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, Revised September 2024

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, in addition to 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 was also part of the activities of SCCER CREST, 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, Revised September 2024
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 context of energy-using durables, where a long-standing policy debate continues on whether taxes or standards are superior for regulating externalities. We derive formulae to quantify welfare effects for each instrument, accounting for misperceptions. Standards have a notable advantage over taxes. They reduce the variance of the misperceived attribute in the choice set, which reduces allocative inefficiencies. In the U.S. appliance market, standards dominate taxes across scenarios, and correctly characterizing misperception heterogeneity is less important for setting optimal standards than for taxes.

Sébastien Houde
Grenoble Ecole de Management
Centre for Energy Policy and Economics
ETH Zurich
sebastien.houde@grenoble-em.com

Erica Myers
Department of Economics
University of Calgary
414 Social Sciences Building
2500 University Drive NW
Calgary, AB T2N 1N4
Canada
and E2e
and also NBER
erica.myers@ucalgary.ca

1 Introduction

Consumers are prone to errors when making choices across various economically important settings, such as health care plans, mutual funds, mortgages, schools, and food choices.[1] As a result, policy-makers increasingly design programs to address consumer misperceptions driving these mistakes.[2] However, to do this well, policy-makers need information on not just the average degree of consumer misperception but also its full distribution (Goldin and Homonoff, 2013; Allcott, Mullainathan, and Taubinsky, 2014; Farhi and Gabaix, 2020; Taubinsky and Rees-Jones, 2018). For example, consider a Pigouvian tax levied on an aspect of product cost that is misperceived heterogeneously by consumers. As they make their consumption decisions, it would be as if the externality itself were heterogeneous, but this cannot be optimally addressed with a single tax rate (Diamond, 1973).

In this paper, we investigate the role of consumer perception heterogeneity in an important context in which a potentially misperceived aspect of product cost has been central to the policy debate: regulating greenhouse gas emissions from energy-using durables. For over 40 years, policy-makers have presumed that consumers undervalue energy costs, implying that minimum efficiency standards could do better than price-based instruments at addressing externalities while protecting consumers from their mistakes. However, if consumers correctly perceive energy costs, market-based instruments such as taxes or cap-and-trade could be superior. Therefore, there has been an active literature assessing consumers' perceptions of energy costs, yet the focus has been almost exclusively on the population average. (e.g., Busse, Knittel, and Zettelmeyer, 2013; Allcott and Wozny, 2014; Rapson, 2014; Sallee, West, and Fan, 2016; Houde, 2018a; Grigolon, Reynaert, and Verboven, 2018; Leard, Linn, and Zhou, 2017; Leard, Linn, and Springel, 2019; Gillingham, Houde, and Van Benthem, 2021; Houde and Myers, 2021). The impact of consumer perception heterogeneity has not been formally considered in evaluating the relative merits of taxes and minimum efficiency standards for climate policy.

This paper extends our understanding of the design of policies with behavioral agents; this has been recently investigated in general settings (e.g., Taubinsky and Rees-Jones, 2018; Farhi and Gabaix, 2020) and in more specific contexts such as health (Handel, Kolstad,

[1]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, Kuminoff, and Powers, 2019), mutual funds (Barber, Odean, and Zheng, 2005), mortgages (Allen, Clark, and Houde, 2019; Guiso et al., 2021), schools (Jensen, 2010), and nutrition (Bollinger, Leslie, and Sorensen, 2011).

[2]In this paper, we will use the term misperception to refer to any biases, heuristics, information frictions, or biased beliefs that induce consumers to undervalue or overvalue a product cost aspect.

and Spinnewijn, 2019), nutrition (Allcott, Lockwood, and Taubinsky, 2019) and energy (Allcott, Mullainathan, and Taubinsky, 2014; Gerster and Kramm, 2024; Rodemeier and Löschel, 2023), among others. We add to this literature by evaluating the role of energy cost perception heterogeneity in the design and performance of taxes versus minimum efficiency standards for addressing pollution externalities from energy-using durables. We first derive a behavioral welfare decomposition for each instrument as a function of sufficient statistics. We demonstrate that standards have a notable welfare advantage over taxes in the presence of heterogeneous misperceptions. This result hinges on a novel insight: because standards reduce the variance of the potentially misperceived attribute in the choice set, they reduce allocative inefficiencies induced by misperceptions. Conversely, a Pigouvian tax exacerbates such inefficiencies. We also demonstrate that an optimal standard minimizes the variance in product costs, which allows for the internalization of misperceptions. The optimal tax, on the other hand, can only internalize misperceptions if it is differentiated and perfectly targets each consumer's degree of misperception. In practice, such perfect targeting is unlikely to be possible. We show that the deadweight loss of a second-best non-targeted tax grows with heterogeneity in misperceptions.

We demonstrate these insights using a behavioral welfare analysis based on empirically recovered distributions of consumer misperceptions. To date, there are few estimates of such distributions. Goldin and Homonoff (2013) and Taubinsky and Rees-Jones (2018) are two notable examples of studies that recovered heterogeneous misperceptions of sales taxes and investigated the equity and efficiency issues such heterogeneity raises. We use a structural econometric model coupled with rich administrative data on refrigerator purchases, consumer-specific demographics, and county-level energy prices to estimate heterogeneity in consumer perception of energy costs. In particular, we use a two-step semi-parametric estimator, which we adapt from Fox, Kim, Ryan, and Bajari (FKRB) (2011). We also present results from a more conventional parametric mixed logit model for comparison. To estimate demand responsiveness to price, we exploit product-specific idiosyncratic variation in price induced by the retailer's national pricing algorithm, which we demonstrate is exogenous to local demand shocks. We estimate demand responsiveness to local energy operating costs using variations in electricity prices across regions and time. To account for any variation in perceived quality that may be correlated with the purchase price or electricity cost, we include a rich set of controls that account for interactions between detailed product attributes and transaction-specific consumer demographics.

We evaluate perceptions of energy cost as the ratio of consumers' demand response to changes in energy costs and to changes to the potentially more salient and correctly perceived purchase price. This approach has been widely used in the literature to quantify the *average* degree of consumer misperception of various product costs, such as shipping and handling fees (Hossain and Morgan, 2006), sales tax (Chetty, Looney, and Kroft, 2009), highway tolls (Finkelstein, 2009), as well as energy operating costs (e.g., Busse, Knittel, and Zettelmeyer, 2013; Allcott and Wozny, 2014; Grigolon, Reynaert, and Verboven, 2018; Sallee, West, and Fan, 2016; Myers, 2019). In Houde and Myers (2021), we use the same empirical setting and a reduced-form approach to demonstrate that on *average* consumers demonstrate little undervaluation of local energy costs when purchasing appliances. Our present estimation exploits the same rich variation in the data to recover the full distribution of this ratio.

Our estimates reveal substantial heterogeneity in consumer perceptions of energy costs, which suggests significant under- and over-valuation of energy costs. Although about half of consumers have modest to no misperceptions, a large share undervalue energy costs with implied discount rates greater than the sample period's average APR rate for credit cards, 12%, and a smaller share (about 20%) have implied discount rates of less than the sample period's 3-year return on Treasury bonds, 2%, potentially reflecting overvaluation. The patterns we find are robust to flexible controls for unobserved variation in preferences for the ENERGY STAR® label and other energy-related attributes as well as idiosyncratic product-specific preferences. We also show that, although lower-income groups are more prone to undervaluation than higher-income groups, the patterns of heterogeneity are qualitatively similar across the income distribution. Therefore, income and, in particular, explanations related to credit availability are not primary drivers of the prevalence of undervaluation.

Assumptions about the expected refrigerator lifetime and discount rates are important for quantifying misperceptions. There may be substantial heterogeneity in these parameters across consumers, driving some of the heterogeneity we recover in energy cost responsiveness. To account for such heterogeneity, we define a broad range of behaviors as "rational" in our optimal policy analysis. In our main specification, we assume there is no misperception associated with implied discount rates between 2% and 12%.

In addition, we explore the sensitivity of our optimal policy estimation to how we define misperceptions by exploring several alternatives. First, we allow any discount rate below 2% to be rational, reflecting that the seeming overvaluation of energy costs may not be driven by misperception, but instead by a so-called warm glow from choosing a greener product. Second, we assume that only an implied discount rate of 5%, used by the U.S. Department of

Energy (DOE) to design refrigerators' minimum standards (Office of Energy Efficiency and Renewable Energy, 2010), is rational and all other behaviors are driven by misperceptions. This demonstrates the welfare impacts in our policy analysis of a stricter misperception categorization. Finally, we consider a scenario where, like our preferred approach, discount rates of 2-12% are rational, but with an expected lifetime of 12 years rather than the lifetime utilized by the DOE of 18 years.

Our optimal policy analysis reveals two important takeaways.[3] First, across all scenarios, standards induce larger welfare gains than the Pigouvian tax for addressing carbon externalities. The ability of a standard to address misperceptions by reducing the variance in energy costs while simultaneously addressing the externality provides a clear advantage. Second, the optimal tax level is more sensitive to the distribution of misperceptions than the optimal level of any of the types of standards we consider. The optimal tax is negative (-\$25.70/ton) for our benchmark approach using the full, non-parametric distribution of consumer perceptions from the FKRB estimator, whereas it is positive if we impose symmetry on the distribution of misperceptions using a random coefficient logit model (\$111.30), or if we ignore heterogeneity and assume all consumers display the average degree of misperception (\$75.70). The high degree of sensitivity of the optimal tax level is largely driven by its responsiveness to the part of the distribution that overvalues energy costs. These consumers have a higher tax elasticity than consumers who undervalue energy cost, such that the downward adjustment in the optimal tax for overvaluers will outsize the upward adjustment for less elastic undervaluers and can result in a negative optimal tax. In contrast to the sensitivity we see in the optimal tax levels according to how misperception heterogeneity is characterized, the levels of the optimal standards are strikingly similar, often the same, across scenarios. Because there are fewer allocative inefficiencies induced by a standard relative to a tax, the optimal policy is less sensitive to the exact distribution of misperceptions. For ambiguity-averse policy-makers this is an important advantage of a standard because the welfare loss of setting policy based on a distribution of mischaracterized misperceptions would be lower than for a tax.[4]

[3]In the energy context, others have performed related behavioral welfare analysis of policy instruments. For instance, both Allcott (2013) and Leard, Linn, and Springel (2020) considered the role of consumers' heterogeneous misperceptions in estimating the welfare impacts of vehicles' efficiency standards. They focused on the passenger vehicle market and a single policy instrument, the U.S. fuel economy standard. Allcott and Taubinsky (2015) used an artefactual field experiment to recover the distribution of consumer misperceptions for light bulbs to estimate the welfare impacts of an incandescent ban.

[4]We refer to ambiguity-averse policy-makers as agents that are uncertain about the true distribution of misperceptions, and thus have a certain prior over such distribution, and evaluate welfare with respect to such prior such that they are averse to such uncertainty (i.e., ambiguity in this setting).

The fact that minimum efficiency standards outperform Pigouvian taxes for regulating negative externalities when agents have heterogenous misperceptions is important because implementing carbon pricing policies can be less politically feasible compared to establishing minimum energy efficiency standards. For instance, as of 2022, only 47 countries had either implemented carbon taxes or cap-and-trade programs, while 120 countries had established energy efficiency standards and labeling programs (IEA (2021), World Bank (2022)). Our work also contributes to the broader literature on second-best climate mitigation approaches, which demonstrates that Pigouvian taxes may not be optimal in certain circumstances. For example, previous studies have identified incomplete markets (Holland, 2012), fiscal interactions (Goulder, Hafstead, and Williams III, 2016), technology and innovation spillovers (Fischer and Newell, 2008; Fischer, Preonas, and Newell, 2017), and behavioral considerations (Fischer, Preonas, and Newell, 2017; Chan, 2023) as factors that can affect the effectiveness of Pigouvian taxes.

Our study focuses on energy-using durables and energy-policy design. However, our insights have implications for other economically important settings with potentially shrouded attributes, such as in purchasing products on online platforms and in making banking, investment, health care, housing, or schooling decisions (Gabaix and Laibson, 2006). If consumers frequently choose options that are suboptimal due to a shrouded attribute, policy-makers can address this issue through taxation or regulation. For example, policy-makers could impose taxes on products with undesirable levels of the shrouded attribute or establish standards that limit consumers' choices for those attributes, such as by regulating the types of bank accounts, investment options, healthcare plans, or schooling curricula available in the market. To the extent that consumers exhibit heterogeneity in their misperceptions of the shrouded attributes, standards will have a welfare advantage over taxes due to their ability to better reduce allocative inefficiencies with a single instrument. In theory, differentiated taxes targeted to groups of consumers with systematically varying degrees of misperception could improve the efficiency of a tax mechanism relative to the standard. However, in most contexts, it is difficult to identify and target subgroups of consumers in this way (Sallee, 2019). For example, as we demonstrate empirically, within the groupings one might consider for differentiating a tax, such as income, there is still plenty of misperception heterogeneity, which makes perfect targeting infeasible. Finally, the fact that optimal policy levels and their resulting welfare effects are less sensitive to the distribution of misperceptions for standards than taxes is particularly relevant in policy contexts in which it may not be possible to recover precisely estimated fine-grained heterogeneity in consumer perceptions. Or, like in our setting, it may be feasible, but estimation techniques and assumptions, about discount rates, for instance, affect the estimated distributions and add uncertainty to the characterization of misperceptions. In these cases, minimum efficiency standards offer a welfare advantage over taxes for policy-makers who are averse to making mistakes when choosing policies.

The remainder of the paper is organized as follows. In the next section, we derive welfare decomposition formulae for energy tax and energy standards in the presence of heterogeneous consumer misperceptions and externalities and investigate the design of optimal policies in this context. In Section 3, we discuss the data and choice environment for our empirical investigation. In Section 4, we present our empirical framework for recovering heterogeneous perceptions. In Section 5, we present the results of our estimation. In Section 6, we investigate the design and evaluation of behavioral environmental policies, and in Section 7, we offer our conclusion.

2 Behavioral Environmental Policy Design

2.1 Set-Up

The test we use for characterizing misperception of energy costs compares whether consumers respond the same way to a one-dollar change in the salient, correctly perceived purchase price (P_j) of energy durable j versus a one-dollar change in the present value of its lifetime energy operating costs. This latter is a complex and, perhaps, less salient attribute relative to the purchase price. It is the product of the energy price (P^e) and the energy consumed over the lifetime of the appliance (E_j).[5] Consumer i trades off these attributes by deriving utility for each product j as follows:

$$U_{ij} = \gamma_j - \eta P_j - \theta P^e \cdot E_j + \epsilon_{ij}, \tag{1}$$

where γ_j is the vertical quality of product j and ϵ_{ij} captures consumer $i's$ idiosyncratic preferences for product j. The formal test of misperception is whether the ratio of the energy cost parameter to the price parameter equals one: $\theta/\eta = 1$. Hereafter, we will refer to this ratio as the misperception parameter: $m = \theta/\eta$. Misperception might arise from errors in perceiving any of the determinants of energy costs such as energy price, expected lifetime, or annual energy consumption. Or, if information is costly to acquire and (mentally)

[5] In practice, uncertainty and consumer-specific heterogeneity in the product lifetime, future energy prices, utilization, depreciation, and discount factor make it difficult to empirically characterize E_j. We discuss these assumptions for our empirical exercise in what follows. For expository purposes, assume for now that $P^E \cdot E_j$ is homogeneous in the population and captures the exact measure of expected energy operating costs discounted with a normal rate of return.

process, some consumers might simply ignore this attribute, which would imply that $\theta = 0$, and $m = 0$.

Given the possibility of several behavioral phenomena affecting consumer perception, the ratio $m = \theta/\eta$ may be heterogeneous in the population. To capture this heterogeneity in a flexible manner and rationalize various underlying behavioral models, we will assume there are K different consumer types with k-specific misperception parameters, m^k, and define α_k as the fraction of the population subject to each level of misperception. We place no restriction on the discrete distribution, except that $\sum_k^K \alpha_k = 1$ and $\alpha_k \geq 0$.

The alternative-specific utility for each consumer type is expressed as an explicit function of the k-specific misperception parameter (m^k) as follows.

$$U_{ijk} = \gamma_j + \eta^k(-P_j - m^k P^e \cdot E_j) + \epsilon_{ijk}. \tag{2}$$

Equation 2 corresponds to consumers' decision utility for each type. However, the welfare associated with a particular choice is not governed by consumers' ex-ante decision utility, which could be subject to misperceptions about energy costs, but rather by their ex-post *experienced* utility, which is a function of correctly perceived realized energy costs.

To derive an expression for consumer welfare that accounts for the gap between decision and experienced utility, we borrow from the approach first proposed by Leggett (2002).[6] We first assume that the idiosyncratic preferences, ϵ_{ijk} in Equation 2, are i.i.d. and follow an extreme value distribution. The choice probabilities that characterize decision utility thus take the standard multinomial logit form. The parameter q_{jk} denotes the probability that a type-k consumer purchases product j. In this framework, Leggett (2002) has shown that the standard expression for measuring the consumer surplus, first derived by Small and Rosen (1981), can be adapted to account for imperfect information. The modified formula

[6] Recently, similar approaches have been developed by Allcott (2013), Dubois, Griffith, and O'Connell (2017), Houde (2018a), Ketcham, Kuminoff, and Powers (2019), and Allcott and Knittel (2019), among others, to conduct applied welfare analysis in contexts where behavioral biases are present.

can be expressed as[7]:

$$CS^k = \frac{1}{\eta^k} \cdot \left[ln \left(\sum_j^J e^{V_{jk}} \right) + \sum_j^J q_{jk} \cdot (V_{jk}^E - V_{jk}) \right]. \tag{3}$$

where V_{jk}^E and V_{jk} are, respectively, the components of experienced and decision utility without the idiosyncratic preferences, ϵ_{ijk}. The first part of the expression in Equation 3, $\frac{1}{\eta^k} \cdot ln \left(\sum_j^J e^{V_{jk}} \right)$, corresponds to the standard welfare measure for the multinomial logit. The second part of the expression, $\frac{1}{\eta^k} \cdot \sum_j^J q_{jk} \cdot (V_{jk}^E - V_{jk})$, is the correction term for misperceptions, which quantifies the expected (private) cost consumers incur because of the discrepancy between what they perceive they will experience and what they actually experience. In our setting, the gap between experienced utility and decision utility is a function of the misperception parameter (m^k), marginal utility of income (η^k), and lifetime energy operating costs ($P^e \cdot E_j$), namely $V_{jk}^E - V_{jk} = \eta^k \cdot (m^k - 1) \cdot P^e \cdot E_j$.

2.2 Behavioral Welfare Decomposition of a Pigouvian Tax

We begin by deriving an expression for the welfare change associated with imposing a Pigouvian tax relative to a no-policy baseline. We consider a negative externality associated with the lifetime energy consumption of each product (E_j), where the marginal externality, denoted ϕ, is constant. If we impose a Pigouvian tax, τ, on the price of energy with a lump-sum redistribution of the tax revenues, the change in welfare for each type k, denoted ΔSW_τ^k, can be computed as the difference between the welfare under the tax (SW_τ^k) and the baseline welfare (SW_0^k):

$$\Delta SW_\tau^k = SW_\tau^k - SW_0^k = CS_\tau^k - CS_0^k + \tau \cdot \sum_j q_{jk}(\tau) \cdot E_j - \phi \cdot \left(\sum_j q_{jk}(\tau) \cdot E_j - \sum_j q_{jk} \cdot E_j \right). \tag{4}$$

In the above expression, the term $\sum_j q_{jk} \cdot E_j$ is the expected energy consumption induced by the choices of a type-k consumer before the tax, which we denote $\overline{E}_0^k \equiv \sum_j q_{jk} \cdot E_j$. After the imposition of tax, we denote this quantity $\overline{E}_\tau^k \equiv \sum_j q_{jk}(\tau) \cdot E_j$. Finally, we define the policy-induced change in expected energy consumption for a type-k consumer as follows: $\Delta \overline{E}_\tau^k \equiv \overline{E}_\tau^k - \overline{E}_0^k$.

[7] Note that in the formula we have omitted an integration constant $\mathcal{I}$, which cancels out when computing the change in consumer surplus if the distribution of the idiosyncratic preferences (ϵ_{ijk}) remains constant across policy changes.

Substituting the expression for the consumer surplus given in Equation 3 into Equation 4 and rearranging, the change in welfare from the imposition of a tax, τ, can be decomposed into three components:

$$
\begin{aligned}
\Delta SW^k_\tau = \; & \underbrace{\frac{1}{\eta^k} \cdot \ln\left(\frac{\sum_j e^{V_{jk}(\tau)}}{\sum_j e^{V_{jk}}}\right)}_{\text{Change in decision utility}} \\
- \; & \underbrace{\phi \cdot \Delta \overline{E}^k_\tau}_{\text{Change in externality}} \\
+ \; & \underbrace{(m^k - 1) \cdot P^e \cdot \Delta \overline{E}^k_\tau + m^k \cdot \tau \cdot \overline{E}^k_\tau .}_{\text{Correction for misperceptions}}
\end{aligned}
\tag{5}
$$

To sign each of these components, we will assume $m > 0$.[8] The first component in Equation 5 is negative, corresponding to the perceived tax burden, i.e., the change in decision utility due to the tax. The second component is the change in externality. The tax induces consumers to choose more energy-efficient products, i.e., $\Delta \overline{E}^k_\tau \leq 0$. This reduction in energy consumption reduces the externality and is, therefore, a source of welfare gain.

