

The Efficiency Consequences of Heterogeneous Behavioral Responses to Energy Fiscal Policies

Sébastien Houde and Joseph E. Aldy

1616 P St. NW
Washington, DC 20036
202-328-5000 www.rff.org

The Efficiency Consequences of Heterogeneous Behavioral Responses to Energy Fiscal Policies

Sébastien Houde, Joseph E. Aldy*

December 1, 2017

The behavioral responses to taxes and subsidies are often subject to various behavioral biases and transaction costs—what we define as “microfrictions.” We develop a theoretical framework to show how these microfrictions—and their heterogeneity across the population and policy instruments—affect the design of Pigouvian policies. Standard Pigouvian pricing still holds with transaction costs, but requires adjustment with behavioral biases. We use transaction-level data from the US appliance market to estimate the heterogeneous behavioral responses to an array of energy fiscal policies and to quantify microfrictions. We then assess optimal fiscal policies and find that it is rarely optimal to couple a Pigouvian tax on energy with an investment subsidy in this context. We also find that energy labels—intended to increase the salience of energy information—can interact in perverse ways with both taxes and subsidies.

Key Words: energy fiscal policies, behavioral taxation, demand estimation, durables.

JEL Codes: Q4, Q48, Q58, H31.

*Houde: Department of Agricultural and Resource Economics, University of Maryland, ETH Zürich, Center for Economic Research, E2e Project, shoude@umd.edu. Aldy: John F. Kennedy School of Government, Harvard University, Resources for the Future, National Bureau of Economic Research, and Center for Strategic and International Studies, Joseph_Aldy@hks.harvard.edu. We thank Blake Barr and Carlos Paez for excellent research assistance. We would like to thank Erica Myers, Catherine Wolfram, Matthias Rodemeier, Andreas Löschel, Nicolai Kuminoff, and Massimo Filippini, for their comments and suggestions on previous drafts. We would also like to thank participants at NBER workshops on the energy policy trade-offs between economic efficiency and distributional equity, the AERE Summer Conference, and several seminar participants. This project benefited from support received from the Alfred P. Sloan Foundation through the NBER Energy Policy Efficiency and Equity Tradeoffs project. This research is also part of the activities of SCCER CREST, which is financially supported by the Swiss Commission for Technology and Innovation (CTI).

1. Introduction

Policymakers have frequently deployed a variety of tax and subsidy instruments to promote socially beneficial behavioral changes: cigarette excise taxes, subsidies for energy-efficient appliances, and subsidies for health insurance. These policies interact with other fiscal instruments—such as sales taxes, energy taxes, and income taxes—resulting in a complicated mix of overlapping fiscal instruments that ultimately impact behaviors. There are several rationales that justify why policy makers do and should combine fiscal instruments. This includes trading-off equity and efficiency concerns, accounting for heterogeneous externalities, and correcting for pre-existing tax distortions. The recent literature in public finance has also found that heterogeneity in the behavioral responses to fiscal instruments could justify using multiple instruments to achieve a single policy goal (Allcott et al. 2014; Farhi and Gabaix 2015). That is, consumers’ responses to fiscal instruments may vary over preferences and income, but also as a function of behavioral biases and transaction costs associated with a given instrument. These “microfrictions” encompass a range of phenomena, such as a lack of salience and misperception of prices and policies as well as the opportunity cost of the time and effort to respond to a fiscal policy. The presence of such microfrictions have important implications for the design of a welfare-maximizing fiscal policy.

The roles of heterogeneity and microfrictions are particularly important in the energy context, where governments have relied on an array of fiscal instruments to mitigate the negative externalities associated with energy use. For example, numerous state and national governments simultaneously subsidize the investment in energy-efficient technologies and raise energy prices to increase the returns on energy savings, through an energy tax or a cap-and-trade scheme. Each of these policies is associated with particular microfrictions, which in turn induce heterogeneous behavioral responses across policy instruments. The primary goal of this paper is to develop a framework to show how the heterogeneous behavioral responses created by microfrictions associated with energy fiscal policies affect the effectiveness and design of policies.

Our theoretical framework builds upon Allcott et al. (2014) and Farhi and Gabaix (2015), who derive two important results pertaining to the design of corrective policies. First, they show how a Pigouvian tax should be adjusted if consumers are subject to various biases when responding to the tax. In both frameworks, it is always optimal to set a corrective tax higher than the externality cost if a single instrument is used and all consumers under-respond to the tax, and vice versa. Second, heterogeneity in behavioral biases motivates a departure from the Pigouvian targeting principle—combining multiple instruments to address a single (homogeneous) externality may be

optimal. In particular, Allcott et al. (2014) show that an energy tax combined with a product subsidy should always perform better than a single bias-adjusted tax, and the level of the tax should always be below the externality cost when a subsidy is offered. In our framework, both results can be overturned due to three distinct mechanisms.

First, we define the concept of microfriction to generalize the various phenomena that lead to behavioral under- (or over-) responses to fiscal instruments. We distinguish two categories of microfrictions: behavioral biases and transaction costs. We argue that this distinction is important because, although both types of microfrictions can be observationally equivalent, they have different implications for the design of a corrective tax. Only the presence of behavioral biases justifies a departure from the standard prescription that the optimal corrective tax should be set equal to the marginal external cost.

Second, we highlight the role of heterogeneity in microfrictions across consumers and policy instruments in determining whether it is optimal to combine fiscal instruments. We show that it is optimal to combine multiple instruments, a tax with a subsidy for instance, only if microfrictions associated with each instrument lead to heterogeneous behavioral responses that follow very specific and empirically rare patterns.

Third, we show that observed and unobserved heterogeneity in microfrictions across consumers have different implications for policy design. In particular, unobserved heterogeneity, which arises when we cannot perfectly segment the population in estimating microfrictions, has an ambiguous effect on the direction of the adjustment of a Pigouvian tax that accounts for behavioral biases.

We apply our framework to the context of policies promoting the purchase of energy-efficient appliances in the US market. In this context, the question of how we should design overlapping fiscal instruments to account for negative externalities associated with energy use is particularly relevant and timely. Currently, federal, state, and local governments use different types of investment subsidies for energy-efficient products, specifically, appliance purchase rebates, sales taxes, and tax credits, which interact with energy taxes and costs of environmental regulations imposed on power producers. The efficiency consequences of this mix of instruments have received little attention. Moreover, given the absence of a federal policy that effectively puts a price on carbon dioxide emissions, and some states introducing regional climate policies, the role of energy efficiency subsidies to compensate the lack of or to complement carbon pricing needs to be clarified. This paper is a first step in this direction.

