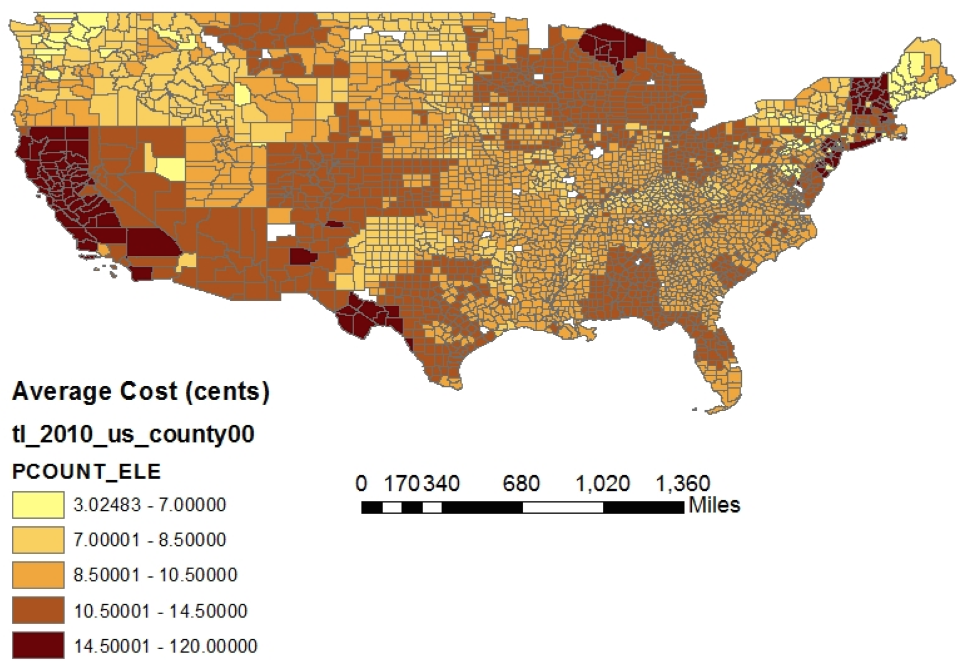

Figure 2. Average County Electricity: 2010

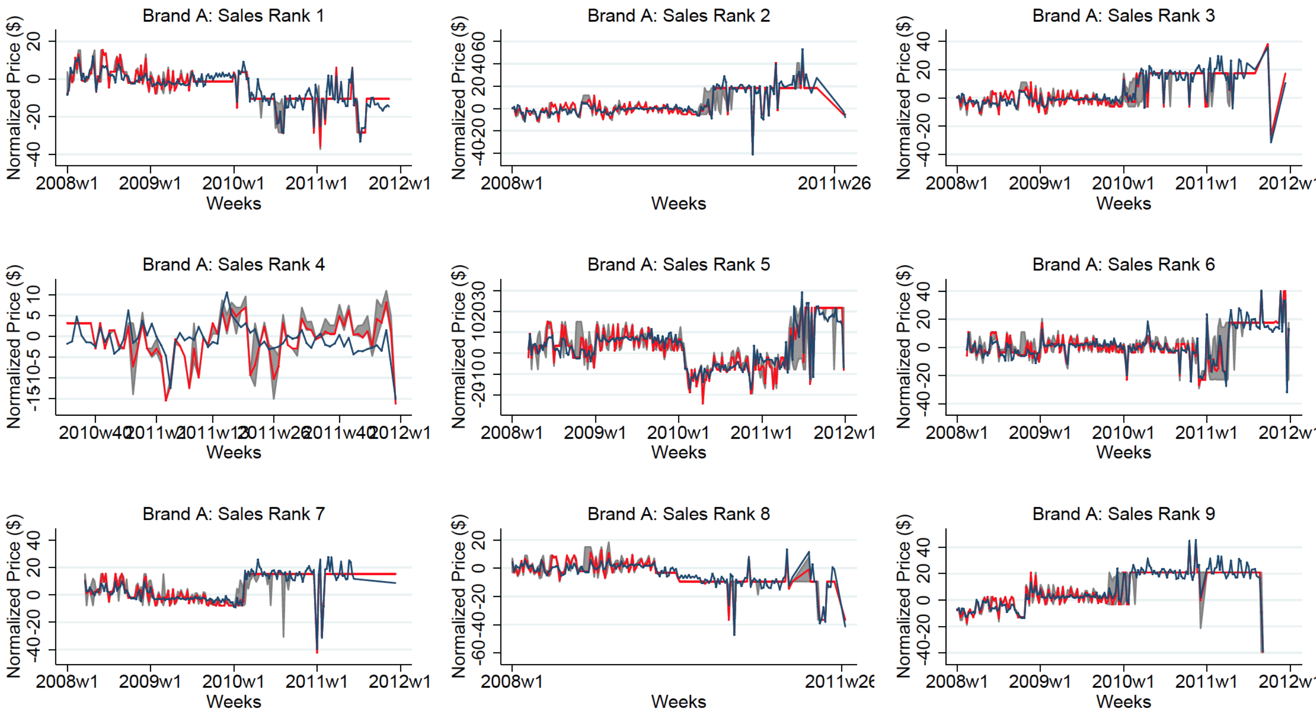

FIGURE 3. Price Variation

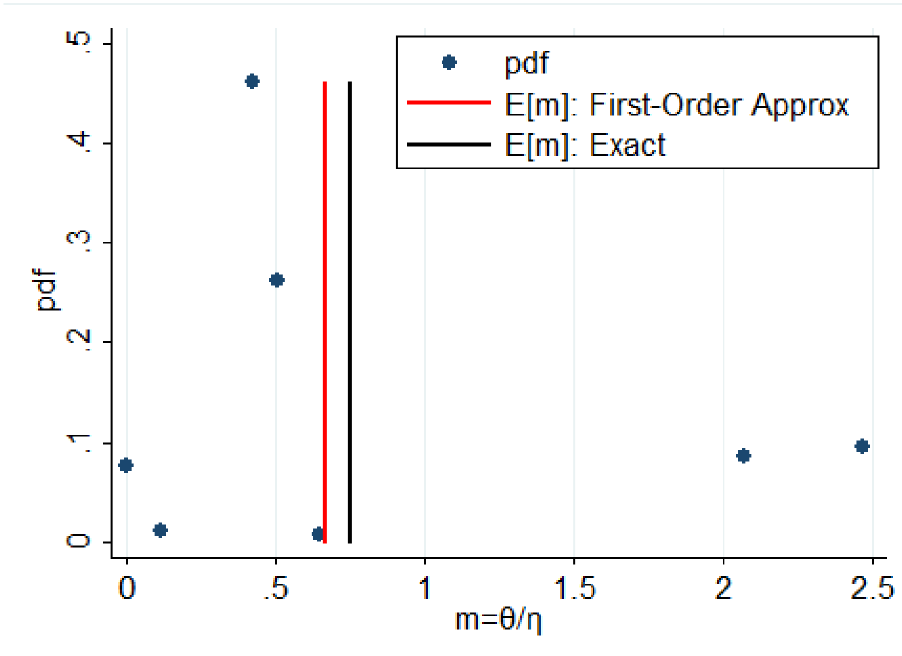

(a) Misperception Parameter

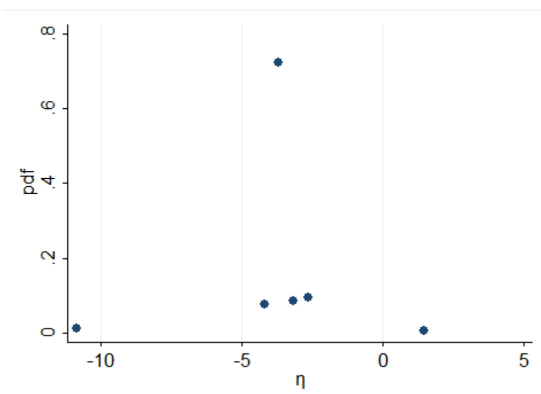

(b) Coefficient on Price

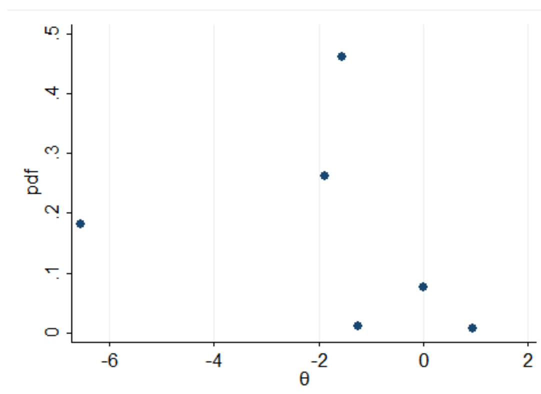

(c) Coefficient on Energy Costs

FIGURE 4. FKRB Joint and Marginal Probability Densities for Parameters η and θ

Notes: The first panel plots the estimated probability density for the ratio θ/η. The distribution identifies three consumer types and shows that the first-order approximation leads to a downward bias in evaluating $E[m]$. Panels b) and c) are the estimated marginal distributions for the parameters η and θ, respectively.

