NBER WORKING PAPER SERIES

HETEROGENEOUS (MIS-) PERCEPTIONS OF ENERGY COSTS: IMPLICATIONS FOR MEASUREMENT AND POLICY DESIGN

Sébastien Houde
Erica Myers

Working Paper 25722
http://www.nber.org/papers/w25722

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2019

We would like to thank Meredith Fowlie, Christopher Knittle, Richard Newell, Ryan Kellogg, Frank Wolak, Thomas Colvert, Kochiro Ito, Anna Alberini, Akshaya Jha, Tomomichi Amano, Tatyana Deryugina, Kenneth Gillingham, Arthur van Benthem, Joseph Aldy, in addition of numerous seminar participants. The U.S. Department of Energy and the National Science Foundation provided funding for this project through the NBER Economics of Energy Market program. This research is also part of the activities of SCCER CREST, which is financially supported by the Swiss Commission for Technology and Innovation (CTI) / Innosuisse. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Heterogeneous (Mis-) Perceptions of Energy Costs: Implications for Measurement and Policy Design
Sébastien Houde and Erica Myers
NBER Working Paper No. 25722
April 2019
JEL No. D12,D83,L15,Q41,Q50

ABSTRACT

Quantifying heterogeneity in consumers' misperceptions of product costs is crucial for policy design. We illustrate this point in the energy context and the design of Pigouvian policies. We estimate non-parametric distributions of perceptions of energy costs in the U.S. appliance market using a revealed preference approach. We show that the average degree of misperception is misleading— while the largest share of consumers correctly perceives energy costs, a significant share undervalues them, and smaller shares either significantly overvalues or completely ignores them. We show that setting a tax based on mean misperception deviates substantially from the optimal tax that accounts for heterogeneous misperceptions. While correctly characterizing misperception is crucial for setting optimal Pigouvian taxes for externalities, it is less important for setting optimal standards. We find that standards can largely outperform taxes. Standards' advantage is they reduce variance in energy operating costs relative to taxes, which internalizes distortionary effects from misperceptions.

Sébastien Houde
ETH Zurich
Department of Management, Technology, and Economic
Centre for Energy Policy and Economics
Zurichbergstrasse 18, 8092 Zurich
shoude@ethz.ch

Erica Myers
University of Illinois at Urbana-Champaign
and E2e
ecmyers@illinois.edu

1. Introduction

Behavioral economists have long pointed out that limited cognition, biases, and heuristics have important implications for a wide range of economic decisions. However, it is only recently that empirical evidence of such behavior has been demonstrated in naturally-occurring markets.[1] There is now a large and growing literature demonstrating that consumers are prone to making mistakes across various economically important markets, such as in choosing health care plans, mutual funds, mortgages, which foods to consume, and in accounting for sales taxes and various shrouded attributes like bank fees, mini bar fees, and shipping and handling expenses among others.[2] For policy design, it is crucial to understand which types of behavior can be attributable to mistakes and how to empirically detect and quantify them. For example, taxation, one of the most commonly used policy levers, has dramatically different effects on behavior if consumers misperceive aspects of product costs targeted by taxes than if they do not. Recent work in behavioral public finance has shown that the presence of misperceptions, in particular the degree of heterogeneity in misperceptions, is a critical input for optimal policy design.[3] The distribution of misperception in the population not only influences the optimal level of a particular policy instrument, but also can determine the type and/or combination of instruments that should be used (Farhi and Gabaix 2018; Allcott et al. 2014).

While characterizing heterogeneous misperceptions is crucial for policy design, it is challenging to do in practice. The researcher needs: 1) a rich description of consumer behavior in the "naturally occurring" environment, where they make mistakes, 2) a rich description of consumer behavior in the "welfare-relevant" environment, where they do not make mistakes, and 3) the mapping between them (Bernheim and Taubinsky 2018). Previous work has relied on two approaches to measure misperceptions. The first approach is to use "artefactual field experiments" (Harrison and List

[1]The series "Anomalies" by Richard Thaler and his co-authors published in the Journal of Economic Perspectives (1987-2006) is probably the first sytematic effort to report real-world manifestations of biases and heuristics that behavioral economists previously documented in laboratory settings.

[2]See for example mistakes in healthcare plan choice (Abaluck and Gruber 2011; Kling et al. 2012; Handel and Kolstad 2015; Heiss et al. 2016; Ketcham et al. 2016a), mutual funds (Barber et al. 2005), mortgages (Allen et al. 2014; Guiso et al. 2018), schools (Jensen 2010), and nutrition (Bollinger et al. 2011). In addition see the following for inattention to sales taxes (Chetty et al. 2009), Taubinsky and Rees-Jones (2018) and other shrouded attributes (Ellison 2005; Gabaix and Laibson 2006; Ellison and Ellison 2009; Hossain and Morgan 2006).

[3]In this paper, we will use the term "misperception" to refer to any biases, heuristics, or biased beliefs that induce consumers to undervalue/overvalue an aspect of product cost.

2004) to create linkages between the naturally occurring and welfare-relevant environment with experimental choice environments (Allcott and Taubinsky 2015; Taubinsky and Rees-Jones 2018). The second approach is to rely on revealed preference and estimate a "behavioral econometric model", (DellaVigna 2018) which explicitly or implicitly characterizes a particular departure from the neo-classical rational model. The model is then used to simulate choices where the behavioral departure is turned off to create a mapping between decision utility and the welfare-relevant experienced utility.

This paper makes several contributions to measuring and demonstrating the policy impacts of consumers' misperceptions. First, we develop a revealed preference estimator to recover heterogeneity in perceptions for one dimension of product costs: energy operating costs. This is the first attempt that we are aware of—not only in the energy context, but more broadly—to use a revealed preference approach that explicitly focuses on recovering heterogeneity in cost perception. Our estimator consists of a semi-parametric estimator that recovers a fully non-parametric distribution of perceptions to energy costs. The advantage of our approach is that it does not require to take a stand on rational behavior, and allows to simply depict heterogeneity in preferences. Post-estimation, different analysts can then use different assumptions to judge if the distribution of preferences is consistent with a fully rational model. We apply our estimator to the U.S. appliance market, in particular to refrigerator purchase decisions. We find that while the largest share of consumers appears to correctly perceive energy costs, a significant share undervalues them, and smaller shares either significantly overvalues or does not pay attention to them. These patterns are strikingly similar across income groups and robust to various normative assumptions used to determine the true experienced utility of consumers.

Our second contribution is to the nascent literature investigating behavioral policy design. We show how heterogeneous misperceptions can impact the design of Pigouvian policies. Our framework builds on the work of Farhi and Gabaix (2018); Allcott et al. (2014) and Houde and Aldy (2017b), deriving expressions for the optimal *behavioral* policies that aim to address negative externalities while accounting for misperceptions. We show that while heterogeneity can significantly impact the level of optimal price instruments, it has small and possibly no effect on the design of quantity instruments, such as efficiency standards. When we use our estimated demand model to compare different types of efficiency standards and Pigouvian taxes, we find that standards can largely outperform taxes. We show that standards, unlike taxes, can internalize externalities and misperceptions at once. The key advantage of standards is that they may reduce variance in energy operating costs in the choice set, which reduces the distortionary effects of misperceptions in trading off energy costs with other attributes. When the variance in energy costs within a choice set is low, misperceptions of energy costs have little effects on substitution patterns and ultimately welfare.

Third, our work makes a methodological contribution by developing an empirical strategy to recover unobserved heterogeneity non-parametrically, which is well-suited for Big Data demand analysis in other contexts like ours, where the number of observations and products are both high dimensional. We first show that focusing on estimating non-parametric distribution is crucial in our setting. Using Monte Carlo evidence, we demonstrate that relying on parametric assumptions to specify the distribution of heterogeneity is likely to yield large biases about the true nature of the heterogeneity presents in a market. This finding is particularly important given that characterizing heterogeneity is the main objective of this paper. We then propose a two-step estimator, which we adapt from Fox, Kim, Ryan, and Bajari (2011) (FKRB, thereafter) to recover a multi-dimensional non-parametric distribution of preferences while addressing the curse of dimensionality. The FKRB estimator has been used in few applications (e.g., Nevo et al. 2016; Blundell et al. 2018). To our knowledge, we are the first to use it for a demand analysis in a context where there are several hundred of options in the choice set and there is a need to account for a large number of product fixed effects. We also provide Monte Carlo evidence that our estimator can achieve consistency in datasets typically used by applied researchers. For instance, we use administrative data and perform the estimation with a sample of close to 200,000 consumers and a choice set that spans more than 400 options.

Finally, our work contributes to the debate about the existence of a so-called Energy Efficiency Gap,[4] which refers to the apparent under-adoption of energy-efficient technologies. Several reasons have been proposed to explain this phenomenon but one of the most highly debated question is whether consumers' systematically fail to properly account for energy operating costs in their purchase decisions. The seminal paper that started the debate about misperceptions of energy costs is Hausman (1979), which found estimates suggesting that consumers were discounting too steeply

[4]The literature has also used the term Energy Paradox in this context.

future operating costs. Following Hausman (1979), a large number of studies have performed a similar exercise and found a wide range of implied discount rates, but that all tend to exceed normal rate of return (Train 1985). One shortcoming of Hausman (1979) and other early studies is that most estimators did not control for unobserved product attributes correlated with energy costs.

This is in contrast with most recent studies (Busse et al. 2013; Allcott and Wozny 2014; Grigolon et al. 2018; Sallee et al. 2016b; Myers 2018; Houde and Myers 2019), which have paid close attention to unobservables that might bias the estimate of misperceptions. 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. In contrast to earlier work, these studies find modest, or no, undervaluation of energy operating costs in the context of cars, housing, and appliances. Until now, the above literature has focused, however, solely on estimating the average degree of misperception.

In this paper, we show that the average masks substantial heterogeneity and can be a very misleading statistic. Underlying the modest level or absence of misperceptions that we found in Houde and Myers (2019) are heterogeneity patterns that suggest that consumers are prone to making mistakes in this decision context. We also assess the plausibility of the severe misperceptions that we observe by investigating two mechanisms explicitly: search frictions and biased beliefs. First, we examine the role of search frictions with a fully non-parametric estimator that exploits the existence of pairs of identical refrigerator models that differ only with respect to energy efficiency and price. While the more efficient model is almost always more expensive, there are sales events when it is less expensive, dominant option. We find that 1/3 to 1/2 of consumers still buy the dominated option, suggesting that search frictions could rationalize inattention to energy costs. Second we assess the role of biased beliefs using a survey designed to assess energy literacy. We find that a significant share of consumers have biased beliefs, which can rationalize the severe misperception we observe particularly in overvaluing energy costs.

The remainder of the paper is organized as follows. In the next section, we discuss the data and choice environment for our empirical investigation. In Section 3, we present our empirical framework for recovering heterogenous perceptions. In Section 4, we present the results of our estimation and in Section 5 we examine the role of search frictions and biased beliefs in driving the heterogeneity patterns that we observe. Section 6 follows and investigates the design and evaluation of behavioral environmental policies.

2. 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 and have little variation in utilization across consumers. Although refrigerator energy costs could be subject to idiosyncratic variation across households, the characteristics of a refrigerator such as its size, door design, and presence of ice maker are the main determinants of its energy costs. Therefore, it is not necessary to explicitly model the endogeneity of the utilization and purchase decisions, 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 of interest. 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 (Gertler et al. 2016). Contributing to the design of policies that improve the energy efficiency of refrigerators is thus important to reduce the negative externalities associated with household long-term 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 2008-2012. 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 includes: 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.

One particular feature of the U.S. appliance market is that appliance retailers, such as ours, have a national pricing policy and retail prices are subject to large and frequent changes. 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. In Appendix A1, we provide more details on this variation and the idiosyncratic price variation at the model level.[5]

To construct our main dataset, 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. If more than one utility serves a county, we take the average of those utilities' prices. Prices are highest in New England (14-20 cents/kWh) and the lowest in the Midwest and South (6-10 cents/kWh). There is also variation over time in price with some states experiencing price increases and others price decreases over the study period.[6]

We estimate the annual energy costs for each model-year-store location by multiplying the annual kwh consumption reported by the manufacturer by the energy price in the store's county. Almost all models sold in our sample are less than $2000, though there are some much higher priced models available. Using an expected lifetime of 18 years and a 5% discount rate the lifetime costs range from $555 for the 10th percentile of energy price to $1000 for the 90th percentile, and the ratio of lifetime cost to purchase price ranges from .44 at the 10th percentile of energy price to .79 at the 90th percentile. This variation in energy costs for particular models across time and space identify the coefficient on energy costs.[7]

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)

[5]Houde and Myers (2019) discuss this variation and demonstrate that this rich variation can be exploited, with few controls, to identify the coefficient on price. Using extensive robustness checks and an instrumental variable approach they demonstrate that the variation in the retail price is not driven by demand shocks and credibly exogenous.

[6]See Appendix A1 for a graphical display of electricity price variation across regions.

[7]See Appendix A1 for a graphical display of the price and energy cost variation used in the analysis.

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. We construct a measure of average rebate at the county-week level using SEEARP data collected by Houde and Aldy (2017a) and utility rebates from the DSIRE database. In the estimation, we do not explicitly distinguish between SEEARP and utility rebates. Houde and Aldy (2017b) show that both types of rebates have a small influence on the adoption of Energy Star certified products.

We carry the estimation using a large subsample of transactions. Each model that we estimate uses a large random subsample of approximately 200,000 transactions. We restrict the sample to transactions made by households owning their housing unit[8] with the goal of focusing on transactions made by consumers who are likely to pay the energy operating costs of their appliances.

3. Recovering Heterogeneous Misperceptions

3.1. Framework

In this section, we present the standard framework that has been used to estimate consumers' average degree of misperception of energy operating costs. We then discuss how we apply it to recover heterogeneous misperceptions.

The test for misperception of energy costs compares whether consumers respond the same way to a one dollar change in the purchase price of an energy durable versus a one dollar change in its energy operating costs. If not, this is taken as evidence of potential misperceptions. To implement this test, the two main variables to consider in the purchasing decision are thus the price (capital cost) of product j, denoted P_j, and the future energy operating costs over the entire expected lifetime of the product, denoted E_j. For expository purposes, we begin by abstracting away from uncertainty and consumer-specific heterogeneity in the product lifetime, future energy

[8]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 two refrigerators during the period 2008-2012 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.

prices, utilization, depreciation, and discount factor.[9] 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:

$$U_{ij} = \gamma_j - \eta P_j - \theta E_j + \epsilon_{ij} \tag{1}$$

where γ_j is the vertical quality of product j, capturing all of a product's attributes and ϵ_{ij} captures consumer $i's$ idiosyncratic preferences for product j. Numerous studies have estimated variants of this model with the goal of identifying the preference parameters for price and energy cost: η and θ, respectively. The formal test of misperception is whether the ratio of the two preference parameters equals one: $\theta/\eta = 1$. This statistic, which will thereafter call m, is now commonly referred as the valuation ratio (Allcott 2013).[10]

3.2. The Role of Preference Heterogeneity

When the preference parameters η and θ are heterogeneous, there are three important challenges for quantifying potential misperceptions. The first challenge is to consistently estimate the misperception measure, $m = \theta/\eta$. As we demonstrate in Appendix A2, if both η and θ are random parameters, quantifying the mean of their ratio (θ/η) requires an estimation of higher moments of the joint distribution of the two parameters and is not equivalent to the ratio of their independent means. A closed form solution for the distribution of m only exists for a few specific distributions (e.g., two lognormals). For the general case, moments of that distribution can be approximated given any distributions of η and θ. For the first moment, we show that: $E[\theta/\eta] \approx E[\theta]/E[\eta] - cov(\eta, \theta)/E[\eta]^2 + Var(\eta)E[\theta]/E[\eta]^3$.[11]

[9] We describe the assumptions we make and the heterogeneity in discount factor that we explore in what follows.

[10] An alternative, but equivalent approach that has been widely used in the literature used yearly energy operating costs in Equation 1 and then solves, post-estimation, for the implicit discount rate, r, such that the ratio $\theta \frac{1-\delta(r)}{\delta(r)(1-\delta(r)^L)} \eta = 1$, where $\delta(r) = 1/(1+r)$ and L is the expected lifetime (see Houde (2018b) for derivation of this result). Using this approach, an implicit discount rate that is markedly above the normal rate of return is seen as a sign that consumers undervalue energy costs.

[11] To our knowledge, we are the first to report this approximation, which has important implications for reconciling the various estimates of valuation ratios that have been recently reported in the literature. Some estimates (e.g., Busse et al. 2013; Grigolon et al. 2018; Houde and Myers 2019) rely on frameworks that yield

The second challenge that arises due to preference heterogeneity is the difficulty of consistently recovering higher moments of a distribution of preferences in contexts where we do not observe repeated choices by the same decision makers. This problem is well-established in the discrete choice literature. For instance, in a Monte Carlo experiment, Train (2009) shows that in a context where there is only one choice situation per decision-maker, the parametric random coefficient logit (RCL) model recovers only approximately 40% of the variation in population preferences.[12] This issue is particularly salient with energy-using durables given that households purchase these products infrequently.

The third challenge arises when consumers' perceptions of product quality are heterogeneous and possibly correlated with preferences for prices and/or energy costs. If heterogeneous quality is not properly accounted for, this would, of course, bias the estimate of the joint distribution of η and θ.