The institutional features of the US appliance market combined with our rich micro-data allow us to study an exhaustive set of energy fiscal policies and highlight issues associated with microfrictions. Our empirical strategy extends our previous work on consumer rebates and energy labels (Houde and Aldy 2017; Houde 2017) and consists of a parsimonious structural demand model that captures heterogeneity in the responses to energy fiscal instruments. We focus on the US refrigerator market over 2008-2012, when local, state, and national policymakers implemented a rich array of fiscal instruments intended to promote energy-efficient purchases. For example, the states launched their Cash for Appliances (CFA) rebate programs at various times during 2010 and these programs varied in terms of rebate amount, appliance coverage, and other characteristics. Sales taxes and sales tax holidays targeting energy-efficient products were also subject to large variations across regions and time during this period. Electricity prices also varied widely across states, and in some regions over time as well, which allows us to identify the consumer response to energy operating costs. We estimate discrete choice models for the refrigerator purchasing decision that exploit this temporal and regional variation in the features of various US energy efficiency subsidy policies as well as in energy operating costs (i.e., electricity prices). Our estimation sample draws from millions of transactions from a large US retailer, which have matched household demographic data. This rich microdataset allows us to recover heterogeneous behavioral responses by segmenting the population using observable information (e.g., household income), while controlling for region-specific unobservables, consumer sorting, and time trends. We also account for the role of energy labels that serve as certification and disclosure policies, and show that they are an important source of unobserved heterogeneity in microfrictions.

Our empirical results show that the three mechanisms that can overturn the results of Allcott et al. (2014) and Farhi and Gabaix (2015) are important in the US appliance market. We find substantial heterogeneity across policy instruments and consumers, which follows systematic patterns across income groups. In particular, the response to energy operating costs increases with income. Therefore, higher income households would respond more to taxes and regulation that impact energy prices. In contrast, consumers' responses to rebates decrease with income. For the sales tax, the pattern is less clear; the highest income group tends to respond more to this instrument. But across all income groups, we find that consumers systematically respond less to a dollar change in energy-efficiency subsidies than to a dollar change in the purchase price. These results suggest that microfrictions taking the form of behavioral biases and transaction costs both play an important role in this context. Behavioral biases might have a larger effect, in relative terms, on the purchasing decision of lower-income households, but tangible hassle costs would have a relatively lower

effect on these households. Finally, the Energy Star certification has the unintended consequence of reducing the behavioral response to energy operating costs across all income groups—that is, higher energy prices have no impact on the energy efficiency of an appliance conditional on its receiving an Energy Star rating. This arises because some consumers, which are not observed *ex ante*, may pay more attention to the coarse Energy Star label, but dismiss other information about energy operating costs. Finally, we find that some observed demographic characteristics (e.g., education, age of the head of the household, and family structure) influence the estimated response to energy operating costs, but we find few other systematic and robust effects of demographics, other than income, on the responses to energy efficiency subsidies.

Using the estimated demand model, we investigate the optimal design of fiscal policies to address the carbon dioxide emission externality. We first find that the bias-adjusted Pigouvian tax could be large. In one scenario, it corresponds to about a five-fold increase of the marginal external cost. Whether this adjustment is desirable from a welfare standpoint, however, crucially depends on how we interpret the effect of the Energy Star certification on consumers' decisions. If we assume that the Energy Star label biases consumers' perception of quality and/or energy operating costs, the optimal Pigouvian tax could be close to zero, or even negative—that is, it could be a subsidy for energy-intensive products. We also find that there is substantial heterogeneity in the behavioral responses across the population, which in theory could justify combining a Pigouvian tax with a subsidy. It is, however, almost never optimal to do so in our context. Given the pattern of behavioral biases pertaining to energy operating costs across income groups, the bias-adjusted tax is too low for lower income groups, but too high for more affluent households. From the standpoint of economic efficiency, offering a subsidy, in addition to the tax, is thus desirable for the former groups, but not the latter. However, the behavioral responses to subsidies, rebates or sales tax, is such that high- and low-income households will take advantage of subsidies when offered, although at a different rate. The take-up of subsidies at the lower-end of the income distribution increases economic efficiency, but these gains are offset by households at the upper-end of the distribution who also take advantage of the subsidies.

The next section discusses the related literature on behavioral response to fiscal instruments with a special emphasis on the energy context. In Section 3, we describe our framework for the design of corrective policies in the presence of microfrictions. In Section 4, we discuss our empirical setting and data, and in Section 5, we present our empirical strategy. We present the estimation results in Section 6 and we conduct our policy analysis in Section 7. Conclusions follow in Section 8.

2. Related Literature on Behavioral Response to Fiscal Instruments

The recent literature on behavioral public finance has raised important questions about the design of tax policy in light of various behavioral phenomena. One important take away from this literature is that different instruments are associated with different biases, and that heterogeneity in behavioral biases exists across segments of the population, types of instruments, and market environments. For instance, take the case of salience, one phenomenon that has received much attention and has been shown to be important in several settings. Finkelstein (2009) shows how electronic toll payments—in lieu of paying tolls at tollbooths—have reduced the salience of tolls and effectively made the response of driving to these taxes less elastic. Likewise, Sexton (2015) shows that automatic bill payment for monthly electricity bills reduces attentiveness to and salience of electricity prices. As a result, automated payment has increased residential and commercial electricity consumption. In an influential study, Chetty et al. (2009) show that salience—or, more precisely, the lack thereof, affects the response to sales taxes. They find that presenting tax-inclusive posted prices reduces demand relative to the case where taxes are applied at the register. Moreover, they document that consumers appear not to lack information about sales tax rates, which suggests that salience itself is an important psychological phenomenon to consider in the design of tax policy. Goldin and Homonoff (2013) expand on this work by examining how the salience of register taxes on cigarettes varies across population focussing on the role of income. Taubinsky and Rees-Jones (2016) examine heterogeneity in consumers' inattentiveness to sales taxes, and demonstrate how heterogeneity among the population impacts welfare. Baker et al. (2017) also find that sales taxes lack salience, but the behavioral responses to these taxes are highly heterogeneous, and some consumers are sophisticated enough to optimize intertemporally to take advantage of change in sales tax rates over time.