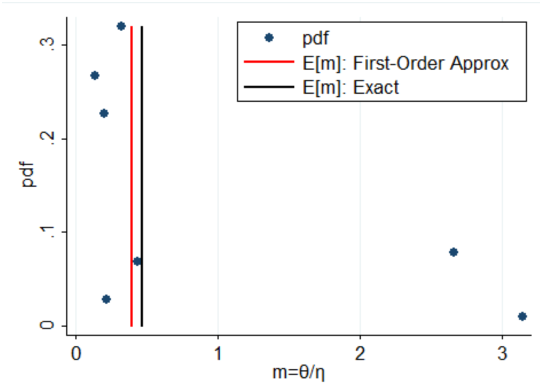

(a) <$30k

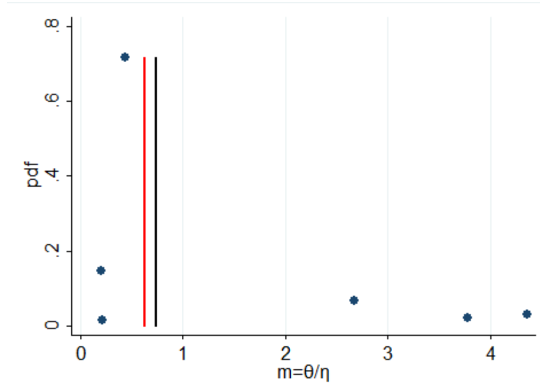

(b) $30k-$50k

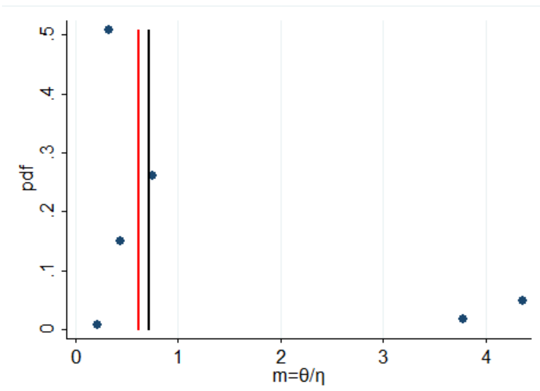

(c) $50k-$75k

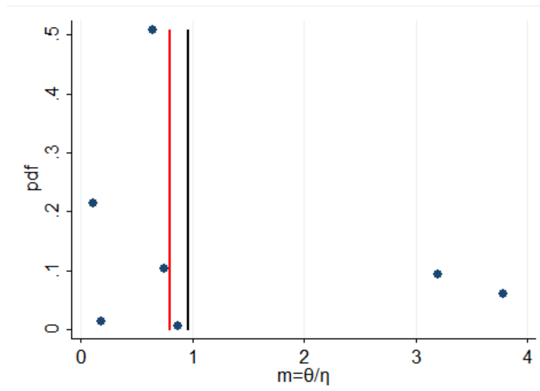

(d) $75k-$100k

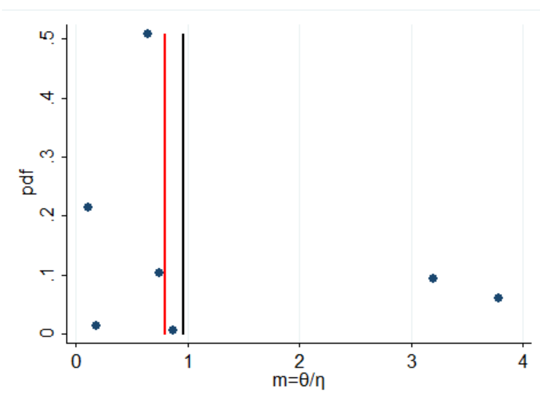

(e) $100k-$150k

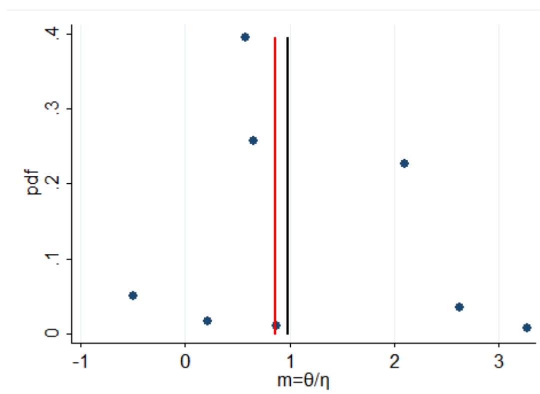

(f) >150k

FIGURE 5. FKRB Probability Density Estimation by Income Group

Notes: Each panel plots the estimated probability density for the ratio θ/η. For all income groups, we compute the ratio assuming a discount rate of 3%, $r = 3\%$.