The third component is a correction term that arises because of misperceptions. It reflects how experienced utility changes relative to the change in decision utility from choosing more efficient models as a result of the tax. The term $(m^k - 1) \cdot P^e \cdot \Delta \overline{E}^k_\tau$ represents the correction for the effect of misperceptions of energy costs. The term: $m^k \cdot \tau \cdot \overline{E}^k_\tau$ corrects for consumers' misperception of the tax, which we assume is redistributed to consumers via a lump-sum. To the extent energy costs are misperceived, the tax will also be misperceived, as it is included in the price of energy. Given $\Delta \overline{E}^k_\tau < 0$, the correction for misperceptions of consumer energy cost will be positive for $0 < m^k < 1$. The intuition for this positive welfare effect is the following. By pushing consumers to more efficient products, the tax ameliorates the distortionary effects of consumers' energy cost undervaluation, thus narrowing the gap between decision and experienced utility. If consumers overvalue energy costs ($m^k > 1$), the sign of the misperception correction is negative, making the overall sign of the third component ambiguous. In this case, the tax would exacerbate rather than ameliorate the

[8] We can also sign the expression for the cases where $m = 0$ or $m < 0$. Note that if $m = 0$, the tax would not affect consumers' perceptions of energy costs and thus have no effect on decisions. Severe misperceptions, $m < 0$, would mean that consumers prefer higher energy costs, all else equal, and therefore would be uncommon.

effects of misperceptions on consumer choice, driving a larger wedge between decision and experienced utility.

Calculating the overall change in welfare requires integrating Equation 5 over the distribution of consumer types. Focusing on the expectation of the third component of the expression offers further insights into how heterogeneous misperceptions affect the performance of a tax. For a given distribution of misperceptions, this expectation is as follows:

$$\mathbf{E}\left[P^e \cdot (m^k - 1) \cdot \Delta \overline{E}^k_\tau + \tau \cdot m^k \cdot \overline{E}^k_\tau\right] = \tag{6}$$

$$\underbrace{P^e \cdot (\overline{m} - 1) \cdot \Delta \overline{E}_\tau + \tau \cdot \overline{m} \cdot \overline{E}_\tau}_{\text{Correction for mean level of misperception}}$$

$$\underbrace{P^e \cdot \mathbf{corr}(m^k, \Delta \overline{E}^k_\tau) \cdot \mathbf{sd}_{m^k} \cdot \mathbf{sd}_{\Delta \overline{E}^k_\tau} + \tau \cdot \mathbf{corr}(m^k, \overline{E}^k_\tau) \cdot \mathbf{sd}_{m^k} \cdot \mathbf{sd}_{\overline{E}^k_\tau}}_{\text{Correction for heterogeneity in misperceptions}},$$

where **corr** denotes the correlation coefficient and **sd** the standard deviation. This expectation can be decomposed into a correction at the mean level of misperception and a correction for heterogeneity in misperceptions.

The correction at the mean level of misperception will be positive when the tax is reducing the expected energy costs and consumers, on average, undervalue energy costs: $0 < \bar{m} < 1$. But the correction term at the mean could also be negative if $\bar{m} > 1$.[9]

However, the sign of the second component, attributable to heterogeneity, is unambiguously negative and is dictated by the correlation between misperceptions and the change in energy costs. This correlation is negative given that the more weight consumers place on lifetime energy costs relative to purchase price (i.e., the higher the m), the more important the energy savings would be. Mathematically, higher m means $\Delta \overline{E}^k_\tau$ is more negative (see Appendix A). Likewise, consumers who undervalue energy costs (i.e., m close to 0) will be the least responsive to the tax, and the energy savings will be smaller. Therefore, both

[9] It could also be negative if a significant proportion of consumers have severe misperceptions such that $m < 0$. This scenario is highly unlikely, but for completeness, it is worth mentioning it could lead to an increase in expected energy consumption and, thus, a negative sign for the correction at the mean level of misperception.

the correlation between m^k and $\overline{E}_\tau^k$ and the correlation between m^k and $\Delta\overline{E}_\tau^k$ are negative. Heterogeneous misperceptions, therefore, always work against the welfare effects of a tax.

As a result of these negative correlations, the variance in energy consumption across the set of chosen products also increases in response to an increase in the energy price induced by the tax, i.e., $\mathbf{var}_{\overline{E}_0^k} < \mathbf{var}_{\overline{E}_\tau^k}$, which further exacerbates the distortionary effects of the tax. This can readily be seen in Equation 6 given that the negative correlation terms are multiplied by the standard deviation of the distribution of both m^k and $\overline{E}_\tau^k$.

2.3 Behavioral Welfare Decomposition of an Energy Standard

We next consider the welfare effect of an energy standard. A standard impacts consumers via two channels. First, it induces movements in the product space. Second, it affects product prices. For the decomposition below, we will first assume that all products remain in the choice set after the standard is imposed whereby any non-compliant products are redesigned to meet the standard at a cost that is passed through to product prices.[10]

Using analogous expressions for consumer surplus and the benefits from the externality correction as described for the tax, we can decompose the welfare effect of an energy standard relative to a no-policy baseline as follows:

$$\begin{aligned}\Delta SW^k = \quad & \underbrace{\frac{1}{\eta^k} \cdot \ln\left(\frac{\sum_j^J e^{\widetilde{V}_{jk}}}{\sum_j^J e^{V_{jk}}}\right)}_{\text{Change in decision utility}} \\ - \quad & \underbrace{\phi \cdot \Delta\overline{E}_s^k}_{\text{Change in externality}} \\ + \quad & \underbrace{(m^k - 1) \cdot P^e \cdot \Delta\overline{E}_s^k.}_{\text{Correction for misperceptions}}\end{aligned} \tag{7}$$

where $\tilde{V}_{jk}$ is the alternative-specific utility associated with product j under the standard where there is a new upfront price and potentially a new energy consumption level, and $\Delta\overline{E}_s^k \equiv \overline{E}_s^k - \overline{E}_0^k$ is the change in expected energy consumption, for type k consumer, induced by the standard.

To sign the components of the decomposition, we will begin by considering a minimum

[10]For our policy analysis, we will also assume that only the efficiency of the appliance is adjusted so that non-energy dimensions of quality are fixed. Our welfare decomposition formula, however, holds for a more general setting where standards induce product movements in the non-energy dimensions.

efficiency standard, which moves the efficiency of models with the highest energy consumption levels in the choice set such that they become compliant with the standard. Further, as with the tax, we will assume $m > 0$. Depending on the marginal cost of adjusting the energy level, it is possible that a minimum efficiency standard provides a (perceived) net private benefit where the (perceived) reduction in lifetime energy costs dominates the increase in upfront prices. However, this is unlikely to be the case for consumer types that undervalue energy costs. Therefore, for most distributions of energy level-adjustment costs, the sign of the first component will be negative.

As with the tax, the sign of the second component will be positive, because a minimum efficiency standard induces consumers to choose more efficient products, i.e., $\Delta\overline{E}_s^k < 0$. Likewise, the sign of the third component is positive if $0 < m < 1$, i.e., the standard reduces the gap between decision and experienced utility. By removing the lowest levels of energy efficiency from the choice set, the standard is particularly effective at reducing the distortionary effects of the highest levels of undervaluation. For consumers who undervalue energy costs the most, products with the lowest level of energy efficiency have the largest gap between decision and experienced utility, and thus are more likely to be chosen by them all else equal. For consumer types that overvalue energy costs ($m > 1$), the misperception correction term will be negative to the extent that the standard affects consumer choice. For these consumers, the least efficient products actually have a higher experienced utility than decision utility (because they are biased toward purchasing too energy-efficient products).

As with the tax, we can take the expectation over this third component to garner insights into how heterogeneity in misperceptions impacts the welfare effect of a standard:

$$\mathbf{E}_k\left[(m^k - 1)\cdot P^e \cdot \Delta\overline{E}_s^k\right] = \underbrace{P^e \cdot (\overline{m} - 1)\cdot \Delta\overline{E}_s}_{\text{Correction for mean level of misperception}} + \underbrace{P^e \cdot \mathbf{corr}(m^k, \Delta\overline{E}_s^k)\cdot \mathbf{sd}_{m^k}\cdot \mathbf{sd}_{\Delta\overline{E}_s^k}}_{\text{Correction for heterogeneity in misperceptions}} . \tag{8}$$

Again, this expectation has two terms. The first term, the correction at the mean level of misperception, will be positive for $0 < \overline{m} < 1$. The second term, the correction for heterogeneity in misperceptions, particularly the sign of the correlation between m^k and the change in energy consumption, is the distinguishing feature between a tax and a minimum efficiency standard. Under the tax instrument, this correlation is always negative, while it is always positive under a minimum efficiency standard. This difference in sign comes from the fact that the tax moves demand toward energy-efficient products more for consumers that undervalue energy costs the least; whereas the standard moves demand more for consumers that undervalue energy costs the most. Therefore, unlike the tax, the minimum efficiency standard addresses externalities while at the same time ameliorating the distortionary effects of misperception on product choice, irrespective of the distribution of misperceptions. To gain an intuition for this, we can show how misperceptions impact the change in energy consumption under a standard by taking the derivative of $\Delta\overline{E}_s^k$ with respect to m^k:

$$\frac{\partial \Delta\overline{E}_s^k}{\partial m_k} = \eta^k \cdot P^e \cdot (\mathbf{var}_{\overline{E}_0^k} - \mathbf{var}_{\overline{E}_s^k}). \tag{9}$$

The above expression will be positive, and thus also the correlation $\mathbf{corr}(m^k, \Delta\overline{E}_s^k)$, because the minimum efficiency standard decreases the variance in energy consumption of the chosen products. Therefore, the presence of heterogeneity in misperceptions will work in favor of a minimum efficiency standard relative to a tax.

While we have focused on a minimum efficiency standard in this section, we note the sign of these welfare effects will hold for any type of standard that: 1) induces a reduction in mean energy consumption ($\Delta\overline{E}_s^k < 0$) and 2) reduces the variance of the energy consumption of the chosen products. However, there may be certain types of attribute-based standards (ABS), where these conditions do not hold. For example, if minimum efficiencies are assigned according to product features such as the size or location of the freezer door, an ABS could conceivably lead to an increase in energy consumption. This could happen if, for example, the regulation had the perverse effect of encouraging larger models to be manufactured. An ABS could also increase the ex-post variance in energy costs, making the sign for Equation 9 negative.[11]

[11]In 2014, the DOE updated their ABS for refrigerators to include more attributes, thus expanding categories with specific efficiency requirements. In Appendix B, we provide suggestive evidence the new standard had little effect on the energy consumption variance of available products. The proliferation of categories may have served to increase the energy consumption variance, thus offsetting any variance-reducing effects from tightening the energy efficiency requirements.

2.4 Implications for Optimal Policies

2.4.1 Pigouvian Tax

Heterogeneity in misperceptions not only has implications for the direction of a welfare change associated with a policy, but also for the design of the optimal policies. Allcott, Mullainathan, and Taubinsky (2014), Farhi and Gabaix (2020), Gerster and Kramm (2024), and Rodemeier and Löschel (2023) have derived formulas for the optimal Pigouvian instruments considering misperceptions in different frameworks. We extend these results to a discrete choice framework and approximate the deadweight loss of a second-best uniform tax. In our discrete choice setting, we show (see Appendix A) that an optimal uniform Pigouvian tax is:

$$\tau^* = \phi \frac{\sum_k \alpha_k \mathcal{E}_\tau^k}{\sum_k \alpha_k m_k \mathcal{E}_\tau^k} + P^e \frac{\sum_k \alpha_k (1 - m_k) \mathcal{E}_\tau^k}{\sum_k \alpha_k m_k \mathcal{E}_\tau^k} \tag{10}$$

with

$$\mathcal{E}_\tau^k = \sum_j \frac{\partial q_j^k}{\partial \tau} E_j.$$

If it is possible to target each type with a tax, the first-best solution is a type-specific differentiated tax where:

$$\tau_k^* = \frac{\phi}{m_k} + P^e \frac{1 - m_k}{m_k}. \tag{11}$$

In practice, it is unlikely that policy-makers will be able to perfectly target consumer types subject to different levels of misperceptions; instead, a second-best tax should be used, which will induce a deadweight loss. The following proposition shows that the size of this loss is proportional to the variance in misperceptions, variance in the energy levels of the products being chosen, and the distance between the first- and second-best tax.

Proposition 1 *The deadweight loss associated with the second-best tax, τ^*, relative to the first-best tax, τ_k^*, is proportional to:*

$$\overline{\boldsymbol{var}}(E) \cdot \boldsymbol{var}(m_k) \cdot E_k[(\tau^* - \tau_k^*)^2], \tag{12}$$

where $\boldsymbol{var}(m_k)$ is the variance of the distribution of misperceptions, $\overline{\boldsymbol{var}}(E)$ is the variance in energy levels E_j for products chosen under the tax, and $E_k[(\tau^* - \tau_k^*)^2]$ is the distance between *the second-best differentiated tax and the first-best fully differentiated taxes t_k^*, respectively.*

Proof. See Appendix A.

Proposition 1 confirms the intuition of the welfare decomposition formula and shows that a (non-differentiated) tax instrument performs poorly with higher degrees of misperception heterogeneity. As in the welfare decomposition formula, two dimensions of heterogeneity matter. The one driven by misperceptions, captured by the term $\mathbf{var}(m_k)$, and the one driven by the choice set, captured by the term $\overline{\mathbf{var}}(E)$.[12] For quantity-based instruments, we show next that demand heterogeneity, especially misperceptions, has less of a role to play.

2.4.2 Optimal Standard

There is a large theoretical literature on minimum quality standards (Leland, 1979; Ronnen, 1991; Crampes and Hollander, 1995) that is relevant to the design of energy standards. In a realistic setting with multiple products and firms, it is, however, difficult to have tractable solutions. The cost function and strategic interactions between firms are two important elements to consider.

Therefore, to highlight how heterogeneous misperceptions and the design of an optimal standard interact, we are making two simplifying assumptions. First, we assume that there is no correlation in the cost function between adjusting energy level and other attributes. Second, we rule out strategic interactions between firms. Under these assumptions, we can define the purchase prices of product j as follows:

$$P_j(E_j) = c(E_j) + \omega_j \tag{13}$$

where $c(E_j)$ is the product's manufacturing 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. Under these assumptions, we are modeling the scenario discussed above where all non-compliant products are considered marginal to the regulation. We, therefore, rule out

[12]The variance $\overline{\mathbf{var}}(E)$ is not solely determined by the supply-side. This variance is also determined by demand-side parameters given that the market share of each product j for each type k enters the computation. To see this, note that we have $\overline{\mathbf{var}}(E) = E_k[\mathbf{var}^k(E)]$, where $\mathbf{var}^k(E) = \sum_j q_j^k \cdot E_j^2 - \left(\sum_j q_j^k \cdot E_j\right)^2$, and the share q_j^k is the demand function, which is function of the misperception parameter m_k and other preference parameters because it is a choice probability obtained from the type-specific and alternative-specific decision utility specified in Equation 2.

product exit and movement in the non-energy dimension.

As before, we will assume a linear and additive externality cost and k different consumer types subject to a misperception m_k. The following proposition derives the expression for the optimal standard.

Proposition 2 *When $P_j(E_j) = c(E_j) + \omega_j$ and with a constant marginal externality cost, ϕ, the optimal standard is a uniform standard, denoted $\bar{E}^*$. The uniform standard is the solution of the following equation, where c' denotes the derivative of the product cost with respect to energy level and P^e is the marginal price of energy:*

$$-c'(\bar{E}^*) = P^e + \phi. \tag{14}$$

Proof. See Appendix A.

The important takeaway from Proposition 2 is the optimal standard fully internalizes the distortionary effects of misperceptions. A uniform standard drives the variance of energy usage to zero. Because there is no substitution across products along the energy usage dimension, misperceptions do not affect consumer choice regardless of their distribution. The only element that matters 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 price of energy P^e inclusive of the externality cost ϕ.

The preceding result is specific to a restrictive type of cost function where there is no correlation between attributes. However, we note that even with a correlation between energy consumption and other attributes of the cost function, the impact of misperceptions can still be ameliorated, or, in some cases, fully internalized with an ABS if it is designed to reduce the variance of energy usage within product classes.[13] Note also, that in a context where energy prices and externalities vary across regions, as in the United States, a uniform standard determined at the Federal level will be second-best. A first-best uniform standard corresponds to a solution of $-c'(\bar{E}^*_r) = P^e_r + \phi_r$, where r defines a region with a specific

[13]Consider a case where the energy level, E_j, is correlated with another attribute in the cost function, a_j. In this case, the optimal standard is no longer uniform but attribute-based like in Ito and Sallee (2018). The extent to which misperceptions can be addressed depends on the elasticity of demand with respect to the attribute a_j. Suppose that a_j represents a categorical variable, size for instance, taking two values: $a_j = \{Small, Big\}$. If consumers are inelastic with respect to size, the distortionary effects of misperceptions can still be fully internalized with an ABS, where the marginal cost for each level of a_j would determine the optimal standard: $-c'(\bar{E}^*_j, a_j) = P^e + \phi$.

energy price-externality cost combination. But, even in this case, demand heterogeneity is not the source of a deadweight loss between the first and second-best optimal standard, unlike for the tax.

2.5 Uncertainty in the Distribution of Misperceptions

In many policy contexts obtaining precise estimations of heterogeneity in consumer perceptions may prove challenging. Further, estimation techniques and assumptions about discount rates or lifetimes may affect the estimated distributions and add uncertainty to the characterization of misperceptions. Therefore, it is worth considering the relative advantages of a tax versus a standard instrument if policy-makers are not fully informed about consumer misperceptions. Specifically, suppose policy-makers do not perfectly know the variance in misperceptions and their correlation with energy savings. Moreover, suppose they are averse to making mistakes in designing policies, i.e., they are ambiguity averse.[14] In this setting, minimum efficiency standards have a key advantage relative to a tax. Because they reduce the magnitude of the variance in energy savings, i.e., the term $\mathbf{sd}_{\Delta \bar{E}^k_s}$ in Equation 8, minimum efficiency standards reduce the importance of heterogeneous misperceptions in the evaluation of welfare. Therefore, all else equal, ambiguity-adverse policy-makers would prefer variance-reducing standards over variance-increasing taxes.

2.6 Discussion: Empirical Requirements

A key challenge in applying our framework is that information about the whole distribution of misperceptions must be recovered. Specifically, we want to estimate m_k, the magnitude of the misperception for each type k, together with α_k, the fraction of the population subject to such misperception. Moreover, to measure the welfare of each component of the decomposition formula it is also important to have a rich demand model in which the market share elasticities with respect to the different policy instruments are also estimated. In practice, this requires estimating the marginal utility of income (i.e., the coefficient on price) and characterizing own- and cross-price elasticities.

Our welfare decomposition formulae, however, offer a way to quickly assess the role of heterogeneous misperceptions as a function of a few sufficient statistics. If analysts have prior information on such a distribution it is possible to determine if the impact of misperceptions could be economically important.

[14] Berger and Bosetti (2020) provide evidence that policy-makers are ambiguity-averse in the climate context.

In what follows, we describe our data and revealed preference approach to recovering heterogeneity in perceptions of energy operating costs. Our empirical strategy recovers distributions (the weights α_k) of m_k for different income groups. We also estimate heterogeneity in the marginal utility of income and quality in the non-energy dimension in order to obtain a full demand system with rich substitution patterns.

3 Data and Environment

Our empirical investigation focuses on the U.S. refrigerator market, which offers several advantages. First, 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 vary idiosyncratically across households, the characteristics of a refrigerator, such as its size, door design, and presence of an 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; this 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, which allows us to identify the demand parameters of interest. Third, the refrigerator market is important in the United States and elsewhere, and it is expected to grow particularly fast in developing countries in the upcoming decades (Gertler et al., 2016).

Our main data source is transaction-level data from a large U.S. appliance retailer. The sample includes all in-store transactions involving refrigerator purchases during the period between 2008 and 2012.[15] 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's model number, and a consumer identifier. For a large subset of transactions, the identifier is matched with household demographics collected by a data aggregator (Table C.1). Detailed attribute information for each manufacturer's model number is also available and includes the manufacturer's reported energy use; refrigerator's dimensions (width, height, and depth); whether a product is ENERGY STAR certified; the presence of an ice maker; and color, brand, door design, and several other features pertaining to design and technology options. One particular feature of the U.S. appliance market is 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 the same retailer is subject to weekly variations that can exceed 20%, and those are model-specific variations and not

[15]Online transactions for refrigerators during that period were less than 5% and excluded from the sample.

perfectly correlated within and across brands. We further discuss the exogeneity of this variation in what follows and provide more details in Appendix E.

To construct our main dataset, we match the transaction data with local energy prices and rebate information. We construct energy prices from the Energy Information Administration's (EIA's) Form 861, which contains revenue and quantity of kWh consumed by residential consumers. Together, these variables provide a measure of the average annual electricity price for each electric utility operating in the United States. The EIA also provides information about which utility is operating in each county; this allows us to compute average electricity prices at the county level. If more than one utility serves a county, we take the sales-weighted average of those utilities' prices. Prices are highest in New England ($0.14 to $0.20/kWh) and lowest in the Midwest and South ($0.06 to $0.10/kWh). There is also variation over time in price: some states experienced price increases and others price decreases over the study period (See Appendix C for a graphical display of electricity price variation across regions).

We estimate the annual energy costs for each model-year-store location by multiplying the annual kWh consumption reported by the manufacturer by the average energy price for the year in the store's county. The energy price variation is largely driven by the variation in fuel costs different electric utilities face. However, we consider and carefully address potential bias arising from correlations between energy prices and energy efficiency preferences in section 4. Higher energy prices create larger differences in operating costs between high- and low-efficiency models. In Appendix C, we show a graphical display of the price and energy cost variation used in the analysis. Almost all models sold in our sample are less than $2,000, though there are some much higher-priced models available. Using the average life expectancy for refrigerators of 18 years and the average discount rate of 5% used by the DOE for appliance standards, the lifetime costs range from $555 for the 10th percentile of energy price to $1,000 for the 90th percentile. Energy costs are thus a large proportion of total product costs: the ratio of lifetime cost to purchase price ranges from .44 at the 10th percentile of energy price to .79 at the 90th percentile.

Both state governments and electric utilities offered rebates for energy-efficient appliances during the sample period. The State Energy Efficiency Appliance Rebate Program (SEEARP) was funded as part of the stimulus package of the American Recovery Act. This program led to generous rebates for ENERGY STAR-certified products during the years 2010 and 2011. Several electric utilities also offered rebates for ENERGY STAR-certified refrigerators. Both rebate programs varied across time and regions. We construct a measure of average rebate at the county-week level using SEEARP data collected by Houde and Aldy (2017b) and utility rebates from the Database of State Incentives for Renewables & Efficiency (DSIRE). In the estimation, we do not explicitly distinguish between SEEARP and utility rebates. Houde and Aldy (2017a) show that utility rebates have a small influence on the adoption of ENERGY STAR-certified products in comparison to SEEARP rebates. Houde and Aldy (2017b), which relies on similar data, also finds that the generous but short-lived SEEARP rebates did little to delay or pull forward purchasing decisions. In this context, dynamics in the purchasing decision thus appear to be limited, which motivated our static model.