In the energy context, salience has also been found to be an important phenomenon. For instance, Li et al. (2014) evaluate how drivers respond differentially to fuel taxes and fuel prices. They find a demand elasticity with respect to the tax about three times larger in magnitude than the demand elasticity with respect to gasoline prices. They ascribe this difference to the extensive media coverage of tax changes, in contrast to the lesser coverage of fuel prices.¹ Davis and Kilian (2011) also show larger responses to gasoline taxes than to prices, which they attribute to salience of tax changes. Rivers and Schaufele (2015) find a larger response to the British Columbia carbon tax in gasoline consumption relative to the demand elasticity with respect to prices. They likewise

¹They also describe how consumers may perceive a fuel tax increase as more persistent than a fuel price increase, and hence may be more likely to make changes in consumption behavior.

attribute this to the salience of the carbon tax, which received substantial media exposure upon its implementation in 2008. Baranzini and Weber (2013) also find a more pronounced demand response to the imposition of a 1993 tax in Swiss gasoline markets than what they estimate with respect to prices. For vehicle tax, Huse and Koptuyug (2017) show the opposite—consumers tend to respond less to such tax than variation in fuel prices, but there is heterogeneity in the response to the tax due to salience. Increasing the salience to a Pigouvian tax might have, however, unintended consequences. For instance, Lanz et al. (2017) provide experimental evidence that the simple framing of a change in relative prices might impact behavioral responses. In their setting, they find that informing consumers that a price change for food products was induced by a Pigouvian charge on embedded CO_2 lead to less substitution across products than an equally-salient change in price with no particular frame. They discuss how different psychological mechanisms such as self-image motivation (Ariely et al. 2009) and biased prior beliefs about external costs (Gneezy and Rustichini 2000) could explain this result. Their result also brings support to motivation crowding-out theory (Frey and Oberholzer-Gee 1997), which stipulates that financial incentives might reduce intrinsic motivation to change behaviors.

The welfare implications of salience, framing, and more generally behavioral biases, have been discussed in many recent papers, and several of the classic results of public finance have been challenged. In the context of externality-correcting instruments, the conventional view, dating back to Pigou (1920), has been to set the tax on the externality equal to its marginal social damage. This view of assigning one instrument to correct one market failure may result in socially suboptimal outcomes if at least some populations behave in ways that do not appear to be rational. Allcott et al. (2014) describe types of consumers who undervalue energy costs and, as a result, bear an “internality” by purchasing privately suboptimal energy-consuming durable goods. They define an internality as a broad concept that captures different behavioral biases such as inattentiveness, bounded rationality, or lack of salience. They show that when consumers are heterogeneous with respect to internalities, not only is it socially optimal to adjust the Pigouvian tax to account for the internalities, but also a subsidy could be justified, and important interactions could exist among the two types of instruments. For example, if a tax on energy is less effective in reducing the impact of the internality than a well-designed subsidy, then the optimal tax would fall below the marginal damage.

The lack of salience of some instruments relative to their alternatives also suggests that the statutory design of a tax could have important implications on the economic incidence of the tax (Chetty et al. 2009). Gallagher and Muehlegger (2011) provide evidence that this could be

an economically important phenomenon to consider. They investigate consumer responses to tax deductions, tax credits, and sales tax exemptions for hybrid vehicle purchases. They find that a dollar reduction in sales tax has 10 times the effect of a dollar of value in tax credits in increasing hybrid sales. Since the sales tax waiver delivers immediate savings, but a taxpayer must wait until she files a tax return to claim the benefit of a deduction or tax credit, they note that the difference could be attributed to discounting. The seasonal variation in the response to these two tax instruments, however, is inconsistent with discounting. Instead, it appears more in line with car buyers becoming more knowledgeable about income tax credits during tax filing season, but less knowledgeable during other times of the year.²

Finally, the heterogeneity across the population in behavioral responses indicates an opportunity for targeting in tax design (Congdon et al. 2009). This may also be analogous to the importance of targeting of social programs in ways to enhance take-up (Currie 2004). High administrative burden in signing up for various social programs may discourage some take-up, although this could also enable selection by those who most benefit from the program (i.e., households with lower income or lower opportunity cost of time). Designing some energy fiscal policies that may require paperwork for claiming the benefits—such as completing forms for a tax credit as in Gallagher and Muehlegger (2011)—could deter take-up. But these kinds of microfrictions should be considered different from the behavioral microfrictions described above. The cost of completing a form is a transaction cost that offsets the value of a tax credit. It is not a behavioral bias like the lack of salience. As we formally show next, the nature of the microfrictions consumers are subject to when they respond to various fiscal instruments—and especially the distinction between behavioral biases and tangible transaction costs—has important design implications.

3. Optimal Pigouvian Instruments in the Presence of microfrictions

The recent literature on behavioral public finance highlights the existence of behavioral biases that should inform the design of fiscal instruments. Our work builds on this foundation by developing a unified framework using the concept of microfrictions. In this section, we propose a formal definition of microfrictions and then use our theoretical framework to derive formulas that characterize

²De Groot and Verboven (2016) find evidence that discounting can influence the design of energy taxes and subsidies. They show that adopters of residential solar panels steeply discount future energy prices, which are subject to generous government subsidies with the goal of making solar panels more attractive. They show that policies that subsidize the up-front investment cost would require much less government funds than subsidized tariffs to achieve similar outcomes.

the optimal Pigouvian fiscal instruments in the presence of these microfrictions. Our analysis illustrates the difference between behavioral biases and tangible transaction costs in their impact on optimal instrument design. In addition, we show the important role of heterogeneity—both heterogeneity across policy instruments and heterogeneity across consumers. Finally, we describe how our framework can motivate empirical work.

3.1. Setup

We consider a context where consumers have unit demand and can choose among J technologies. A technology j requires E_j units of energy to operate, and there is a negative externality associated with the use of energy. The marginal external cost associated with each unit of E_j is denoted ϕ . The technology could be a new vehicle, water heater, or refrigerator, which uses gasoline, natural gas, or electricity, respectively. We make an abstraction of the utilization decision and assume that E_j is entirely determined by manufacturers prior to consumers' purchase decisions. This assumption is motivated by our empirical application that focuses on the appliance market, in particular refrigerators for which utilization should be very inelastic with respect to the price of energy (denoted P^e). We assume that the technology lasts only one time period and consumers must pay $P^e \cdot E_j$ in energy operating cost over this period.

Consumer k values technology j as function of its quality, V_{kj} ; purchase price, P_j ; and the energy operating cost, $P^e \cdot E_j$. We assume quasilinear preferences, and the marginal utility of income is denoted η_k . The alternative-specific utility for technology j is thus given by

$$(1) \quad U_{kj} = V_{kj} - \eta_k(P_j + P^e \cdot E_j).$$

The quality term, V_{kj} , captures all vertical and horizontal dimensions of quality valued by consumer k . We assume that these two dimensions of quality are additive and separable: $V_{kj} = \nu_k + \epsilon_{kj}$. To derive our main results with the case of $J > 2$ technologies, it is useful to assume that the horizontal component of quality, ϵ_{kj} , corresponds to idiosyncratic preferences and follows a generalized extreme value distribution. Under this assumption, the probability that consumer k chooses technology j , denoted σ_{kj} , takes the well-known multinomial logit form.³

³The two-technology case where the distribution of ϵ_{kj} is left unspecified is discussed by Allcott et al. (2014).