3.3. **Solutions & Estimator**

There is a straightforward solution to the first challenge. As first suggested by Revelt and Train (1998), we can transform Equation 1 by factoring out the marginal utility of income, the parameter η in our model, and by estimating the parameter m or its distribution directly. For instance, the distribution of m can be estimated using the model:

$$U_{ij} = \gamma_j + \eta_i(-P_j - m_i E_j) + \epsilon_{ij}. \tag{2}$$

the approximation: $E[\theta]/E[\eta]$, whereas other frameworks (e.g., Allcott and Wozny 2014; Sallee et al. 2016a; Myers 2018) exactly estimate the quantity: $E[\theta/\eta]$.

[12] In our own Monte Carlo experiment (Appendix A3), we reach a similar conclusion to Train (2009) regarding the RCL and its ability to recover higher moments of the true distribution of preferences. We investigate the role of the number of observations (i.e., decision-makers), the number of alternatives, and mis-specification in the parametric distribution of preferences. Across the different scenarios, we find that the RCL performs well at recovering the mean values of the true population parameters. However, the estimates of the covariance matrix of the distribution of preferences are imprecise and subject to large biases. Scenarios where the parametric distribution is mis-specified are particularly discouraging. For instance, if the true distribution is a mixture of normals, but we specify that preferences follow an unimodal normal, the estimated covariance matrix would suggest a much wider distribution of preferences compared to the true data generating process.

The advantage of this approach is that the estimation directly yields an estimate of $E[m]$, or a distribution of m if it is specified as a random parameter.

A solution to the second challenge is to relax parametric assumptions and to rely on a non-parametric estimator to flexibly recover a distribution for the parameter m. We do so by adapting the semi-parametric estimator of Fox, Kim, Ryan, and Bajari (2011) (FKRB), which allows us to estimate a non-parametric distribution of m. The intuition of the FKRB estimator is that a continuous distribution of random parameters can be approximated by estimating population weights over a discretized support. We adapt this estimator to recover the joint distribution of $F(\eta, m)$ as follows.

We first discretize the support of η and m into K grid points: $\beta^k = \{\eta^k, m^k\}$, $k \in K$. We then compute the choice model for each β^k using a parametric model where the probability of choosing product j given β^k is noted $P_j(\beta^k) = P_j^k$. For the parametric model, we are using the conditional logit. The choice probability, P_j, is then a mixture of K conditional logit models:

$$P_j = \int P_j(\eta, m) dF(\eta, m) \approx \sum_k^K \alpha^k P_j^k \tag{3}$$

where $\sum_k^K \alpha^k = 1$ because the weights are a discrete probability density function that approximates the true underlying continuous function. By choosing a parametric form for the choice model, each P_j^k can be first be computed for each grid point and then treated as data in the estimation. The estimator is thus semi-parametric and the estimation can proceed by running a linear regression with P_j as the dependent variable, P_j^k as regressors, and α^k, $\forall k \in K$ as coefficients 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 main weakness of the FKRB estimator is that it suffers from the curse of dimensionality, which is also related to the third challenge, determining how to account for quality. For example, product fixed effects, denoted γ_j in Equation 1, control for a products' mean unobserved quality, and are important in obtaining unbiased estimates of the distributions of η and θ, since appliances that are nicer along unobserved dimensions might also have higher prices and be more efficient. However estimating each of these fixed effects as random coefficients discretized with 100 + grid points in the FKRB framework would rapidly become intractable. Further, we have rich data on consumer demographics and product attributes that could be used to control for other dimensions of perceived quality that might be correlated with preferences for prices or energy costs.

To overcome the curse of dimensionality and obtain a consistent set of estimates for quality and the joint distribution of η and m, we follow a suggestion in FKRB to perform a two-stage estimation, which allows us to incorporate a rich set of controls for quality. In the first step we use simulated maximum likelihood to estimate the full model with a parametric random coefficient logit where η and m are assumed to follow a multivariate normal with an unknown mean and co-variance matrix. In the second step, we estimate the non-parametric joint distribution of η and m following the FKRB procedure outlined above, fixing all of the other regressors at their means.

To investigate the consistency of this two-step approach, we extended our Monte Carlo experiment. In particular, we were interested in whether a mis-specified first-stage has implications for the estimation of the non-parametric distribution in the second-stage. In Appendix A3, we show that the FKRB estimator implemented with a two-step approach performs well when the quality terms are first estimated with a parametric random coefficient logit. This is true even when the parametric distribution of the random coefficient logit is mis-specified.

The intuition behind this result is that while the parametric random coefficient logit model has difficultly recovering higher moments of the distributions of parameters when it is mis-specified, it still performs well at recovering the mean values. The mean of a mixture of several distributions is simply a weighted average of the mean of each distribution. The higher moments of the mixture are, however, not a simple weighted average of the higher moments of the individual distributions entering the mixture, which leads to bias if these are mis-specified.[13] One added advantage of the two-step approach is that the parametric distribution from the random coefficient logit can be readily compared to the non-parametric distribution of the FKRB estimator, which allows us to show the bias induced by the parametric model.

[13]We also show that it is important to use a model with parameter heterogeneity in the first stage. A model that does have heterogeneity, namely a conditional logit, leads to biases in the estimates of quality, which in turn lead to a downward bias in the estimated non-parametric distribution of the preference parameters of interest, θ and η.

To construct the choice probabilities, we infer a zip code-trimester-specific choice set, i.e., all models offered in a given zip code during a given trimester are considered to be in a consumer's consideration set.[14] We focus on modeling the purchase decision conditional on the fact that a consumer has decided to buy a refrigerator at time t and in location r. The timing decision and the choice of the retail chain store are thus not explicitly modeled. The marginal effects of the coefficients on price and energy costs thus capture the substitution across different models offered.

To recover the standard error in the second-stage and to account for uncertainty due to the two-step estimation, we implement the subsampling bootstrap (Politis et al. 1999). This consists of slicing the dataset in S subsamples, without replacement, and performing the estimation for each subsample. The estimated joint pdfs are averaged across all subsamples. The standard errors of each discrete weight of the pdf capture the variation across subsamples.[15] When applying the subsampling bootstrap, we should be careful in making inference for parameters situated at the edge of the support of the distribution, which often arise with the FKRB' estimator. The choice of the grid points, especially the end of the support should therefore be done with caution. We perform various robustness checks with respect to the choice of the grid. We also use our Monte Carlo analysis to investigate how the choice of the grid impacts the estimates. We provide additional details about the estimation procedure in Appendix A4.

Given that the two-stage approach can consistently recover a non-parametric joint distribution of our parameters of interest, η and m, while accommodating high-dimensional fixed effects, it allows us to address the third challenge by controlling for many aspects of heterogenous quality. The parametric choice model for each grid point $k \in K$, modeled as a conditional logit with alternative-specific utility given by:

$$U_{ijrt}^{k} = \eta^{k}(P_{jrt} + m^{k} E_{jrt}) + \tau ES_{jt} + \phi Rebate_{jrt} + \gamma_j + Demo_i \times Att_t + \epsilon_{ijrt}, \tag{4}$$

[14]Almost all zip codes have only one store. Our choice sets are thus store specific. 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.

[15]The subsampling bootstrap is recommended by FKRB because the estimated distribution may bunch at the edge of the support.

where all coefficients except η^k, m^k are fixed at their mean values estimated in the first stage. We control for Energy Star certification (ES_{jt}), Energy Star rebates ($Rebate_{jrt}$), product fixed effects (γ_j), and demographic information interacted with a subset of attributes ($Demo_i \times Att_t$). Specifically, we interact income level, education, age of the head of the household, family size, and political orientation with refrigerator's overall size, and freezer location (i.e., top-freezer, bottom-freezer, or side-by-side). One concern in our setting is that preferences for energy-related attributes could be correlated with local electricity prices. This correlation could exist for a variety of reasons. For example, high income households might prefer larger refrigerators and live in regions subject to higher electricity prices. Demographic information interacted refrigerator attributes is included to control for any of these types of preferences for specific refrigerator features that may systematically vary with consumer characteristics.

Another way to address heterogeneity in consumers' perceptions of product quality that might be correlated with preference for prices and/or energy costs, is to allow the coefficients on important dimensions of quality to be a random parameter in the two stage estimator. As a robustness test, we estimate the a non-parametric joint distribution for η, m, τ, allowing the coefficient on the Energy Star certification to be a random parameter as well. The reason we focus on the Energy Star certification is that consumers perceive certified models as being of higher quality, irrespective of their energy costs (Houde 2018b). Second, we also consider that the product fixed effects could be random parameters, which essentially allows for measurement error in the estimates of mean quality.[16]

In addition, we present estimates where we perform the estimation separately for each of six different income groups: <\$30k, \$30k-\$50k, \$50k-\$75k, \$75k-\$100k. \$100k-\$150k, and >\$150k. This allows us to have income-specific controls for quality. We implement the estimator on six different samples using the same approach described above. Each sample is drawn from the universe of transactions made by households in a particular income group.

[16]To implement this estimator, we assume that the coefficient for each product fixed effect follows a normal distribution, where the mean is equal to the estimates obtained in the first stage and the standard deviation is 25% of the mean. We then integrate the choice probabilities over these distributions when we perform the second-stage estimation.

3.4. **Identification**

The identification of the non-parametric distribution in the FKRB estimator is induced by variation that leads to substitution across products. Variation in prices, energy operating costs, and entry and exits of products are the main sources of variation that we exploit. Given that the estimator is semi-parametric and the kernel of the model is the conditional logit, the joint distribution of the parameters η and m will allow for the relaxation of the independence of the irrelevant alternative (IIA) assumption that the conditional logit imposes. The identification of the distribution thus comes from the variation in prices and energy costs that induces relative movement in market shares for products that are close substitutes.

For the identification, it is crucial that we control for quality and use variation in prices and energy costs that are not correlated with product attributes. Given our rich controls for quality the identification of the distribution of the parameters η and m works as follows. Consider two products j and k that differ with respect to energy use and are subject to changes in prices that are not perfectly correlated. Each product may also differ along several dimensions, which will be captured by the controls for quality discussed above. Once quality is controlled for, the relative market shares of products j and k will be function of the differences in energy operating costs and prices. If consumers pay attention to energy costs and prices with equal weight, both sources of variation will lead to similar changes in relative market shares.

If consumers pay more attention to prices relative to energy costs, it would lead to lower value of m. In fact, complete inattention to energy costs could even lead to negative values for m. This could happen if consumers choose an option with high energy costs but its quality and price alone cannot rationalize this choice, i.e., it appears the option is dominated in the quality-price dimension. The preference parameter for energy costs then needs to take on a negative value to explain this choice. Flexibly controlling for quality and price is thus crucial to capture instances where consumers may make mistakes by choosing a dominated option.[17] Large values of m can happen when consumers pay more attention to energy costs relative to prices. Again controlling for quality, especially for attributes related to energy consumption is important. In particular, controlling for refrigerator

[17]Having a random coefficient on price is particularly important to identify the tail of the distribution of the parameter m, especially the potential negative values. Consumers will choose dominated option if they pay little attention to price, i.e., for values η close to zero, or they put a very high weight on price, i.e., very high values of η corresponding to very high price elasticities, relative to quality.

size is particularly important as it is positively correlated with energy use, and presumably, is a desirable attribute.

As in any demand estimation, we also need to ensure that variation in price is exogenous to local market conditions. To the extent that the price changes from retailer's national pricing algorithm are driven demand shocks, it would create endogeneity between price and demand that could bias our estimations. Using the conditional logit, we explore whether changes in demand over time might affect our coefficient on price. The main concern is that contemporaneous demand shocks are correlated with prices due to brand-specific weekly promotions and time trends. We thus investigate the impact of week-of-sample fixed effects interacted with brand dummies.

Table 2 presents the results for a simple conditional logit. We first present a very parsimonious model where we only control for product fixed effects, the Energy Star certification status, and Energy Star rebates (Specification I). We then include interaction terms between demographic and product attributes (Specification II). In Specification III, we interact the Energy Star dummy with county fixed effects to account for region-specific preference for energy-efficient refrigerators that could be correlated with local electricity prices. In Specification IV, we include week-of-sample interacted with brand dummies—the coefficient estimates on price and energy costs change little. In Appendix A5, we further control for temporal demand shocks correlated with preferences for energy efficiency by considering week-of-sample fixed effects interacted with various attributes related to energy usage, such as the Energy Star certification status, refrigerator size, and freezer location. This set of controls has little effect on the price coefficient, indicating that the large model-specific price variation we observe from the retailer's pricing model is is unlikely to be driven by demand-side shocks.

One worry could be that areas with high energy prices also have residents with preferences for energy efficiency, since policies to promote renewable energy generation and utility efficiency programs can raise rates. Specifications IV-VII in Table 2 show the results for the conditional logit estimates where we interact various attributes correlated with a refrigerator energy usage with county fixed effects. In addition to the Energy Star certification status, we also consider a dummy for large refrigerators (overall volume larger than the mean: 22.5 cu. ft.), and a dummy for refrigerator with a top freezer (i.e., the most efficient models relative to bottom-freezer or side-by-side). If regions with high energy prices also have preferences for energy-related attributes, the addition of these controls will effect on the coefficient on energy costs. Again, these controls change the coefficients on price and energy costs little. This suggests that it is unlikely that a correlation between preferences for energy efficiency and local energy prices is biasing our results. In addition Houde and Myers (2019), provide evidence that the variation in local energy prices during this period is largely being driven by exogenous variation in fuel costs. They instrument for energy costs using average annual electricity fuel prices weighted by the local utility's pre-sample capacity shares of coal, oil and gas-fired power plants. The two-staged least squares (2SLS) estimates using this instrumental variables approach are indistinguishable from OLS estimates indicating that the variation in electricity price is driven mostly from the underlying fuel price variation rather than demand side factors.

The fact that the conditional logit model is relatively unchanged with the addition of many fine-grained controls for quality suggests that our relatively parsimonious model used in the 2-stage structural estimation is capturing the important aspects of heterogenous preferences for quality.

3.5. **Assumptions: Lifetime Energy Costs**

Quantifying misperceptions requires making assumptions about consumers' beliefs about future fuel prices, the lifespan of the refrigerator, and the discount rate. In this section, we describe the assumptions made in our analysis. Though given our estimation, analysts could apply different assumptions ex post and evaluate their effects on the distribution of misperceptions. For our analysis, we assume that consumers believe that annual electricity prices follow a no-change forecast, so that contemporaneous prices are the best predictor of future prices. In our main specification, we assume the average life expectancy for refrigerators of 18 years for all models and the average discount rate used by the DOE for appliance standards of 5% (the market rate of return) for all consumers.

In interpreting our results, we also show how heterogeneous discount rates can impact our interpretation. We consider that discount rates between 2% and 12% could be "rational". This is motivated by the return on 3 year U.S. Treasury bonds, which is close to 2% during our sample period. This low discount rate thus represents a market returns for a risk averse consumer with no credit constraints. On the other extreme, we use the average APR rate for credit cards during our period, which is around 12%. This represents a cost of funds for consumers carrying credit card debt.

4. Results

4.1. Non-Parametric Distribution of η and m

We first present the results of the FKRB estimator graphically. Figure 1 shows the marginal pdf of η (Panel a), and the marginal pdf of m (Panel b) estimated with the FKRB estimator (blue dots) and the parametric random coefficient logit (red stars). In each panel, each pdf weight corresponds to an average of the weights estimated across the bootstrap iterations.[18] The 95% confidence interval of each estimated weight is depicted by the vertical lines.

The estimated distributions for both parameters show substantial heterogeneity. The distribution of misperceptions, m, is especially dispersed (panel b) and has several modes and statistically significant weights at grid points located between zero and one. However, there are also statistically significant weights above one and below zero. A value of m greater than one suggests that there is share of consumers that place a very high weight on energy operating costs in their purchase decision. Put another way, some consumers have an implied discount rate below 5%. In fact, a value of $m = 1.54$, is equivalent to having a negative implied discount rate under our assumptions. A negative value of m implies not only that (some) consumers do not value energy costs in their purchase decision and they choose energy-inefficient models over cheaper more energy-efficient models of similar quality. We provide evidence of existence of search frictions that might lead to this extreme undervaluation of energy costs in Section 5.

Compared to the parametric random coefficient logit, the FKRB estimator suggests larger tails for the distributions and a larger average degree of misperception: 0.77 instead of 0.64. The estimator used to recover heterogeneity plays thus an important role. In the present case, the parametric model underestimates both the amount of heterogeneity and the average degree of misperception. This conclusion also holds for the parameter η.

[18]The pdf displayed in each panel is effectively a point-by-point average of sixteen different pdfs. Therefore, the sum of the weights may not sum to 1.

The marginal pdf of η is left-skewed and has most of its mass between its mean value (-5.43) and zero. Some consumers have a very high price elasticity and others are price inelastic. We also find a positive and statistically significant mass for positive values of the price coefficient, which suggests that there is a share of consumers that are inattentive to the purchase price and/or have high search costs. Interestingly, consumers that have a negative value for m are also the ones that either have a positive or have a largely negative price coefficient. This suggests that these consumers are either focusing on dimensions of the product other than price or energy costs, or put a very strong emphasis on price alone, and dismiss energy costs completely.