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 comprise a consumer's consideration set).[16] We focus on modeling the purchase decision conditional on the fact a consumer has decided to buy a refrigerator at a given week and in a given store. Thus, we do not explicitly model the timing decision and the choice of the retail chain store. The marginal effects of the coefficients on price and energy costs, therefore, capture the substitution across different models offered.

We carry out the estimation using a large subsample of several million transactions. Each model estimated uses a large random subsample of approximately 200,000 transactions. We restrict the sample to transactions made by households owning their housing unit with the goal of focusing on transactions made by consumers who are likely to pay for the energy-operating costs of their appliances.[17]

4 Recovering Heterogeneous Misperceptions

4.1 Model Estimation

We begin by estimating the joint distribution of the parameters η and m for a discrete choice model of the general form:

[16] In almost all zip codes, there is 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.

[17] Due to principal-agent problems, landlords may underinvest in energy efficiency in instances where they purchase appliances and tenants pay utility bills.(Levinson and Niemann, 2004; Gillingham, Harding, and Rapson, 2012; Myers, 2020) While this could be another rationale for minimum efficiency standards, this study is aimed at uncovering the role of homeowner perceptions of energy operating costs.

$$U_{ijkrt} = \gamma_{ijrt} + \eta_k(-P_{jrt} - m_k P^e_{rt} \cdot E_{jrt}) + \epsilon_{ijkrt}, \tag{15}$$

where the subscript r denotes the region (zip code) where consumer i of type k makes a purchase of product j, and the subscript t denotes the time (i.e., week). Again, we denote the share of each type k by α_k, which corresponds to the discrete probability density we want to estimate. The purchase price, P_{jrt}, and the lifetime energy costs, $P^e_{rt} \cdot E_{jrt}$, are the two main regressors. In the above equation, we make it explicit that the marginal price of energy varies across regions and time. The term γ_{ijrt} denotes product- and consumer-specific controls for perceived quality, which we later describe in detail. Finally, ϵ_{ijkrt} represents idiosyncratic preferences.

We use two different estimators to characterize this distribution: 1) our preferred semi-parametric estimator, as developed by Fox et al. (2011) (FKRB), and 2) a parametric random coefficient logit model. We use FKRB as our primary approach because it does not impose any functional form on the joint distribution of η and m. The second approach is more standard in the literature, but it comes at a cost because it imposes a specific functional form for the distribution of misperceptions.

The intuition behind the FKRB approach is a continuous distribution of random parameters can be approximated by estimating population weights over a discretized support. We first divide the support of η and m into K grid points: $\beta_k = \{\eta_k, m_k\}$, $k \in K$. These can be thought of as the K different types of consumers in the population, which corresponds to the heterogeneity specified in our theoretical model. Assuming ϵ_{ijkrt} is independent and identically distributed (IID) and follows a generalized extreme value distribution, we can use a parametric multinomial logit to compute the choice model for each β_k, where the probability of choosing product j given β_k is denoted $q_j(\beta_k) = q_j^k$, and the subscripts r and t are omitted to simplify the exposition. The choice probability, q_j, is then a mixture of K multinomial logit models:

$$q_j = \int q_j(\eta, m) dF(\eta, m) \approx \sum_k^K \alpha_k q_j^k \tag{16}$$

where $\sum_k^K \alpha_k = 1$ because the weights are a discrete probability density function (pdf) that approximates the true underlying continuous function. By choosing a parametric form for the choice model, each q_j^k can 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 q_j as the dependent variable, q_j^k as regressors, and α_k, $\forall k \in K$ as coefficients to be estimated. To ensure the weights α_k sum to one, we estimate a linear least squares model with the constraint: $\sum_k^K \alpha_k = 1$.

The FKRB estimator suffers from the curse of dimensionality, meaning not all parameters can be estimated non-parametrically. As we describe in more detail later, we wish to control for several high-dimensional fixed effects to account for important consumer-specific, unobserved differences in perceived quality. However, estimating each of these quality indexes as random coefficients discretized with over 100 grid points in the FKRB framework would rapidly become intractable. To overcome this, we perform the two-stage estimation suggested by FKRB. In the first step, we use simulated maximum likelihood to estimate all the parameters of the model with a parametric random coefficient logit where η and m are assumed to follow a multivariate normal with an unknown mean and covariance matrix. We directly estimate η and m, as specified in Equation 15, where η is factorized to multiply all the money-metric variables. In the second step, we estimate the non-parametric joint distribution of η and m following the FKRB procedure previously outlined, fixing all of the other regressors at their means estimated in the first stage.

To account for uncertainty due to the two-step estimation, we implement the subsampling bootstrap, such that the standard errors of each discrete weight of the pdf capture the variation across subsamples (Politis, Romano, and Wolf, 1999).[18] When applying the subsampling bootstrap, it is important to be careful regarding making inferences for parameters situated at the edge of the support of the distribution, which often arises with the FKRB estimator. The choice of the grid points, especially at the end of the support, should therefore be cautiously chosen; therefore, we perform various robustness checks with respect to the grid choice. We provide additional details about the estimation procedure in Appendix D.

Conveniently, the first stage of the two-step procedure recovers a joint distribution of η and m using a parametric random coefficient logit. Therefore, we can readily compare these estimates with those from the FKRB estimator.

[18]The subsampling bootstrap is recommended by FKRB because the estimated distribution may bunch at the edge of the support. It consists of slicing the dataset in S subsamples, without replacement, and performing the estimation for each subsample. We average the estimated joint pdfs across all subsamples.

4.2 Identification

The key challenge in identifying our parameters of interest, η and m (as well as their joint distribution), is to control for all aspects of perceived quality that might be correlated with the purchase price or electricity cost of each refrigerator model. Our identification strategy relies on using high-dimensional fixed effects that capture rich dimensions of product quality.

In our main model, for each grid point $k \in K$ (effectively consumer type), the alternative-specific utility for individual i of type k for product j in zip code r in week-of-sample t is as follows:

$$U_{ijkrt} = \eta_k(P_{jrt} + m_k P^e_{rt} \cdot E_{jt}) + \tau ES_{jt} + \phi Rebate_{jrt} + \gamma_j + Demo_i \times Att_{jt} + \epsilon_{ijkrt}, \quad (17)$$

We first control for perceived quality using time-invariant product fixed effects that are constant across regions and consumers; these are denoted γ_j for each j and capture the attributes of model j. One concern is that some attributes might be perceived differently by different groups of consumers. For attributes correlated with energy usage, this might be particularly problematic for our identification. To control for such potential confounders, we use demographic information interacted with a subset of attributes that are strongly correlated with energy use, denoted $Demo_i \times Att_{jt}$ in Equation 17. Specifically, we interact income level, education, age of the head of the household, family size, and political orientation with the refrigerator's overall size and freezer location (i.e., top-freezer, bottom-freezer, or side-by-side refrigerator-freezer).[19]

Finally, we include an indicator for the ENERGY STAR certification (ES_{jt}) because consumers perceive ENERGY STAR-certified models as being of higher quality, irrespective of their energy costs (Houde, 2018a). We also include variables for ENERGY STAR rebates ($Rebate_{jrt}$) and the government incentives to purchase certified refrigerators, which vary over time and region.

For both the parametric random coefficient logit and the FKRB approach, identification of the joint distribution of η and m is induced by variation that leads to substitution across products. Prices, energy operating costs, and entry and exit of products are the main sources of variation that we exploit. Given that the kernel of the discrete choice model is the

[19]These two attributes are highly correlated with a refrigerator's reported energy use. They explain as much as 70% of the variation in energy use across models (Houde and Myers, 2021).

multinomial logit, the joint distribution of the parameters η and m will allow for relaxing the independence of the irrelevant alternative (IIA) assumption the multinomial logit imposes.

As described in the data section, the variation in refrigerator model price, P_{jrt}, is driven by the retailer's national pricing algorithm. The retailer frequently dictates model-specific price changes with weekly variations that can vary by as much as 20%. In Appendix E, we demonstrate the national pricing algorithm appears to provide credible exogenous variation in retail prices, which is uncorrelated with shifts in demand. Specifically, we show controlling for brand-by-week fixed effects, ENERGY STAR-by-week fixed effects, and other energy-related attribute-by-week fixed effects (i.e., size-by-week and freezer location-by-week) removes little of the variation observed in the normalized prices. This confirms the variation is driven by supply-side or idiosyncratic features and is independent of demand shocks for particular brands, ENERGY STAR status, or particular refrigerator features. We also demonstrate although there is some cross-sectional variation in model price, departures from the national pricing policy are small and infrequent, i.e., local store managers follow the policy closely. Finally, we note that during the sample period, several other large appliance retailers were also following a national price policy. This suggests price competition at the national level would serve to dampen local price competition and selection effects.

The lifetime energy cost, $P^e_{rt} \cdot E_{jt}$, is a function of the average appliance lifetime (T), the discount factor (δ^t), expected future energy prices at time t ($E[P^e_r|t]$), and the annual energy consumption of product j (e_j), where i indexes year t as follows.

$$P^e_{rt} \cdot E_{jt} = \sum_{i=t}^{t+T} \delta^{t-i} E[P^e_r|t] e_j. \tag{18}$$

We construct $P^e_{rt} \cdot E_{jt}$ pre-estimation based on several baseline assumptions: 1) consumers believe annual electricity prices follow a no-change forecast, 2) consumers rely on the county-level average electricity price corresponding to the year and zip code in which they made their purchase decisions, 3) the average life expectancy of a refrigerator is 18 years for all models, and 4) the average discount rate used by consumers is 5%. The first assumption implies consumers believe current electricity prices are the best predictor of future prices.[20]

[20]An alternative could be consumers are attentive to trends in futures prices for the major electricity-generation fuels. However, if so, their beliefs would not differ substantially from this no-change forecast. Forward curves for oil and natural gas rarely deviated substantially from spot prices for our sample period (Myers, 2019). Further, coal, the other major electric-generation fuel, is not traded with enough volume in

The baseline lifetime and discount rate assumptions were chosen to be in line with the latest revision of the DOE's minimum efficiency standards for refrigerators (Office of Energy Efficiency and Renewable Energy, 2010).

Based on the assumption of a no-change forecast of electricity prices, the identifying variation in $P^e_{rt} \cdot E_{jt}$ is driven by variation in contemporaneous local electricity prices across space and over time. As seen in Equation 18 the assumptions about lifetime and discount rates only change the multiplicative factor on electricity price. Therefore, as we describe in more detail later, we can readily accommodate sensitivity analysis to alternative parameter assumptions by rescaling our estimates of m post estimation.

Because the identifying variation is driven by local electricity prices, one might worry that even conditional on our controls, there is some remaining unobserved or uncontrolled-for correlation between those prices and preferences for energy efficiency, which might be biasing our results. We do several robustness tests to probe whether this is the case.

First, we implement an estimator by which we expand the dimension of the non-parametric joint distribution to include a third parameter for the ENERGY STAR label. This allows us to specifically control for any remaining unobserved heterogeneity in preferences for an environmental signal that might be correlated with preferences for energy cost and/or price. Second, we provide estimation results of our main model without the demographic-by-energy attribute controls ($Demo_i \times Att_{jt}$). If there is an important variation in perceived quality correlated with energy usage, the inclusion of these controls would substantially affect our estimates. We provide estimates of these alternative versions of our main model in section 5.2.

In addition, in Appendix E, we provide results from an exercise by which we estimate the effect of additional fine-grained controls on the relationships between market share and both energy costs and purchase price using the more restrictive but easier-to-implement conditional logit model. We estimate the model using aggregate market share data with our previously described preferred controls. We then include two sets of additional controls. First, we control for county-specific energy efficiency preferences by including indicators for the energy-related refrigerator attributes (ENERGY STAR indicator, size, freezer location) each interacted with county-fixed effects. Second, we control for trends in preferences for these features over time by including indicators both for the same energy-related refrigerator attributes each interacted with week-of-sample fixed effects and for brand-by-week-of-sample

futures markets for them to be particularly informative.

fixed effects. We demonstrate the coefficients on price and energy costs change little with the inclusion of either set of these controls. The fact the estimates of the conditional logit models are relatively unchanged with the addition of many fine-grained controls for quality suggests unobserved or uncontrolled-for preferences for energy efficiency are unlikely to be a significant biasing factor in the relatively parsimonious model used in the two-stage structural estimation.

Given the plausibly exogenous variation in P_{jrt} and $P^e_{rt} \cdot E_{jt}$ conditional on our preferred controls previously described, 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 not perfectly correlated changes in prices. Each product may also differ along several dimensions, which will be captured by the controls for perceived quality previously discussed. The relative market shares of products j and k will be a 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 values of m. In fact, complete inattention to energy costs could lead to negative m values. This could happen if consumers choose an option with high energy costs, but quality and price alone cannot rationalize this choice (i.e., the option appears to be dominated in the quality-price dimension). Flexibly controlling for perceived quality and price is thus crucial to capture instances where consumers may make mistakes by choosing a dominated option.[21] Large values of m can occur when consumers pay more attention to energy costs relative to prices.

4.3 Defining Consumer Misperceptions

In constructing $P^e_{rt} \cdot E_{jt}$, assumptions about the expected lifetime and discount rate are important for interpreting whether values of $m \neq 1$ correspond to misperceptions. Our baseline assumptions are in line with those used by the DOE in their latest revision of refrigerators' minimum standards (Office of Energy Efficiency and Renewable Energy, 2010). However, we acknowledge substantial heterogeneity in these parameters could exist as well.

To account for such heterogeneity, we allow for a broad range of behaviors to be considered

[21] Having a random coefficient on price is particularly important to identify the tail of the distribution of the parameter m, especially negative values. Consumers will choose a dominated option if they pay little attention to both price and energy efficiency. See Appendix F for evidence of such mistakes.

rational in our optimal policy analysis. In Table 1, we define four consumer types. We categorize consumers as subject to little or no misperceptions if the parameters recovered in the estimation could be rationalized by a discount rate ranging between 2% and 12% for the preferred 18-year lifetime. The low end of the range reflects the return on 3-year U.S. Treasury notes during our sample period and thus represents a market return for a risk-averse consumer with no credit constraints. The upper end reflects the average APR rate for credit cards during our period and, therefore, represents a cost of funds for consumers carrying credit card debt.

Note that for any values of $m_k \neq 1$ obtained under our baseline assumptions, we can solve for a discount rate or an expected lifetime such that m would be rationalized (i.e., equal to one). The first row in Table 1 displays the range of m under our baseline assumptions that could be rationalized under the corresponding discount rates or lifetimes in rows two and three for each consumer type. For instance, consider consumers with little or no misperceptions. Under a discount rate of 5% and an expected lifetime of 18 years, a value of $m = 0.62$ is equivalent to a value of $m = 1$ if we were to use a discount rate of 12% for an expected lifetime of 18 years. Alternatively, $m = 0.62$ could also be rationalized (i.e., rescaled to one) using a 5% discount rate and an expected lifetime of nine years.[22] A value of $m = 1.28$ under our baseline assumptions can be rationalized by a discount rate of 2% and an expected lifetime of 18 years or a 5% discount rate and a 28-year lifetime.

We define consumers as undervaluers of energy costs if they have implied discount rates above 12% for the preferred 18-year lifetime. We categorize consumers as overvaluers of energy costs if their choice can only be rationalized with discount rates lower than 2% for an 18-year lifetime. Finally, we categorize any consumers with negative values of m that cannot be rationalized by a positive discount rate or expected lifetime as having severe misperceptions.

For quantifying the welfare effects in our policy simulation, we assume attention is not endogenous, and therefore, m does not change in response to the policy.[23] Further, we treat

[22]Values of $0 < m < 1$ could be driven by an (expected) behavioral mistake of future home buyers in being inattentive to or myopic about appliance efficiency, rather than the home seller themselves making a behavioral mistake at the appliance store. The implications of either of these types of behavioral mistakes would then be similar in our welfare framework–they create a gap between decision and experienced utility either on the part of a homeowner using the appliance over its lifetime or a home buyer overvaluing inefficient appliances being conveyed with the home.

[23]If consumers did respond to a Pigouvian tax by exerting more effort and attention to energy costs it would have associated costs and benefits, which we do not explicitly model here (e.g., Gabaix, 2014; Matejka and McKay, 2015; Mackowiak, Matejka, and Wiederholt, 2023).

consumers categorized in column three in Table 1 as having no misperceptions (i.e., we set $m_k = 1$). To be conservative in how we define misperceptions, we calculate m_k for the other consumer categories using the extremes of the implied discount rates we consider to be "rational" along with the preferred 18-year lifetime. Therefore, we use a 2% discount rate to calculate m_k for consumers who overvalue energy costs. Likewise, we use a 12% discount rate to calculate m_k for consumers who undervalue energy costs. Finally, for any values of m_k below zero, we set $m_k = 0$ to correct for the previously discussed mistaken choices that led to an apparent positive valuation of energy inefficiency.

We provide results using several alternatives to explore the sensitivity of our optimal policy estimation of how we define misperceptions. First, we consider the possibility consumers who appear to overvalue energy costs are not actually misperceiving energy costs. Rather they are experiencing a so-called warm glow from or positive association with choosing a greener product, which is not being fully accounted for by controls for perceived quality in the estimation. For this scenario, we set $m = 1$ for all consumers with implied discount rates below 12%. We also consider the effects of using the distribution of m recovered using our baseline assumptions of a 5% discount rate and 18-year lifetime directly in the optimal policy analysis. In this scenario, rather than allowing for a range of behavior to be rational, as in our preferred approach, any deviation from $m = 1$ under the baseline assumptions is treated as a misperception. Finally, we define misperceptions as in our preferred approach, except the range of consumers treated as displaying little to no misperceptions (2 to 12% implied discount rate) is calculated based on a 12-year lifetime rather than an 18-year lifetime.[24]

5 Results

5.1 Non-Parametric Distribution of η and m

Our estimation results reveal substantial heterogeneity in consumers' perceptions. Figure 1 presents the marginal probability density functions (pdfs) of two key parameters η, and m, for both our preferred FKRB estimator and the random coefficient logit model. The parameter η captures the effect of energy costs in consumers' choices and m measures the relative degree of responsiveness to energy costs versus purchase price. Our preferred esti-

[24]Estimates based on home inspections and surveys of consumer beliefs by consumer advocacy and trade association groups (e.g., the National Association of Home Builders (National Association of Home Builders/Bank of America (NAHB), 2007), the International Association of Certified Home Inspectors (The International Association of Certified Home Inspectors (NACHI), 2020), and the Consumer Reports (Consumer Reports, 2019) report lifespans somewhat lower than those used by the DOE, ranging from nine to thirteen years.

mator (FKRB) suggests larger tails for the distributions of both parameters compared to the random coefficient logit (RCL), indicating greater heterogeneity among consumers.

For both estimators, the distribution of misperceptions, m, is highly dispersed, with statistically significant weights at grid points located between zero and one, as well as above one and below zero. A value of m greater than one suggests that some consumers place a high weight on energy operating costs in their purchasing decisions, with implied discount rates below 5%. On the other hand, values of m less than one imply that some consumers value a dollar of operating costs less than a dollar of purchase price. Negative values of m suggest that some consumers not only do not value energy costs in their purchasing decisions but also choose energy-inefficient models over cheaper, more energy-efficient ones of similar quality.

Our analysis in section 5.3 provides evidence that rationalizes some mass in the negative part of the support for m. Specifically, we observe that consumers regularly select dominated options when two otherwise identical models are available, and the more efficient one happens to be less expensive due to promotional pricing.

The left-skewed marginal pdf of η 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 of positive values of the price coefficient for a small share of consumers; this suggests they are inattentive to the purchase price and/or have high search costs. Large values of m are correlated with lower magnitude coefficients on price (see Appendix Figure D.1).

In Table 2, we summarize these results using the four broad types of consumers defined in Table 1. The first row of Table 2 reports the estimated shares of each of these types of consumers for our main model: 48.4% of consumers have no or little misperceptions, 37.5% undervalue energy operating costs, 17.0% overvalue them, and the remaining 12.6% are subject to severe misperceptions.[25] In Table D.1 in the appendix, we investigate the sensitivity of the distribution of consumer types to our baseline assumptions about the discount rate (5%) and expected lifetime (18 years). Specifically, we use the ranges of m values from row one of Table 1 to classify consumer types for a span of alternative baseline discount rates (2-18%) and expected lifetimes (4-24 years). In general, we find significant mass for each type across all scenarios. Undervaluation or overvaluation persists even for

[25] We estimate a discrete probability distribution function (pdf) where each pdf weight has a standard error. The percentages we report are average weights, and each weight has its confidence interval. Thus, the sum of these averages may not sum to one.

extreme cases where the discount rate and lifetime are very low/high.

5.2 Robustness

We now assess the robustness of our estimated distribution along several dimensions. First, it is crucial we have properly specified the grid for the FKRB estimator so we recover the empirical distribution of the parameters of interest at the relevant density and over the relevant range, particularly in the tails of the distribution.[26] We investigate the robustness of the estimated pdfs by re-estimating the model using alternative grids. We report the results in Table 2. In Model II, we perform the estimation with a smaller support for the grid.[27] We find this increases the bunching at the lower end of the support, but it does not affect the mean and the overall patterns for the joint or marginal pdf distributions (See Figure D.2, Appendix D.2). In Model III, we kept the same span for the support, but we increased the number of grid points. With a denser grid, we find more mass in the range $m \geq 0$, $m < 0.62$, but it has little impact on the mean and the qualitative patterns (See Figure D.2, Appendix D.2). This suggests our chosen grid characterizes heterogeneity in the parameters of interest well.