3.2. Behavioral Biases versus Transaction Costs

Fiscal instruments can influence the consumers' choice of technology, by impacting either the investment margin (through P^j) or the operating margin (through P^e). The efficiency of fiscal instruments in driving behavioral change could be undermined by microfrictions, which we formally define here.

Definition 1. *microfrictions are any phenomena impacting the behavioral response to a fiscal instrument such that a one dollar variation induced by the fiscal instrument does not have the same effect than a one dollar variation in relative prices.*

In our model, a fiscal instrument is subject to microfrictions when its marginal effect on utility is not equal to the marginal utility of income.⁴ microfrictions capture a range of phenomena: behavioral biases, such as lack of salience, inattention, misperceptions, and bounded rationality, and transaction costs, such as the time and cost to search for tax information and to complete the relevant tax paperwork.⁵

In our framework, we represent a behavioral bias with a term m that drives a gap between decision and experienced utility. For example, consumer k 's perception of the policy I is subject to behavioral biases if his decision utility is given by $U_{kj} = V_{kj} - \eta_k(P_j + P^e \cdot E_j + m \cdot I)$, but the utility he actually experiences after a purchase decision is given by $U_{kj} = V_{kj} - \eta_k(P_j + P^e \cdot E_j + I)$. When a consumer is subject to behavioral biases, the welfare effects of a policy can still be measured but need to account for the discrepancy between decision and experienced utility. Leggett (2002) and, more recently, Allcott (2013), Ketcham et al. (2016), and Houde (2017) show that the expression for the change in consumer surplus in a discrete choice framework can be adapted to account for behavioral biases.⁶ For an instrument change $\mathcal{I} \rightarrow \tilde{\mathcal{I}}$, the (expected) change in consumer surplus

⁴To illustrate this, consider the effect of a small change in the price P_j on utility: $\frac{\partial U_{kj}}{\partial P_j} = -\eta_k$. In the absence of microfrictions, the following equality must hold for a small tax, τ , levied on technology j : $\frac{\partial U_{kj}}{\partial \tau} = \frac{\partial U_{kj}}{\partial P_j} = -\eta_k$.

⁵In the literature on health plan choice, Handel and Kolstad (2015) distinguish between the concepts of information frictions and hassle costs. Our definition of microfrictions encompass these two concepts. We consider information costs as being one source of behavioral biases and hassle costs as being one type of tangible transaction cost.

⁶Recently, this expression has also been used to quantify the welfare effects of advertising (Dubois et al. 2017), firms' strategic gaming of attribute information (Reynaert and Sallee 2016), and fuel economy standards and nudges in the presence of inattentive consumers (Allcott and Knittel 2017).

(CS) for consumer k is given by

$$(2) \quad \Delta CS_k = \frac{1}{\eta_k} \cdot \left[\ln \sum_j^J \exp(\tilde{U}_{kj}) + \sum_j^J \tilde{\sigma}_{kj}(\tilde{U}_{kj}^E - \tilde{U}_{kj}) \right] - \frac{1}{\eta_k} \cdot \left[\ln \sum_j^J \exp(U_{kj}) + \sum_j^J \sigma_{kj}(U_{kj}^E - U_{kj}) \right],$$

where the terms with a tilde are evaluated at the postpolicy change, U_{kj}^E denotes experienced utility and U_{kj} corresponds to decision utility where behavioral biases m influences the behavioral response to the policy. The above expression differs from the standard welfare measure derived by Small and Rosen (1981) due to the terms $\sum_j^J \sigma_{kj}(U_{kj}^E - U_{kj})$, which we refer as the Leggett (2002) correction. This represents the expected difference between experienced and decision utility and quantifies the size of the welfare losses due to behavioral biases.

Abstracting from the gap between decision and experienced utility, we find that behavioral biases do not impose additional costs on welfare. This interpretation is consistent with the concept of internality proposed by Allcott et al. (2014) and the model of behavioral biases proposed by Farhi and Gabaix (2015). In contrast, we represent transaction cost microfrictions as an *economic* cost C that the consumer must bear to respond (denoted $R = 1$) to the policy I . Consumer k then faces the following decision problem:

$$\max_{R=\{0,1\}} (1 - R) \cdot \left[\max_j V_{kj} - \eta_k(P_j + P^e \cdot E_j) \right] + R \cdot \left[\max_j V_{kj} - \eta_k(P_j + P^e \cdot E_j + I) - C \right]$$

In this decision problem, if the consumer does not pay C , then she will optimize as if the policy I was not present. If she takes up the policy, then her welfare will change by the amount $-\eta_k I - C$.⁷

3.3. Optimal Pigouvian Instruments

Let us now turn to the subset of fiscal instruments I that price negative externalities—that is, a Pigouvian tax. Let ϕ represent the marginal external cost of energy. In our framework, consumers must pay $(P^e + \tau)E_j$ for operating technology j in the presence of tax τ . If consumers face microfrictions in computing the energy costs inclusive of the tax, we can derive an expression for the optimal Pigouvian tax.

⁷Chetty et al. (2007) propose a similar model to model boundedly rational agents that do not take into account a small tax. They refer to the cost C as a small cognitive cost. The difference in our framework is that we consider that C is a tangible economic cost that must be accounted for in social welfare.

Proposition 1. *If microfrictions are transaction costs, Pigou holds that $\tau = \phi$.*

If microfrictions are behavioral biases represented by m , the bias-adjusted Pigouvian tax is

$$(3) \quad \tau = \frac{\phi}{m} + P^e \cdot \frac{(1-m)}{m}$$

All proofs are in Appendix A.

There is thus a stark difference in the level of the Pigouvian tax that must be charged whether microfrictions are considered tangible economic costs or behavioral biases. In the presence of behavioral biases, the expression for the optimal Pigouvian tax departs from the classic result in two ways because it aims to internalize both the externality and the behavioral biases. The marginal external cost is thus scaled by the size of the bias.⁸ If consumers' biases with respect to energy costs lead to an undervaluation (i.e., $m < 1$), the tax must be adjusted upward, and vice versa. The price of energy also enters the expression because with a price-inclusive tax, the gap between decision and experienced utility is a function of not only the tax, but also the price of energy. Therefore, the adjustment could be quite large even for a modest level of m , and, as pointed out by Allcott et al. (2014), it could be optimal to levy a tax when the marginal external cost is zero.