Figure 2 shows the joint pdf of η and m. The size and the color of the markers represent the value of the pdf. The uncertainty of each estimated weight is depicted by a circle or a star, where a star corresponds to an estimate that has a t-statistic of 1.1 or larger.[19] The joint pdf shows that large values of m are correlated with lower coefficients on price. As we show below, this inverse correlation is partly, but not entirely, driven by the difference in the distribution of misperception across income groups. Several phenomena can explain overvaluation of energy costs. For instance, some consumers might pay much higher electricity prices to the ones that we have assumed in building the estimator, might have beliefs about current or future energy prices that are upward biased, may overestimate the energy consumption of refrigerators, or simply value energy efficiency highly because of environmental motives.

We use the 2%-12% range in discount rates to classify consumers in four broad types. First, there are the consumers who undervalue energy costs, which corresponds to a discount rate above 12%.[20] Second, there are consumers who respond to energy costs and are not prone to major misperceptions, i.e., no under or overvaluation. We classify the mass of the distribution for values of m that correspond to an implied discount rate of 2%-12% into that category. Third, there are consumers that appear to overvalue energy costs, i.e., consumers for whom choice can only be rationalized with discount rates smaller than 2%. Finally, there are consumers with negative values of m, which we interpret as strong signal that behavioral biases are at play. Table 3 reports the

[19]We favor this approach to show uncertainty in the estimates over confidence intervals simply to make the Figure more readable. Given that we use sixteen slices for the subsampling bootstrap, the value 1.1 corresponds to the critical value of a one-tail t-test with ten degrees of freedom and a significance level of approximately 15%.

[20]Under our assumptions, a value of $m = 0.62$ corresponds to an implied discount rate of approximately 12%. The mass of the distribution where $m \geq 0$, $m < 0.62$ thus corresponds to undervaluation.

estimated shares of consumers that belong to these four different types, which are 40.2%, 37.9%, 18.9%, and 12.6% respectively.

4.2. **Robustness**

We investigate the robustness of the estimated pdfs by first considering alternative grid size and grid density. In Table 3, we report that result for estimators where we first reduced the span of the support and reduced the size of the grid (Model II).[21] We find that this increases the bunching at the lower end of the support, but it does not affect the mean and the overall patterns for the pdf distributions (joint and marginal) (Figure A5). In another specification (Model III), we kept the the same span for the support, but increased the number of grid points. With denser grid, we find more mass in the range $m \geq 0$, $m < 0.62$, but it has little impact on the mean (Table 3) and the qualitative patterns (Figure A5, Appendix A6).

To assess the robustness of the estimates with respect to our controls for quality we consider three alternative estimators. First, we implement an estimator (Model IV) with a non-parametric joint distribution for the parameters η, m, and the parameter for the ENERGY STAR label: τ. Doing so allows us to further control for unobserved heterogeneity in preferences for overall quality, signaled by the ENERGY STAR label, which might be correlated with preferences for energy costs. This has little effect on the joint and marginal distributions of η and m. Second, we also consider that each product fixed effect, denoted γ_j, could be a random parameter. This specification introduces idiosyncratic preferences, in addition to the i.i.d. extreme value error term, for each product in the choice set. Given the large number of such fixed effects, we avoid the curse of dimensionality by assuming a parametric distribution for each fixed effect. We assume a normal distribution where the mean is determined by the first-stage estimates and the standard deviation is proportional to the estimated mean. Again, this is a little impact on the results (Model V in Table 3). Finally, we show the impact of having fewer controls for quality. In Model VI, we do not control for demographics interacted with energy-related attributes when we estimate the joint non-parametric distribution of η and m. The results show that the estimated joint distribution is similar to the specification

[21]We also attempted to implement estimator with a very large minimum value and maximum value for the support of the distribution, but we could not find a numerical solution for these estimators. The specification used for Model 1 in Table 3 is the largest span of the support for which we could recover estimates.

with richer controls for quality. Altogether, this suggests that unobserved quality not likely to be a major source of bias in our setting.

4.3. **Distributions by Income**

Table 3 summarizes the distribution of m across our four broad categories of energy cost perception for each of the six income groups. As discussed above, values of m smaller than zero indicate severe misperception. We find that this share of consumers is largest for the lowest income group, almost 30%. For other income groups, the estimated distribution suggests that this share ranges from 12% to 19%. Values of m between zero and 0.62, correspond to consumers that consider energy operating costs in their purchase decision with an implied discount rate larger than 12%. The share of consumers in this category is larger among the lowest income groups and decreases with income. For values m that correspond to discount rates between 2% and 12%, i.e., consumers subject to no or modest misperception, the share of consumers is the largest among the three highest income groups. Finally, for values of m greater than 1.28, which implies overvaluation, the share of consumers increases with income and is especially large, higher than 30%, for income greater than $100k.[22]

Importantly, for all income groups we observe evidence of heterogeneous perceptions. Although income is one driver of heterogeneity, it does not explain extreme values that suggest overvaluation, or complete inattention. The existence of potential severe misperceptions is therefore not driven by differences across income groups such as access to credit, or systematic difference in the perception of quality, which we account for.

5. **The Role of Search Frictions and Biased Beliefs**

There are several potential behavioral mechanisms that might lead consumers to misperceive energy costs. These include biased beliefs (Allcott and Knittel 2018), lack of salience (Chetty et al. 2009), present bias (Laibson 1997) or bias toward concentration (Koszegi and Szeidl 2013).[23]. Or consumers

[22]Figure A6 (Appendix A6) also shows the joint distribution of η and m for each income group. The overall patterns are strikingly similar. Consumers that either respond strongly to energy costs are or are completely inattentive to energy costs also tend to be less price sensitive (i.e., have a value of η close to zero).

[23]If consumers have bias toward concentration, they underweight future cash flows that accrue in small increments over time, relative to upfront costs.

could be subject to high search and information acquisition costs and be rationally inattentive (Gabaix 2014; Sallee 2014). Each of these mechanisms create bias in consumer perception of energy costs, which would result in under-valuation (or in some cases, over-valuation) of energy costs (Allcott 2016).

In this section, we explore two potential mechanisms behind the severe misperception that we observe: search frictions and lack of information. We show that search frictions may play a role, especially in explaining negative values of m. We also show that some consumer beliefs about key pieces of energy information are severely biased and follow patterns consistent with values of m that largely exceed one.

5.1. **Search Frictions**

In our discrete choice framework, a negative value for m implies that some consumers prefer energy-inefficient products over efficient ones, holding all other dimensions of quality constant. Therefore, in order to rationalize the probability mass of m over a negative support, our controls for the various dimensions of quality are crucial. In our estimator, the consumer-specific mean quality for attributes other than price and energy cost, γ_{ij}, are characterized by the product fixed effects and interaction terms between demographics and product attributes. The consumer-specific idiosyncratic quality is characterized by the error term that follows an extreme value distribution. Negative values of m arise when a consumer chooses an energy-inefficient product over other available options that have lower energy costs but higher quality, as suggested by their estimated mean quality, and given the dispersion of the idiosyncratic error term. Put simply, it appears that some consumers choose a dominated option both in terms of price and quality, and the only way the model can rationalize this decision is by having a negative value for m, which implies that energy costs were a desirable attribute.

In practice, search frictions could be one of the underlying behavioral mechanisms that induces consumers to choose a less efficient product over a more efficient one of higher quality. For instance, if consumers dismiss energy costs in their purchase decision, and focus primarily on other dimensions of quality, this could lead to instances where they would make what appears to be a mistake by selecting a dominated option. It is also possible that consumers allocate attention to a restricted set of options or sales people introduce them to a limited set of models, which results in having a dominant option excluded from the consideration set.

The empirical challenge in identifying these various types of search frictions, and the resulting mistakes it induces, is that preferences for all dimensions quality need to be correctly specified. In our setting, the negative values for m in the FKRB estimator would accurately capture search frictions only if our characterization of the mean quality, γ_{ij}, and the parametric assumption of the distribution of the error term provide a good approximation of the true preferences.

To investigate the role of search frictions in our setting we propose an alternative empirical strategy that requires minimal assumptions about how we specify preferences for quality. We exploit a natural experiment in the U.S. refrigerator market, which allows us to compare choices among pairs of nearly identical refrigerator models, where a clear dominant and dominated option exists. In the U.S. refrigerator market, manufacturers commonly make strategic product line decisions to meet the Energy Star certification (Houde 2018a). This often results in product line with several different models, where some models only differ with respect to their energy consumption, but are otherwise identical from the point of view of the consumers. In our sample, we were able to identify 52 pairs of such refrigerators using a conservative matching process. To identify these pairs we first used the rich attribute information to ensure that paired refrigerators where identical along observable dimensions of quality, except energy use. In particular, we focus on product lines where manufacturers models with the same refrigerator volume, freezer volume, width, height, depth, overall design, color, and technology options but consume different amount of electricity. For all paired refrigerators identified by this matching process, we then manually verified the accuracy of the matches using information from various online marketplaces.

A second institutional detail that offers an ideal setting to detect search frictions is the fact that the retailer has a national pricing strategy and offers large and frequent promotions. Therefore, each model is subject to large price variation and it is often the case that the price of the most efficient model within a pair is lower than the price of the less efficient model. When this occurs, this is a clear instance where consumers face a dominant and dominated option.

The existence of the identical pairs and model-specific price variation over time allows us to construct a fully non-parametric estimator that controls for all dimensions of quality so that we can estimate search frictions to energy cost and price. To build such estimator, we first identify instances when the more efficient model was cheaper than the less efficient model for each pair. We refer to this type of event as a dominated price event. In then identify cases where both models within a pair were offered at the same location during a dominated price event. One challenge is that we do not observe inventory, so we must impute product availability from sales. However, the fact that we have transaction level data where we observe the exact location and date of each purchase allows us to address this issue as follows.

We infer that two models of a matched pair were available at the same location if we observe at least one sales for each model during a time period of specific length. We consider different lengths starting from the most conservative, where we only consider instance when both models sold on the same day at a given location, and we gradually increase the length to two, four, six, and twelve days. As we increase the time interval between sales, it becomes more likely that if the dominant product was not purchased, it is because of it was temporarily out-of-stock, rather than consumer search frictions.

Once we have identified the dominate price events for each pair and restrict such events to locations where both models were offered at the same time, our non-parametric estimator of search frictions simply consists of reporting the share of consumer that chose the dominated option.

Table A5 reports summary statistics for the matched pairs compared to the overall sample of models we observe. One take-away is that matched refrigerator models tend to be smaller and cheaper relative to the overall sample and, therefore, are not fully representative of the U.S. market. A second important take-away is that the average price difference within a pair is $22 a small and positive amount, which is consistent with the fact that more efficient models tend to be more expensive. But there are large variations that span negative and positive values throughout the whole sample. The 10^{th} percentile is negative $150 and the 90^{th} percentile is $155.

Table 4 shows the share of consumers that chose the dominated option, along with the price difference and the number of observations during a dominated price event. For the most conservative assumption about product availability, where we limit our definition of a dominated event to locations where we observe at least one sale of each model in the pair sold on the same day, we have only 304 observations. This is a very small fraction of sales of the several million of transactions that we started with. Arguably, this estimator strives for internal validity and we err on the conservative side to truly identify inattentive consumers. We find that the share of inattentive consumers is 48%. As we increase the length of time between sales within a pair, the share of inattentive consumers tends to decrease and stabilize. It, however, remains large in magnitude. Across the various specifications, the share ranges from 48% to 38%.

These results show that consumers systematically fail to find a clearly dominant option in their choice set, which highlights that search frictions are at play in this setting. However, we should be cautious in interpreting these results for the following reasons. First, the estimator captures search frictions for particular options in the choice set, which could be induced by search frictions over energy costs and prices, and by how consumers restrict their consideration set. Although, we do not distinguish between these two mechanisms, we argue that both are the result of imperfect attention allocation. Our estimator is thus a useful diagnostic to show that search frictions are present and important in our setting. Second, although the estimator controls for all dimensions of quality, it does not control for the retailer's product placement strategies. In the absence of attention allocation costs though, product placement strategies should have a minimal impact. Such strategies can thus be interpreted as an equilibrium response to search frictions. Our estimator is therefore still relevant even if we do not identify how the retailer exacerbates or alleviates this behavioral bias. Third, the estimation is carried on a very small subset of the overall sample. The external validity of the results, therefore, is not guaranteed.

5.2. **Biased Beliefs**

Computing the energy cost of a refrigerator requires two key pieces of information: estimates of electricity consumption and electricity price. Beliefs about these pieces of information should therefore play an important role on how consumers respond to energy costs when they purchase an appliance.

In this section, we report data from a survey designed to assess energy literacy among U.S. households. The survey was administered in the Spring of 2017 on a representative sample of 1,512 U.S. households living in 24 different U.S. states. In this paper, we report the results for two questions that were included in the questionnaire and focused on beliefs. In the first question, we asked survey respondents to provide their best estimate of a full-size refrigerator's annual electricity use. In the second question, we asked them for their best estimate of the average electricity price that they pay. For both questions, respondents were asked to report their best estimates and no numerical or other anchors were provided.

Figure 3 reports the distribution of the ratio between beliefs and true value for the two pieces of information. Allcott (2013) refers to such measures as valuation ratios, where correct valuation implies a ratio of exactly one. To construct our valuation ratios, we matched county average electricity prices using information about respondents' zip codes. For refrigerator's annual electricity use, we use the value 450 kWh/year, which approximately corresponds to the average for the year 2017. For both pieces of information, we see a large distribution of beliefs, with significant under- and over-valuation. To construct the histograms, we censored all values greater than 5, which explains the large mass at this value. This shows that a significant share of consumers tend to have beliefs about refrigerators' energy use and local electricity prices that would lead to an overvaluation of energy costs. The pattern in beliefs is u-shaped and also suggests that a large share of consumers undervalue both the level of energy use of refrigerator and prices. Interestingly, Allcott and Knittel (2018) find similar patterns in beliefs in the U.S. car market.

In sum, we find that beliefs are severely biased in our setting and follow patterns consistent with the distribution of misperception our estimator recovers. In particular, the large values of m could be induced by consumers that overestimate energy use and electricity price. However, biased beliefs are just one possible mechanism. For example, environmental values could also play a role in the appearance of over-valuation of energy efficiency. It is also important to note that in the U.S. appliance market, the mandatory label EnergyGuide should help informing consumers about different pieces of energy information. How effective the label is in correcting beliefs remains an open question, but our survey results show that there is a need to provide better information.

6. **Implications for Policy Design**

In this section, we assess how heterogeneous misperceptions impact the design of externality-correcting policies. We focus on simple cases where the planner uses a single policy instrument, a tax or a specific type of product standard, and we illustrate the role that heterogeneity plays.

That is, our main counterfactuals show how optimal instruments and their welfare effects differ if the social planner only uses the average degree of misperception versus the full distribution. Given that our empirical exercise focuses on the demand side, we do not explicitly model firms' responses in our policy simulations. Though, as we discuss at the end of the section, the supply side response is unlikely to change the welfare rankings of the policies considered in our simulations.

The first step in conducting our policy analysis is to define the type of behavioral response that constitutes consumer bias, or misperception. Our demand estimation yields preference estimates that characterize "decision utility", or the utility consumers thought they would experience from purchasing a good. To the extent that consumers' biases affect their decisions, their decision utility might differ from their "experienced utility." The gap between decision and experienced utility determines the magnitude of the misperceptions. We make the following two assumptions for defining experienced utility.:

(1) For any m_k corresponding to discount rate below 2% or larger than 12%, we consider that consumer of type k misperceives energy costs and decision utility differs from experienced utility.
(2) If $\eta_k > 0$, consumer of type k misperceived the product price and decision utility differs from experienced utility. To measure experienced 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 implies that if consumers do not perceive a one dollar change in the present value of future energy operating costs, the same way they perceive a dollar change in product price (for a reasonable range of discount rates), they are prone to misperceptions. The second assumption aims to make a tractable representation of experienced utility for parts of the distribution with positive coefficients on purchase price, which could be a manifestation of various consumers' biases such as search frictions.[24]

The second step to implement the policy simulation is to derive a measure of welfare that accounts for the gap between decision and experienced utility. For this purpose, we follow Allcott (2013), Ketcham et al. (2016b), and Houde (2018b), among others, and use Leggett (2002)'s formula

[24]Less than 3% of the estimated joint density for η and θ has positive price coefficients.

to measure welfare in a discrete choice framework in the presence of imperfect information or alternatively consumers' biases. The expression for the change in consumer surplus is:

$$\Delta CS_k = \frac{1}{\eta_k} \cdot \left[ln \sum_j^J e^{(\tilde{U}_{kj})} + \sum_j^J \tilde{\sigma}_j^k \cdot (\tilde{U}_{kj}^E - \tilde{U}_{kj}) \right] - \frac{1}{\eta_k} \cdot \left[ln \sum_j^J e^{(U_{kj})} + \sum_j^J \sigma_j^k \cdot (U_{kj}^E - U_{kj}) \right]. \tag{5}$$

where the terms with a tilde are evaluated after the policy change, U_{kj}^E denotes experienced utility, U_{kj} corresponds to decision utility for consumer of type k and σ_j^k refers to the probability that consumer type k will choose product j. The expression in Equation 5 corresponds to the standard measure of welfare for the multinomial logit (Small and Rosen 1981) plus the term $\sum_j^J \sigma_{jtr}^k \cdot (U_{kjtr}^E - U_{kj})$, which we refer to as the correction term. The correction term arises because of the discrepancy between what consumers perceive that they will experience and what they actually experience, and represents the expected (private) cost that consumers incur because of their misperceptions.