We also address whether our model is well-specified to capture idiosyncratic tastes for product quality. If these tastes are not properly accounted for and are correlated with prices of energy costs, it could bias the estimates of the joint distribution we aim to recover. For example, if consumers have a high m because they value energy efficiency highly simply due to environmental motives, we would not want to correct for that type of preference in our optimal behavioral policy. To assess the robustness of our estimates of the possibility some people may have environmental values that might be correlated with living in areas with high energy costs, we implement an estimator (Table 2, Model IV) with which we expand the dimension of the non-parametric joint distribution to include a third parameter for the ENERGY STAR label: τ. Doing so allows us to control for any remaining unobserved heterogeneity in preferences for an environmental signal (the ENERGY STAR label), which might be correlated with preferences for energy costs and/or price. This has little effect on the joint and marginal distributions of η and m; this suggests unobserved environmental

[26]Note the parametric normal distribution estimated in the first stage is not affected by the choice of the grid. Therefore, an additional way to assess the robustness of our results is to assess whether the results obtained with the parametric distribution are qualitatively similar.

[27]We also attempted to implement the estimator with both a very large minimum value and maximum value for the support of the distribution; however, we could not find a numerical solution for these estimators. The specification used for Model I in Table 2 is the largest span of the support for which we could recover estimates.

preferences are not a biasing factor in our estimation.

Next, we consider the possibility there may be heterogeneity in consumer preferences for particular models our main specification does not fully capture. Therefore, we estimate a version of our model where each product fixed effect, denoted γ_j, is also a random parameter. This specification introduces idiosyncratic preferences, in addition to the IID extreme value error term, for each product in the choice set. Note we still control for consumer-specific perceived quality indexed with interaction terms between demographics and product attributes. Given a large number of random product 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. The inclusion of these controls has little impact on the results (Model V in Table 2), which suggests unobserved heterogeneity for specific models is also not an important biasing factor in the distributions we recover.

We also look at the impact of having fewer controls for quality. Specifically, in Model VI in Table 2, we do not control for demographics interacted with energy-related attributes when we estimate the joint non-parametric distribution of η and m. The estimated joint distribution in Model VI is very similar to the specification with richer controls for quality. This again shows unobserved heterogeneity in perceived quality is not likely to be a major source of bias in our setting.

Our preferred model controls for refrigerator size interacted with demographics. However, size could also have an impact on choices beyond its impact on perceived quality. For a lot of households, kitchens are designed in such a way that it limits households' consideration set to refrigerators of specific range sizes. To account for such a size constraint, we estimate a model with household-specific consideration sets where the set of possible refrigerator models household i can choose from is determined by the size of the chosen model we observe plus or minus some variation (2.5 Cubic Feet), in overall size. This model cuts the size of the consideration set of each household by more than 60%, on average. Compared to the main model and other specifications, we now find a larger share of consumers that overvalue energy costs (Model VII in Table 2). However, undervaluation and severe misperceptions persist at similar levels to our main, unconstrained model. Accounting for consumer-specific consideration sets does not reduce the substantial heterogeneity in misperceptions we recover.

In Model VIII in Table 2, we use just the cross-sectional variation in electricity price (i.e., the average county-level electricity price across the sample), removing the changes in price over time. The 95% confidence intervals of the shares of each consumer type overlap with those of our main model. This is consistent with Houde and Myers (2021), where we demonstrated for the same time period that county-level electricity prices were highly correlated with capacity shares of coal, oil and gas-fired power plants, and exploiting just this variation had little effect on our estimates of consumers' average responsiveness to energy operating costs. The identifying electricity price variation is, therefore, largely driven by the variation in fuel costs different electric utilities face, and electricity price dynamics do not appear to be a significant biasing factor in estimating consumer responsiveness.

Finally, we consider whether misperceptions could potentially be confounded by credit constraints. Two long-standing questions are: first, is the undervaluation of energy costs more pronounced among low-income households[28], and second, if so, are those differences driven by credit constraints. To shed light on these questions, we use our estimator to recover a non-parametric distribution of η and m for six different income groups. Table 3 summarizes the distribution of m across our four broad categories of energy cost perception for each of the six income groups. We find the largest share of consumers with severe misperceptions, at almost 30%, is the lowest income group. For other income groups, the estimated distribution suggests this share ranges from 12% to 19%. The share of consumers that undervalue energy costs is larger among the lowest-income groups and decreases with income. The share of consumers subject to no or very modest misperception is the largest among the three highest income groups. Finally, the share of consumers overvaluing energy costs increases with income and is especially large, higher than 30%, for incomes greater than \$100k.[29]

Importantly, the patterns in heterogeneous misperception persist when we carry out the estimation by income groups. There is evidence of complete inattention to energy costs for some portion of the population, even in the highest income group. Therefore, the existence of severe misperceptions is not entirely attributable to income and, in particular, to explanations that relate to access to credit. Moreover, the results also provide an important robustness test that demonstrates systematic differences in the perception of quality of each product by different income groups is not the main driver of the heterogeneity patterns.

[28] Hausman (1979) was also the first to provide evidence that lower-income households could be subject to much more severe misperceptions. His heterogeneous estimates of implied discount rates relied, however, on a small sample size (n < 10).

[29] Figure D.3 (Appendix D) shows the joint distribution of η and m for each income group. The overall patterns are strikingly similar. Consumers who either respond strongly to energy costs or are completely inattentive to energy costs also tend to be less price sensitive (i.e., have a value of η close to zero).

5.3 Supporting Evidence

In what follows, we consider potential mechanisms that could rationalize the extreme values of m that we observe. Negative values of m imply some consumers must be choosing energy-inefficient models over less expensive, more energy-efficient models of similar quality. In Appendix F, we examine whether we observe this type of behavior in the data. We take advantage of the existence of pairs of identical refrigerator models that differ only with respect to energy efficiency and price. Although the more efficient model is almost always more expensive, sales events occur, making it the least expensive than the least efficient model. We thus have instances of a pair of products where there is a clear dominated option. In those cases, we find 1/3 to 1/2 of consumers still buy the dominated option. This could be for multiple reasons. For example, search frictions may make it difficult for consumers to compare all dimensions of all models in the store. Or, consumers may take price as a signal of (partially unobservable) quality. Either of these explanations could lead to a dominated option being chosen and could rationalize the severe misperceptions of energy costs that we observe. In testing for the role of search frictions, we have also paid a particular attention to stock-outs, which could also rationalize severe misperceptions. But even when we are very conservative in defining the consideration set such that we rule out stock-outs, severe misperceptions persist, which gives us confidence that they are driven by demand-side behavioral phenomena.[30]

Second, we assess the role of biased beliefs using a survey designed to assess energy literacy (see Appendix F for further details about the survey procedure). Focusing on questions about beliefs pertaining to the average usage of a refrigerator and the local average electricity price, we find that a significant share of consumers largely under- or over-value two key pieces of information consumers must use to determine a refrigerator's energy operating costs. The patterns in biased beliefs we found can rationalize some, but not all, of the heterogeneity in misperceptions we estimate. In particular, the fact we find a significant mass of consumers overvalue energy usage and energy price can explain the overvaluation of energy costs. Biased beliefs, however, are one of many other behavioral phenomena contributing to misperceptions, so we would not expect the distributions of biased beliefs and misperceptions to match perfectly.

[30]We consider several ranges of time around the purchase date in defining the consideration set. For our most conservative definition, we define the consideration set at the day-level as the products we observed sold on that day; products not sold that day are not in the consideration set.

6 Implications for Policy Design

We now turn to assessing how heterogeneous misperceptions affect the design of externality-correcting policies. First, we use our welfare decomposition formula to quantify the impact of misperceptions on the performance of a Pigouvian tax and an environmentally equivalent standard. Second, we use our framework to derive optimal behavioral taxes and standards. Again, we focus on showing the role of heterogeneous misperceptions in making welfare statements, but also on their influence on the design of policies. Accounting for heterogeneous misperceptions has an important impact on the stringency of externality-reducing policies.

For our policy simulations, we use the county-level average annual electricity prices constructed for our demand estimation and allow emissions factors to vary by 20 different electricity market regions in the U.S. For our main results, we use our preferred approach to defining misperceptions as described in section 4.3. However, we provide sensitivity analyses of our optimal policy evaluation to several alternative approaches to defining misperceptions. Finally, we use a subsampling bootstrap technique to account for uncertainty in our demand estimation. See Appendix G for additional details on the policy simulations.

6.1 Welfare Decomposition of a Policy Change

In Table 4, we use our framework to decompose the welfare impacts of a \$50 tax per ton of CO_2 and compare the results to a minimum efficiency standard designed to achieve similar reductions in carbon emissions. As of the writing of this paper, the social cost of carbon used by the U.S. Environmental Protection Agency is \$51 per ton. An upward revision has, however, been proposed. We thus present a second scenario where we use the recent estimate of \$185 per ton recommended by Rennert et al. (2022). For the standard, we assume manufacturers can adjust the energy efficiency level of any product to meet the revised standard, making all non-compliant products marginal to the regulation. In particular, we assume adjusting a product's efficiency increases manufacturing costs ($c(E_j)$), and markups (ω_j) are fixed such that compliance costs are fully passed through to the retail price ($P_j(E_j)$) as follows[31]: $P_j(E_j) = c(E_j) + \omega_j$. We rely on an estimate of the manufacturer's cost of providing energy efficiency from Houde (2018b), which uses the same data and covers the same sample period as this analysis. The cost function takes the following parametric form.

[31] In an oligopoly setting, firms would adjust prices strategically in response to a carbon tax or a more stringent standard. For this policy analysis, we exclude this margin to isolate the role of demand. This is the first step of a broader research agenda aiming to demonstrate the importance of heterogeneous biases in the design of environmental policies. In future research, it would be interesting to account for firms' strategic responses and investigate if market power magnifies or mitigates the impact of heterogeneous misperceptions.

$$c(E_j) = exp(\frac{\psi}{E_j} + \beta_j) \tag{19}$$

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

To find the standard equivalent to a \$50 carbon tax, we do a numerical search to find the standard that would achieve comparable carbon reductions by simulating the carbon emissions effects of incrementally strengthening the existing attribute-based standards in place during the sample period.

The first row in Table 4 indicates the percentage increase in the stringency of the existing standards' energy consumption. Existing standards are attribute-based and vary by refrigerator size and freezer location (i.e., top, bottom, or side-by-side refrigerator-freezer). An increase in stringency of 4% and 12% within each attribute-based refrigerator category would be needed to achieve an emissions outcome similar to an electricity tax of \$50/ton CO_2 and \$185/ton CO_2 respectively. The second row indicates the change in the damages from the externality (\$/capita), which is similar across the two policies, by construction. The third indicates the change in consumer surplus, which we decompose into two components: the change in decision utility and the change in the misperception correction. The misperception correction can be further decomposed into two components using the sufficient statistics approach outlined in Section 2.4.2: one that is attributable to the mean level of misperception and one that is driven by heterogeneity in misperceptions (i.e., the variance). Finally, in the last row of the top panel, we report the cumulative change in social welfare from the policy.

Several important results stand out. First, relative to the change in decision utility, the adjustment for misperception is large in magnitude for both types of instruments. In fact, this component dominates in the standard scenario. Second, heterogeneity is driving most of the adjustments attributable to misperceptions. If we were to focus only on average misperception, the gap between decision and experienced utility in the welfare evaluation would be quite marginal. Third, the direction of the welfare effect due to heterogeneous misperceptions differs for the tax versus the standard. As discussed in the theory section, this is to be expected. A tax exacerbates the effect of heterogeneous misperceptions, and, thus, contributes to a loss in welfare. A minimum efficiency standard, on the other hand, partly internalizes heterogeneous misperceptions, which causes a welfare gain. Finally, the impact of \$50 carbon tax or equivalent ABS that is 4% more stringent is modest, but the welfare impact could be larger if we were to consider a higher social cost of carbon. In particular, the change in externality responds non-linearly to an increase in the carbon tax or standard stringency. Overall, the results are qualitatively similar. Heterogeneous misperceptions are an important component of overall welfare and impact welfare in a different direction for a tax versus a standard.

6.2 Optimal Policies

We now simulate the demand model to find the optimal tax level and the optimal level for three different types of standards: a uniform kWh/year requirement for all models, a minimum performance standard, requiring all models to consume a certain kWh/year or less, and an attribute-based minimum performance standard, which closely resembles existing U.S. refrigerator standards.

For each policy, we report the level of the optimal policy using both the full distribution of misperceptions and the average degree of misperception.[32] In our simulations, we consider heterogeneity in the share of each income group across counties. Therefore, we incorporate spatial heterogeneity in the degree of misperception to energy costs due to differences 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. Because our empirical exercise focuses on the demand side, we do not explicitly model firms' strategic responses (e.g., the entry and exit of new products) in our policy simulations.

6.2.1 Result I: Average versus Heterogeneous Misperceptions

Table 5 presents results from our main demand model to define misperceptions along with alternative distributions used to characterize misperceptions. For all policy instruments, the benchmark for the level of the optimal Pigouvian tax without misperceptions is $50/ton.

The first important result is the difference in the optimal tax when we consider the full distribution of misperceptions ($F(m_k)$) versus the average ($E(m_k)$). If we consider the full distribution of misperceptions, we find a *negative* optimal tax of -$25.70/ton, using the FKRB estimator to quantify misperceptions. A negative tax means we should subsidize the carbon externality in this market. This surprising and counterintuitive result is driven by the large

[32] For both scenarios, we account for heterogeneity across income groups by taking a weighted average of the demand model estimated for each income group.

share of consumers who overvalue energy costs, which induces a large downward adjustment to the tax. If, instead, we ignore heterogeneity and set the tax based on the average value of m, the optimal tax is \$75.70/ton, which is about 50% more than the externality cost. Accounting for heterogeneous misperceptions has a dramatic effect on the level of the optimal tax. These results are caused by the asymmetric effect under- and over-valuation have on dictating the optimal level of the tax. This asymmetry arises because consumers who overvalue energy costs are also more sensitive to the tax. The average energy level induced by their choice is thus more tax-elastic. Given that the optimal tax weights each type as a function of the elasticity of energy with respect to the tax (see the term $\mathcal{E}_\tau^k = \sum_j \frac{\partial q_j^k}{\partial \tau} E_j$) in Equation 10), consumers with the largest tax elasticity will tend to dominate. In particular, overvaluation will induce a larger downward adjustment in the optimal tax formula compared to the upward adjustment induced by undervaluation, even if the shares of consumers who over/under value are the same. Therefore, in setting the optimal tax, overvaluing consumers have a disproportional effect.

A second important result is accounting for heterogeneity has very little effect on the levels of the three different types of standards we consider. For the uniform standard, we would expect this to be true, as shown in Proposition 2. 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. 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 to make it 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 55% of the existing standard (i.e., it is 45% more stringent). However, even with this more flexible type of standard, we find accounting for heterogeneous misperceptions has a negligible effect on the optimal level of standard.

6.2.2 Result II: Taxes vs. Standards

In Table 5, we also observe a third major result of our policy analysis: across all scenarios, standards induce welfare gains much larger than the gains under the Pigouvian tax. The ability of a standard to address misperceptions by reducing the variance in energy costs while simultaneously addressing the externality provides another clear advantage for this type of instrument.

We note that in this policy analysis, the optimal uniform standard is still second-best because we consider heterogeneity in energy prices and CO_2 emission factors across regions, which reflects the U.S. context. The uniform tax is also second-best both for these reasons and because it cannot target the misperception level for each consumer type. The analytical expression for the welfare loss associated with a non-differentiated behavioral Pigouvian tax relative to a fully differentiated tax is derived in Appendix A. When we quantify it numerically, we find that the deadweight loss of a second-best versus a first-best tax is approximately \$28 per consumer. For the standard case, this deadweight loss is just \$4.5 per consumer.

In our setting, the convexity of our cost function, which is determined by the functional form we use in Equation 19, impacts the welfare results. Because the convexity of this expression is low–on the order of $1e-6$–it is worth considering how sensitive our welfare comparison is to our cost functional form assumptions.

For illustrative purposes, consider a scenario in which the cost of providing energy efficiency is unrealistically highly convex and takes a value of four.[33] A more convex cost function disadvantages a standard relative to a tax. However, even with this extreme scenario, the variance in differentiated uniform standards would need to be almost twice as large ($28/4 = 7 > 4.5$) to make the tax more desirable. Therefore, in our context, the welfare ranking does not appear to be sensitive to functional form assumptions about compliance costs.

6.2.3 Result III: Sensitivity to the Distribution of Misperceptions

Our first result demonstrated using heterogeneous versus an average degree of misperception can have a large effect on the optimal tax, but not on optimal standards. We now explore the sensitivity of optimal policy to how misperception heterogeneity is characterized.

We first consider the effect of the estimator. The main advantage of the FKRB estimator is it flexibly recovers heterogeneity without imposing symmetry on the distribution, which an RCL approach would. However, an RCL approach is potentially easier for policy-makers to implement. In the appendix, Table G.2 presents results using the first-stage estimates obtained with a parametric mixed logit estimated via maximum likelihood with a joint normal distribution for the coefficient on price and energy costs. Using this distribution, the

[33]Numerically, we found such a convex cost function for providing energy efficiency would be unrealistically high. It would mean decreasing the energy consumption of a particular refrigerator model by 1 kWh/year at the margin will increase the marginal cost by a factor of four.

optimal tax based on the full distribution of misperceptions is \$111.30/ton. This is about twice the level of the no-misperception benchmark, and it is much lower than the optimal tax computed with the average degree of misperception: \$146.30/ton. Qualitatively, the direction of the bias is the same as before: using the full distribution lowers the adjustment to the tax relative to the benchmark. However, the tax is no longer negative, which was induced by the share of consumers who overvalue energy costs. By imposing symmetry on the distribution of misperceptions, overvaluation no longer dominates.

In contrast to the dramatic differences we see in the optimal tax when using the distribution recovered by FKRB versus RCL, the optimal standards are quite similar. The uniform and minimum efficiency standards are exactly the same, and the stringency of the attribute-based standard is only two percentage points higher with the RCL estimates. This is again driven by the RCL distribution having less mass in the share of consumers who overvalue energy costs than the FKRB distribution. However, clearly, the welfare loss of setting policy based on a distribution with imposed symmetry rather than a semi-parametric estimator that better recovers the tails of the distribution will be much smaller for standards than taxes.[34]

We also investigate the role of overvaluation bias (i.e., values of m_k greater than one). Until now, we have assumed these values are entirely due to misperceptions. However, consumers may value energy costs beyond their pure monetary value due to environmental preferences that are not being otherwise controlled for in our preferred model. In particular, purchasing energy-efficient appliances could induce a so-called warm glow that should be considered as part of their true experienced utility. We consider what would happen if we were to assume there was no overvaluation biases (i.e., we set $m = 1$ for all consumers with an implied discount rate less than 12% for an 18-year lifetime). Appendix G's Table G.3 shows this drives the level of the optimal tax to a much higher level.[35] This is intuitive. We discussed how values of m_k greater than one were inducing downward adjustments and even driving the tax to a negative value. If we do not allow for m_k greater than one, such adjustments are not made, and only undervaluation bias is driving the adjustments. This leads to a tax much larger than the externality cost. Comparing the optimal standards from

[34] Farhi and Gabaix (2020) also illustrate how heterogeneous misperceptions affect the calculation of a Pigouvian tax. They implicitly consider a symmetric distribution of misperceptions. Qualitatively, the adjustment due to heterogeneous misperceptions that Farhi and Gabaix (2020) propose is thus similar to ours in Table G.1, but they are different than the results with the FKRB estimator.

[35] See also Appendix G for estimates of optimal behavioral policy allowing for a so-called warm glow using the RCL distribution. There is also an upward adjustment in the tax when allowing for warm glow, though the change is not as dramatic as for FKRB-based estimates.

this exercise versus our benchmark approach in column 1 of 5 shows accounting for warm glow has almost no effect on optimal policy. The stringency of the attribute-based standard is less than half a percentage point higher than the benchmark.

In Appendix G, we further explore the sensitivity of our results to assumptions about lifetimes and discount rates. We also report the full set of results for each scenario where we report the change in consumer surplus and externality costs and the standard errors for each result.

We notably consider the impact of a stricter definition of misperceptions by using the distribution of m recovered using our baseline assumptions for quantifying welfare, rather than allowing for a range of values to be rational (Tables G.5 and G.4). The optimal tax is still negative, but it is almost three times larger in magnitude (-$76.90) than in our benchmark. Conversely, the magnitude of the optimal standard is minimally affected. We also consider the effects of defining misperceptions as in our preferred approach, except we use a 12-year lifetime rather than an 18-year lifetime for categorizing each group of consumers (Table G.6). Here, again, this change in underlying assumptions about misperceptions increases the magnitude of the tax more than three times to -$87.50, relative to our main estimation, but it has a minimal impact on the optimal standard. Relative to our preferred approach, both of these alternatives result in a broader range of consumers being characterized as having misperceptions. The optimal tax adjusts downward due to the asymmetric effect of having more consumers with overvaluation relative to undervaluation.

The fact standards are much more robust than taxes with respect to the way we characterize heterogeneous misperceptions is an important advantage, especially in a context where policy-makers are averse to making mistakes in designing policies. More precisely, suppose policy-makers were to assign a probability that each column of Table G.2 could be true, the standard instruments will provide an additional benefit compared to the tax, depending on policy-makers' degree of ambiguity aversion.

6.2.4 Discussion: Supply-Side Factors

In evaluating the welfare costs of different types of standards, our implicit assumption is all the compliance costs were reflected in the marginal cost of providing energy efficiency. However, the way we model the supply-side response is unlikely to change the welfare rankings of the policies considered in our simulations. First, we consider the role of the convexity of the marginal cost function, which is determined by the parameter ϕ. For our preferred specification, we rely on the cost function estimated in Houde (2018b) ($\phi = 191$), which uses the same data as the current analysis. However, even if we allow for much more convex marginal cost functions, where $\phi = 300$ or $\phi = 400$, the welfare effects of optimal uniform and minimum efficiency standards still dominate the optimal tax (see Appendix Table H.1). For the attribute-based standard, it takes an almost doubling of the baseline ϕ, i.e., $\phi = 400$, for there to be little welfare improvement, thus making it comparable to a tax. At this extreme value, DOE's existing attribute-based standards maximize our social welfare function—i.e, welfare cannot be improved by changing the stringency of the policy.[36] We consider this an upper-bound on ϕ given there is ample evidence environmental regulations tend to overestimate compliance costs Harrington, Morgenstern, and Nelson (2000), especially when manufacturers self-report these costs as with appliance regulations.