Table 1. Summary Statistics

	Mean	SD
Price ($)	1252.6	627.0
kWh/y	514.9	78.4
County Elec. Price (cents)	11.4	3.7
State Elec. Price (cents)	12.3	3.3
County Elec. Cost/y ($)	58.5	20.6
State Elec. Cost/y ($)	63.2	18.9
Rebate Amount ($)	25.9	68.8
% Energy Star	68.5	
% w Ice-Maker	76.0	
Overal Size (cu. ft.)	22.5	3.4
% w Top Freezer	30.3	
Demographics		
% of Households	67.6	
% w. Demo. Info.	56.6	
% Renters	1.9	
Income distribution		
<$30k	12.2	
$30k-$50k	16.8	
$50k-$75k	25.2	
$75k-$100k	18.2	
$100k-$150k	11.8	
>$150k)	15.7	

TABLE 2. Estimation of the Effect of Purchase Price and Operating Cost on Sales Count

Dependent Variable	sales count	sales count	sales count
purchase price	-0.00190(***)	-0.00196(***)	-0.00186(***)
	(0.00002)	(0.00002)	(0.00002)
annual electric cost	-0.01596(***)	-0.01687(***)	-0.01667(***)
	(0.00067)	(0.00076)	(0.00077)
Product FE	Yes	Yes	Yes
Brand × Week FE	Yes	Yes	Yes
Zip × 6 month FE	Yes	No	Yes
Zip × month FE	No	Yes	No
EStar × State FE	Yes	Yes	No
EStar × County × Year FE	No	No	Yes
N	1090865	1090865	1090865
Implied Discount Rate	9.6%	9.3%	8.6%

Notes: The dependent variable is the number of units of a particular appliance sold in a given week in a given zip code. The model is estimated according to a Poisson regression, using the algorithm proposed by Guimares and Portugal (2009) to absorb high dimensional fixed effects. The Standard errors are in parentheses. ***, ** and * denote statistical significance at the 1, 5 and 10 percent levels.

Table 3. Estimation of the Effect of Purchase Price and Operating Cost on Sales Count: Sample Limited to Address Measurement Error

Dependent Variable	sales count	sales count
Sub-sample	(counties served by one utility)	(states with no retail competition)
purchase price	-0.00155(***)	-0.00192(***)
	(0.0000467)	(0.0000257)
annual electric cost	-0.0160(***)	-0.0174(***)
	(0.000985)	(0.000927)
Product FE	Yes	Yes
Brand × Week FE	Yes	Yes
Zip × 6 month FE	Yes	Yes
EStar × State FE	Yes	Yes
N	136471	618358
Implied Discount Rate	6.6%	8.5%

Notes: The dependent variable is the number of units of a particular appliance sold in a given week in a given zip code. The model is estimated according to a Poisson regression, using the algorithm proposed by Guimares and Portugal (2009) to absorb high dimensional fixed effects. The Standard errors are in parentheses. ***, ** and * denote statistical significance at the 1, 5 and 10 percent levels.

Table 4. Welfare Effects: Pigouvian Tax

	No Adjustment τ =0.02 \$/kWh	Adjustment w. First-Order Approx. τ =0.087 \$/kWh	Adjustment w. Exact $E[m]$ τ =0.063 \$/kWh
Δ CS w. Heterogeneity	-140.3	-604.8	-441.8
Δ CS Homogeneous	-137.4	-589.1	-431.1
Δ Externality w. Heterogeneity	-0.9	-3.6	-2.7
Δ Externality Homogeneous	-0.9	-3.6	-2.7
Δ Gvt Revenue w. Heterogeneity	138.7	591.2	433.6
Δ Gvt Revenue Homogeneous	139.0	592.2	434.3
Δ SW w. Heterogeneity	-0.7	-9.9	-5.5
Δ SW Homogeneous	2.5	6.7	5.9

Notes: Social Welfare (SW) is the sum of the consumer surplus (CS), externality costs, and government (GVT) revenues. The externality cost is 0.02 $/kWh in all scenarios. All welfare estimates are measured in dollar and compare to the case where there is no tax levied, which is the optimal policy if we account for heterogeneity in determining the level of the tax.