6.1. **Optimal Pigouvian Tax**

In order to demonstrate the importance of effect of using the full distribution of heterogeneous misperceptions versus the mean degree of misperception in the population on the optimal Pigouvian taxation, we follow Houde and Aldy (2017b). Given the above expression as a measure consumer surplus and assumptions that 1) tax revenues are redistributed lump-sum, 2) the tax is included in the marginal price of energy, denoted P_e, and 3) each unit of energy consumed (E_j) is associated with a constant marginal damage cost, denoted ϕ, the optimal Pigouvian tax (τ^*) is as follows:

$$\tau^* = \frac{\phi}{1 - \mathcal{A}} + P_e \frac{\mathcal{A}}{1 - \mathcal{A}} \tag{6}$$

with

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

where $\frac{\partial \sigma_j^k}{\partial \tau}$ is the derivative of demand for product j with respect to the tax for consumers subject to misperceptions m_k.

If instead the tax is set using the average degree of misperceptions, noted $\bar{m}$, the optimal (price-inclusive) Pigouvian energy tax is:[25]

$$\tau^* = \frac{\phi}{\bar{m}} + P_e \frac{1 - \bar{m}}{\bar{m}} \tag{7}$$

The important difference between the two expressions is that in the presence of heterogeneous misperceptions the demand elasticities determine the size of the adjustment associated with each m_k. Because we consider a tax included in the price of energy, misperceptions apply to both energy costs and the tax. The magnitude of the adjustment is thus function of two demand parameters: the degree of misperception and the demand elasticity with respect to the tax. For values of m_k that are well-above one, σ_j^k is more elastic with respect to the tax and the degree of overvaluation creates a downward adjustment to the tax. For m_k close to zero, misperception creates an upward adjustment to the tax, but the demand elasticity with respect to the tax is smaller due to undervaluation of energy costs. The share of consumers that undervalue energy costs thus has a smaller impact on the overall magnitude of the behavioral adjustment compared to consumers prone to overvaluation. In practice, the tails of the distribution of misperceptions can therefore lead to large differences between the expressions 6 and 7, as we show below.

6.2. **Optimal Standard**

Minimum energy-efficiency standards are widely used by governments to manage energy demand. Historically, they have been justified by consumers' undervaluation of energy costs (Hausman and Joscow 1982). In our policy simulations, we consider three types of standards. First, we consider an uniform minimum standard where all products must consume the same amount of energy. This standard represents an extreme case where manufacturers cannot differentiate their products in the energy dimension.[26] Second, we consider a minimum standard where all products must meet

[25] The expression 7 is similar to the results of Farhi and Gabaix (2018) except that the price of energy also enters the expression of the optimal Pigouvian tax. It can also be shown that Proposition 1 of Allcott et al. (2014) yields a similar expression to Equation 6, except that their expression for the optimal tax is function of the degree of misperception in future utilization of the durable, in addition of the price of energy, and misperception m_k.

[26] The government of India has implemented a version of such standard by requiring that only one variant of CFL light bulb be offered. They are currently considering a similar approach for energy-efficient air-conditioners where only one variant would be offered.

or exceed a lower bound with respect to energy-efficiency, or equivalently, an upper bound with respect to energy consumption. Third, we consider a variant of the minimum standard where the standard varies as a function of key attributes. This attribute-based standard mimics the regulations currently in place in the U.S. appliance market, which vary with refrigerator and freezer size, freezer location, and defrost technology.

An expression for the optimal uniform standard can be easily derived in our framework and shows a striking result: a single uniform standard can internalize both externalities and heterogeneous misperceptions with a single instrument. Removing high energy models to address externalities also eliminates the distortionary effects of misperceptions by getting rid of any variance in energy costs.

To derive our main result, we focus on the the demand side and assume a stylized market structure, which abstracts away from the strategic behavior of the firms. We assume that manufacturers' purchase prices are given by:

$$p_j(E_j) = c(E_j) + \omega_j \tag{8}$$

where $c(e)$ is the product cost that varies as a function of the energy level E_j, with $c'(E) < 0$, and ω_j is a product-specific additive markup that does not vary with E_j. As before, we will assume a linear and additive externality cost and the existence of k different types, each subject to a bias m_k. The following proposition is the expression for the optimal uniform standard under this setup.

Proposition 1. *When $p_j(E_j) = c(E_j) + \omega_j$ and with a constant marginal externality cost, ϕ, the optimal uniform standard, denoted $\bar{E}^*$ is the solution of:*

$$-c'(\bar{E}^*) = 1 + \phi. \tag{9}$$

The important takeaway from Equation 9 is that the optimal standard is not function of misperceptions, nor demand parameters. That is, this result holds irrespective of the distribution of m_k—the design of an uniform standard is unaffected by misperceptions.

The intuition is that by setting a uniform standard, all options have the same energy level. Therefore, energy operating costs do not induce substitution across products and misperceptions of energy costs are completely internalized regardless of their distribution. The only factors that matter in setting the optimal standard is the trade-off between the increase in product cost induced by making the standard more stringent (i.e., $c'(E)$) and the externality cost ϕ. In comparison to a price instrument, a standard therefore has a clear advantage given its ability to internalize misperceptions, even if they are heterogeneous, while simultaneously addressing the negative externalities.[27]

The above result is specific to a very restrictive type of standard, but provides important insights on how other types of standards will be affected by heterogeneity in misperceptions. As standards reduce the variance in energy costs between products present in the choice set, this reduces the misallocations induced by misperceptions to energy costs. The uniform standard is the extreme case where there is a no variance in energy cost and all the substitution is ruled out. A minimum standard should induce some variation in energy operating costs, but less than an attribute-based minimum standard, which by design varies in several dimensions of the product space. Therefore, the former should be more robust to the distribution of misperceptions relative to the latter.

Note that although a uniform standard rules out substitution across products, it may affect demand via other margins, namely the decision to adopt a product or not and the timing of the purchase. A very stringent uniform standard could raise the average purchase price of a given technology, which in turn could induce some consumers to reconsider adopting such technology. In markets where those margins are important, it may no longer be true that misperceptions would not affect the level of the optimal standard. In those settings, the impact of misperceptions on the elasticity of the outside option is an important statistic.[28]

6.3. **Setup: Policy Simulations**

We simulate the demand model to find the optimal level of the tax or of the different types of standards discussed above. For each policy, we compare two scenarios: the level of the optimal

[27]This result is reminiscent of Allcott and Knittel (2018)'s result showing that an optimal fuel economy standard for vehicles should be equivalent to a 'pure nudge' that fully internalizes biases.

[28]In the U.S. appliance market, Houde and Aldy (2017a) found that the decision to adopt a new refrigerator, clothes washer, or dishwasher and the timing of purchase for such appliance are very inelastic, at least during the period 2008-2012. In particular, they found very little evidence that generous consumer rebates lead to a market expansion and long-term substitution in replacement decisions. The role of the outside option is thus second order in the present policy exercise.

policy using the full distribution of misperceptions versus the average degree of misperception. For each scenario, however, we take into account heterogeneity across income groups. The overall demand is thus a weighted average of the demand model estimated for each income group. We also consider heterogeneity in the share of each income group across all the U.S. counties. The demand model thus incorporates spatial heterogeneity in the degree of misperceptions to energy costs due to difference in income across regions.

In addition to the level of the optimal policy, we also report total social welfare, consumer surplus, and externality costs, but we abstract away from firms' profits. Across scenarios, we assume that firms' markups are fixed, and prices are set by: $p_j(E_j) = c(E_j) + \omega_j$. For simulating the standards, we need to model how energy consumption affects the product prices. For this purpose, we use the cost function estimated in Houde (2018c), which takes the following parametric form:

$$c(E_j) = \frac{\psi}{E_j} + \beta_j \tag{10}$$

where the parameter ψ is estimated using quasi-experimental variation and β_j is recovered using information about wholesale prices in this market.

For all scenarios, we assume that the only externality comes from carbon emissions with a marginal damage of 50 \$/ton of CO2. We consider heterogeneous local electricity prices and emissions factors, which vary with electricity market areas. Data for the CO2 emission factors comes from the EIA scenarios used to simulate the impact of the Clean Power Plan on electricity prices (EIA 2015). We infer the emission factors for 20 different regions of the U.S. corresponding to different electricity markets.[29] For electricity prices, we use the county averages constructed for our estimation.

To characterize experienced utility, we consider a range of discount rate that spans 2% to 12%. For all values of m_k that correspond to implicit discount rates that fit this range, we assume that

[29]Our simulations thus account for correlation in spatial heterogeneity between misperceptions and externalities. As shown by Farhi and Gabaix (2018), this correlation could have important implications for the choice of policy instruments. However, we find that in our setting accounting for spatial heterogeneity has little effect on our policy ranking compared to a scenario where misperceptions and externalities do not vary across the regions. There results can be provided by the authors.

there is not misperception and then set $m_k = 1$. For values outside this range, we then use 2% as the basis to quantify overvaluation and 12% to quantify undervaluation.

Finally, we account for uncertainty due the estimation by sampling the different demand estimates obtained from the subsampling bootstrap procedure. We sample one set of demand estimates for each income group, simulate the optimal policy, store the results, and repeat. We report the mean and standard errors across a large number of repetitions (N=500).[30]

6.4. **Results**

Table 5 presents the results from our policy simulation. For the tax instrument, the benchmark for the level of the optimal Pigouvian tax without misperceptions is 50 $/ton. If we consider the full distribution of heterogeneity in misperceptions, we found that the optimal tax is large and negative: -91.4 $/ton. A negative tax means that we should subsidize the carbon externality in this market. This surprising result is driven by the large share of consumers that overvalue energy costs, which induces a large downward adjustment to the tax. If instead we ignore heterogeneity, and we set the tax based on the average value of m, the optimal tax is 148.1 $/ton, which is more than two times the externality cost. Accounting for heterogeneous misperceptions thus has a dramatic effect on the level of the optimal tax.

Conversely, we find that accounting for heterogeneity has very little effect on the levels of the 3 types of standards we consider. For the uniform standard, we would expect this to be true, as shown in Proposition 1. For our setting, the optimal uniform standard is 335 kWh/year. Compared to the average energy consumption observed in the sample, 514.9 kWh/year (Table 1), this corresponds to a reduction of about 35% in energy usage. For our setting the minimum standard turns out to be the same as the uniform standard. This is because in our setting, the level of that optimal standard is stringent enough that it is binding for all products, making a uniform standard equivalent to a minimum standard. For the attribute-based minimum standard, which we express as a percentage reduction relative to the existing minimum standard, the optimal standard is set at 56% of the existing standard, i.e., it is 44% more stringent. However, even with this more flexible type of

[30]Note that for each income group, we have used 16 different slices of the data to implement the subsampling bootstrap. Given that there are 6 income groups, there are then 16^6 different combinations of the demand models that we can sample from for the policy simulations.

standard, we find that heterogeneous misperceptions have a negligeable effect on the optimal level of standard.

Overall, the optimal level of all three types of standards are very robust to various definitions of misperception and distributions of heterogeneity. For example, the results are robust to the way we define experienced utility. If we use a constant discount rate of 5% to define unbiased perception (rather than the 2-12% range), the optimal standards are quite close (see Table A6, Appendix A6).

Across all scenarios, standards also induce welfare gains that are much larger than the gains under the Pigouvian tax. One caveat is that we have abstracted away from firms' profits. Since standards directly impose costs on manufacturers they will negatively impact firms' profits relative to a tax. However, it is still possible to draw conclusions about the distributional effects of each policy instrument for consumers. Standards are clearly more beneficial to consumers than taxes. The ability of a quantity instrument to address misperceptions by reducing the variance in energy costs while simultaneously addressing the externality provides a clear advantage for this type of instrument and is an important take-away from our policy exercise. Without careful consideration of the effect of standards on producers' profits, the results from our exercise can be thought of as an upper bound of the relative benefits of standards relative to a tax. However, as we discuss below, compliance costs would have to be very high to change the welfare ranking of standards relative to a tax.

6.5. **Discussion: Supply-Side**

In what follows, we turn to previous analyses of the market impact of appliance standards to understand how compliance costs might affect the benefits of standards relative to a tax. Houde and Spurlock (2015), Spurlock (2013) and Brucal and Roberts (2017) are examples of recent studies that use a retrospective approach to show how past revisions in minimum appliance standards have impacted firms. These three analyses show that more stringent standards had a modest impact on prices and may in fact have lead to a reduction in equilibrium prices, and point toward the role of firms' strategic behavior as a possible explanation. Standards could have induced unexpected cost-efficiency and/or a change in market structure that lead to lower prices.[31]

[31]Houde (2018c) finds that the Energy Star certification, a voluntary energy efficiency standard, has very little impact on firms' profits in equilibrium, although it may lead to important changes in markups of a subset of products. The strategic reallocation of products in attribute space and optimal pricing gives rise

Given our policy simulation, compliance costs would have to be high, more than $100 per refrigerator sold, to change the welfare ranking of our policies. The ex post analyses of the market impact of previous changes appliance standards suggest that historical compliance costs have been much lower (Houde and Spurlock 2015; Spurlock 2013; Brucal and Roberts 2017). This suggests that appliance standards dominate a tax instrument for the historical range of compliance costs. One caveat is that we find that the optimal standards could increase stringency by 35% to 45% relative to the current standard, a larger increase than has been historically considered. One way to minimize potentially large compliance costs would be to gradually increase standards' stringency. This could provide firms additional flexibility to strategically react as past standard changes suggest.

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 therefore been used as a sufficient statistic for the average degree of consumer misperception.

This paper shows that the average degree of misperception can be a misleading statistic. From a policy design standpoint, ignoring heterogeneity and focusing on the average, can lead to very different conclusions. We illustrate these points using the U.S. appliance market and perception of energy operating costs as a case study. We show that the valuation of energy operating costs are very heterogeneous and suggest large misperceptions ranging from undervaluation to large overvaluation. The average degree of misperception, however, suggests a very modest amount of bias.

to substitution patterns that make firms almost as well-off with or without a voluntary standard. Houde (2018c)'s results do not readily apply to present context because minimum standards are likely to change the nature of the strategic responses of firms compared to a voluntary standard. Nonetheless, they suggest that firms' strategic behaviors have important implications for reducing the compliance costs of minimum standards.

We use these results to show how different policies that aim to address the carbon externality associated with energy use while accounting for misperceptions are affected by heterogeneity. We show heterogeneous misperceptions have a large impact on the level of the Pigouvian tax, but very little effect on a quantity instrument such as a standard. In fact, we show that a special type of standard, uniform standard, could even internalize misperceptions and externalities all at once. Empirically, we find that different types of standards performs relatively well from a welfare standpoint and dominate a tax for a wide range of compliance costs.

References

Abaluck, Jason and Jonathan Gruber, "Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program," *American Economic Review*, 2011, *101* (4), 1180–1210.

Allcott, Hunt, "The Welfare Effects of Misperceived Product Costs: Data and Calibrations from the Automobile Market," *American Economic Journal: Economic Policy*, 2013, *5* (3), 30–66.

———, "Paternalism and Energy Efficiency: An Overview," *Annual Review of Economics*, 2016, *8*, 145–176.

——— **and Christopher Knittel**, "Are Consumers Poorly Informed about Fuel Economy? Evidence from Two Experiments," *American Economic Journal: Economic Policy*, 2018, *forthcoming.*

——— **and Dmitry Taubinsky**, "Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market," *American Economic Review*, 2015, *105* (8), 2501–2538.

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

———**, Sendhil Mullainathan, and Dmitry Taubinsky**, "Energy Policy with Externalities and Internalities," *Journal of Public Economics*, 2014, *112* (April), 72–88.

Allen, Jason, Robert Clark, and Jean-François Houde, "Search Frictions and Market Power in Negotiated Price Markets," Technical Report, National Bureau of Economic Research 2014.

Barber, Brad M., Terrance Odean, and Lu Zheng, "Out of Sight, Out of Mind: The Effects of Expenses on Mutual Fund Flows," *Journal of Business*, 2005, *78* (6), 2095–2120.

Bernheim, B. Douglas and Dmitry Taubinsky, "Behavioral Public Economics," *NBER Working Paper*, 2018, (24828).

Blundell, Wesley, Gautam Gowrisankaran, and Ashley Langer, "Escalation of Scrutiny: The Gains from Dynamic Enforcement of Environmental Regulations," Technical Report, National Bureau of Economic Research 2018.

Bollinger, Brian, Phillip Leslie, and Alan Sorensen, "Calorie Posting in Chain Restaurants," *American Economic Journal: Economic Policy*, 2011, *3* (1), 91–128.

Brucal, Arlan and Michael Roberts, "Do Energy Efficiency Standards Hurt Consumers? Evidence from Household Appliance Sales," Technical Report, Grantham Research Institute on Climate Change and the Environment 2017.

Busse, Meghan, Christopher Knittel, and Florian 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.

DellaVigna, Stefano, "Structural Behavioral Economics," Technical Report, National Bureau of Economic Research 2018.

EIA, US, "Analysis of the impacts of the Clean Power Plan," *US Energy Information Administration. Washington, DC*, 2015, *20585.*

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

——— **and Sara Fisher Ellison**, "Search, Obfuscation, and Price Elasticities on the Internet," *Econometrica*, 2009, *77* (2), 427–452.

Farhi, Emmanuel and Xavier Gabaix, "Optimal Taxation with Behavioral Agents," *NBER Working Paper*, 2018, (21524).

Gabaix, Xavier, "A Sparsity-Based Model of Bounded Rationality," *Quarterly Journal of Economics*, 2014, *129*, 1661–1710.