Second, we consider the role of any sunk and fixed costs associated with compliance, which are not included in our stylized representation of the supply side. Given our policy simulation, fixed 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 to appliance standards indicate historical fixed compliance costs have been much lower.[37] Therefore, if the fixed component of compliance costs fell into the historical range, it would not overturn our welfare comparison.

For standards, we have shown the main source of welfare loss may come from a reduction in production variety. If a more stringent standard were to dramatically reduce the number of options, our previous welfare decomposition exercise suggests that this effect could dominate. However, we find little evidence that past standard changes lead to a smaller choice set. If anything, the refrigerator choice set in the U.S. context has expanded despite more stringent standards.

7 Conclusion

Recent work in behavioral public economics has demonstrated if consumers misperceive an aspect of product costs, optimal policies for addressing consumption externalities must ac-

[36]This can readily be seen in Table H.1) where the value of the optimal attribute-based standard for $\phi = 400$ is 100%, which means that no scaling of the existing standards is required to maximize social welfare.

[37]Spurlock (2013) and Brucal and Roberts (2019) show more stringent standards had a modest impact on prices and may, in fact, have led to a reduction in equilibrium prices. This work points 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 led to lower prices.

count for the *full* distribution of consumer misperceptions; the average degree of misperception is not sufficient. We extend this work to the context of regulating negative externalities from energy-using durables. While there has been a long-going policy debate on the relative merit of standard versus tax instruments in this context, the impact of heterogeneity in consumers' perceptions of energy costs on optimal policy design has not been formally considered.

This paper fills this gap using a behavioral welfare analysis to derive optimal standards and Pigouvian taxation in the presence of heterogeneous misperceptions of energy costs. Using a welfare decomposition, we demonstrate minimum efficiency standards confer a major advantage to policy-makers over Pigouvian taxation. Because minimum efficiency standards reduce the variance of the potentially misperceived attribute in the choice set, they reduce allocative inefficiencies induced by misperceptions. Conversely, a Pigouvian tax exacerbates such inefficiencies. Standards have thus the potential to internalize heterogeneous misperceptions, at least partly, in addition to addressing the externality. This insight about the advantage of a variance-reducing standard has also broader implications. Because there are fewer allocative inefficiencies induced by a minimum efficiency standard relative to a tax, the optimal standard can be less sensitive to the exact distribution of misperceptions. Thus, the preferred instrument to use may also be driven in part by the regulator's aversion to uncertainty in the distribution of misperceptions.

We demonstrate these insights using a behavioral welfare simulation based on empirically recovered distributions of consumers' perceptions of energy costs in the appliance purchasing decision. We use a large sample of transaction-level data from a U.S. appliance retailer. We recover the distribution using both a semi-parametric estimator and a more conventional parametric mixed logit model. Our behavioral welfare simulations demonstrate: 1) the level of the optimal tax varies greatly depending on whether policy-makers have access to the mean level of misperception versus the full distribution, 2) The optimal tax is highly sensitive to the assumptions we use to define misperceptions, such as consumers' discount rates and expected lifetimes, and to the estimation technique used to recover heterogeneity. Whereas, the level of optimal standards varies much less across the scenarios we consider.

References

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

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

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

Allcott, Hunt and Christopher Knittel. "Are consumers poorly informed about fuel economy? Evidence from two experiments". *American Economic Journal: Economic Policy* 11 (2019), 1–37.

Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky. "Regressive sin taxes, with an application to the optimal soda tax". *The Quarterly Journal of Economics* 134 (2019), 1557–1626.

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

Allcott, Hunt and Dmitry Taubinsky. "Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market". *American Economic Review* 105 (2015), 2501–38.

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

Allen, Jason, Robert Clark, and Jean-François Houde. "Search Frictions and Market Power in Negotiated-Price Markets". *Journal of Political Economy* 127 (2019), 1550–1598.

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* 78 (2005), 2095–2120.

Berger, Loic and Valentina Bosetti. "Are policymakers ambiguity averse?" *The Economic Journal* 130 (2020), 331–355.

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

Brucal, Arlan and Michael J Roberts. "Do Energy Efficiency Standards Hurt Consumers? Evidence from Household Appliance Sales". *Journal of Environmental Economics and Management* 96 (2019), 88–107.

Busse, Meghan, Christopher Knittel, and Florian Zettelmeyer. "Are Consumers Myopic? Evidence from New and Used Car Purchases". *American Economic Review* 103 (2013), 220–256.

Chan, Nathan W. "Pigouvian policies under behavioral motives" (2023).

Chetty, Raj, Adam Looney, and Kory Kroft. "Salience and Taxation: Theory and Evidence". *American Economic Review* 99 (2009), 1145–1177.

Consumer Reports. "How Long Will Your Appliances Last?" `https://www.consumerreports.org/appliances/how-long-will-your-appliances-last/` (Accessed on December 9, 2020). 2019.

Crampes, Claude and Abraham Hollander. "Duopoly and quality standards". *European Economic Review* 39 (1995), 71–82.

Diamond, Peter A. "Consumption Externalities and Imperfect Corrective Pricing". *Bell Journal of Economics* 4 (1973), 526–538.

Dubois, Pierre, Rachel Griffith, and Martin O'Connell. "The Effects of Banning Advertising in Junk Food Markets". *The Review of Economic Studies* 85 (2017), 396–436.

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

Farhi, Emmanuel and Xavier Gabaix. "Optimal Taxation with Behavioral Agents". *American Economic Review* 110 (2020), 298–336.

Finkelstein, Amy. "E-ztax: Tax Salience and Tax Rates". *Quarterly Journal of Economics* 124 (2009), 969–1010.

Fischer, Carolyn and Richard G Newell. "Environmental and technology policies for climate mitigation". *Journal of environmental economics and management* 55 (2008), 142–162.

Fischer, Carolyn, Louis Preonas, and Richard G Newell. "Environmental and technology policy options in the electricity sector: are we deploying too many?" *Journal of the Association of Environmental and Resource Economists* 4 (2017), 959–984.

Fox, Jeremy T, Kyoo Il Kim, Stephen P Ryan, and Patrick Bajari. "A Simple Estimator for the Distribution of Random Coefficients". *Quantitative Economics* 2 (2011), 381–418.

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

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

Gerster, Andreas and Michael Kramm. "Optimal Internality Taxation of Product Attributes". *American Economic Journal: Economic Policy* (2024).

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

Gillingham, Kenneth, Matthew Harding, and David Rapson. "Split incentives in residential energy consumption". *The Energy Journal* 33 (2012), 37–62.

Gillingham, Kenneth, Sébastien Houde, and Arthur Van Benthem. "Consumer Myopia in Vehicle Purchases: Evidence from a Natural Experiment". *American Economic Journal: Economic Policy* 13 (2021), 207–238.

Goldin, Jacob and Tatiana Homonoff. "Smoke gets in your eyes: cigarette tax salience and regressivity". *American Economic Journal: Economic Policy* 5 (2013), 302–36.

Goulder, Lawrence H, Marc AC Hafstead, and Roberton C Williams III. "General equilibrium impacts of a federal clean energy standard". *American Economic Journal: Economic Policy* 8 (2016), 186–218.

Grigolon, Laura, Mathias Reynaert, and Frank Verboven. "Consumer Valuation of Fuel Costs and Tax Policy: Evidence from the European Car Market". *American Economic Journal: Economic Policy* 10 (2018), 193–225.

Guiso, Luigi, Andrea Pozzi, Anton Tsoy, Leonardo Gambacorta, and Paolo Emilio Mistrulli. "The cost of steering in financial markets: Evidence from the mortgage market". *Journal of Financial Economics* (2021).

Handel, Benjamin R and Jonathan T Kolstad. "Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare". *American Economic Review* 105 (2015), 2449–2500.

Handel, Benjamin R, Jonathan T Kolstad, and Johannes Spinnewijn. "Information frictions and adverse selection: Policy interventions in health insurance markets". *Review of Economics and Statistics* 101 (2019), 326–340.

Harrington, Winston, Richard D Morgenstern, and Peter Nelson. "On the accuracy of regulatory cost estimates". *Journal of Policy Analysis and Management: The Journal of the Association for Public Policy Analysis and Management* 19 (2000), 297–322.

Hausman, Jerry A. "Individual Discount Rates and the Purchase and Utilization of Energy Using Durables". *Bell Journal of Economics* 10 (1979), 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. National Bureau of Economic Research, Oct. 2016.

Holland, Stephen P. "Emissions taxes versus intensity standards: Second-best environmental policies with incomplete regulation". *Journal of Environmental Economics and management* 63 (2012), 375–387.

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

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

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

— "Bunching with the stars: How firms respond to environmental certification". *Management Science* 68 (2022), 5569–5590.

Houde, Sébastien and Joseph E Aldy. *The efficiency consequences of heterogeneous behavioral responses to energy fiscal policies*. Tech. rep. National Bureau of Economic Research, 2017.

— "Consumers' Response to State Energy Efficient Appliance Rebate Programs". *American Economic Journal: Economic Policy* 9 (2017), 227–255.

Houde, Sebastien and Erica Myers. "Are consumers attentive to local energy costs? Evidence from the appliance market". *Journal of Public Economics* 201 (2021), 104480.

IEA. *Achievements of Energy Efficiency Appliance and Equipment Standards and Labelling Programmes*. License: Creative Commons Attribution CC BY-SA 3.0 IGO. Paris: IEA, 2021.

Ito, Koichiro and James M. Sallee. "The Economics of Attribute-based Regulation: Theory and Evidence from Fuel Economy Standards". *Review of Economics and Statistics* 100 (2018), 319–336.

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

Ketcham, Jonathan D, Nicolai V Kuminoff, and Christopher A Powers. "Estimating the heterogeneous welfare effects of choice architecture". *International Economic Review* 60 (2019), 1171–1208.

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

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

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

Leard, Benjamin, Joshua Linn, and Katalin Springel. "Pass-through and welfare effects of regulations that affect product attributes". *Resources for the Future Discussion Paper* (2019).

— "Have US Fuel Economy and Greenhouse Gas Emissions Standards Improved Social Welfare?" *Report. Resources for the Future, Washington, DC* (2020).

Leard, Benjamin, Joshua Linn, and Yichen Christy Zhou. "How much do consumers value fuel economy and performance? Evidence from technology adoption". *The Review of Economics and Statistics* (2017), 1–45.

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

Leland, Hayne E. "Quacks, lemons, and licensing: A theory of minimum quality standards". *Journal of political economy* 87 (1979), 1328–1346.

Levinson, Arik and Scott Niemann. "Energy use by apartment tenants when landlords pay for utilities". *Resource and energy economics* 26 (2004), 51–75.

Mackowiak, Bartosz, Filip Matejka, and Mirko Wiederholt. "Rational inattention: A review". *Journal of Economic Literature* 61 (2023), 226–273.

Matejka, Filip and Alisdair McKay. "Rational inattention to discrete choices: A new foundation for the multinomial logit model". *American Economic Review* 105 (2015), 272–298.

Myers, Erica. "Are Home Buyers Inattentive? Evidence from Capitalization of Energy Costs". *American Economic Journal: Economic Policy* 11 (2019), 165–88.

— "Asymmetric information in residential rental markets: Implications for the energy efficiency gap". *Journal of Public Economics* 190 (2020), 104251.

National Association of Home Builders/Bank of America (NAHB). "Study of Life Expectancy of Home Components". https://www.interstatebrick.com/sites/default/files/library/nahb20study20of20life20expectancy20of20home20components.pdf (Accessed on December 9, 2020). 2007.

Office of Energy Efficiency and Renewable Energy. "Final Rule Technical Support Document". https://www1.eere.energy.gov/buildings/appliance_standards/pdfs/refrig_finalrule_tsd.pdf (Accessed on December 9, 2020). 2010.

Politis, Dimitris N, Joseph P Romano, and Michael Wolf. *Subsampling*. Springer Science & Business Media, 1999.

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

Rennert, Kevin, Frank Errickson, Brian C Prest, Lisa Rennels, Richard G Newell, William Pizer, Cora Kingdon, Jordan Wingenroth, Roger Cooke, Bryan Parthum, et al. "Comprehensive evidence implies a higher social cost of CO2". *Nature* 610 (2022), 687–692.

Rodemeier, Matthias and Andreas Löschel. *Information Nudges, Subsidies, and Crowding Out of Attention: Field Evidence from Energy Efficiency Investments*. Tech. rep. Institute of Labor Economics (IZA), 2023.

Ronnen, Uri. "Minimum quality standards, fixed costs, and competition". *The RAND Journal of economics* (1991), 490–504.

Sallee, James M. "Pigou creates losers: On the implausibility of achieving pareto improvements from efficiency-enhancing policies". *NBER Working Paper 25831* (2019).

— "Rational Inattention and Energy Efficiency". *Journal of Law and Economics* 57 (2014), 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* 135 (2016), 61–73.

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

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

Taubinsky, Dmitry and Alex Rees-Jones. "Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment". *The Review of Economic Studies* 85 (2018), 2462–2496.

The International Association of Certified Home Inspectors (NACHI). "InterNACHI's Standard Estimated Life Expectancy Chart for Homes". https://www.nachi.org/life-expectancy.htm (Accessed on December 9, 2020). 2020.

World Bank. *State and Trends of Carbon Pricing 2022*. State and Trends of Carbon Pricing. License: CC BY 3.0 IGO. Washington, DC: World Bank, 2022.

Tables and Figures

Table 1: Categorization of Consumer Type

	Severe Misperceptions	Undervalue Energy Costs	Little or No Misperceptions	Overvalue Energy Costs
Misperception Parameter (m) 18 Year Lifetime 5% Discount Rate	< 0	[0,0.62)	[0.62,1.28)	> 1.28
Implied Discount Rate (%) For 18 Year Lifetime, $m = 1$	N/A	$(12,\infty]$	(2,12]	(0,2]
Implied Lifetime (Years) For 5% Discount Rate, $m = 1$	N/A	[0,9)	[9,28)	28+

Notes: This table provides definitions of the four broad categories of consumer types identified in the analysis. The first row provides the ranges of misperception parameters that belong to each category. The second row defines the implied discount rates associated with each range of of m from the first row, assuming an 18 year lifetime. The third row defines the implied lifetimes associated with each range of m from the first row, assuming a 5% discount rate.

Table 2: Estimated Share of Each Consumer Type (%):

	Severe Misperceptions	Undervalue Energy Costs	Little or No Misperceptions	Overvalue Energy Costs
Model I: Main Model	**12.63** **(2.27)**	**37.53** **(5.37)**	**48.42** **(5.70)**	**17.00** **(3.80)**
Model II: Smaller Grid	12.94 (2.20)	35.31 (5.22)	45.08 (5.84)	19.13 (3.16)
Model III: Denser Grid	13.56 (2.15)	42.47 (6.06)	39.00 (7.31)	18.67 (4.03)
Model IV: Random τ	15.83 (2.67)	30.33 (7.83)	48.00 (8.29)	19.36 (5.46)
Model V: Random τ & γ_j	15.38 (1.98)	35.94 (3.37)	32.06 (2.91)	16.63 (2.05)
Model VI: No Demographic $\times$ Attribute Controls	13.94 (1.51)	22.81 (3.11)	36.63 (5.16)	26.7 (4.52)
Model VII: Size-Constrained Consideration Set	10.20 (4.71)	39.00 (5.11)	19.86 (7.29)	46.00 (3.99)
Model VIII: Cross Sectional Elec. Price	13.56 (2.36)	34.94 (5.17)	37.56 (6.14)	13.75 (2.48)

Notes: The first four columns report the marginal PDF of m computed from the joint PDF of η and m and aggregated over ranges of m corresponding to the four consumer types we defined. Models II through VI correspond to different implementations of the FKRB estimator. We varied how we specified the discrete grid (Models II and III), the coefficients that are specified as random (Model IV and V), and the controls for quality (Model VI). Finally, Model VII uses information about the size of the refrigerator chosen to define consumer-specific consideration sets. Standard errors were computed through the subsampling bootstrap procedure and are in parentheses.

Table 3: Estimated Shares of Consumer Types By Income

Income Category	Consumer Type Severe Misperceptions	Undervalue Energy Costs	Little or No Misperceptions	Overvalue Energy Costs
< \$30k	27.43 (6.71)	49.00 (10.16)	17.86 (8.72)	6.67 (2.54)
\$30-50k	17.81 (3.82)	56.63 (6.44)	25.75 (7.84)	7.14 (1.26)
\$50-75k	12.75 (2.70)	43.88 (6.99)	31.85 (6.76)	18.80 (2.79)
\$75-100k	15.31 (2.05)	35.73 (5.39)	32.80 (6.56)	21.93 (2.46)
\$100-125k	19.56 (4.61)	28.92 (6.15)	41.07 (5.78)	23.00 (3.91)
≥\$125k	13.38 (2.48)	31.46 (6.34)	34.81 (5.76)	26.25 (4.96)

Notes: This table reports the marginal PDF of m computed from the joint PDF of η and m for the estimated distribution for each income group. The estimation is implemented separately for each income group using the FKRB estimator described in the main text. Standard errors were computed through the subsampling bootstrap procedure and are in parentheses.

Table 4: Welfare Decomposition: Pigouvian Tax versus Attribute-Based Standard (ABS)

	Pigouvian Tax 50$/$CO_2$	ABS: equiv to 50$/$CO_2$	Pigouvian Tax 185$/$CO_2$	ABS: equiv to 185$/$CO_2$
Δ Standard Stringency (%)	-	4	-	12
Δ Externality	-1.0	-0.9	-12.3	-11.5
Δ Consumer Surplus	-1.5	3.1	-10.2	9.8
Δ Decision Utility	-0.6	1.4	-7.0	4.5
Δ Misperception Correction	-0.9	1.7	-3.2	5.3
Δ Correction for Mean Level of Misperception	0.2	0.2	0.7	1.4
Δ Correction for Misperception Heterogeneity	-1.2	1.5	-4.0	4.0
Δ Social Welfare	-0.5	4.0	2.1	21.4

Notes: The first row in Table 4 indicates the percentage increase in the stringency of the existing standards' energy consumption. Standards are set based on refrigerator size and freezer location (i.e., top, bottom, or side-by-side refrigerator-freezer). The first row is the increase in standard stringency of the minimum efficiency standards that were in place during the sample period. A 4% increase means that minimum efficiency requirements would have to be 4% more efficient, relative to the baseline, to reduce the externality cost by the same amount as what is achieved by a 50$/$CO_2$ imposed on electricity prices. All the other numbers represent the welfare effect, measured in $/capita, for the different components of welfare. The change in consumer surplus is decomposed into the components due to decision utility and misperceptions. This latter component is further decomposed, using the sufficient statistic approach, to distinguish the effect of mean misperception versus heterogeneous misperceptions.

Table 5: Optimal Behavioral Policies

		FKRB and Heterogeneous r	RCL and Heterogeneous r	FKRB, Heterogeneous r, and Warm Glow	FKRB, Homogeneous r
Pigou Tax					
With $F(m_k)$:	\$/ton of CO_2	-25.7	111.3	92.5	-76.9
	Δ SW (\$/capita)	0.2	1.8	1.2	1.6
With $E[m_k]$:	\$/ton of CO_2	75.7	146.3	205.0	55.4
	Δ SW (\$/capita)	1.4	3.5	5.5	0.8
Uniform Standard					
With $F(m_k)$:	kWh/year	335.0	335.0	335.0	335.0
	Δ SW (\$/capita)	114.2	98.1	109.4	122.0
With $E[m_k]$:	kWh/year	335.0	335.0	335.0	335.0
	Δ SW (\$/capita)	101.0	94.7	101.0	101.0
Minimum Standard (kWh/year)					
With $F(m_k)$:	kWh/year	335.0	335.0	335.0	335.0
	Δ SW (\$/capita)	114.2	98.1	109.4	122.0
With $E[m_k]$:	kWh/year	335.0	335.0	335.0	335.0
	Δ SW (\$/capita)	95.6	94.7	100.0	97.0
Attribute-Based Minimum Standard					
With $F(m_k)$:	% of Existing Standard	55.6%	57.9%	56.0%	54.9%
	Δ SW (\$/capita)	91.3	76.0	89.1	94.7
With $E[m_k]$:	% of Existing Standard	56.1%	57.3%	56.1%	56.1%
	Δ SW (\$/capita)	78.3	73.6	80.3	78.9

Notes: The table reports the optimal policies evaluated using the full distribution of misperceptions (denoted $F(m_k)$) or only the average misperception (denoted $E[m]$). For each optimal policy, we report the change in social welfare relative to a baseline scenario with no policy (denoted Δ SW). Each column defines misperceptions using a different estimator and/or assumptions. The results with the non-parametric distribution of the FKRB estimator are denoted by FKRB. The second column uses the estimates from a Mixed Logit. The term "Heterogeneous r" means that we define misperceptions assuming that consumers could have an implicit discount rate that spans 2% to 12%. For the results with "Homogeneous r", we assume all consumers have a discount rate of 5%. For the results with "Warm Glow", we assume that overvaluation of energy costs, i.e., $m > 1$, corresponds to preferences (experienced utility), not a bias.