If there are additional fiscal instruments that impact the purchase decision, this will also impact the expression for the optimal tax. To illustrate, consider that an ad valorem sales tax, denoted T^s , is levied on the price of each technology, and consumers' response to T^s is scaled by d , which may capture the lack of tax salience or other biases. The expression for the bias-adjusted Pigouvian tax then becomes

$$(4) \quad \tau = \frac{\phi}{m} + P^e \cdot \frac{(1-m)}{m} - T^s \cdot \frac{d}{m} \cdot \frac{\sum_j \frac{\partial \sigma_{kj}}{\partial \tau}}{\sum_j \frac{\partial e_j}{\partial \tau}}$$

In the absence of biases—that is, $d = 1$ and $m = 1$ —the optimal Pigouvian tax must be adjusted to account for the tax interaction effect, a well-known result related to the double dividend of environmental taxation (Bovenberg and Goulder 2002; Parry and Bento 2000; Goulder and Williams III 2003). The size of this adjustment is, holding everything else constant, decreasing when the sales

⁸Farhi and Gabaix (2015) consider a simple setting where the price of energy is zero. They show that the bias-adjusted tax is simply $\tau = \phi/m$. Allcott et al. (2014), on the other hand, consider a more general setting where, after some rearranging of their Proposition 1 (p. 76), the bias-adjusted tax is equal to the expression we obtain and an additional term that captures misperceptions in the future benefits associated with utilization. Our main contribution is thus to highlight that the interpretation of microfriction matters. To explicitly show that the price of energy matters for the adjustment is our secondary contribution.

tax is undervalued ($d < 1$) and, in the limit, for a very large degree of undervaluation $d \rightarrow 0$, can even go to zero. But ultimately the size of the adjustment is determined by interacting biases: the ratio d/m in our framework. If consumers undervalue the sales tax, but less than tax-inclusive energy costs, this will amplify the tax interaction effect, for instance. In sum, in the presence of multiple fiscal instruments, heterogeneity in microfrictions across instruments impacts the tax interaction effect, and it is the relative size of the biases that matters.

3.4. Heterogeneity Across Types and Instruments

We now extend the analysis for the case with K different consumer types, where each type $k \in K$ has a different share in the population, denoted α_k , and a different level of behavioral biases, denoted m_k . In our setting, the expression for the optimal Pigouvian tax with heterogeneous consumers is

$$(5) \quad \tau = \frac{\phi}{1 - \frac{\sum_k \alpha_k (1-m_k) \Delta_k^{\tau,e}}{\sum_k \alpha_k \Delta_k^{\tau,e}}} + P^e \cdot \frac{\frac{\sum_k \alpha_k (1-m_k) \Delta_k^{\tau,e}}{\sum_k \alpha_k \Delta_k^{\tau,e}}}{1 - \frac{\sum_k \alpha_k (1-m_k) \Delta_k^{\tau,e}}{\sum_k \alpha_k \Delta_k^{\tau,e}}},$$

where the expression $\Delta_k^{\tau,e}$ is the net change in energy use due to a small tax τ :

$$(6) \quad \Delta_k^{\tau,e} = \sum_j \frac{\partial \sigma_{kj}}{\partial \tau} E_j.$$

Although the above expression is more complex than Equation 3, it can still be signed.

Lemma 1. *If all consumer types undervalue tax-inclusive energy costs $m_k \leq 1, \forall k$, the bias-adjusted Pigouvian tax is larger than the marginal external cost—that is, $\tau > \phi$. The opposite also holds.*

Heterogeneity in behavioral biases has several implications. If we cannot perfectly target each consumer with a type-specific tax, the bias-adjusted Pigouvian tax will be too high for some consumers and too low for others. Such heterogeneity can motivate the use of multiple instruments, a tax combined with a subsidy, for instance (Allcott et al. 2014), or the use of quantity instruments, such as technology standards (Farhi and Gabaix 2015; Allcott and Knittel 2017). In the framework of Allcott et al. (2014), one reason it is optimal to combine an energy tax with a product subsidy, is that the bias-adjusted energy tax distorts the utilization decision of the energy-intensive durable (a car, in their empirical illustration), which is not the case for the subsidy. If the change in utilization induced by an energy tax is marginal or zero, as in our setting, it becomes harder to justify a subsidy.

When subsidies are also subject to microfrictions, it can only be optimal to combine a tax and a subsidy, if the behavioral responses to each instrument are inversely correlated. The intuition is simple. Ideally, we should target the subsidy to consumers with the larger biases with respect to the tax. If these consumers have lower microfrictions with respect to subsidies, they will respond more to the subsidy. Heterogeneous microfrictions associated with subsidies, thus act as an implicit targeting mechanism. Therefore, offering a subsidy, in addition to the tax, is only optimal if this implicit targeting allocates the subsidy to the consumers with the larger biases. Note that an inverse correlation in the behavioral responses to a tax and a subsidy is not a sufficient condition to motivate combining these instruments. In the present framework, every consumer for whom the bias-adjusted tax is too high and ends up taking advantage of a subsidy is at the source of a welfare loss. Therefore, combining a tax and a subsidy is only optimal when the share of such consumers is small enough relative to the share of consumers for which the bias-adjusted tax is too low.

3.5. Observed versus Unobserved Heterogeneity

Applying the expression in 5 in an empirical setting requires segmenting the population into K observable segments and quantifying the microfrictions for each segment. Within each segment, however, additional heterogeneity—not observed by the researcher—might remain. The implication of such heterogeneity for the measurement (estimation) of m_k is that the level of the behavioral bias might become technology-specific: $m_k \rightarrow m_{kj}$. This will arise when heterogeneity in microfrictions within a segment k leads to consumers sorting into specific regions of the product space. In this case, the behavioral response to a fiscal instrument across different technologies, but within a population segment, will then differ. We refer to heterogeneity within consumer type as unobserved heterogeneity because it arises when the econometrician is not able to perfectly segment the population based on observable information.