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

Gertler, Paul J., Orie Shelef, Catherine D. Wolfram, and Alan Fuchs, "The energy-efficiency gap What does it mean?," *American Economic Review*, 2016, *106* (6), 1366–1401.

Grigolon, Laura, Mathias Reynaert, and Frank Verboven, "Consumer Valuation of Fuel Costs and the Effectiveness of Tax Policy: Evidence from the European Car Market," *American Economic Journal: Economic Policy: Forthcoming*, 2018.

Guiso, Luigi, Andrea Pozzi, Anton Tsoy, Leonardo Gambacorta, and Paolo Mistrulli, "The Cost of Steering in Financial Markets: Evidence from the Mortgage Market," 2018.

Handel, Benjamin R and Jonathan T Kolstad, "Health insurance for" humans": Information frictions, plan choice, and consumer welfare," *American Economic Review*, 2015, *105* (8), 2449–2500.

Harrison, Glenn W and John A List, "Field Experiments," *Journal of Economic literature*, 2004, *42* (4), 1009–1055.

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

Hausman, Jerry A. and Paul L. Joscow, "Evaluating the Costs and Benefits of Appliance Efficiency Standards," *American Economic Review Papers and Proceedings*, 1982, *72* (2), 220–225.

Heiss, Florian, Daniel McFadden, Joachim Winter, Amelie Wuppermann, and Bo Zhou, "Inattention and Switching Costs as Sources of Inertia in Medicare Part D," Working Paper 22765, National Bureau of Economic Research October 2016.

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, Sébastien, "Bunching with the Stars: How Firms Respond to Environmental Certification," *Economics Working Paper Series*, 2018, *18* (292).

———, "How Consumers Respond to Product Certification and the Value of Energy Information," *RAND Journal of Economics*, 2018, *49* (2), 453–477.

———, "The Incidence of Coarse Certification: Evidence from the Energy Star Program," *Economics Working Paper Series*, 2018, *18* (290).

——— **and Cecilia A. Spurlock**, "Do Energy Efficiency Standards Improve Quality? Evidence from a Revealed Preference Approach," *Working Paper: University of Maryland*, 2015.

——— **and Erica Myers**, "Are Consumers Attentive to Local Energy Costs? Evidence from the Appliance Market," *CER Working paper*, 2019, *19* (313).

——— **and Joseph E. Aldy**, "Consumers' Response to State Energy Efficient Appliance Rebate

Programs," *American Economic Journal: Economic Policy*, 2017, *9* (4), 227–255.

——— **and** ———, "The Efficiency Consequences of Heterogeneous Behavioral Responses to Energy Fiscal Policies," Working Paper 24103, National Bureau of Economic Research December 2017.

Jensen, Robert, "The Perceived Returns to Education and the Demand for Schooling," *Quarterly Journal of Economics*, 2010, *125* (2), 515–548.

Ketcham, Jonathan D, Nicolai V Kuminoff, and Christopher A Powers, "Estimating the Heterogeneous Welfare Effects of Choice Architecture: An Application to the Medicare Prescription Drug Insurance Market," Working Paper 22732, National Bureau of Economic Research October 2016.

Ketcham, Jonathan D., Nicolai V. Kuminoff, and Christopher A. Powers, "Estimating the Heterogeneous Welfare Effects of Choice Architecture: an Application to the Medicare Prescription Drug Insurance Market," *NBER Working Paper*, 2016, (22732).

Kling, Jeffrey, Sendhil Mullainathan, Eldar Shafir, Lee Vermeulen, and Marian Wrobel, "Comparison Friction: Experimental Evidence from Medicare Drug Plans," *Quarterly Journal of Economics*, 2012, *127* (1), 199–235.

Koszegi, Botond and Adam Szeidl, "A Model of Focusing in Economic Choice," *Quarterly Journal of Economics*, 2013, *128*, 53–107.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 1997, *112*, 443–477.

Leggett, Christopher G, "Environmental Valuation with Imperfect Information the Case of the Random Utility Model," *Environmental and Resource Economics*, 2002, *23* (3), 343–355.

Myers, E., "Are Home Buyers Inattentive? Evidence From Capitalization of Energy Costs," *American Economic Journal: Economic Policy: Forthcoming*, 2018.

Nevo, Aviv, John L Turner, and Jonathan W Williams, "Usage-Based Pricing and Demand for Residential Broadband," *Econometrica*, 2016, *84* (2), 411–443.

Politis, Dimitris N, Joseph P Romano, and Michael Wolf, "Subsampling Springer-Verlag New York," 1999.

Revelt, D. and K. Train, "Mixed Logit with Repeated Choices: Households' Choices of Appliance Efficiency Level," *Review of Economics and Statistics*, 1998, *8* (4), 647–657.

Sallee, James M., "Rational Inattention and Energy Efficiency," *Journal of Law and Economics*, 2014, *57*, 781–820.

Sallee, James M, Sarah E West, and Wei 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.

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.

Small, K.A. and H.S. Rosen, "Applied Welfare Economics with Discrete Choice Models," *Econometrica*, 1981, *49* (1), 105–130.

Spurlock, C. Anna, "Appliance Efficiency Standards and Price Discrimination," Working Paper, Lawrence Berkeley National Laboratory May 2013.

Taubinsky, D. and A. Rees-Jones, "Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment," *Review of Economic Studies: Forthcoming*, 2018.

Train, K., "Discount Rates in Consumers' Energy-Related Decisions: A Review of the Literature," *Energy*, 1985, *10* (12), 1243–1253.

Train, Kenneth E, *Discrete Choice Methods with Simulation*, Cambridge university press, 2009.

TABLE 1. Summary Statistics

	Mean	SD
Attribute Information		
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 Information		
% 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	-

Notes: The retailer's data do not explicitly identify transactions made by households. We observe a unique identifier for each customer's credit card. We classify as "non-household" cases where we observe more than one purchase of a full-size refrigerator during the period: 2008-2012 for a single identifier (credit card).

TABLE 2. Conditional Logit Results

	I	II	III	IV	V	VI	VII
Purchase Price	-0.00348***	-0.00348***	-0.00348***	-0.00337***	-0.00338***	-0.00331***	-0.00338***
	(0.00007)	(0.00007)	(0.00007)	(0.00007)	(0.00007)	(0.00007)	(0.00007)
Energy Cost	-0.02901***	-0.03045***	-0.03088***	-0.03081***	-0.02730***	-0.03026***	-0.02871***
	(0.00318)	(0.00318)	(0.00318)	(0.00320)	(0.00408)	(0.00323)	(0.00444)
m	0.7131	0.7485	0.7591	0.782	0.691	0.781	0.726
Product FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demo × Attributes	No	Yes	Yes	Yes	Yes	Yes	Yes
Brand × Week	No	No	No	Yes	Yes	Yes	Yes
County × EStar	No	No	Yes	Yes	No	No	Yes
County × TopFreezer	No	No	No	No	Yes	No	Yes
County × Size	No	No	No	No	No	Yes	No

Notes: All specifications include controls for energy star certification and rebates offered. Standard errors (in parentheses) clustered at the zip code level. The average m is computed assuming a discount rate of 5% and a refrigerator lifetime of 18 years.

TABLE 3. Marginal PDF of m: Robustness and Heterogeneity w.r.t. Income

(Mis-) Perception Parameter	m<0	m≥0, m< 0.62	m≥ 0.62, m< 1.28	m≥ 1.28	E[m]	E[m]
Implied Discount Rate		$r > 12\%$	$r > 2\%$, $r \leq 12\%$	$r \leq 2\%$	r = 5%	r =[2%, 12%]
Model I:	12.63	37.53	48.42	17.00	0.77	0.68
Main Specification	(2.27)	(5.37)	(5.70)	(3.80)	(0.08)	(0.07)
Robustness Tests						
Model II:	12.94	35.31	45.08	19.13	0.73	0.66
Smaller Grid	(2.20)	(5.22)	(5.84)	(3.16)	(0.07)	(0.06)
Model III:	13.56	42.47	39.00	18.67	0.76	0.69
Denser Grid	(2.15)	(6.06)	(7.31)	(4.03)	(0.08)	(0.07)
Model IV:	15.83	30.33	48.00	19.36	0.78	0.67
τ Random	(2.67)	(7.83)	(8.29)	(5.46)	(0.10)	(0.08)
Model V:	15.38	35.94	32.06	16.63	0.74	0.67
τ & γ_j Random	(1.98)	(3.37)	(2.91)	(2.05)	(0.06)	(0.06)
Model VI:	13.94	22.81	36.63	26.75	0.88	0.68
No Demo × Att	(1.51)	(3.11)	(5.16)	(4.52)	(0.06)	(0.05)
Heterogeneity w.r.t. Income						
< \$30k, Model I	27.43	49.00	17.86	6.67	0.35	0.33
	(6.71)	(10.16)	(8.72)	(2.54)	(0.07)	(0.11)
\$30-50k, Model I	17.81	56.63	25.75	7.14	0.48	0.51
	(3.82)	(6.44)	(7.84)	(1.26)	(0.07)	(0.07)
\$50-75k, Model I	12.75	43.88	31.85	18.80	0.73	0.66
	(2.70)	(6.99)	(6.76)	(2.79)	(0.10)	(0.09)
\$75-100k, Model I	15.31	35.73	32.80	21.93	0.82	0.64
	(2.05)	(5.39)	(6.56)	(2.46)	(0.07)	(0.07)
\$100-125k, Model I	19.56	28.92	41.07	23.00	0.94	0.66
	(4.61)	(6.15)	(5.78)	(3.91)	(0.12)	(0.12)
≥\$125k, Model I	13.38	31.46	34.81	26.25	1.00	0.80
	(2.48)	(6.34)	(5.76)	(4.96)	(0.12)	(0.11)

Notes: Marginal PDF of m computed from the joint PDF of η and m. The two-step estimation was performed for model and for each income group separately. Standard errors (in parentheses) are obtained using subsampling bootstrap. The first four columns refer to different bins of the pdf, which each maps into a different type of misperception. Each column corresponds to a type based on a range of implied discount rates that the parameter m translates into. The last two columns is the average amount of misperception implied by the full distribution assuming $r = 5\%$ for all consumers. For the last column, all values of m that correspond to a discount rate between 2% and 12% were set to $m = 1$.

Table 4. Matched Pairs: Share of Inattentive Consumers

	Time Interval Between Sales Within Pair Same Day	+/- 1 Day	+/- 2 Days	+/- 4 Days	+/- 6 Days	+/- 12 Days	+/- 24 Days
Dominated Option	48.03%	45.85%	44.44%	43.31%	43.18%	41.27%	38.27%
Δ Price ($)	-19.70	-21.86	-25.05	-29.38	-33.13	-39.24	-45.61
# Obs	304	325	360	441	528	773	1304

Notes: The percentage in the first row is the fraction of consumers that chose the dominated option. The second row is the mean size of the price difference during the events where the most efficient model was offered at a lower price. Each column represents the time interval between the sales of two models of the same pair in the same store. For instance, "+/- 2 Days" means that the sales between two models of the same pair were at most two days apart.

Table 5. Optimal Behavioral Policies

	Misperceptions with Heterogeneous r Policy	Δ SW $/capita	Δ CS $/capita	Δ Ext $/capita
Pigou Tax ($/ton of CO_2)				
With $F(m_k)$	-91.4	3.9	208.7	1.4
	(0.38)	(0.06)	(0.89)	(0.01)
With $E[m_k]$	148.1	3.9	-312.4	-2.6
	(1.80)	(0.08)	(3.76)	(0.03)
Uniform Standard (kWh/year)				
With $F(m_k)$	335.3	128.0	90.8	-37.1
	(0.00)	(0.15)	(0.14)	(0.02)
With $E[m_k]$	335.3	103.2	67.4	-35.8
	(0.00)	(0.11)	(0.09)	(0.03)
Minimum Standard (kWh/year)				
With $F(m_k)$	335.3	128.0	90.8	-37.1
	(0.00)	(0.15)	(0.14)	(0.02)
With $E[m_k]$	335.3	103.2	67.4	-35.8
	(0.00)	(0.11)	(0.09)	(0.03)
Attribute-Based Minimum Standard (% of Existing Standard)				
With $F(m_k)$	55.0	99.5	63.6	-35.8
	(0.03)	(0.10)	(0.08)	(0.04)
With $E[m_k]$	56.1	80.9	47.8	-33.1
	(0.01)	(0.10)	(0.10)	(0.02)

Notes: The table reports the optimal policies evaluated using the full distribution of misperceptions (with $F(m_k)$) or only the average misperception (with $E[m]$). For each of the two cases, the change in social welfare (SW), consumer surplus before redistribution of the tax revenues (CS), and externality costs (Ext) are computed relative to the case without a policy. The mean and standard errors of the optimal taxes, standards, and welfare metrics are computed using 500 random draw of the estimated distributions.

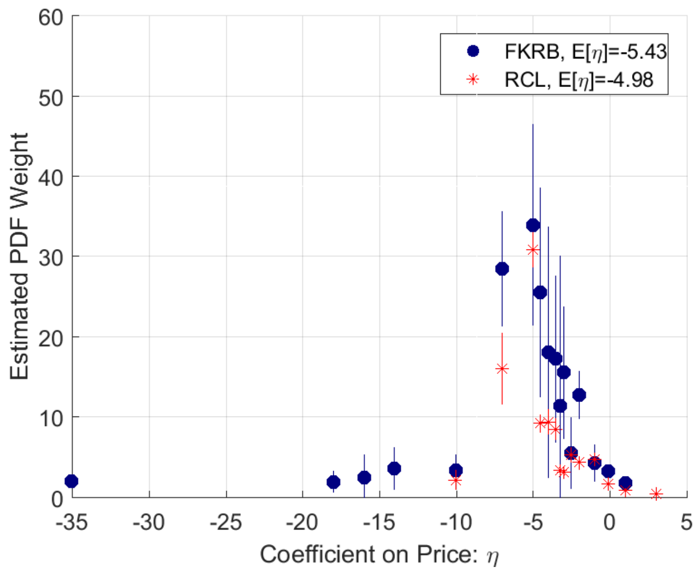

(a) Marginal PDF of η: $f(\eta)$

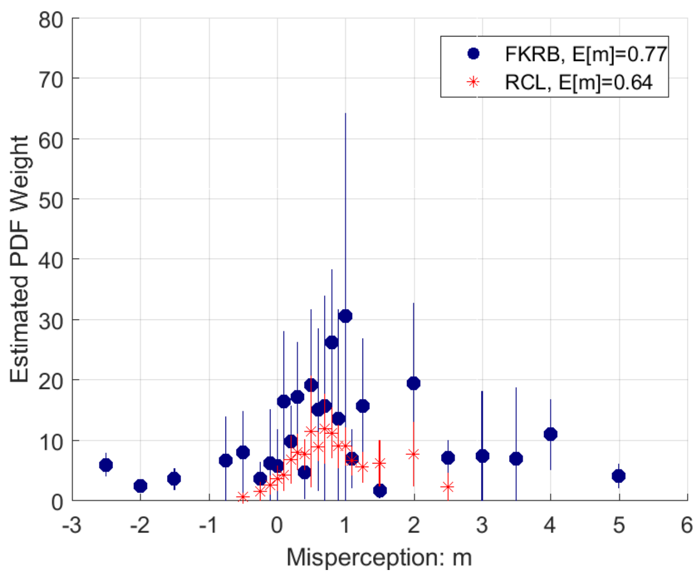

(b) Marginal PDF of m: $f(m)$

Figure 1. Marginal Probability Distributions for η and m
Notes: : The blue dots represent the estimated pdf with the FKRB estimator and the 95% confidence interval for each estimated pdf weight. The red stars represent the estimated pdf implied by the first-stage estimation with the parametric random coefficient logit.

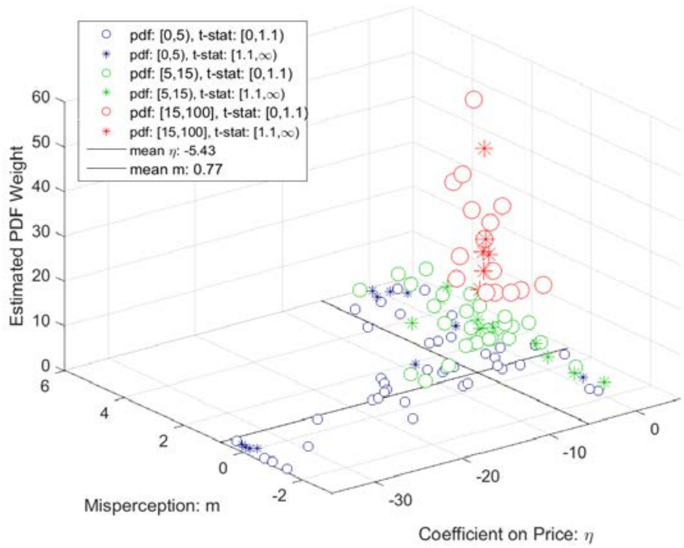

Figure 2. Joint Probability Distribution for η and m
Notes: : Stars represent pdf weights with t-statistics greater than 1.1 and dots represent pdf weights with t-statistics less than 1.1. Color and marker size represent the level of the pdf weights, where bigger markers capture larger weights.