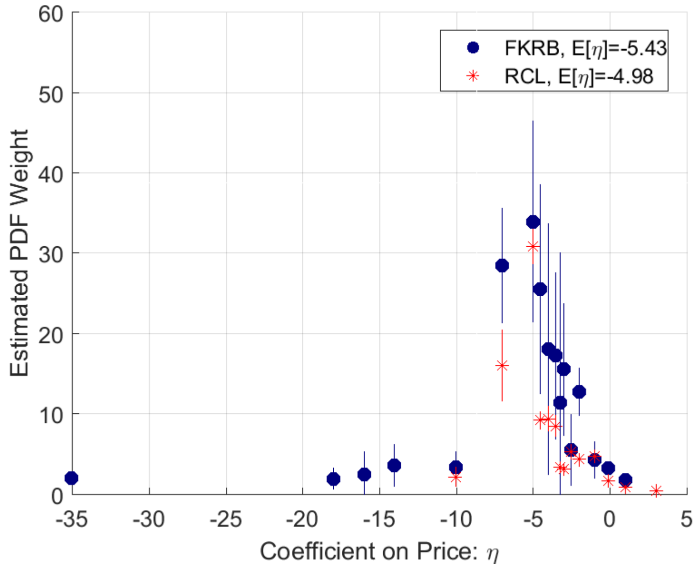

(a) Marginal PDF of η: $f(\eta)$

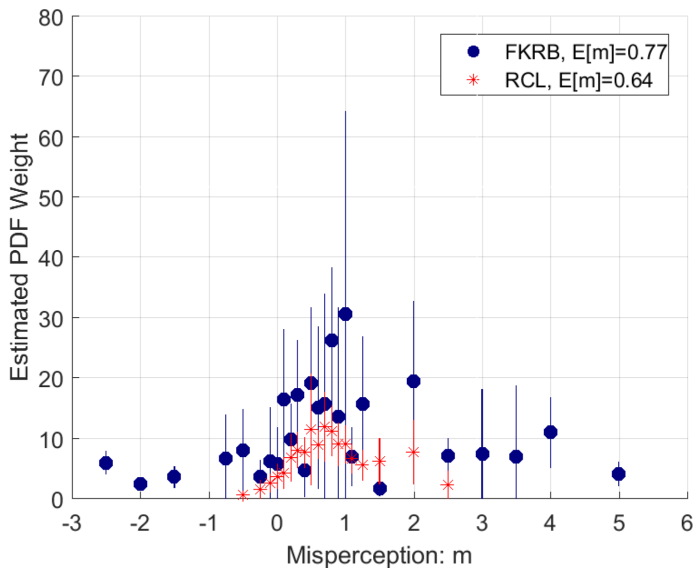

(b) Marginal PDF of m: $f(m)$

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 from the first-stage estimation with the parametric random coefficient logit. Each pdf weight corresponds to an average of the weights estimated across the bootstrap iterations. The pdf displayed in each panel is effectively a point-by-point average of 16 different pdfs. Therefore, the sum of the weights may not total 1. The support of the RCL is more narrow than the FKRB, which is why the pdf weights appear somewhat lower. m is calculated using the baseline parameters of a 5% discount rate and 18-year lifespan.

For Online Publication

A Theory: Proofs and Additional Results

Impact of Misperceptions on Average Energy Consumption For the multinomial logit model, the expression for the derivative $\frac{\partial \overline{E}^k}{\partial m_k}$ is calculated as follows:

$$\frac{\partial \overline{E}^k}{\partial m_k} = \frac{\partial}{\partial m_k} \sum_j q_j^k \cdot E_j = \sum_j \frac{\partial q_j^k}{\partial m_k} \cdot E_j = \sum_j \frac{\partial}{\partial m_k} \frac{e^{V_j^k}}{\sum_l e^{V_l^k}} \cdot E_j = \eta^k \cdot P^e \cdot \mathbf{var}_{E^k}, \quad \text{(A.1)}$$

where $\mathbf{var}_{E^k}$ denotes the variance in energy usage induced by the choices of a type-k consumer. The important take-away from the above expression is that this expression is always positive. The misperception parameter and expected energy usage are thus positively correlated.

As stated in the main text: higher m means $\Delta \overline{E}_\tau^k$ is more negative.

Optimal Behavioral Pigouvian Tax To derive the expression of the optimal Pigouvian tax with misperceptions, we start with a social welfare function consisting of the sum of consumer surplus, government revenue, and externality costs. The expression for social welfare after a tax τ is levied and included in the price of energy, P^e, is given by

$$\begin{aligned} SW(\tau) = \frac{1}{\eta} \cdot & \left[ln \sum_j^J e^{(V_j(\tau))} + \sum_j^J q_j (V_j^E(\tau) - V_j(\tau)) \right] \\ & + (\tau - \phi) \sum_j q_j(\tau) \cdot E_j, \end{aligned}$$

where we have used the expression for consumer surplus in Equation 3 and the fact that the gap between decision utility: $V_j(\tau) = \gamma_j - \eta(P_j + m(P^e + \tau)E_j)$ and experienced utility: $V_j^E(\tau) = \gamma_j - \eta(P_j + (P^e + \tau)E_j)$ is a 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 the expression for social welfare is given by:

$$SW(\tau) = \sum_k \alpha_k CS^k + (\tau - \phi) \sum_k \sum_j \alpha_k q_{jk} E_j. \tag{A.2}$$

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:

$$\tau^* = \phi \frac{\sum_k \alpha_k \mathcal{E}_\tau^k}{\sum_k \alpha_k m_k \mathcal{E}_\tau^k} + P^e \frac{\sum_k \alpha_k (1 - m_k) \mathcal{E}_\tau^k}{\sum_k \alpha_k m_k \mathcal{E}_\tau^k} \tag{A.3}$$

with

$$\mathcal{E}_\tau^k = \sum_j \frac{\partial q_j^k}{\partial \tau} E_j.$$

The Deadweight Loss of a Non-Differentiated Tax: Proof of Proposition 1

To approximate the deadweight loss associated with the non-differentiated tax relative, we compare the welfare under a first-best fully differentiated instrument with the welfare under a second-best non-differentiated instrument. To obtain tractable closed-form expressions, we follow Farhi and Gabaix (2020)'s approach and use a second-order Taylor approximation of the welfare function around the first-best tax. The proof proceeds in three steps.

- **Step 1**: Deriving the second-order Taylor expansion of the welfare function for the tax.

For a given type k, the social welfare function is:

$$SW^k = CS^k + (\tau - \phi) \sum_j q_{jk} E_j. \tag{A.4}$$

Replacing for the expression of the consumer surplus (CS), we obtain:

$$SW^k = \frac{1}{\eta_k} \cdot \left[ln \sum_j^J e^{(V_{jk}(\tau))} + \sum_j^J q_{jk} (V_{jk}^E(\tau) - V_{jk}(\tau)) \right] + (\tau - \phi) \sum_j q_{jk} \cdot E_j. \tag{A.5}$$

After simplifications, the first derivative of the above expression with respect to a tax τ is:

$$\frac{\partial SW^k}{\partial \tau} = -(m_k \cdot \tau + (m_k - 1) \cdot P_k^e \cdot \phi^k) \cdot m_k \cdot \eta_k \cdot \mathbf{var}^k(E), \tag{A.6}$$

where we have used the fact that

$$\sum_j \frac{\partial q_j^k}{\partial \tau} \cdot E_j = m_k \cdot \eta_k \cdot \left[\left(\sum_j q_j^k \cdot E_j \right)^2 - \sum_j q_j^k \cdot E_j^2 \right] = m_k \cdot \eta_k \cdot \mathbf{var}^k(E). \tag{A.7}$$

The second derivative of Equation A.6 is:

$$\frac{\partial^2 SW^k}{\partial \tau^2} = -(m_k \cdot \tau + (m_k - 1) \cdot P_k^e \cdot \phi^k) \cdot m_k \cdot \eta_k \cdot \frac{\partial \mathbf{var}^k(E)}{\partial \tau}(E) - \mathbf{var}^k(E) \cdot m_k^2 \cdot \eta_k. \tag{A.8}$$

- **Step 2**: Deriving an expression for the deadweight loss of the second-best tax (τ^*) using a Taylor approximation around the first-best fully differentiated tax (τ_k^*).

We want to compute the loss function: $\mathcal{L}(\tau^*) = SW(\tau^*) - SW(\tau_k^*)$. For each type k, using the second-order Taylor expansion, we have

$$\begin{aligned} & SW^k(\tau^*) - SW^k(\tau_k^*) \approx \\ & SW^k(\tau_k^*) + \frac{\partial SW^k(\tau_k^*)}{\partial \tau} \cdot (\tau^* - \tau_k^*) + 1/2 \cdot \frac{\partial^2 SW^k(\tau_k^*)}{\partial \tau^2} \cdot (\tau^* - \tau_k^*)^2 - SW^k(\tau_k^*) = \\ & 1/2 \cdot \frac{\partial^2 SW^k(\tau_k^*)}{\partial \tau^2} \cdot (\tau^* - \tau_k^*)^2 \end{aligned} \tag{A.9}$$

Note that we have simplified the above expression using the first-order condition that determines the tax τ_k^*: $\frac{\partial SW^k(\tau_k^*)}{\partial \tau} = 0$.

Using the second-order Taylor expansion derived in Step 1 and the expression for τ_k^* we have:[38]

$$1/2 \cdot \frac{\partial^2 SW^k(\tau_k^*)}{\partial \tau^2} \cdot (\tau^* - \tau_k^*)^2 = -1/2 \cdot \mathbf{var}^k(E) \cdot m_k^2 \cdot \eta_k \cdot (\tau^* - \tau_k^*)^2. \tag{A.10}$$

The loss function can then be approximated by:

$$\mathcal{L}(\tau^*) \approx -1/2 \cdot \sum_k \alpha_k \cdot \mathbf{var}^k(E) \cdot m_k^2 \cdot \eta_k \cdot (\tau^* - \tau_k^*)^2. \tag{A.11}$$

[38] In Equation A.8, the first term on the RHS cancels to zero once we replace the tax τ by the optimal tax τ_k^*.

- **Step 3**: Deriving an expression for the deadweight loss as a function of the variances.

The deadweight loss in Equation A.11 is an expectation over four random variables, η_k, m_k, $(\tau^* - \tau^*_k)^2$ and $\mathbf{var}^k(E)$, which vary with the type k and, thus, the discrete distribution defined by the weight α_k where $\sum_k \alpha_k = 1$. If we assume independence across those four variables, we have our main result:

$$\begin{aligned}\mathcal{L}(\tau^*) \propto &- E_k[\mathbf{var}^k(E)] \cdot E_k[m_k^2] \cdot E_k[\eta_k] \cdot E_k[(\tau^* - \tau^*_k)^2] = \\ &- \overline{\mathbf{var}}(E) \cdot (\mathbf{var}(m_k) + \bar{m}^2) \cdot \bar{\eta} \cdot E_k[(\tau^* - \tau^*_k)^2]\end{aligned} \tag{A.12}$$

where the expectations are taken over the discrete distribution of types k. If the four random variables are correlated, the covariance matrix will also matter in determining the size of the deadweight loss. Nonetheless, the product of the variances and the distance between the second-best and first-best taxes: $\overline{\mathbf{var}}(E) \cdot \mathbf{var}(m_k) \cdot E_k[(\tau^* - \tau^*_k)^2]$ will always be of the same sign and increases the deadweight loss induced by a non-differentiated tax, which is the main result of Proposition 1.

Optimal Uniform Standard: Proof of Proposition 2 The proof proceeds in two steps. The first step consists of showing that the optimal standard, from the first-best perspective with only a constant externality, is a uniform standard. The second step consists of showing that the first-best solution can be attained in a second-best setting.

- **Step 1**: A uniform standard is optimal in the first-best setting.

In the first-best setting, there is no misperception, and the social planner has perfect information about preferences and consumer types. Considering that there is a constant externality ϕ, the social planner's optimal standard consists of maximizing the quantity E_j for each consumer accounting for ϕ: i.e.,

$$\max_{E_j} \frac{1}{\eta^k} \cdot U^k_{ij} - \phi \cdot E_j, \tag{A.13}$$

where $U^k_{ij} = \gamma_j - \eta^k \cdot (P_j - P^e \cdot E_j) + \epsilon_{ij}$. Replacing P_j by $P_j = c(E_j) + \omega_j$, the first-order condition is:

$$-c'(E_j) = P^e + \phi. \tag{A.14}$$

Given that P^e and ϕ are constant, the optimal standard is uniform for all consumers in this setting.

- **Step 2**: A uniform standard is optimal in the second-best setting.

In the second-best setting, consumers are prone to misperceptions and the social planner does not have perfect information about preferences and consumer types. For each type k, the optimal standard consists of solving the following problem:

$$\max_{E_j,\ \forall j} CS^k - \phi \sum_j q_j^k E_j \tag{A.15}$$

Replacing the CS^k by the expression in Equation 3, taking the first-order condition for each E_j, and rearranging, we obtain:

$$-c'(E_j) = P^e + \phi + \frac{1}{q_j^k} \sum_l \frac{\partial q_l^k}{\partial E_j} \cdot \left(\phi - (m^k - 1) \cdot P^e\right). \tag{A.16}$$

When the derivative $\frac{\partial q_l^k}{\partial E_j}$ equals zero, we recover the solution of the first-best setting. Note that $\frac{\partial q_l^k}{\partial E_j} = 0$ only when all consumers are completely inattentive to E_j or all products have the same level of energy consumption. This latter case corresponds to a uniform standard.

B Choice Set Variance Pre- Versus Post-2014 Standard Change

In order to gain some insight into the effect of a real world standard change on the variance of the annual energy costs in the choice of refrigerator models available to consumers, we consider the DOE's 2014 standard change for refrigerators in the United States. In 2011, the DOE established the prevailing regulations for refrigerators, and these regulations were implemented in 2014. These standards specify the maximum annual energy consumption values according to the specific product type and the product's refrigerated volume. Compared to the preceding standards enforced from 2001 to 2014, the current standards were designed to yield energy savings of 25% for the majority of refrigerators.[39] Relative to the 2001-2014 period, the number of specific product types with their own minimum efficiency standards went up (from 18 product distinctions to 42 product distinctions).

We use data collected and provided by the Lawrence Berkeley National Laboratory (LBNL) on existing appliances and their features. LBNL wrote a "data crawling" script to collect information about refrigerators offered on the main U.S. retailers' websites. The data provided to us were collected during the period 2012 to 2016. To create "cohorts" of available models in each year, we take advantage of two variables that are used to indicate when the model entered the data set "date first collected" and when the model left the data set "date last collected." We use these variables as proxies for when the model entered and left the market. The dataset contains entries for 2,911 models available in 2013 and 3,559 available in 2015. A model was defined as available in year 20XX if it had a first collection date before or in 20XX and a last collection date after or in 20XX. However, information on kwh consumed per year is only available for a subset of those data, 594 for those available in 2013 and 1021 for those available in 2015.

Table B.1 shows summary statistics for both annual kWh consumption and annual kWh consumption by unit volume for 2013 and 2015. Assuming that the models for which the energy consumption information is available is a random sample of all refrigerator models, we can test whether the introduction of the 2014 standard had an impact on the variance of annual energy consumption of available models. In Table B.2 we report results of a ratio test for homogeneity of variances between models available in 2013 and models available in 2015. For both annual energy consumption and annual energy consumption per unit volume, we fail to reject that the variances are the same (i.e. the null hypothesis that the ratio of variances is equal to one). Whereas, tightening an existing standard would be expected to

[39]For more information, see the following from the Appliance Standards Awareness Project: `https://appliance-standards.org/product/refrigerators-and-freezers`, accessed January 2024.

reduce the variance of energy consumption for available models, the 2014 standard appears to have little effect. It is possible that having many more categories of appliance types served to increase the variance of annual energy consumption, thus offsetting some of the variance-reducing effects from tightening the efficiency standard.

Table B.1: Summary Statistics by Year Available

	Mean	SD	Min	Max	Count
2013 kWh/year	458.14	159.30	158	778	594
2015 kWh/year	449.27	163.72	158	855	1021
2013 kWh/year/ft^3	40.21	33.77	2	185	594
2015 kWh/year/ft^3	44.90	35.44	12	198	1021

Source: LBNL Refrigerator Model and Attribute Data

Table B.2: Test for Homogeneity of Variance

	SD 2013	SD 2015	F-Statistic	P-value	Observations
kWh per year	159.3	163.72	.947	.7705	1615
kWh per year/adjusted volume	33.77	35.44	.908	.9044	1615

Notes: This table reports the results of a ratio test for homogeneity of variances between models in the data set as of 2013 and models in the data as of 2015. We report the F-Statistic and upper one-sided p-value for the alternative hypothesis that the ratio of the standard deviation of 2013 models to 2015 models is greater than one.

C Additional Information about Data

Table C.1 summarizes the data used in the analysis.

Table C.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 between 2008 and 2012 for a single identifier (credit card).

Figure C.1 is reproduced from Houde and Myers (2021), gives a sense of the variation in electricity prices across regions and time. It depicts the 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 nine 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.

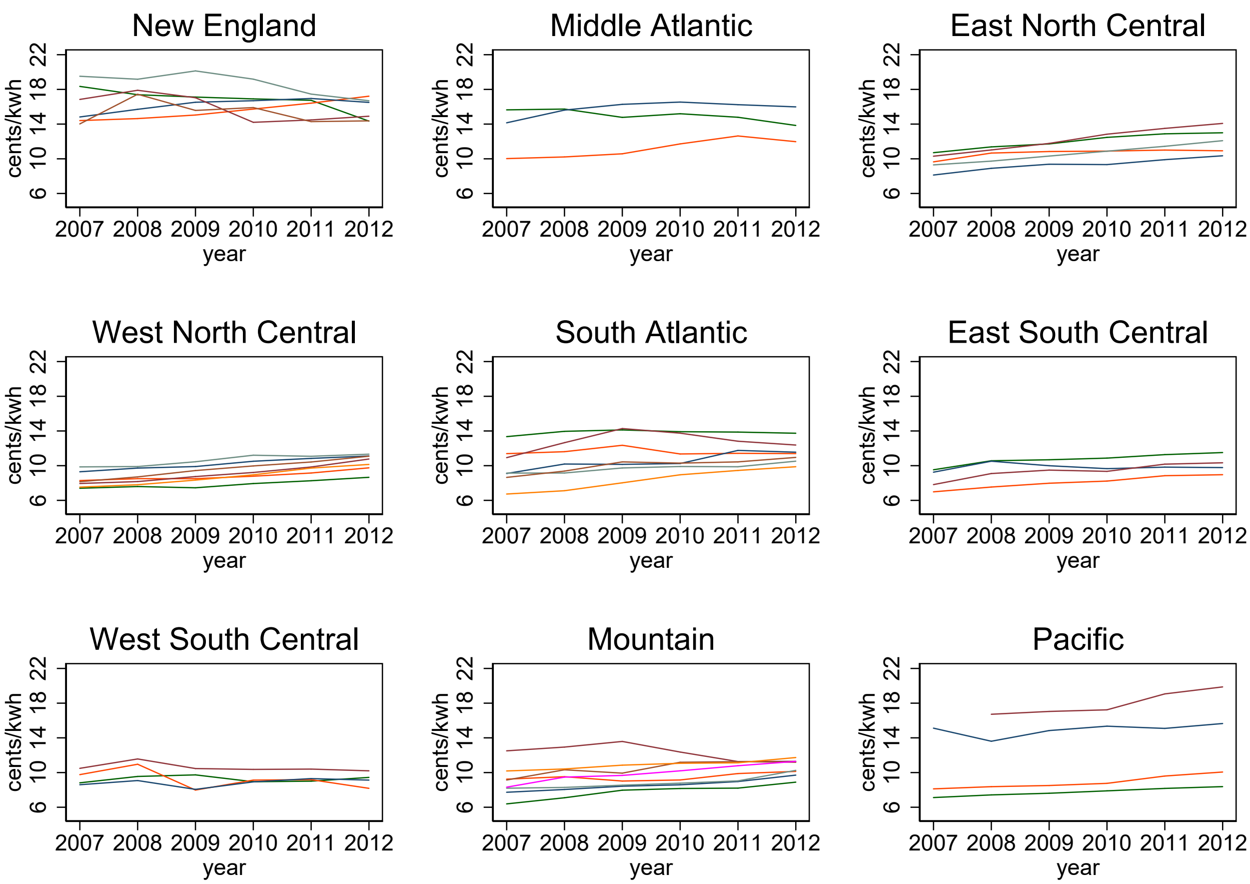

Figure C.1: Average Electricity Prices for Each State in a Census Division

Notes: This Figure is reproduced from Houde and Myers (2021). Each line represents the appliance sales-weighted average electricity price for a particular state within each census division.

Figure C.2 displays density plots of model prices and lifetime energy costs. The first panel plots the distribution of prices paid. Almost all models sold in our sample are less than \$2,000, though there are some much-higher-priced models available. The second panel shows how the distribution of lifetime energy cost of the models sold varies for the 10th, 50th, and 90th percentiles of electricity price. This figure differs slightly from Figure 3 in Houde and Myers (2021) in that this one uses an expected lifetime of 18 years rather than 12 years. We apply a 5% discount rate. The mean lifetime costs range from \$555 for the 10th percentile of energy price to \$1,000 for the 90th percentile. The third panel shows the distribution of the ratio of lifetime cost to purchase price; the means range 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.

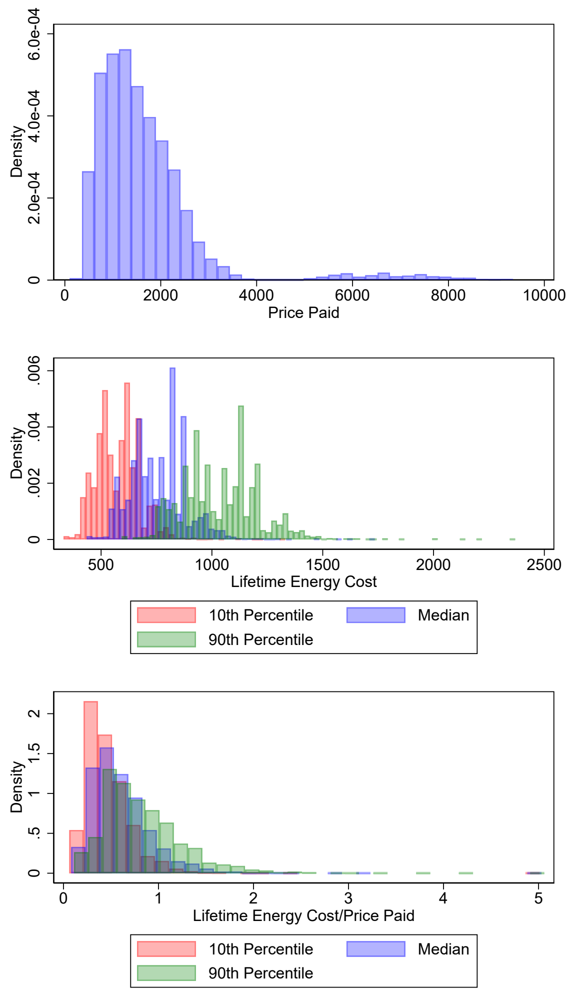

Figure C.2: Distributions of Prices and Lifetime Energy Costs

Notes: The lifetime energy costs are calculated using an expected lifetime of 18 years and a discount rate of 5% and the manufacturer's reported annual energy consumption for each model at the 10th, 50th, and 90th percentile of electricity prices.