The presence of unobserved heterogeneity further complicates the expression for the bias-adjusted Pigouvian tax and makes the direction of the adjustment ambiguous even if biases induce systematic undervaluation ($m_{kj} \leq 1, \forall k$ and j) or overvaluation ($m_{kj} \geq 1, \forall k$ and j). To gain more intuition on this result, consider the expression of the optimal tax targeting consumers of type k for which $m_{kj} \neq m_{ki}$ if $j \neq i$:

$$(7) \quad \tau_k = \phi \cdot \frac{\sum_j \frac{\partial \sigma_{kj}}{\partial \tau} E_j}{\sum_j m_{kj} \frac{\partial \sigma_{kj}}{\partial \tau} E_j} + P^e \cdot \frac{\sum_j (1 - m_{kj}) \frac{\partial \sigma_{kj}}{\partial \tau} E_j}{\sum_j m_{kj} \frac{\partial \sigma_{kj}}{\partial \tau} E_j}.$$

Given that the tax τ will induce substitution from the most energy-intensive technologies to the less energy-intensive ones, the tax will induce a decrease in market share for the most energy-intensive technologies—that is, $\frac{\partial \sigma_{kj}}{\partial \tau} < 0$. The opposite will occur for the less energy-intensive technologies. The net effect of a positive tax should reduce, at least weakly, the net amount of energy use: $\sum_j \frac{\partial \sigma_{kj}}{\partial \tau} E_j < 0$. This allows us to sign the direction of the adjustment. But when the bias m_{kj} for each technology j (or $1 - m_{kj}$) weights the partial derivative of the market, the sign of the following expression, $\sum_j m_{kj} \frac{\partial \sigma_{kj}}{\partial \tau} E_j$ (respectively, $\sum_j (1 - m_{kj}) \frac{\partial \sigma_{kj}}{\partial \tau} E_j$), is ambiguous even if $\tau > 0$. Therefore, the whole expression in Equation 7 cannot be signed. The optimal Pigouvian tax could be below the marginal external cost, ϕ , or even positive (i.e., could be a subsidy) in some cases.

3.6. Implications for Empirical Work

The above results have four important implications for our empirical investigation. First, we need to simultaneously estimate the response to the purchase price of an energy-intensive durable and the responses to fiscal instruments in order to detect and quantify the magnitude of microfrictions. In particular, the empirical strategy needs to exploit variation in prices to recover the marginal utility of income, which is crucial to assess whether microfrictions are present.

Second, it is important to empirically distinguish the type of microfrictions consumers are subject to, which has important implications for policy design. However, the econometrician faces an important challenge in identifying microfrictions when evaluating consumers' decisions. Specifically, if behavioral biases lead to an underresponse to a policy and the econometrician has data only on quantities, then the two types of microfrictions can be observationally equivalent. This arises because by simply observing the effect of policy I on equilibrium quantities, the econometrician does not know if only a fraction of the consumers responded to the policy or if all of them responded, but misperceived the policy (see Appendix A for a formal proof). In such a case, the econometrician must take a stand on whether the observed behavioral response to a policy is subject to transaction costs or behavioral biases.

Third, the empirical strategy must account for heterogeneity across policy instruments and consumers. In practice, this requires us to pay close attention to how we segment the population and capture unobserved heterogeneity. Fourth, the potential presence of unobserved heterogeneity means that a full demand system that captures substitution across products is required.

4. Empirical Setting and Data

4.1. Empirical Setting: US Refrigerator Market

Our empirical setting is the US refrigerator market. Consumers purchase more than 10 million refrigerators annually, making it the largest appliance market in the United States—with an economic value on the order of \$10 billion. Nearly 100 percent of households live in dwellings with a refrigerator and in recent years, the share of housing units with two refrigerators increased to 30 percent. The energy consumed by residential refrigerators exceeds 1 percent of total US energy consumption.

The features of the refrigerator market and policy environment make it particularly well suited to study consumers' responses to energy fiscal policies. First, the federal government implements two labeling schemes that provide appliance customers information on the energy efficiency of the options in their choice set. The EnergyGuide label provides estimates of the expected annual energy consumption and energy expenditures for a given product, as well as a comparison to similar products. The Energy Star (ES) label is a more coarse source of information—only products that meet certain certification criteria have the right to display the simple ES logo. The ES certification plays an important role in branding, marketing, and policy design for energy-efficient products.

Second, policymakers at federal, state, and local levels of government have deployed an array of instruments influencing the investment cost and the energy operating cost of refrigerators. State governments and electric utilities have subsidized investment in energy-efficient refrigerators by offering rebates for purchases of qualifying refrigerators. In most programs, the ES certification serves as the eligibility condition. In addition, states subsidize refrigerator purchases through sales tax exemptions and holidays, some of which specifically target ES-certified products. The cost of operation—the price of electricity—depends in part on a number of state and federal energy and environmental regulations. The increasing stringency of power plant regulations on local air pollutants, including fine particulates, ozone precursors, sulfur dioxide, and air toxics, the emerging state carbon dioxide cap-and-trade programs in California and the northeastern states, and the extensive renewable power mandates over a majority of American consumers can each increase power prices. And any future carbon pricing policy would likely do the same.

Third, refrigerators are unique among energy-consuming durables: the energy consumption for a given model is predominantly determined by the manufacturers' design decisions. Thus, we

can abstract from consumer utilization decisions when examining how changes in operating costs influence consumer behavior.

4.2. Data

We have obtained the universe of refrigerator transactions from a large US appliance retailer over 2008-2012. Each transaction contains information about the refrigerator model purchased, which we have matched to detailed attribute information, including the expected annual electricity consumption based on the model’s EnergyGuide label. Each transaction also contains information about the transaction date, the exact price paid, the total amount of sales taxes paid, and the store location. Through the retailer, we also have transaction-specific demographics from the data aggregator Acxiom for approximately 50% of the transactions. This includes information on household size, income, education, homeownership, housing type, political orientation, and age of the head of the household.

We have limited our sample for analysis along several dimensions. First, we focus only on transactions that contain demographic information. Second, we consider only transactions made by homeowners who bought no more than one refrigerator over 2008-2012. Our estimation samples thus exclude transactions made by contractors making bulk purchases and renters who may not pay for their electricity bills.⁹ For our demand estimation, we draw large random samples ($N \approx 55,000$) of refrigerator purchase transactions for each of six different income groups that are a function of the income category classifications in the source data: <\$30k, \$30k-\$50k, \$50k-\$75k, \$75k-\$100k, \$100k-\$150k, and >\$150k.

We have drawn from multiple sources to construct our data on investment subsidies. We have compiled data characterizing utility rebate programs—including rebate amounts and eligibility criteria—over 2008-2012 from the DSIRE database.¹⁰ We estimate the average utility rebate available for qualifying refrigerators on a county-by-year level by mapping utility program rebate

⁹Renters represent less than 3 percent of transactions for which demographic data are available. Although contractors are not explicitly identified in the data, we observe that about 32 percent of the transactions are made by entities that purchased more than two refrigerators during the period 2008-2012.

¹⁰The Database of State Incentives for Renewables and Efficiency tracks national, state, local, and utility policies that impact renewable power and energy efficiency investments. Refer to: <http://www.dsireusa.org/>.

amounts to the counties served by a utility based on the EIA Form 861 database.¹¹ If a county is served by multiple utilities, then we take a nonweighted average of the rebate amounts.