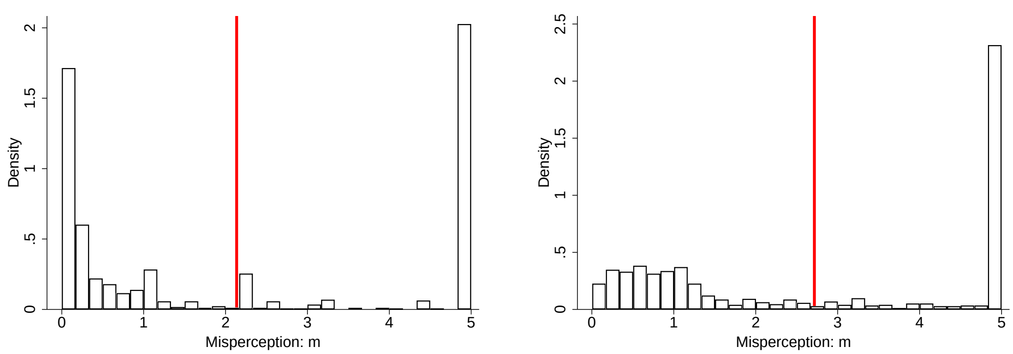

(a) Beliefs: Refrigerator Usage

(b) Beliefs: Electricity Price

Figure 3. Beliefs Valuation Ratios: Refrigerator Usage and Electricity Price
Notes: The red bar identifies the mean valuation ratio. All the values above 5 were rounded at 5 for exposition purposes only. The correct energy usage of refrigerator is assumed to be approximately the national average for 2017: 450 kWh/year. The correct electricity price is assumed to be the average price at the county level for the corresponding zip code a survey respondent lived at the time of answering the survey.

Appendix

For Online Publication

A1. Figures of Price and Energy Cost Variation

Figure A1 displays an example of the variation in weekly price for over the study period for the top 9 sales ranked models from a particular refrigerator brand. The red line shows the median weekly price change across all zip codes with the 25th and 75th percentile depicted with gray bands. The blue line shows the remaining variation in weekly price change after controlling for week-of-sample fixed effects interacted with brand dummies. Even with these controls, large variation in in price remains suggesting that the model-specific variation we observe is highly idiosyncratic and is generated by the retailer's dynamic pricing algorithm. We exploit this variation to identify the coefficient on price.

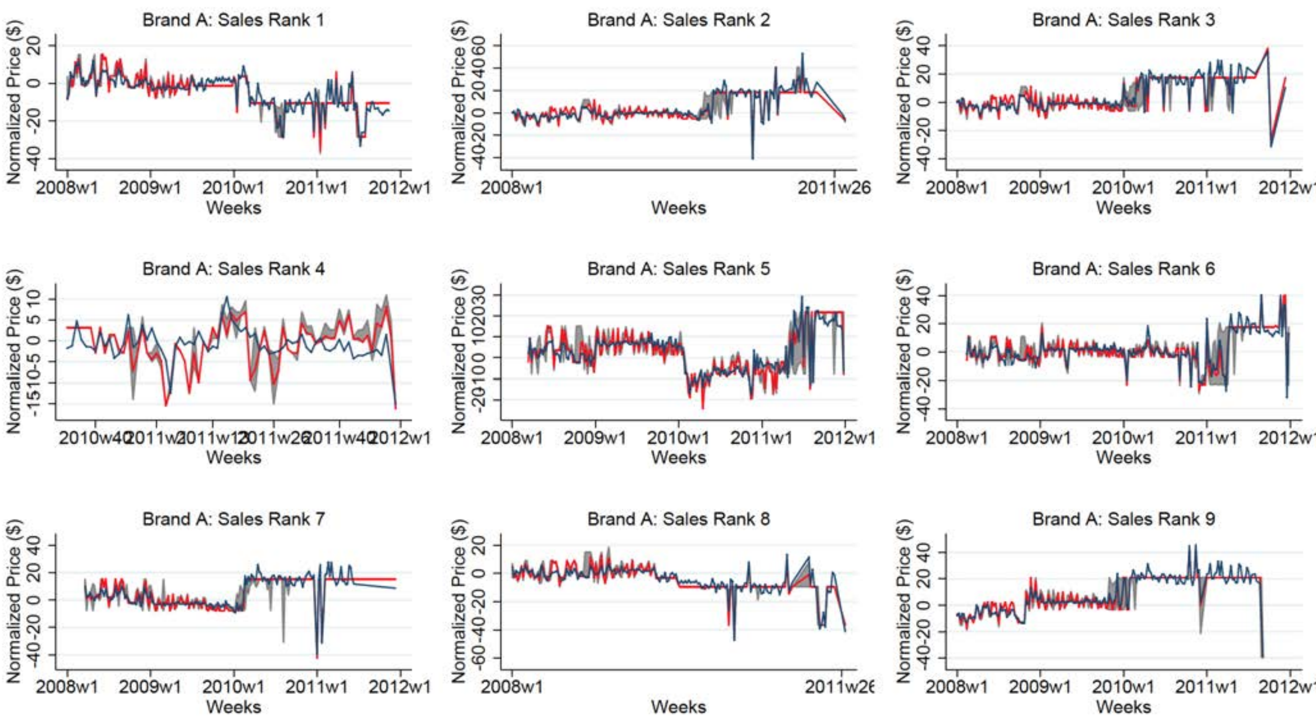

Figure A1

Figures A2 gives a sense of the variation in electricity price across regions and time. It depicts mean sales-weighted annual electricity price for each state in a U.S. census division. Each line represents the sales weighted average price for a state plotted for each of the 9 U.S. census divisions. Prices vary quite a bit regionally, with the highest prices in New England and the lowest prices in the Midwest and South. There is also variation over time in price with some states experiencing price increases and others price decreases over the study period.

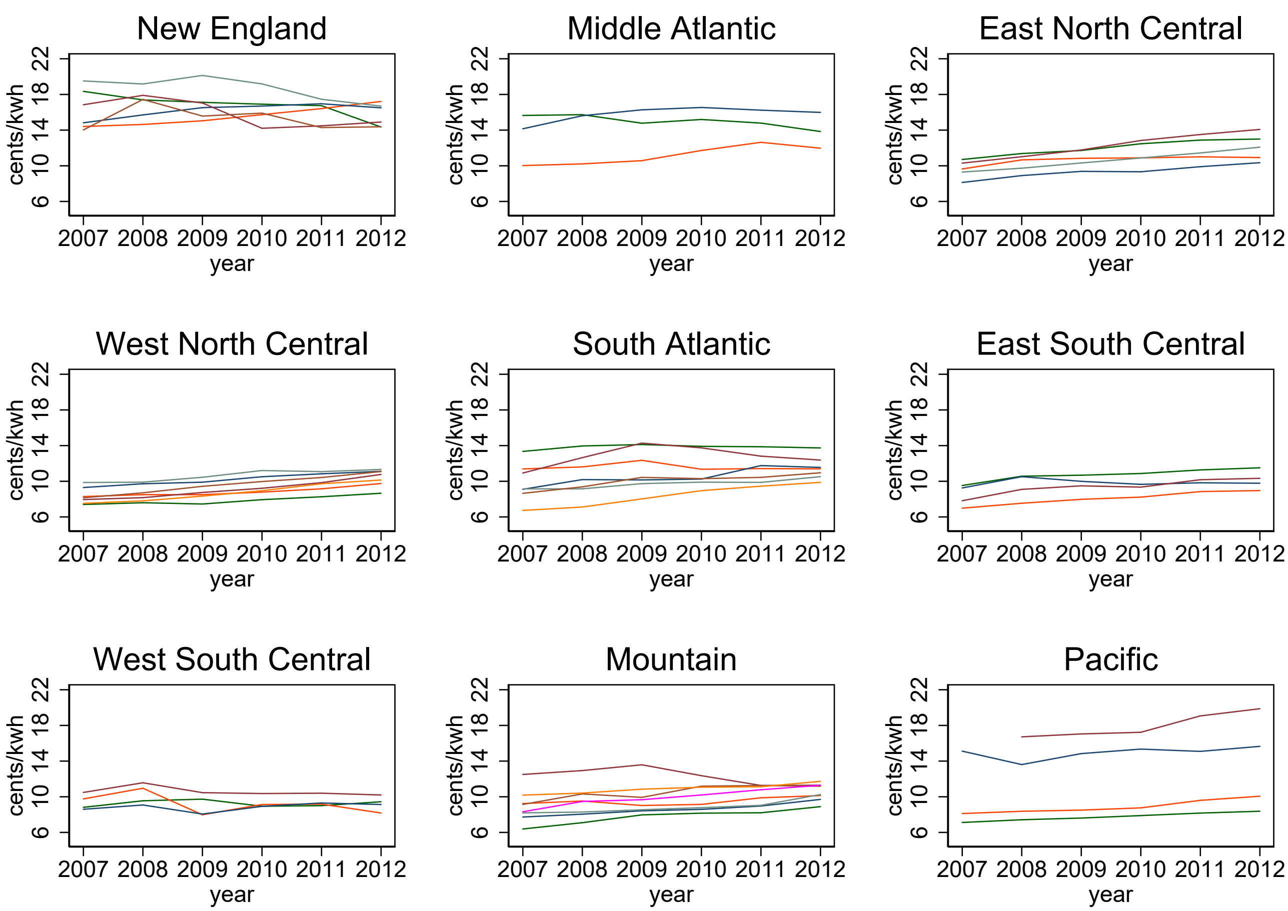

Figure A2. Average Electricity Prices for Each State in a Census Divsion

Figure A3 displays density plots of model's prices and lifetime energy costs. The first panel plots the distribution of prices paid. Almost all models sold in our sample are less than $2000, though there are some much higher priced models available. The second panel shows how the the distribution of life time energy cost of the models sold varies for the 10th, 50th, and 90th percentile of electricity price. Using an expected lifetime of 18 years and a 5% discount rate the mean lifetime costs range from \$555 for the 10th percentile of energy price to \$1000 for the 90th percentile. The third panel shows the distribution of the ratio of lifetime cost to purchase price, with means ranging from .44 at the 10th percentile of energy price to .79 at the 90th percentile. This variation in energy costs for particular models across time and space identify the coefficient on energy costs.

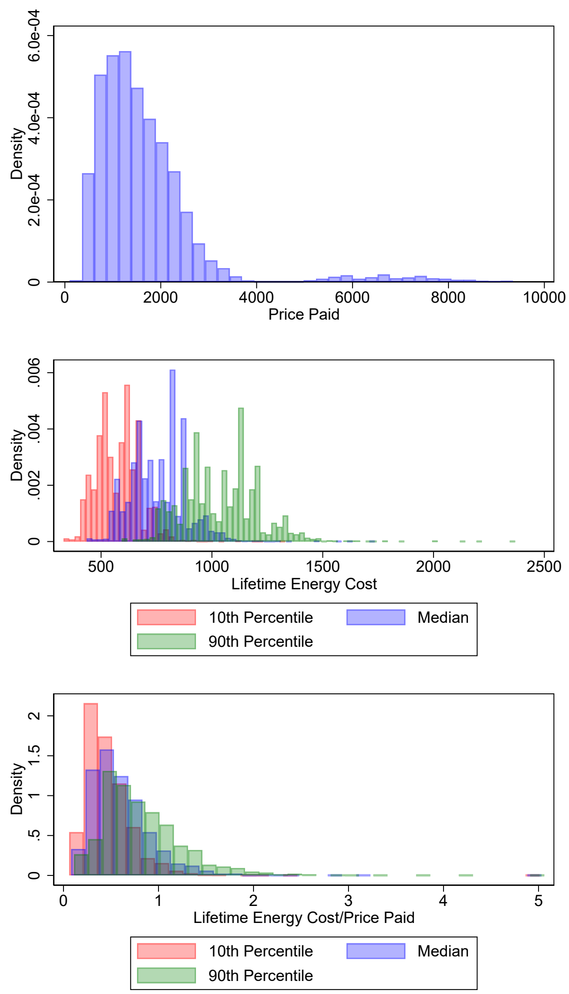

FIGURE A3. Distributions of Prices and Lifetime Energy Costs

A2. **Approximation of the Distribution of Misperceptions**

Suppose that in choice model given by Equation 1, the parameters η and θ vary across consumers such that:

$$U_{ij} = \gamma_j - \eta_i P_j - \theta_i E_j + \epsilon_{ij} \tag{A1}$$

where η_i and θ_i each follows a joint distribution $F(\eta, \theta)$ with well defined first of second moments for the marginal and joint distributions. The following proposition shows that it is possible to approximate the first moment of the distribution of the statistics: $m_i = \theta_i/\eta_i$.

Proposition 2. *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:*

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

Proof. The ratio m is a function of the two variables: θ and η, which we define as $m = f(\theta, \eta) = \theta/\eta$. The proof of the proposition is obtained by simply taking a first-order and second-order, respectively, Taylor series expansion around the mean of θ and η, which we denote: $E[\theta]$ and $E[\eta]$. Remember that for any bivariate function: $f(\theta, \eta)$, a second-order Taylor series expansion around $(E[\theta], E[\eta])$ is given by:

(A2)

$$\begin{aligned} f(\theta,\eta) \approx & f(E[\theta],E[\eta]) + (\theta - E[\theta])f_\theta(E[\theta],E[\eta]) + (\eta - E[\eta])f_\eta(E[\theta],E[\eta]) + \\ & \frac{1}{2}\Big((\theta - E[\theta])^2 f_{\theta\theta}(E[\theta],E[\eta]) + 2(\theta - E[\theta])(\eta - E[\eta])f_{\theta\eta}(E[\theta],E[\eta]) + (\eta - E[\eta])^2 f_{\eta\eta}(E[\theta],E[\eta])\Big) \end{aligned}$$

where the subscripts denote the first and second derivatives of the function $f()$.

We can readily see that the approximation of $E[m] = E[f(\theta, \eta)]$ at $(E[\theta], E[\eta])$ using the first-order expansion yields:

(A3)

$$\begin{aligned} E[f(\theta,\eta)] \approx & E[f(E[\theta],E[\eta])] + E[(\theta - E[\theta])] \cdot E[f_\theta(E[\theta],E[\eta])] + E[(\eta - E[\eta])] \cdot E[f_\eta(E[\theta],E[\eta])] \\ & \equiv E[f(E[\theta],E[\eta])] + 0 \cdot E[f_\theta(E[\theta],E[\eta])] + 0 \cdot E[f_\eta(E[\theta],E[\eta])] \\ & \equiv E[f(E[\theta],E[\eta])] \\ & \equiv \frac{E[\theta]}{E[\eta]} \end{aligned}$$

For the second order expansion, we need the terms: $f_{\theta\theta}(\theta, \eta) = 0$, $f_{\theta\eta}(\theta, \eta) = -\frac{1}{\eta^2}$, and $f_{\theta\eta}(\theta, \eta) = \frac{2\theta}{\eta^3}$. Replacing these terms in A2, and using the results of A3, the approximation of $E[m] = E[f(\theta, \eta)]$ at $(E[\theta], E[\eta])$ using the second-order expansion is:

(A4)
$$\begin{aligned} E[f(\theta,\eta)] \approx & \frac{E[\theta]}{E[\eta]} + \frac{1}{2}E\left[-2(\theta - E[\theta])(\eta - E[\eta]) \cdot \frac{1}{E[\eta]^2} + (\eta - E[\eta])^2 \cdot \frac{2E[\theta]}{E[\eta]^3}\right] \\ & \equiv \frac{E[\theta]}{E[\eta]} - \frac{Cov(\theta,\eta)}{E[\eta]^2} + \frac{Var(\eta)E[\theta]}{E[\eta]^3} \end{aligned}$$

Note that we use the fact that $E[(\theta - E[\theta])(\eta - E[\eta])] = Cov(\theta, \eta)$ and $E[(\eta - E[\eta])^2] = Var(\eta)$.

□

Proposition 2 shows that, even if one is only interested in the average degree of misperception, heterogeneous responses to both prices and energy costs and how these responses are correlated must be considered. Note that several studies that have quantified misperceptions of energy costs only report the first-order Taylor approximation shown in Proposition 2. This approximation induces a bias of the order of:

(A5)
$$-\frac{Cov[\eta,\theta]}{E[\eta]^2} + \frac{Var[\eta] \cdot E[\theta]}{E[\eta]^3} + \mathcal{O}(n^{-1})$$

where $\mathcal{O}(n^{-1})$ represents the higher order terms that are not captured by the second-order Taylor approximation. To characterize the distribution of misperceptions, Proposition 2 can be extended to approximate the higher moments of the distribution of m_i, which will also be subject to approximation biases.

A3. **Monte Carlo**

A3.1. **Performance of the Random Coefficient Logit: One-Choice Situation Case**

We investigate the performance of the random coefficient logit for the case where we observe only one choice for each decision-maker. We vary three dimensions: the number of alternatives (3, 6, or 10), the number of decision-makers (1,000, 10,000, or 25,000), and the parametric distribution of preferences. The data generating process is based on the following model. Decision-makers make a discrete choice among J alternatives where the alternative-specific utility of option j for consumer i is function of the price, denoted P_j, a second attribute that may represent energy operating cost, denoted E_j, overall quality, denoted γ_j, and an idiosyncratic taste parameter, denoted ϵ_{ij}:

$$U_{ij} = \eta_i P_j + \theta_i E_j + \delta_j + \epsilon_{ij} \tag{A6}$$

The preference parameters η_i and θ_i are random coefficients that follow a specific parametric distribution. Finally, we assume that ϵ_{ij} is i.i.d. and follows a extreme value distribution. This gives rise to the random coefficient logit.