D Implementation of the FKRB Estimator

We implement the FKRB estimator with a two-step approach, as suggested by Fox et al. (2011). 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 for which 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 we estimate takes the following form for each alternative j chosen by consumer i of type k. The subscripts r and t correspond to the zip code and time of purchase, respectively.

$$U_{ijkrt} = \eta_k(P_{jrt} + m_k P^e_{rt} \cdot E_{jrt}) + \tau ES_{jt} + \phi Rebate_{jrt} + \gamma_j + Demo_i \times Att_{jt} + \epsilon_{ijkrt}. \quad \text{(D.1)}$$

The parameters we include in ξ are τ, ϕ, and 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 the joint distribution of η and m follows a multivariate normal distribution with unknown mean μ and covariance matrix q. The first step estimates μ, q, and ξ. We used simulated maximum likelihood with 50 draws to approximate the distribution of random parameters.

In the second step, we implement the FKRB estimator with a discrete grid over the support of η and m. We use the estimates of the first step to determine the range of the support. We experimented with grids of different sizes and noticed the span of the grid could lead to computation issues. In particular, if the span is too large, the constrained least-square algorithm (implemented in Matlab) fails to converge. Moreover, the number of grid points (i.e., the density of the grid) increases computational time, but it does create convergence issues. Our choice of the grid was thus made to ensure convergence and computational efficiency. The estimates we report in the main text (Table 2) use 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 a robustness check, we experimented with grids that spanned a smaller region of the support. In Table 2, 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 refer to as Model IV in Table 2:

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

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 2):

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

The estimation was carried out via the subsampling bootstrap procedure. We sliced the sample data into 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 the subsample drawn from the transaction data only included transactions for a specific income group.

D.1 Joint Distribution of η and m for FKRB Estimation

Figure D.1 shows the joint pdf of η and m with the FKRB estimator. The mean of each variable is depicted by a solid black line on the grid. The size and the color of the markers represent the pdf weight for each point. The uncertainty of each estimated weight is depicted by a circle or a star. A star corresponds to an estimate with a t-statistic of 1.1 or larger. We favor this approach to show uncertainty in the estimates over confidence intervals simply to make the figure more readable. Given we use 16 slices for the subsampling bootstrap, the value 1.1 corresponds to the critical value of a one-tail t-test with 10 degrees of freedom and a significance level of approximately 15%. The joint pdf shows large values of m are correlated with lower magnitude coefficients on price.

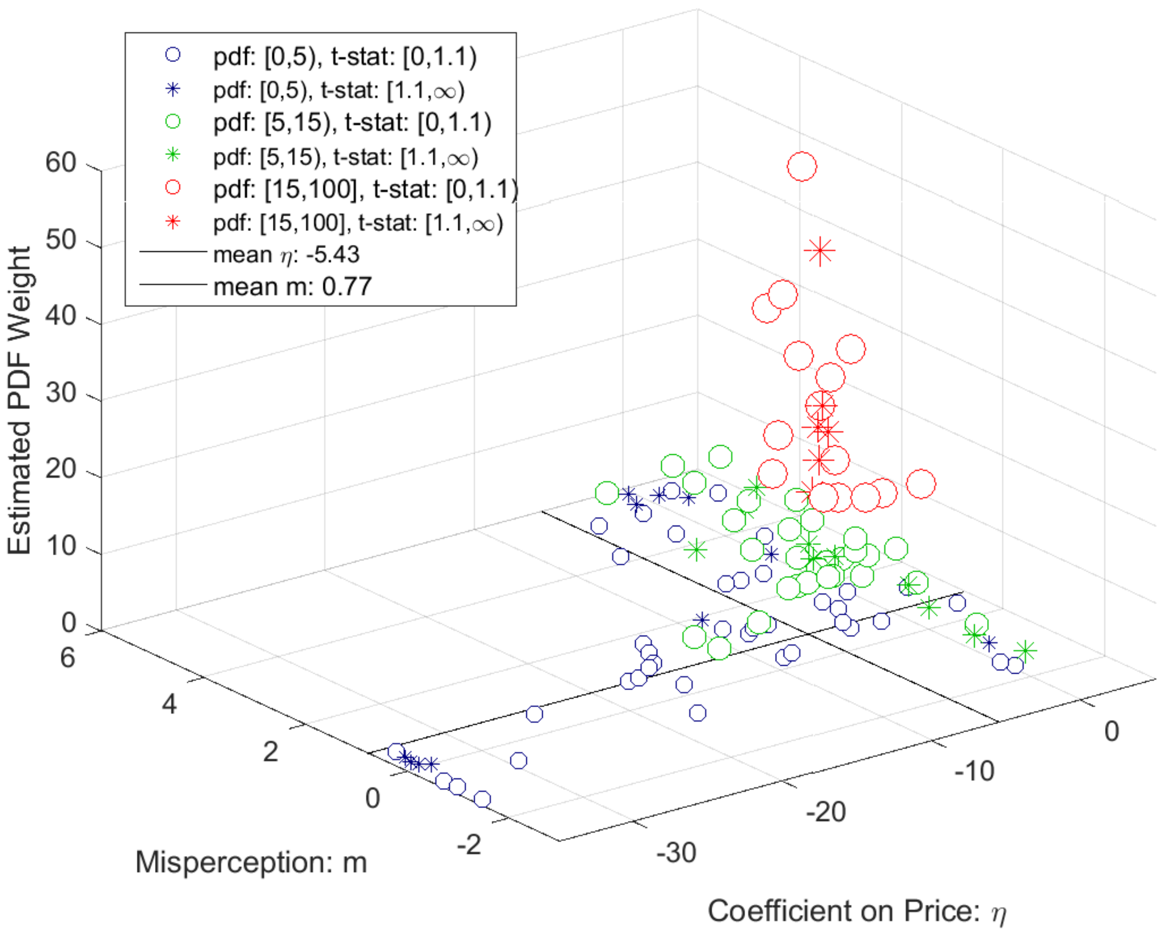

Figure D.1: Joint Probability Distribution for η and m

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.

D.2 Robustness Tests for FKRB Estimation

Table D.1 presents the estimated share of each consumer type for different assumptions about the discount rate and refrigerator's expected lifetime.

Table D.1: Estimated Share of Each Consumer Type (%):

Discount Rate	Expected Lifetime	Severe Misperceptions	Undervalue Energy Costs	Little or No Misperceptions	Overvalue Energy Costs
2%	4 years	12.63	14.21	15.67	66.06
	8 years	12.63	23.14	38.00	29.44
	10 years	12.63	24.50	39.38	26.88
	12 years	12.63	30.47	45.86	18.88
	18 years	12.63	40.19	42.33	16.67
	24 years	12.63	66.44	16.50	11.40
5%	4 years	12.63	11.22	18.42	67.25
	8 years	12.63	23.14	26.58	47.44
	10 years	12.63	24.50	36.88	29.44
	12 years	12.63	24.50	39.38	26.88
	18 years	12.63	37.53	48.42	17.00
	24 years	12.63	40.19	42.33	16.67
12%	4 years	12.63	11.22	18.42	67.25
	8 years	12.63	14.21	19.69	59.00
	10 years	12.63	14.21	25.79	52.31
	12 years	12.63	23.14	21.82	52.31
	18 years	12.63	23.14	26.58	47.44
	24 years	12.63	24.50	36.88	29.44
18%	4 years	12.63	11.22	9.80	74.94
	8 years	12.63	14.21	15.67	66.06
	10 years	12.63	14.21	15.67	66.06
	12 years	12.63	14.21	19.69	59.00
	18 years	12.63	14.21	19.69	59.00
	24 years	12.63	14.21	19.69	59.00

Notes: Each row reports the distribution of consumers across the four types assuming a different value of the discount rate and a refrigerator's expected lifetime. These percentages were computed post-estimation using the estimated distribution of the main model. For each combination of discount rate and expected lifetime, we can re-scale the distribution.

The Figure D.2 shows the estimated joint distribution for η and m and the marginal distribution of m for various specifications of the FKRB estimator. It shows the overall pattern of the joint and marginal distributions are robust by the way the grid is specified and how quality is controlled for.

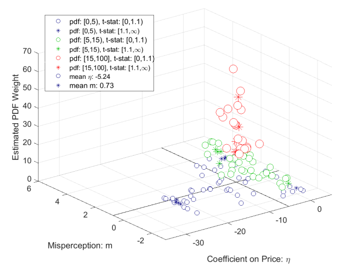

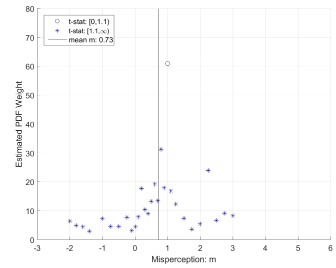

(a) Model II Smaller Span; $f(\eta, m)$ (b) Model II Smaller Span; $f(m)$

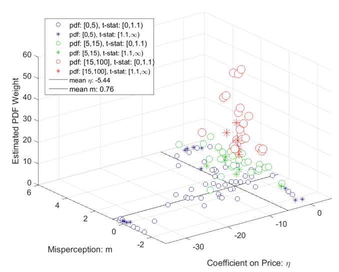

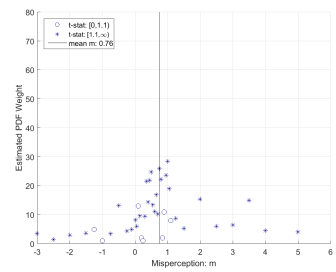

(c) Model III Dense Grid; $f(\eta, m)$ (d) Model III Dense Grid; $f(m)$

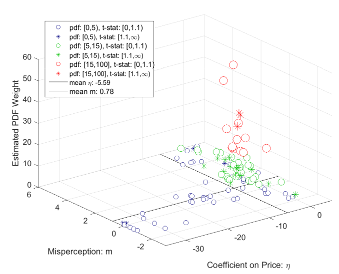

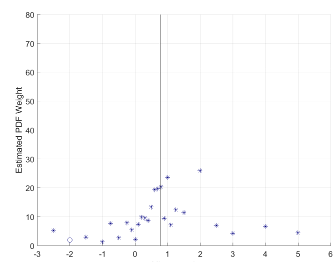

(e) Model IV τ Heterogeneous; $f(\eta, m)$ (f) Model IV τ Heterogeneous; $f(m)$

Figure D.2: Robustness Tests: Joint Probability and Marginal Probability of m as a Function of the Grid

Notes: Each estimated pdf weight is 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 Figure D.3 shows the estimated joint distribution for η and m for different income groups.

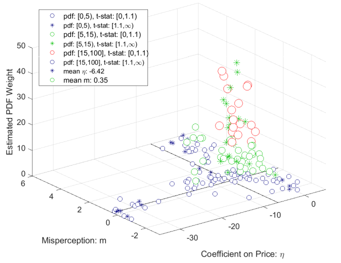

(a) Income: <$30k

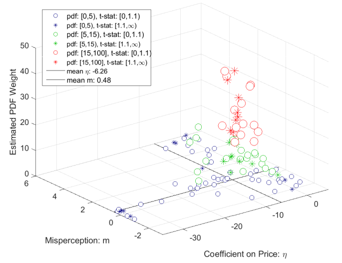

(b) Income: [$30k, $50k)

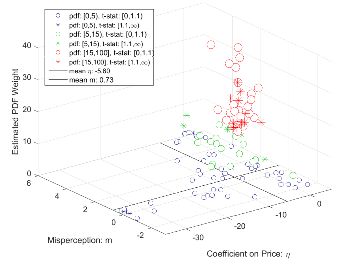

(c) Income: [$50k, $75k)

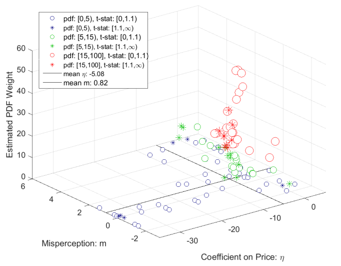

(d) Income: [$75k, $100k)

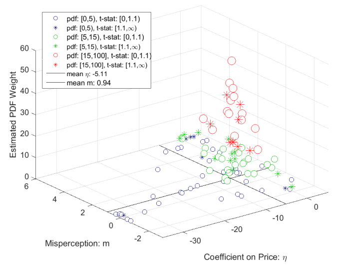

(e) Income: [$100k, $150k)

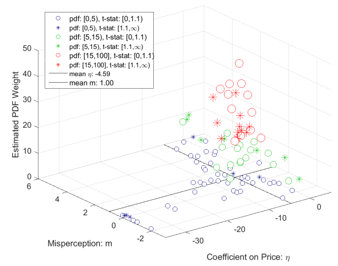

(f) Income: ≥$150k

Figure D.3: Joint Probability Distribution for η and m by Income Group

Notes: : Each estimated pdf weight is 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.

E Exogeneity of Price and Energy Costs

Figure E.1, reproduced from Houde and Myers (2021), displays an example of the variation in weekly price over the study period for the top nine sales-ranked models from a particular refrigerator brand. The red line shows the median weekly price change across all zip codes, and the 25th and 75th percentiles are 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, a large variation in price remains, suggesting the model-specific variation we observe is independent of demand shocks for particular brands.

Table E.1, also reproduced from Houde and Myers (2021), shows the change in price with respect to the mean after controlling for: 1) brand-by-week fixed effects, 2) ENERGY STAR-by-week fixed effects, and 3) other energy-related attribute-by-week fixed effects (i.e., size-by-week and freezer location-by-week fixed effects). These various controls remove little variation compared to the variation observed in normalized prices. This suggests the pricing algorithm provides credible exogenous variation in retail prices, which appears to be independent of demand shocks for particular refrigerator features.

Table E.1: Idiosyncratic Variation in Retail Prices

	(1)	(2)	(3)	(4)	(5)	(6)
Δ Price w.r.t. Mean Price (%)	9.18	8.04	7.88	7.39	7.11	6.91
R^2	0.965	0.973	0.974	0.977	0.979	0.980
Product FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	No	Yes	No	No	No	No
Brand × Week FE	No	No	Yes	Yes	Yes	Yes
EStar × Week FE	No	No	No	Yes	Yes	Yes
Attribute × Week FE	No	No	No	Yes	Yes	Yes
County × Product FE	No	No	No	No	Yes	Yes
County × Brand × Week FE	No	No	No	No	No	Yes

Notes: This table is reproduced from Houde and Myers (2021). The dependent variable is the log of the retail prices. Each column reports the standard deviation of the residuals of a regression of the log of retail prices on various fixed effects (Δ Price with respect to the Mean). The residuals correspond to the percentage variation in retail prices relative to the mean price of each refrigerator model. We consider the attributes other than ENERGY STAR (EStar) are: a dummy variable that identifies small versus large full-size refrigerators and a dummy variable that distinguishes top-freezer versus other types of refrigerators.

Table E.2 presents the results for a simple conditional logit estimated via maximum likelihood. We first present a very parsimonious model in which 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 Table E.3, 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 the large model-specific price variation we observe from the retailer's pricing model is unlikely to be driven by demand-side shocks.

One worry could be areas with high energy prices also have residents with preferences for energy efficiency because policies to promote renewable-energy generation and utility efficiency programs can raise rates. Specifications IV to VII in Table E.2 show the results for the conditional logit estimates where we interact various attributes correlated with 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 refrigerator-freezer models). If regions with high energy prices also have preferences for energy-related attributes, the addition of these controls will affect the coefficient on energy costs. Again, these controls change the coefficients on price and energy costs little. This suggests it is unlikely a correlation between preferences for energy efficiency and local energy prices is biasing our results.

Table E.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) are clustered at the zip code level. The average m is computed by assuming a discount rate of 5% and a refrigerator lifetime of 18 years.

Table E.3: Conditional Logit Results Continued

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

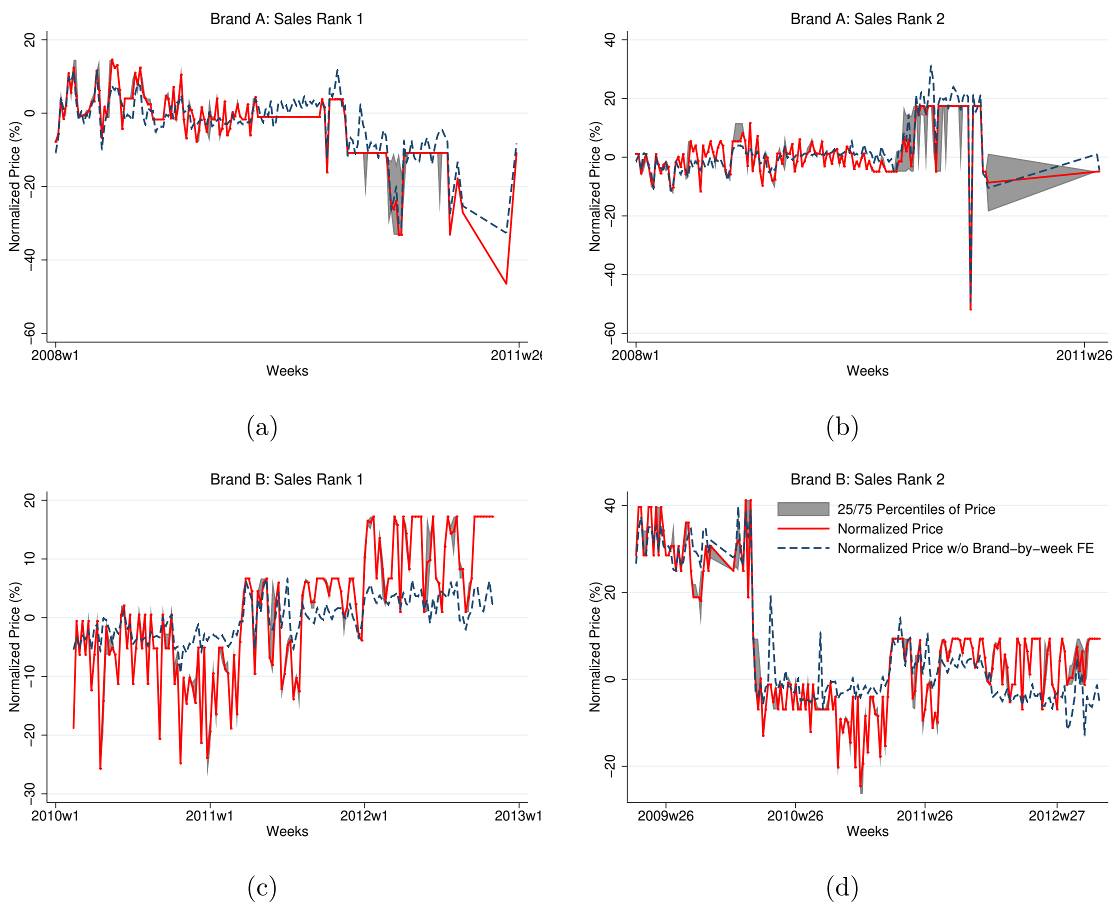

Figure E.1: Price Variation Due to Retailer's National Pricing Algorithm

Notes: This figure is reproduced from Houde and Myers (2021). Each panel corresponds to a particular model offered by a particular brand. Models with sales rank equal to one correspond to the most popular model offered by a given brand. Each panel displays the week-to-week variation in retail price relative to the mean price for a particular model. The plain red line corresponds to the median price across zip codes. The grey band depicts the 25th and 75th percentiles. The dashed blue line is the median price without brand-by-week fixed effects.

F 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, 2019), lack of salience (Chetty, Looney, and Kroft, 2009), present bias (Laibson, 1997), or bias toward concentration (Koszegi and Szeidl, 2013).[40] Additionally, consumers 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 undervaluation (or in some cases, overvaluation) of energy costs (Allcott, 2016).

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

F.1 Search Frictions

In our discrete choice framework, a negative value for m implies 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}, is 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 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 energy costs were a desirable attribute.

In practice, search frictions could be one underlying behavioral mechanism 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

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

appears to be a mistake by selecting a dominated option. It is also possible consumers either 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 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 for which 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, 2022). This often results in a product line with several different models, some of which only differ with respect to their energy consumption and are otherwise identical from the consumer's point of view.

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 paired refrigerators were identical along observable dimensions of quality, but not energy use. In particular, we focused on product lines in which manufacturers' models had the same refrigerator volume, freezer volume, width, height, depth, overall design, color, and technology options, but they consume different amounts 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 retailers having a national pricing strategy that includes offering large and frequent promotions. Therefore, each model is subject to large price variations, and often, the price of the most-efficient model within a pair is lower than the price of the less-efficient model. When this occurs, consumers clearly 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 and to estimate search frictions to energy cost and price. To build such an 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. We then identify cases where both models within a pair were offered at the same location during a dominated price event. One challenge to this is we cannot observe inventory, so we must impute product availability from sales.

With transaction-level data, we can observe the exact location and date of each purchase, which allows us to address this issue as follows. We infer two models of a matched pair were available at the same location if we observe at least one sale for each model during a time period of specific length. We consider different time lengths starting from the most conservative, only considering instances 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, we can attribute the reason for the dominant product not being purchased to it's being temporarily out of stock, not consumer search frictions. As the results will show below, however, the share of consumers that chose a dominant option is quite stable irrespective of the fact that we look at same-day sales or use a longer time interval between sales. The effect of stock-outs in identifying severe misperceptions is thus limited.

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 consumers who chose the dominated option.

Table F.1 reports summary statistics for the matched pairs compared to the overall sample of models we observed. One takeaway from the report 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 takeaway is the average price difference within a pair is \$22, a small and positive amount; this is consistent with the fact that more-efficient models tend to be more expensive. However, there are large variations spanning negative and positive values throughout the whole sample. The 10th percentile is negative \$150, and the 90th percentile is \$155.