In a similar fashion, we have produced a detailed dataset on state Cash for Appliances (CFA) programs, which offer consumers rebates for the purchase of Energy Star-rated refrigerators (and other appliances) (Houde and Aldy 2017). The federal government provided states with approximately \$300 million in grants under the 2009 American Recovery and Reinvestment Act to implement this rebate program. We have compiled data on the rebate amounts offered, appliances covered, eligibility criteria, timing and duration of programs, and mechanisms to claim the rebates.

For sales tax rates, exemptions, and holidays, we have drawn from data produced by the Tax Policy Center, the Sales Tax Institute, and the DSIRE database.¹² We created a day-by-zip code panel of sales tax rates, holidays, and eligibility criteria for holidays and exemptions over 2008-2012.

Finally, we compute the electricity operating cost for each appliance model in the sample using the expected annual electricity consumption reported by the manufacturer multiplied by the average electricity price of the region where each household made a purchase.¹³ We assume that consumers form time-invariant expectations about electricity prices using the current local average price. Time-invariant expectations are consistent with recent evidence in the car market suggesting that consumers' best forecast of future gasoline prices is simply the current prices (Anderson et al. 2011).

4.3. Sources of Variation

Plausibly exogenous sources of variation enable us to identify the heterogeneous behavioral responses to the retail prices of refrigerators as well as to utility rebates, CFA rebates, sales tax policies, and electricity operating costs. We describe below the variation in retail prices and then the variation for each of our policy variables.

¹¹Refer to the Energy Information Administration's detailed data files on electric power sales, revenue, and energy efficiency from Form EIA-861 at <https://www.eia.gov/electricity/data/eia861/>.

¹²For sales tax rates, we have used the transaction-level data and computed average sales tax rates at the zip code-day level. For sales tax holidays and sales tax exemptions targeting energy-efficient appliances, we have used archived webpages for the Sales Tax Institute and DSIRE accessed through the Internet Archive. Detailed data and sourcing are available from the authors upon request.

¹³We do not observe the zip code of each household, but the zip code of the store where each transaction was made. We estimate average annual electricity prices by state using the Energy Information Administration Form EIA-861 database.

Retail Prices. To identify consumers’ responses to retail prices, we exploit a number of institutional features of the appliance market. Our retailer, like most other large US appliance retailers, has a national pricing policy. This implies that a given appliance model has the same retail price across store locations, and the only variation in price is over time. There is, however, substantial temporal variation in prices. This is illustrated in Figure 1, which shows the median prices of the most popular refrigerator models for a major brand.¹⁴ Each panel plots the weekly variation in price for a specific model of a specific brand. For each brand, we show the weekly variation for the four most popular models offered by this brand. We use the sales rank during the period 2008-2012 as our measure of popularity. The red line corresponds to the median change in price relative to the average price over the lifetime of the product, where the median is taken across zip codes. That is, we computed week–zip code-specific changes in price for each model and then plotted the median of the weekly changes for a specific model. The gray band identifies the 25th and 75th percentile of these weekly changes in price. By presenting various quantiles of the distribution of weekly changes, we show that the local store managers typically comply with the national price policy; for most weeks, the 25th and 75th percentiles coincide with the median.

Finally, the blue line plots the median change in price after removing brand dummies interacted with week-of-sample fixed effects. The goal is to show the remaining weekly variation in the price of each model after accounting for seasonal as well as contemporaneous brand-specific shocks. The main takeaway is that after accounting for those temporal shocks, the price time series are smoother, but large and frequent variation persists. These patterns mean that the weekly variation in prices is weakly correlated among models of the same brand, and large price events are model-specific and tend to be idiosyncratic.¹⁵ In sum, there is significant randomness in how the retailer sets prices.

In our estimation, we rely on this high-frequency temporal variation in prices to identify consumers’ sensitivity to prices. The identification argument here is similar to that of Einav et al. (2012): abrupt variations in prices identify price elasticities as long as they are not correlated with slower-moving trends in demand. We show that this exclusion restriction is likely to hold in the present context. For example, we show that controlling for brand-week-specific fixed effects has little impact on the coefficient estimate for the retail price, suggesting that the raw variation in price alone is mostly uncorrelated with demand shocks.

¹⁴The brand name is anonymized to keep the confidentiality of the data. Similar patterns are found for other brands but are not shown here.

¹⁵These patterns are not restricted to the nine most popular models.

Utility Rebates and CFA Rebates. Utilities provide rebates for energy-efficient appliances typically as a result of demand-side management requirements set by their state utility commissions. These requirements provide substantial discretion to utilities on how they promote energy efficiency, and appliance rebate programs are only one of many instruments at their disposal. We observe significant variation over space and time in utility rebate programs' applicability to energy-efficient refrigerators and generosity of rebates. Some programs operate for a few months, while others operate for years, and the rebates range from as little as \$15 to as much as \$1,000 for a refrigerator.

Under CFA, states had sovereignty over the design of several elements of their rebate programs, subject to Department of Energy approval. In practice, the state programs varied by appliance coverage—most but not all states offered rebates for energy-efficient refrigerators—and rebate amounts. The states offered economically significant rebates, on average 12% of refrigerators' sales prices, and these varied significantly among states (Figure 5, Appendix C). The start dates and duration of CFA programs differed considerably as well.

State Sales Tax. In our setting, there are three sources of variation in sales tax rates. First, there is substantial cross-sectional variation across states in the level of the sales tax rate, and it occurs to a lesser extent within states because of local jurisdictions imposing their own sales taxes. This is shown in Figure 4 (Appendix C), from Einav et al. (2014). Second, state and local sales taxes also vary over time, and this variation can be economically important. Tax rates are typically adjusted every year and these changes are usually coordinated with calendar time—that is, new tax rates usually take effect on January 1ST. Finally, sales taxes also vary over time because of tax holidays, some of which specifically target Energy Star-qualified products. As shown in Table 7 in Appendix C, of the 11 states that offered an appliance-oriented sales tax holiday over 2008-2012, 9 of them used Energy Star certification as an eligibility criterion. We compute the sales tax for each product in our sample using the weekly retail price and zip code-specific sales tax rates, accounting for sales tax exemptions and holidays that target ES-certified products.