For the first set of scenarios, we consider that the parameters η_i and θ_i follow a bivariate normal: $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\boldsymbol{\mu} = [\mu_\eta = -1, \mu_\theta = -0.7]$ and $\boldsymbol{\Sigma} = [\Sigma_{1,1} = 0.5, \Sigma_{1,2} = \Sigma_{2,1} = 0.25, \Sigma_{2,2} = 0.75]$. We generate 100 datasets for a given number of alternatives, J, and number of decision-makers, N. The price P_j, the attribute E_j, and quality γ_j are generated with a multivariate standard normal distribution. We assume that there is no correlation between the level of each of those variables across alternatives.

For each simulated dataset, we estimate a random coefficient logit where the preferences are specified with a bivariate normal. We recover the estimates of the mean and the full covariance matrix via simulated maximum likelihood and use 500 draws to simulate the distribution of preferences. Results are presented in Table A1. For a given J and N, we report that average of the point estimates across the 100 datasets. A number of important patterns emerge. As the number of observations increases, the bias in the estimate of the mean vector $\boldsymbol{\mu}$ quickly goes to zero. This is, however, not the case for the components of the covariance matrix: $\boldsymbol{\Sigma}$. Even for a large number of observations, e.g., $N = 25,000$, the estimated covariance matrix differs from the true data generating process. Increasing the number of alternatives helps identifying the components of the covariance matrix. This is intuitive. Adding alternatives generates more substitution events that violate the independence of irrelevant alternative (IIA) assumption. This allows us to pin down the preference parameters of the covariance matrix, which is responsible for the violation of this assumption.

For the second and third sets of scenarios, we consider the effect of mis-specifying the parametric distribution of preferences. In particular, for the data generating process, we consider that preferences are generated by a mixture of two bivariate normal: $\omega_1\mathcal{N}(\boldsymbol{\mu_1}, \boldsymbol{\Sigma_1}) + (1-\omega_1)\mathcal{N}(\boldsymbol{\mu_2}, \boldsymbol{\Sigma_2})$, where $\omega_1 = 0.7$. We first set $\boldsymbol{\mu_1} = [\mu_\eta^1 = -2.0, \mu_\theta^1 = -1.5]$ and $\boldsymbol{\Sigma_1} = [\Sigma_{1,1}^1 = 0.5, \Sigma_{1,2}^1 = \Sigma_{2,1}^1 = 0.25, \Sigma_{2,2}^1 = 0.75]$ and $\boldsymbol{\mu_2} = [\mu_\eta^2 = -0.5, \mu_\theta^2 = -0.05]$ and $\boldsymbol{\Sigma_2} = [\Sigma_{1,1}^2 = 0.15, \Sigma_{1,2}^2 = \Sigma_{2,1}^2 = -0.05, \Sigma_{2,2}^2 = 0.10]$. We then consider a mixture where the two modes have a larger difference and set $\boldsymbol{\mu_1} = [\mu_\eta^1 = -6.0, \mu_\theta^1 = -4.5]$ and $\boldsymbol{\mu_2} = [\mu_\eta^2 = -0.5, \mu_\theta^2 = -0.05]$. For this latter case, all other preference parameters definying the mixture distribution are the same.

Results are presented in A2. The main take-away is that the mis-specified model does well in recovering the mean of the mixture of distribution, but does not perform well in recovering the covariance matrix. This is especially the case if the two modes of the mixture distribution are very different. Adding alternatives or observations improves the estimates somewhat, but large biases remain. Note that the mixing of two normal distributions may not result in a symmetric distribution, but the mean is always a simple weighted average of the mean of each mode. By specifying a random coefficient logit with a normal distribution, it is therefore difficult to capture the skewness in the resulting mixture distribution, but the mean can easily be recovered.

Table A1. Monte Carlo: Performance of the Random Coefficient Logit for the One-Choice Situation Case

	True Value	Estimates N=1,000	Estimates N=10,000	Estimates N=25,000	% Bias N=1,000	% Bias N=10,000	% Bias N=25,000
# Alternatives: 3							
η	-1.00	-1.05	-1.02	-0.99	4.84	1.75	-0.88
θ	-0.70	-0.77	-0.70	-0.69	10.03	-0.05	-0.98
$Cov_{1,1}$	0.50	1.57	0.62	0.46	214.46	23.53	-7.91
$Cov_{2,1}$	0.25	0.30	0.11	0.19	20.64	-55.33	-24.42
$Cov_{2,2}$	0.75	2.38	0.89	0.80	217.28	18.21	6.08
# Alternatives: 6							
η	-1.00	-1.05	-1.00	-0.99	5.00	0.13	-0.52
θ	-0.70	-0.72	-0.70	-0.70	3.10	-0.12	-0.55
$Cov_{1,1}$	0.50	0.92	0.48	0.41	83.87	-3.48	-17.74
$Cov_{2,1}$	0.25	0.27	0.21	0.26	6.32	-17.61	4.47
$Cov_{2,2}$	0.75	1.11	0.69	0.71	48.63	-8.46	-4.71
# Alternatives: 10							
η	-1.00	-1.03	-1.00	-0.99	2.50	-0.06	-0.82
θ	-0.70	-0.70	-0.71	-0.70	-0.52	0.84	-0.08
$Cov_{1,1}$	0.50	0.87	0.43	0.44	74.85	-14.07	-11.55
$Cov_{2,1}$	0.25	0.19	0.25	0.24	-25.71	0.37	-2.24
$Cov_{2,2}$	0.75	1.07	0.65	0.72	42.21	-13.41	-4.31

Notes: The column referring to the true value lists the parameter values used to generate the data. 100 hundred datasets were generated. The columns labeled "Estimates" are the estimates recovered in the estimation. The columns labeled "% Bias" are the percentage bias of the estimated values relative to the true parameter values. The different columns refer to the number of decision-makers in the true data generating process: $N = 1,000$, $10,000$, or $25,000$. The three panels refer to the number of alternatives used in the data generating process: $J = 3$, 6, or 10.

Table A2. Monte Carlo: Performance of the Random Coefficient Logit with Mis-specified Distribution

	True Value			Estimates			% Bias		
				N=1,000	N=10,000	N=25,000	N=1,000	N=10,000	N=25,000
	$\mathcal{N}(\boldsymbol{\mu_1}, \boldsymbol{\Sigma_1})$	$\mathcal{N}(\boldsymbol{\mu_2}, \boldsymbol{\Sigma_2})$	Mixture						
Mixture I. Small difference in means									
ω_i	0.70	0.30							
η	-2.00	-0.50	-1.55	-1.64	-1.55	-1.55	6.06	0.29	-0.31
θ	-1.50	-0.05	-1.07	-1.14	-1.05	-1.06	6.58	-1.47	-0.70
$Cov_{1,1}$	0.50	0.15	0.87	2.03	1.00	0.91	139.40	18.18	7.43
$Cov_{2,1}$	0.25	-0.05	0.62	0.51	0.49	0.58	-15.53	-19.02	-3.87
$Cov_{2,2}$	0.75	0.10	1.00	2.38	1.13	0.98	138.47	12.76	-1.63
Mixture II. Large difference in means									
ω_i	0.70	0.30							
η	-6.00	-0.50	-4.35	-4.68	-4.50	-4.48	7.64	3.47	3.08
θ	-4.50	-0.05	-3.17	-3.44	-3.27	-3.27	8.73	3.26	3.23
$Cov_{1,1}$	0.50	0.15	6.75	11.34	10.04	9.99	68.06	48.81	47.95
$Cov_{2,1}$	0.25	-0.05	5.30	8.36	7.98	7.90	57.84	50.62	49.01
$Cov_{2,2}$	0.75	0.10	4.70	9.48	7.01	6.82	101.66	49.05	45.12

Notes: The column referring to the true value of the mixture distribution corresponds to the mean and covariance matrix of a unimodal distribution fitted on the true mixture distribution of preferences. We refer to it as the true unimodal distribution. The random coefficient logit specified with a bivariate normal is compared to this distribution. The true values of the covariance matrix of the mixture distribution are computed numerically. The columns labeled "Estimates" are the estimates recovered in the estimation. The columns labeled "% Bias" are the percentage bias of the estimated values relative to the true parameter values.

A3.2. **Performance of the FKRB Estimator**

The FKRB estimator is implemented as two-step estimator. In the first step, all the parameters, δ_j, η, and θ of the discrete choice model (Equation A6) are estimated using a conditional logit or a random coefficient logit with bivariate normal for the distribution of η, and θ. In the second step, the product fixed effects, γ_j estimated in the first step are treated as data and the FKRB estimator is implemented with non-parametric joint distribution specified for η, and θ.

Table A3 compares the mean of the estimated non-parametric distribution of the parameters η and θ. In the first panel, the product fixed effects are from the conditional logit. In the second panel, the product fixed effects are from the random coefficient logit. The two approaches yield similar results if the true distribution of preferences follows the simple bivariate normal. However, if the true data generating process follows a more complex mixture distribution, using the product fixed effects from the conditional logit can lead to large bias in the estimated means of η and θ. Using product fixed effects from the random coefficient logit substantially reduces the bias. The difference between the two approaches also impacts the overall distribution.

Figure A4 compares the estimated non-parametric distributions under the two approaches for the case where the true data generating process is a mixture with two modes that have a large difference in means. When the product fixed effects from the random coefficient logit are used, the estimated distribution fits the true distribution well and captures the two modes.

Table A3. Monte Carlo: Performance of the Random Coefficient Logit with Mis-specified Distribution

	True Distribution	True Value	Estimates N=1,000	Estimates N=10,000	Estimates N=25,000	% Bias N=1,000	% Bias N=10,000	% Bias N=25,000
Product Fixed Effects from Conditional Logit								
η	Bivariate normal	-1.00	-1.05	-1.03	-0.99	5.17	2.78	-0.97
θ		-0.70	-0.78	-0.71	-0.70	12.12	1.13	0.27
η	Mixture I.	-1.55	-1.65	-1.54	-1.52	6.17	-0.71	-1.72
θ		-1.07	-1.13	-1.05	-1.05	6.03	-1.77	-1.54
η	Mixture II.	-4.35	-3.72	-3.65	-3.64	-14.44	-16.16	-16.40
θ	Large difference in means	-3.17	-2.71	-2.61	-2.62	-14.31	-17.40	-17.21
Product Fixed Effects from Random Coefficient Logit								
η	Bivariate normal	-1.00	-1.12	-1.05	-1.01	12.12	5.34	1.07
θ		-0.70	-0.85	-0.73	-0.72	21.87	3.98	2.90
η	Mixture I.	-1.55	-1.76	-1.59	-1.56	13.48	2.38	0.92
θ		-1.07	-1.21	-1.08	-1.08	13.82	1.63	1.73
η	Mixture II.	-4.35	-4.99	-4.56	-4.52	14.80	4.78	3.83
θ	Large difference in means	-3.17	-3.71	-3.32	-3.29	17.07	4.83	4.05

Notes: The first panel report results where the fixed effects in the discrete choice model are first estimated using a conditional logit and then treated as data in the FKRB estimator. The second panel reports results where the fixed effects are estimated with a random coefficient logit where the preference parameters follow a bivariate normal. For all scenarios, the number of alternatives if fixed to 3, which means that two product fixed effects need to be estimated. The columns labeled "Estimates" are the estimates recovered in the estimation. The columns labeled "% Bias" are the percentage bias of the estimated values relative to the true parameter values.

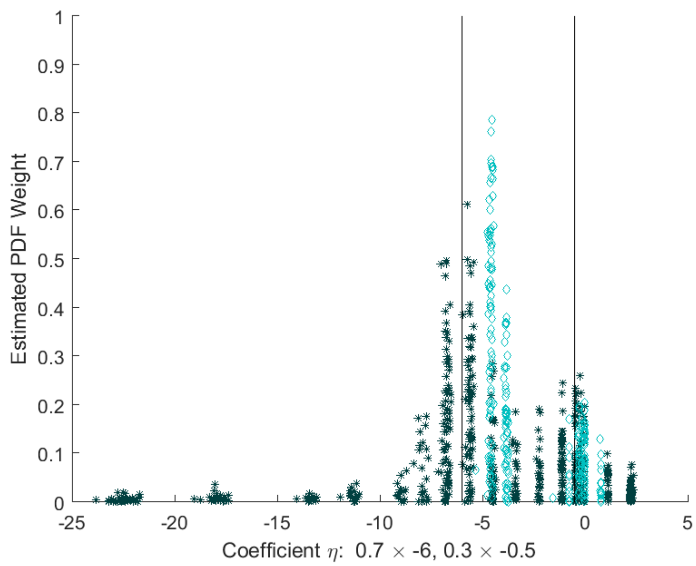

(a) Estimated pdf for η

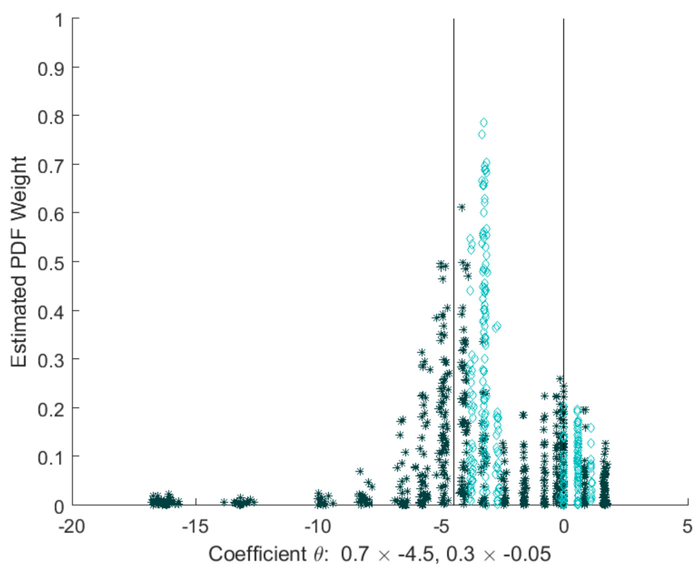

(b) Estimated pdf for θ

FIGURE A4. Performance of FKRB Estimator for a Mixture Distribution
Notes: Each panel shows the estimated marginal distribution of η and θ recovered from the non-parametric joint pdf. Each $\Diamond$ corresponds to the weight of the distribution estimated with the product fixed effects from the conditional logit, and each $*$ corresponds to the weight of the distribution estimated with the product fixed effects from the random coefficient logit. The modes of the mixture distribution of the true data generating process are depicted by the vertical lines. For all simulations, the number of alternatives is $J = 3$ and the number of observations is $N = 25,000$.

A4. **Implementation FKRB Estimator**

We implement the FKRB estimator with a two-step approach. In the first step, we estimate a parametric random coefficient logit to recover estimates for all parameters. In the second step, we recover a joint non-parametric distribution for a subset of parameters. For exposition, we define the set of parameters as $\Omega = \{\xi, \beta\}$, where β is the set of parameters that we aim to estimate a non-parametric distribution. In our application, $\beta = \eta, m$. The remaining parameters are all included in ξ. The choice model that we estimate takes the following form for each alternative j chosen by consumer i:

$$U^k_{ijrt} = \eta^k(P_{jrt} + m^k E_{jrt}) + \tau ES_{jt} + \phi Rebate_{jrt} + \gamma_j + Demo_i \times Att_t + \epsilon_{ijrt}. \tag{A7}$$

The parameters that we include in ξ are τ, ϕ, γ_j for all j, and the coefficients for the interaction terms between demographic information and attributes.

For the parametric random coefficient logit, we assumed that the joint distribution of η and m follows a multivariate normal distribution with unknown mean μ and covariance matrix Σ. The first-step estimates μ, Σ, and ξ. We used simulated maximum likelihood with 50 draws to approximate the distribution of random parameters.