Table F.2 shows the share of consumers who chose the dominated option and the price difference and 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 the several million transactions we started with. Arguably, this estimator strives for internal validity, and we err on the conservative side to truly identify inattentive consumers. We find 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 38% to 48%.

These results show 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 both are the result of imperfect attention allocation. Our estimator is thus a useful diagnostic to show 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, 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 out on a very small subset of the overall sample. The external validity of the results, therefore, is not guaranteed.

Table F.1: 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/year, electricity costs per year, and price are averaged within pairs. 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.

Table F.2: 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 who chose the dominated option. The second row is the mean size of the price difference during the events where the more 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 the sales between two models of the same pair were at most two days apart.

F.2 Biased Beliefs

Computing the energy cost of a refrigerator requires two key pieces of information: electricity consumption estimates and electricity price. Beliefs about these pieces of information should, therefore, play an important role in 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 Spring 2017 on a sample of U.S. households. The survey procedure was reviewed by the Institutional Review Board committee (reference number 1073843-1) at the University of Maryland and was exempt from an IRB review in accordance with federal regulations.

Survey participants were recruited by the survey company SSI. We did not contact the participants and did not collect information about the identity of the participants. The survey company was entirely responsible of selecting the sample, contacting the participants, and administering the online survey. SSI was mandated to recruit participants 18 years or older who are members of households living in 24 U.S. states, we pre-selected. SSI household panels aim to be representative of the population of the United States. However, because we restrict our sampling to a subset of states, our sample is, therefore, not fully representative of the whole U.S. population. Table F.3 reports the distribution of the main sample. Note that although we targeted 24 states, some participants from other states got access to our survey.

The survey was administrated with an online survey tool (Qualtrics). We shared the link with SSI and SSI, in turn, shared the link with its panelists, as well as gathered informed consent and distributed compensation to survey participants. Our recruitment target was 2,000 households; 1,043 households answered the survey, at least in part; and 1,016 of those provided answers to the belief questions we use for this paper.

In this paper, we report the results for two questions 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 they pay. Respondents were asked to report their best estimates for both questions, and no numerical or other anchors were provided.

Table F.3: States Sampled

State State	Number of Participants	Frequency (%)	Cumulative Density (%)
California	2	0.19	0.29
Colorado	1	0.1	0.38
Connecticut	2	0.19	0.58
Delaware	13	1.25	1.83
Florida	3	0.29	2.11
Georgia	1	0.1	2.21
Hawaii	19	1.83	4.03
Illinois	159	15.27	19.31
Indiana	78	7.49	26.8
Kentucky	56	5.38	32.18
Louisiana	68	6.53	38.71
Maryland	76	7.3	46.01
Massachusetts	107	10.28	56.29
Minnesota	1	0.1	56.39
Nevada	1	0.1	56.48
New Hampshire	24	2.31	58.79
New Jersey	102	9.8	68.59
New York	4	0.38	68.97
North Carolina	1	0.1	69.07
Ohio	1	0.1	69.16
Oklahoma	55	5.28	74.45
Oregon	62	5.96	80.4
Pennsylvania	2	0.19	80.6
Tennessee	67	6.44	87.13
Utah	40	3.84	90.97
Vermont	10	0.96	91.93
Virginia	2	0.19	92.12
Washington	80	7.68	99.81
Wyoming	1	0.1	99.9
Total	1,016	100	100

Figure F.1 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 a county's average electricity prices using information about respondents' zip codes. For the 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 overvaluation. To construct the histograms, we censored all values greater than 5, which explains the large mass at this value. This shows 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 a large share of consumers undervalue both refrigerators' level of energy use and their prices. Interestingly, Allcott and Knittel (2019) find similar patterns in beliefs in the U.S. car market.

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

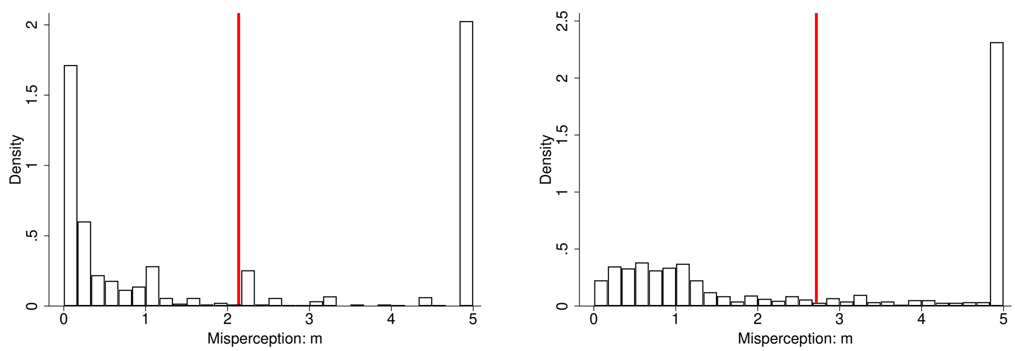

(a) Beliefs: Refrigerator Usage

(b) Beliefs: Electricity Price

Figure F.1: Beliefs Valuation Ratios: Refrigerator Usage and Electricity Price
Notes: The red bar identifies the mean valuation ratio. All the values above 5 were rounded to 5 for exposition purposes only. The correct energy usage of refrigerators 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 in which a survey respondent lived at the time of answering the survey.

G Additional Details: Policy Simulations

G.1 Implementation Details

For all optimal policy scenarios, we assume the only externality comes from carbon emissions with marginal damage of \$50/ton of CO_2. We consider heterogeneous local electricity prices and emissions factors, which vary with electricity market areas. Data for the CO_2 emission factors comes from the EIA scenarios used to simulate the impact of the Clean Power Plan on electricity prices (Energy Information Administration (EIA), 2015). We infer the emission factors for 20 different regions of the United States corresponding to different electricity markets. For electricity prices, we use the county averages constructed for our estimation.

To quantify the degree of misperception in our demand model, it is important we take a clear stand on experienced utility. For our main results, we use our preferred approach to defining misperceptions as described in section 4.3. However, we provide estimates using several alternatives to assess the sensitivity of the results to these assumptions.

We account for parametric uncertainty associated with each estimator by sampling the demand estimates obtained from the estimation procedure. We sample one set of demand estimates for each income group, simulate the policy, store the results, and repeat the process. We report the mean and standard errors across the different slices of the data to implement the subsampling bootstrap.

G.2 Robustness

Table G.1 reproduces the results from our preferred Baseline Scenario in Table 5 with more details. Tables G.2–G.4 display results from two scenarios described in detail in the main text. In Table G.2 we use the RCL approach to estimate the distribution of consumer misperceptions. Tables G.3 and G.4 investigate the impact of a so-called warm glow for the FKRB model and the model with the parametric mixed logit estimator respectively. For these policy scenarios, all values of m greater than 1 are not defined as misperceptions; instead, they are assumed to be due to warm glow. Table G.4 presents the results, which are qualitatively similar to the ones obtained with the FKRB estimator (Table G.3). However, the optimal tax is larger in magnitude.

Table G.5 presents the results from the optimal policy simulation assuming misperceptions are defined using a discount rate of 5%. Under this approach, we are less conservative in defining misperceptions. Compared to our preferred approach, Table G.1, the tax is still negative, but it is almost three times greater in magnitude. Whereas the difference in the optimal standards is zero or minimal.

Finally, we conduct a sensitivity test by revisiting the results of Table G.1 and rescaling the m distribution to make it consistent with a 12-year expected lifetime. Relative to our main simulations where we use an 18-year expected lifetime, a shorter valuation will reduce undervaluation bias, but it increases overvaluation (i.e., the mass of the distribution with values of m greater than one). Table G.6 presents the results and shows an overvaluation bias has an even greater impact on driving the optimal tax to a negative value. Note in this simulation, we do not assume that a so-called warm glow is at play. Values of m greater than one are thus interpreted as misperceptions.

Table G.1: Detailed Results for Optimal Behavioral Policies: Preferred Scenario

	Misperceptions with FKRB Estimator and Heterogeneous r			
	Policy	Δ SW $/capita	Δ CS $/capita	Δ Ext $/capita
Pigou Tax ($/ton of CO_2)				
With $F(m_k)$	-25.7	0.2	56.8	0.4
	(4.4)	(0.1)	(9.8)	(0.1)
With $E[m_k]$	75.7	1.4	-157.9	-1.9
	(9.2)	(0.2)	(18.9)	(0.2)
Uniform Standard (kWh/year)				
With $F(m_k)$	335.0	114.2	78.0	-36.1
	(0.0)	(0.7)	(0.6)	(0.1)
With $E[m_k]$	335.0	101.0	66.2	-34.8
	(0.0)	(0.8)	(0.6)	(0.2)
Minimum Standard (kWh/year)				
With $F(m_k)$	335.0	114.2	78.0	-36.1
	(0.0)	(0.7)	(0.6)	(0.1)
With $E[m_k]$	335.0	95.6	62.2	-33.4
	(0.0)	(0.9)	(0.7)	(0.2)
Attribute-Based Minimum Standard (% of Existing Standard)				
With $F(m_k)$	55.6%	91.3	57.0	-34.3
	(0.2%)	(0.7)	(0.5)	(0.3)
With $E[m_k]$	56.1%	78.3	46.7	-31.7
	(0.1%)	(0.8)	(0.7)	(0.2)

Notes: The table reports the optimal policies evaluated using the full distribution of misperceptions [with $F(m_k)$] or only the average misperception [with $E[m]$]. To define misperceptions, a range of discount rates that spans 2% to 12% is used. 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 (in parentheses) of the optimal taxes, standards, and welfare metrics are computed using 500 random draws of the estimated distributions. The non-parametric distribution of the FKRB estimator is used.

Table G.2: Sensitivity: Optimal Behavioral Policies with Parametric Mixed Logit Estimator and Heterogeneous r

	Misperceptions with Parametric Mixed Logit Estimator and Heterogeneous r			
	Policy	Δ SW $/capita	Δ CS $/capita	Δ Ext $/capita
Pigou Tax ($/ton of CO_2)				
With $F(m_k)$	111.3	1.8	-233.5	-1.8
	(6.3)	(0.1)	(13.0)	(0.1)
With $E[m_k]$	146.3	3.5	-301.3	-3.1
	(11.2)	(0.3)	(22.6)	(0.1)
Uniform Standard (kWh/year)				
With $F(m_k)$	335.0	98.1	63.8	-34.2
	(0.0)	(0.6)	(0.5)	(0.1)
With $E[m_k]$	335.0	94.7	60.9	-33.7
	(0.0)	(0.9)	(0.7)	(0.2)
Minimum Standard (kWh/year)				
With $F(m_k)$	335.0	98.1	63.8	-34.2
	(0.0)	(0.6)	(0.5)	(0.1)
With $E[m_k]$	335.0	94.7	60.9	-33.7
	(0.0)	(0.9)	(0.7)	(0.2)
Attribute-Based Minimum Standard (% of Existing Standard)				
With $F(m_k)$	57.9%	76.0	44.8	-31.2
	(0.1%)	(0.8)	(0.7)	(0.2)
With $E[m_k]$	57.3%	73.6	43.0	-30.6
	(0.2%)	(1.0)	(0.8)	(0.3)

Notes: The table reports the optimal policies evaluated using the full distribution of misperceptions [with $F(m_k)$] or only the average misperception [with $E[m]$]. To define misperceptions, a range of discount rates that spans 2% to 12% is used. 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 (in parentheses) of the optimal taxes, standards, and welfare metrics are computed using 500 random draws of the estimated distributions. The parametric joint normal distribution of the mixed logit estimator is used.

Table G.3: Sensitivity: Optimal Behavioral Policies with FKRB Estimator and Warm Glow

	Misperceptions with FKRB Estimator and Heterogeneous r			
	Policy	Δ SW $/capita	Δ CS $/capita	Δ Ext $/capita
Pigou Tax ($/ton of CO_2)				
With $F(m_k)$	92.5	1.2	-200.1	-1.4
	(3.6)	(0.1)	(7.6)	(0.1)
With $E[m_k]$	205.0	5.5	-425.0	-3.9
	(12.7)	(0.4)	(25.8)	(0.2)
Uniform Standard (kWh/year)				
With $F(m_k)$	335.0	109.4	72.6	-36.8
	(0.0)	(0.7)	(0.6)	(0.1)
With $E[m_k]$	335.0	101.0	66.2	-34.8
	(0.0)	(0.8)	(0.6)	(0.2)
Minimum Standard (kWh/year)				
With $F(m_k)$	335.0	109.4	72.6	-36.8
	(0.0)	(0.7)	(0.6)	(0.1)
With $E[m_k]$	335.0	100.0	64.6	-35.4
	(0.0)	(0.8)	(0.6)	(0.2)
Attribute-Based Minimum Standard (% of Existing Standard)				
With $F(m_k)$	56.0%	89.1	54.9	-34.2
	(0.0%)	(0.7)	(0.6)	(0.1)
With $E[m_k]$	56.1%	80.3	47.7	-32.6
	(0.1%)	(0.8)	(0.7)	(0.2)

Notes: The table reports the optimal policies evaluated using the full distribution of misperceptions [with $F(m_k)$] or only the average misperception (with $E[m]$). To define misperceptions, a range of discount rates that spans 2% to 12% is used. Moreover, all values of m above 1 are assumed to be due to a so-called warm glow. Therefore, there is no overvaluation bias in these simulations. As in the other simulations, 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 (in parentheses) of the optimal taxes, standards, and welfare metrics are computed using 500 random draws of the estimated distributions. The non-parametric distribution of the FKRB estimator is used.

Table G.4: Sensitivity: Optimal Behavioral Policies with Parametric Mixed Logit Estimator and Warm Glow

	Misperceptions with Parametric Mixed Logit Estimator and Warm Glow			
	Policy	Δ SW $/capita	Δ CS $/capita	Δ Ext $/capita
Pigou Tax ($/ton of CO_2)				
With $F(m_k)$	122.5	2.0	-257.0	-2.0
	7.5	0.1	15.5	0.1
With $E[m_k]$	163.8	4.1	-336.9	-3.3
	13.2	0.4	26.6	0.2
Uniform Standard (kWh/year)				
With $F(m_k)$	335.0	98.1	63.8	-34.3
	0.0	0.6	0.5	0.1
With $E[m_k]$	335.0	95.2	61.2	-34.0
	0.0	0.8	0.7	0.2
Minimum Standard (kWh/year)				
With $F(m_k)$	335.0	98.1	63.8	-34.3
	0.0	0.6	0.5	0.1
With $E[m_k]$	335.0	95.2	61.2	-34.0
	0.0	0.8	0.7	0.2
Attribute-Based Minimum Standard (% of Existing Standard)				
With $F(m_k)$	57.9%	76.2	44.9	-31.3
	0.1%	0.7	0.6	0.2
With $E[m_k]$	57.3%	74.1	43.4	-30.7
	0.2%	0.9	0.6	0.3

Notes: The table reports the optimal policies evaluated using the full distribution of misperceptions [with $F(m_k)$] or only the average misperception (with $E[m]$). The parametric joint normal distribution of the mixed logit estimator is used. To define misperceptions, a range of discount rates that spans 2% to 12% is used. Moreover, all values of m above 1 are assumed to be due to a so-called warm glow. Therefore, there is no overvaluation bias in those simulations. As in the other simulations, 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 (in parentheses) of the optimal taxes, standards, and welfare metrics are computed using 500 random draws of the estimated distributions.

Table G.5: Sensitivity: Optimal Behavioral Policies with FKRB Estimator and $r = 5\%$

	Misperceptions with FKRB Estimator and $r = 5\%$			
	Policy	Δ SW $/capita	Δ CS $/capita	Δ Ext $/capita
Pigou Tax ($/ton of CO_2)				
With $F(m_k)$	-76.9	1.6	172.8	1.3
	4.4	0.2	10.1	0.1
With $E[m_k]$	55.4	0.8	-116.2	-1.4
	7.6	0.2	15.7	0.2
Uniform Standard (kWh/year)				
With $F(m_k)$	335.0	122.0	85.4	-36.6
	0.0	0.8	0.8	0.1
With $E[m_k]$	335.0	101.0	66.2	-34.8
	0.0	0.8	0.6	0.2
Minimum Standard (kWh/year)				
With $F(m_k)$	335.0	122.0	85.4	-36.6
	0.0	0.8	0.8	0.1
With $E[m_k]$	335.0	97.0	63.7	-33.3
	0.0	0.8	0.7	0.2
Attribute-Based Minimum Standard (% of Existing Standard)				
With $F(m_k)$	54.9%	94.7	59.3	-35.4
	0.3%	0.8	0.6	0.3
With $E[m_k]$	56.1%	78.9	47.3	-31.6
	0.1%	0.8	0.7	0.2

Notes: The table reports the optimal policies evaluated using the full distribution of misperceptions [with $F(m_k)$] or only the average misperception (with $E[m]$). Misperceptions are determined using a discount rate of $r = 5\%$. 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 (in parentheses) of the optimal taxes, standards, and welfare metrics are computed using 500 random draws of the estimated distributions. The non-parametric distribution of the FKRB estimator is used.

Table G.6: Sensitivity: Optimal Behavioral Policies 12 Year Lifetime

	Misperceptions with FKRB Estimator, $r = 2\% - 12\%$, and 12 Years Lifetime			
	Policy	Δ SW $/capita	Δ CS $/capita	Δ Ext $/capita
Pigou Tax ($/ton of CO_2)				
With $F(m_k)$	-87.5	2.9	196.8	1.9
	4.0	0.3	9.3	0.1
With $E[m_k]$	-1.3	0.2	3.1	0.1
	6.7	0.0	14.2	0.2
Uniform Standard (kWh/year)				
With $F(m_k)$	335.0	120.8	85.5	-35.2
	0.0	0.9	0.9	0.1
With $E[m_k]$	335.0	101.0	66.2	-34.8
	0.0	0.8	0.6	0.2
Minimum Standard (kWh/year)				
With $F(m_k)$	335.0	120.8	85.5	-35.2
	0.0	0.9	0.9	0.1
With $E[m_k]$	335.0	95.5	63.8	-31.6
	0.0	1.0	0.9	0.2
Attribute-Based Minimum Standard (% of Existing Standard)				
With $F(m_k)$	55.3%	93.3	59.1	-34.2
	0.3%	0.9	0.7	0.3
With $E[m_k]$	56.2%	77.1	46.2	-30.9
	0.2%	0.9	0.8	0.2

Notes: The table reports the optimal policies evaluated using the full distribution of misperceptions (with $F(m_k)$) or only the average misperception (with $E[m]$). Misperceptions are determined for a range of discount rate from $r = 2\%$ to $r = 12\%$ and a refrigerator's expected lifetime of 12 years. 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 (in parentheses) of the optimal taxes, standards, and welfare metrics are computed using 500 random draws of the estimated distributions. The non-parametric distribution of the FKRB estimator is used.

H Sensitivity With Respect to Cost Function

Table H.1 displays the results from alternative scenarios, where we vary the convexity of the supply-side cost function defined by Equation 19. We reproduce the results from our preferred approach in the "Baseline Scenario". For our preferred specification, we rely on the cost function estimated in Houde (2018b) ($\phi = 191$), which uses the same data as in the current analysis. Here, we display the results from our optimal policy simulation if we instead allow for more convexity in the cost function. In the "Mid-Case Scenario", we set $\phi = 300$ and in the "Upper-Bound Scenario", we set $\phi = 400$. As shown below, even if we allow for a more convex cost function ($\phi = 300$), the welfare effects of optimal standards still dominate the optimal tax. It takes an almost doubling of the baseline ϕ for there to be little welfare improvement from the attribute-based standard, thus making it comparable to the tax. However, even at this extreme, the uniform and minimum efficiency standards still dominate the tax.

Table H.1: Optimal Behavioral Policies with Different Cost Functions

	Misperceptions with FKRB Estimator and Heterogeneous r			
	Policy	Δ SW $/capita	Δ CS $/capita	Δ Ext $/capita
Pigou Tax ($/ton of CO_2)				
With $F(m_k)$	-25.7	0.2	56.8	0.4
	(4.4)	(0.1)	(9.8)	(0.1)
Baseline Scenario: $\phi = 191$				
Uniform Standard (kWh/year)				
With $F(m_k)$	335.0	114.2	78.0	-36.1
	(0.0)	(0.7)	(0.6)	(0.1)
Minimum Standard (kWh/year)				
With $F(m_k)$	335.0	114.2	78.0	-36.1
	(0.0)	(0.7)	(0.6)	(0.1)
Attribute-Based Minimum Standard (% of Existing Standard)				
With $F(m_k)$	55.6%	91.3	57.0	-34.3
	(0.2%)	(0.7)	(0.5)	(0.3)
Mid-Case Scenario: $\phi = 300$				
Uniform Standard (kWh/year)				
With $F(m_k)$	420.0	45.7	28.1	-17.6
	(0.0)	(0.3)	(0.2)	(0.0)
Minimum Standard (kWh/year)				
With $F(m_k)$	415.1	43.8	24.4	-19.4
	(0.1)	(0.2)	(0.2)	(0.0)
Attribute-Based Minimum Standard (% of Existing Standard)				
With $F(m_k)$	70.0%	15.9	-0.3	-16.2
	(0.0%)	(0.1)	(0.1)	(0.0)
Upper-Bound Scenario: $\phi = 400$				
Uniform Standard (kWh/year)				
With $F(m_k)$	485.0	32.3	28.7	-3.6
	(0.0)	(0.7)	(0.7)	(0.1)
Minimum Standard (kWh/year)				
With $F(m_k)$	468.1	18.8	8.5	-10.4
	(0.6)	(0.4)	(0.3)	(0.1)
Attribute-Based Minimum Standard (% of Existing Standard)				
With $F(m_k)$	100%	0.1	0.0	-0.1
	(0.2%)	(0.0)	(0.0)	(0.0)

Notes: The table reports the optimal policies evaluated using the full distribution of misperceptions [with $F(m_k)$] for different values of the parameter ϕ, which determines the marginal cost of energy efficiency. To define misperceptions, a range of discount rates that spans 2% to 12% is used. 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 (in parentheses) of the optimal taxes, standards, and welfare metrics are computed using 500 random draws of the estimated distributions. The non-parametric distribution of the FKRB estimator is used.