In the second-step, we implemented the FKRB estimator with a discrete grid over the support of η and m. We used the estimates of the first-step to determine the range of the support. We experimented with grids of different size and noticed that the span of the grid could lead to computation issues. In particular, if the span is too large, the constraint least-square algorithm (implemented in Matlab) failed to converge. On the other hand, the number of grids point, i.e., the density of the grid, lead to higher computational time, but it did not create convergence issues. Our choice of the grid was thus made to ensure convergence and computational efficiency. The estimates that we report in the main text (Table 3) uses the following grid:

$$\begin{aligned} m =& [-3.5, -2.5, -2, -1.5, -1.25, -1, -0.75, -0.5, -0.25, -0.1, 0.01, 0.1, 0.2, ... \\ & 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1, 1.1, 1.25, 1.5, 2, 2.5, 3, 3.5, 4, 5] \\ \eta =& [-35, -30, -25, -18, -16, -14, -10, -7, -5, -4.5, -4, -3.5, -3.25, -3, -2.5, -2, -1, -0.1, 1, 2, 3] \end{aligned}$$

As robustness check, we experimented with grids that spanned a smaller region of the support. In Table 3, we report results for the following two grids. Model II refers to:

$$
\begin{aligned}
m =&[-2.5, -2.1, -1.8, -1.5, -1.25, -1, -0.75, -0.5, -0.25, -0.1, 0.01, 0.1, 0.2, ...\\
&0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1, 1.1, 1.25, 1.5, 2, 2.5, 3, 3.5, 4, 5]\\
\eta =&[-30, -26, -22, -18, -16, -14, -10, -7, -5, -4.5, -4, -3.5, -3.25, -3, -2.5, -2, -1, -0.1, 1, 2, 3],
\end{aligned}
$$

and Model III refers to:

$$
\begin{aligned}
m =&[-2.0, -1.8, -1.6, -1.4, -1.2, -1, -0.75, -0.5, -0.25, -0.1, 0.01, 0.1, 0.2, ...\\
&0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1, 1.1, 1.25, 1.5, 1.75, 2, 2.25, 2.5, 2.75, 3]\\
\eta =&[-22, -20, -18, -16, -14, -12, -10, -7, -5, -4.5, -4, -3.5, -3.25, -3, -2.5, -2, -1, -0.1, 1, 1.5, 2].
\end{aligned}
$$

We also experimented with a denser grid, which we refers as Model IV in Table 3:

$$
\begin{aligned}
m =&[-3.5, -3, -2.5, -2, -1.5, -1.25, -1, -0.75, -0.5, -0.25, -0.1, 0.01, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35, 0.4, ...\\
&0.45, 0.5, 0.55, 0.6, 0.65, 0.7, 0.75, 0.8, 0.85, 0.9, 0.95, 1, 1.05, 1.1, 1.15, 1.25, 1.5, 1.75, 2, 2.5, 3, 3.5, 4, 5]\\
\eta =&[-35, -30, -27.5, -25, -22.5, -18, -16, -14, -12, -10, -8, -6, -5, -4.5, -4, ...\\
&-3.5, -3.25, -3, -2.5, -2, -1, -0.1, 1, 2, 3].
\end{aligned}
$$

Finally, we experimented with a three dimensional grid, where the coefficient for the Energy Star label (τ) was also treated as a random coefficient (Model V, Table 3):

$$
\begin{aligned}
m =&[-3.5, -2.5, -2, -1.5, -1.25, -1, -0.75, -0.5, -0.25, -0.1, 0.01, 0.1, 0.2, ...\\
&0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1, 1.1, 1.25, 1.5, 2, 2.5, 3, 3.5, 4, 5]\\
\eta =&[-35, -30, -25, -18, -16, -14, -10, -7, -5, -4.5, -4, -3.5, -3.25, -3, -2.5, -2, -1, -0.1, 1, 2, 3]\\
\tau =&[0 \times \hat{\tau}^{RCL}, \hat{\tau}^{RCL}, 2.5 \times \hat{\tau}^{RCL}]
\end{aligned}
$$

The estimation was carried via the subsampling bootstrap procedure. We sliced the sample data in 16 subsamples of approximately equal size (N $\approx$ 2,000), and performed the two-step estimation on each subsample separately. For the estimation where we account for income, we performed the same procedure, except that the sample drawn from the transaction data only includes transactions for a specific income group.

A5. **Conditional Logit: Robustness**

Table A4 shows the results for demand shocks to attributes related to energy use. We consider week-of-sample fixed effects interacted with with three such attributes: (I) Energy Star certification status, (II) refrigerator size, and (III) freezer location. This set of controls has little effect on the price coefficient, suggesting that the model-specific price variation we observe from the retailer's pricing model is likely driven by supply-side and idiosyncratic factors rather than endogenous demand shocks.

Table A4. Conditional Logit Results

	I	II	III
Purchase Price	-0.00344***	-0.00339***	-0.00349***
	(0.00007)	(0.00007)	(0.00007)
Energy Cost	-0.03056***	-0.02559***	-0.03053***
	(0.00318)	(0.00315)	(0.00320)
m	0.76	0.65	0.75
Product FE	Yes	Yes	Yes
Week × EStar	Yes	No	No
Week × Size	No	Yes	No
Week × Top Freezer	No	No	Yes

Notes: Standard errors (in parentheses) clustered at the zip code level. The average m is computing assuming a discount rate of 5% and a refrigerator lifetime of 18 years.

A6. **Additional Results**

Table A5 provides summary statistics for the refrigerator models that were used to construct the matching pair estimator. We compare them to the whole universe of refrigerator models we observe in our data. Within a pair, the most efficient model sold for an average $22.5 more, with a standard deviation of $151.6. Large negative and positive price differences exist, which shows that the most efficient model is sometimes sold at a lower price.

Table A5. Matched Pairs: Summary Statistics

		Matched Paired	All Models
# Models		102	6,859
MSRP ($)		1,073	1,671
kWh/y		493	575
Adjusted Volume (Cu. Ft.)		24	27
% more efficient than minimum (%)		11	10
Year entered on market		2004.8	2004.1
Δ kWh/y	mean	-69.51	-
	std	27.63	-
	10^th Pctile	-108.00	-
	10^th Pctile	-35.00	-
Δ Elec Cost $/y	mean	-8.80	-
	std	4.57	-
	10^th Pctile	-14.08	-
	10^th Pctile	-3.34	-
Δ Price $	mean	22.54	-
	std	151.64	-
	10^th Pctile	-150.00	-
	10^th Pctile	155.00	-

Notes: The change in kWh/y, electricity costs per year, and price are average within pair. These averages are computed for all weeks where the two models of a pair were offered. The most efficient model is always compared to the least efficient model within a pair.

The following figure shows the estimated joint distribution for η and m and the marginal distribution of m for various specifications of the FKRB estimator. It shows that the overall pattern of the joint and marginal distributions are robust to the way the grid is specified and how quality is controlled for.

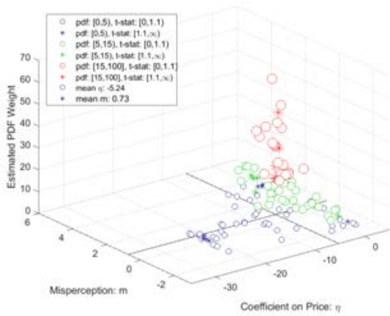

(a) Model II Smaller Span; $f(\eta, m)$

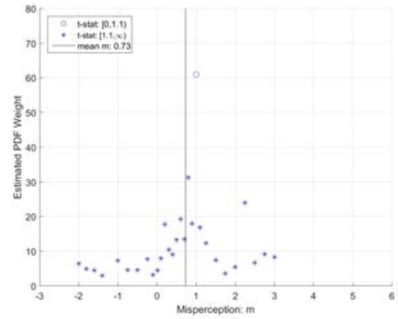

(b) Model II Smaller Span; $f(m)$

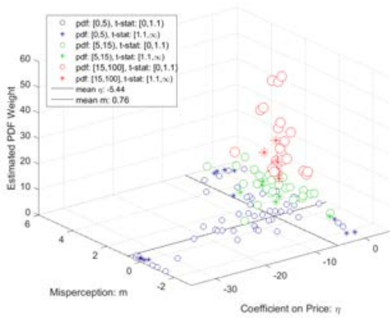

(c) Model III Dense Grid; $f(\eta, m)$

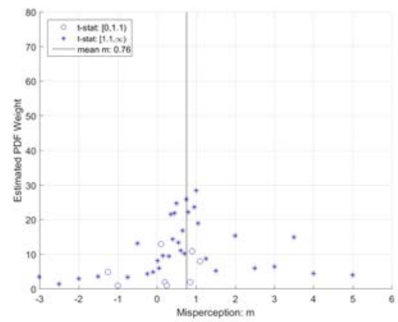

(d) Model III Dense Grid; $f(m)$

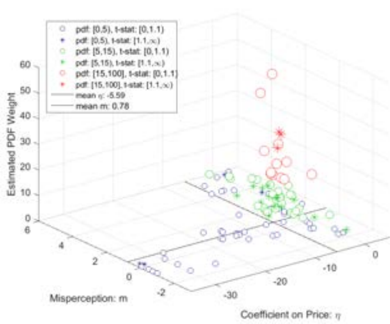

(e) Model IV τ Heterogeneous; $f(\eta, m)$

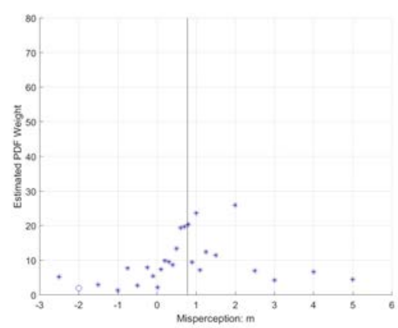

(f) Model IV τ Heterogeneous; $f(m)$

FIGURE A5. Robustness Tests: Joint Probability and Marginal Probability of m as a Function of the Grid

Notes: Each estimated pdf weight represented by stars or dot. Stars represent weight with t-statistics greater than 1.1. For the joint distribution, the color and market size represent the level of the pdf weights, where bigger markers capture larger weights.

The following figure shows the estimated joint distribution for η and m for different income groups.

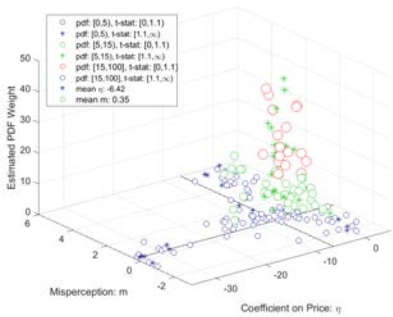

(a) Income: <\$30k

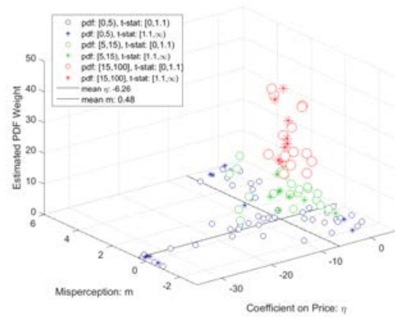

(b) Income: [\$30k, \$50k)

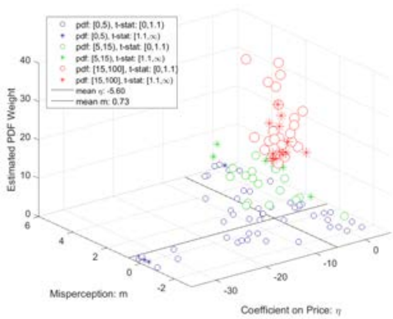

(c) Income: [\$50k, \$75k)

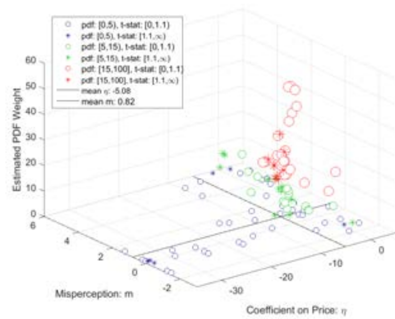

(d) Income: [\$75k, \$100k)

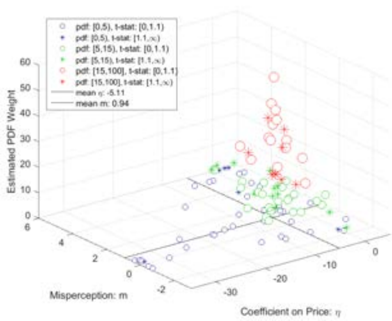

(e) Income: [\$100k, \$150k)

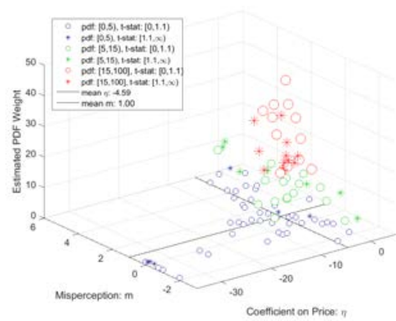

(f) Income: ≥\$150k

Figure A6. Joint Probability Distribution for η and m by Income Group
Notes: : Each estimated pdf weight represented by stars or dot. Stars represent weight with t-statistics greater than 1.1. For the joint distribution, the color and market size represent the level of the pdf weights, where bigger markers capture larger weights.

Table A6 presents the results assuming that misperceptions are defined using $r = 5\%$. Under this approach, there is a broader ranges of estimated values for m that corresponds to misperceptions. The table, however, show that is has little effect on the results. This makes the point that the tails of the distribution of misperceptions drive the results.

Table A6. Optimal Behavioral Policies

	Misperceptions at $r = 5\%$				Misperceptions with Heterogeneous r			
	Policy	Δ SW $/capita	Δ CS $/capita	Δ Ext $/capita	Policy	Δ SW $/capita	Δ CS $/capita	Δ Ext $/capita
Pigou Tax ($/ton of CO_2)								
With $F(m_k)$	-91.7	3.5	209.4	1.3	-91.4	3.9	208.7	1.4
	(0.38)	(0.05)	(0.90)	(0.01)	(0.38)	(0.06)	(0.89)	(0.01)
With $E[m_k]$	113.1	2.6	-237.7	-2.2	148.1	3.9	-312.4	-2.6
	(2.05)	(0.07)	(4.27)	(0.03)	(1.80)	(0.08)	(3.76)	(0.03)
Uniform Standard (kWh/year)								
With $F(m_k)$	335.3	128.3	91.0	-37.3	335.3	128.0	90.8	-37.1
	(0.00)	(0.13)	(0.13)	(0.01)	(0.00)	(0.15)	(0.14)	(0.02)
With $E[m_k]$	335.3	100.8	65.9	-34.9	335.3	103.2	67.4	-35.8
	(0.00)	(0.10)	(0.09)	(0.03)	(0.00)	(0.11)	(0.09)	(0.03)
Minimum Standard (kWh/year)								
With $F(m_k)$	335.3	128.3	91.0	-37.3	335.3	128.0	90.8	-37.1
	(0.00)	(0.13)	(0.13)	(0.01)	(0.00)	(0.15)	(0.14)	(0.02)
With $E[m_k]$	335.3	100.8	65.9	-34.9	335.3	103.2	67.4	-35.8
	(0.00)	(0.10)	(0.09)	(0.03)	(0.00)	(0.11)	(0.09)	(0.03)
Attribute-Based Minimum Standard (% w.r.t. existing standard)								
With $F(m_k)$	54.7	98.2	61.9	-36.3	55.0	99.5	63.6	-35.8
	(0.03)	(0.11)	(0.09)	(0.04)	(0.03)	(0.10)	(0.08)	(0.04)
With $E[m_k]$	56.1	79.7	47.1	-32.6	56.1	80.9	47.8	-33.1
	(0.01)	(0.11)	(0.10)	(0.02)	(0.01)	(0.10)	(0.10)	(0.02)

Notes: The table compares a scenario where experienced utility is defined using a 5% discount rate versus the scenario that considers heterogeneous discount rates that range from 2% to 12%. The table reports the optimal policies evaluated using the full distribution of misperceptions (with $F(m_k)$) or only the average misperception (with $E[m]$) as in the main text.

A7. **Theory: Proofs and Additional Results**

Optimal Behavioral Pigouvian Tax. To derive the expression of the optimal Pigouvian tax with misperceptions, we closely follow the framework presented in Houde and Aldy (2017b). We start with a social welfare function that consists of the sum of consumer surplus, government revenue, and externality costs.

In the presence of homogenous misperception, the expression for social welfare after a tax τ is levied and included in the price of energy, P^e, is given by

$$SW(\tau) = \frac{1}{\eta} \cdot \left[ln \sum_j^J e^{(U_j(\tau))} + \sum_j^J \sigma_j (U_j^E(\tau) - U_j(\tau)) \right] + (\tau - \phi) \sum_j \sigma_j(\tau) \cdot E_j,$$

where the gap between decision utility: $U_j(\tau) = V_j - \eta(P_j + m(P^e + \tau)E_j)$ and experienced utility: $U_j^E(\tau) = V_j - \eta(P_j + (P^e + \tau)E_j)$ is function of the size of the misperception m. Taking the derivative of $SW(\tau)$ with respect to τ and rearranging, we obtain

$$\tau = \frac{\phi}{m} + P^e \cdot \frac{(1-m)}{m},$$

which is the expression of the optimal Pigouvian tax for the average degree of misperception $\bar{m}$.

The expression for the case where misperceptions are heterogeneous is derived using the same approach, except that the expression for social welfare is given by:

$$SW = \sum_k \alpha_k CS_k + (\tau - \phi) \sum_k \sum_j \alpha_k \sigma_j^k E_j. \tag{A8}$$

The gap between decision and experienced utility is now indexed by the degree of misperception for each type k: m_k. Taking the derivative with respect to τ and rearranging, we obtain:

$$\tau^* = \frac{\phi}{1-\mathcal{A}} + P_e \frac{\mathcal{A}}{1-\mathcal{A}} \tag{A9}$$

with

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

Optimal Uniform Standard. Proof of proposition 1:

Under the assumption that the price of each product is the sum of unit cost and a constant markup:

$$p_j(E_j) = c(E_j) + \omega_j$$

for a uniform standard where $E_j = \bar{E}$ for all j, we can factorize $e^{(-m_k\bar{E})}$ and $e^{(-\eta_k c(\bar{E}))}$ from the choice probabilities σ_j^k. In other words, the level of $\bar{E}$ will not affect the choice probabilities given that all options have the same values for $m_k\bar{E}$ and $\eta_k c(\bar{E})$. Note also that the expression: $ln\sum_j^J e^{(U_{kjtr})}$ in the expression for the consumer surplus becomes:

$$ln\sum_j^J e^{(U_{kj})} = -m_k\bar{E} - \eta_k c(\bar{E}) + ln(\sum_j e^{(\omega_j+\gamma_j)}).$$

The difference between decision and experienced utility simply becomes: $(m_k - 1)\bar{E}$ for a given type k. The expression for the social welfare is thus:

$$SW = \sum_k \alpha_k[-m_k\bar{E} - c(\bar{E}) + (m_k - 1)\bar{E}] - \phi\bar{E}.$$

Taking the derivative with respect to $\bar{E}$ and rearranging directly yields:

$$-c'(\bar{E}^*) = 1 + \phi. \tag{A10}$$