Electricity Operating Costs. Electricity operating costs vary substantially as a function of the energy efficiency of refrigerator models as well as retail electricity prices across the country. Each margin is approximately equally responsible for the maximum operating cost to be four times larger than the minimum operating cost in our sample. Since the first margin is a function of the consumer's purchase choice, let us focus on the second margin, which is exogenous to the consumer (conditional on making a residential location decision). Average electricity prices vary across the country as a function of historical investments in different types of power generating technologies,

the change in natural gas prices over our study period, and the significant changes in environmental regulations and renewable power subsidies that contribute to an evolving generation stock and associated power prices. We estimate electricity operating costs by multiplying manufacturers' reported expected annual electricity consumption and average state electricity prices.

5. Empirical Strategy

We represent the consumer's appliance purchasing decision by a discrete choice model, where consumer (household) i values refrigerator model j at time t in region r as follows

$$(8) \quad U_{ijtr} = \gamma_{ij} - \eta_i P_{jrt} - \alpha_i T_{jrt}^S + \nu_i ES_{jt} + \psi_i S_{rt}^{DSM} \times ES_{jt} + \phi_i S_{rt}^{CFA} \times ES_{jt} - \theta_i Elec_{jrt} + \epsilon_{ijtr}$$

The term γ_{ij} represents a consumer-product-specific fixed effect, which is modeled with a product fixed effect, $\tilde{\gamma}_j$, and interaction terms between product attributes (X_j) and demographic information ($Demo_i$)—that is, $\gamma_{ij} = \tilde{\gamma}_j + \beta X_j Demo_i$. The variable P is the retail price before tax and T^S is the sales tax. The variable ES takes a value of 1 if product j is certified at time t and zero otherwise. The variables S_{rt}^{DSM} and S_{rt}^{CFA} are subsidies for ES-certified refrigerators offered by utilities through their demand-side management (DSM) programs and the CFA rebate program, respectively. $Elec$ is the annual electricity cost of operating product j in region r in year t . Finally, ϵ_{ijtr} represents idiosyncratic taste parameters, which we assume follow a Type-1 extreme value distribution and are i.i.d. This leads to the conditional logit.

To capture heterogeneity, we interact the preference parameters with observable demographic information to identify consumer-specific behavioral responses. We focus primarily on income, but we have also investigated heterogeneity with respect to education, age, family structure, and political orientation.

5.1. Interpretation of Model Parameters

In this framework, the coefficient estimate on retail price, η_i , corresponds to the marginal utility of income.¹⁶ We use our estimate of η_i to detect and quantify the magnitude of the microfrictions. For example, if consumers have perfect information about the sales tax rate and this information is equally salient as the retail price, the coefficient estimate for α_i should exactly equal the coefficient estimate for η_i . The lack of sales tax salience (Chetty et al. 2009) or consumers lacking full

¹⁶We let the coefficient η_i vary across income groups, but we rule out an income effect within each group.

information about changes in sales tax rates (Baker et al. 2017) would imply $\alpha_i/\eta_i < 1$, in other words, evidence of microfrictions.

Similarly, by comparing the coefficient estimates on rebates ψ_i and ϕ_i to the coefficient estimate for retail price η_i , we can measure the microfrictions that influence the responses to various types of rebates. In our setting, we do not observe whether a consumer claims a rebate, but we show in Appendix B that the ratios ψ_i/η_i and ϕ_i/η_i each can be interpreted as an approximation of the probability that consumers claim a rebate.¹⁷ microfrictions associated with claiming a rebate offered by utilities or under the CFA program exist if $\psi_i/\eta_i < 1$ or $\phi_i/\eta_i < 1$, respectively. The greater the time and effort to learn about and file a claim for an energy efficiency rebate, the smaller the ratio.

Finally, the relative magnitude of the coefficient estimates on annual electricity cost, θ_i , and retail price, η_i , can reveal microfrictions that influence the perceptions of energy operating costs. The coefficients η_i and θ_i can be used to infer a value of an implicit discount rate r_i that rationalizes consumers trade-off between retail prices and lifetime energy operating costs. To estimate the discount rate r_i , we assume that consumers form time-invariant expectations about the annual operating electricity expenditure, do not account for the effect of depreciation, and assume an appliance lifetime of 18 years (typical for a refrigerator). The lifetime energy operating cost (LC_j) for a refrigerator j is simply

$$(9) \quad LC_{ij} = \sum_{t=1}^L \rho(r_i)^t \cdot Elec_{ij} = \rho(r_i) \cdot \frac{1 - \rho(r_i)^L}{1 - \rho(r_i)} \cdot Elec_{ij},$$

where L is the lifetime of the durable, $\rho(r_i) = 1/(1 + r_i)$ is the discount factor, and $Elec_{ij}$ is the annual electricity operating cost for product j . Using the estimates for η_i and θ_i , the implicit discount rate r_i is then the solution to the following equation:

$$(10) \quad \theta_i = \eta_i \cdot \rho(r_i) \cdot \frac{1 - \rho(r_i)^L}{1 - \rho(r_i)}.$$

A number of papers have taken a similar approach to estimating implicit discount rates by comparing the price of a durable purchase and the expected lifetime operating costs, including Hausman (1979), Gallagher and Muehlegger (2011), Busse et al. (2013), and Allcott and Wozny (2014). These

¹⁷This interpretation is exact in a linear estimation framework, where we can interpret ψ_i and ϕ_i as reduced form intent-to-treat estimators. Given that we use a non-linear framework, we rely on a linear approximation of the choice model to interpret the ratios ψ_i/η_i and ϕ_i/η_i . In our framework, these ratios correspond to probabilities to claim a rebate under the assumption that these probabilities are constant as a function of the rebate amount.

papers reveal some heterogeneity in discount rates across consumers, but in several cases, the implicit discount rates are consistent with rates a consumer may face for credit card financing. This imputed r_i could reflect both consumers’ pure rate of time preference and behavioral biases, such as myopia. To quantify the magnitude of these microfrictions, we compare the implicit discount rate with an “appropriate discount rate,” denoted \bar{r} , in the context of the appliance purchasing decision.

The literature presents two lines of thought for determining the “appropriate discount rate,” \bar{r} . Some have argued that \bar{r} should reflect the rate that consumers face in their daily financial decisions. We have two reservations about this approach. First, financial market frictions may ration the supply of credit and bias upward our estimated discount rate relative to consumers’ pure rate of time preference. Indeed, Hausman (1979) argued that the implicit discount rates should be benchmarked with the normal rate of return found in well-functioning financial markets. Second, we take the perspective in this analysis of the social planner and believe that the “appropriate discount rate” should be the same as the one used to discount the externalities associated with energy use (i.e., carbon dioxide emissions). Throughout the analysis, we use a discount rate of 5% as our “appropriate discount rate” benchmark, with the caveat that it corresponds to the upper end of discount rates employed in the climate change economics literature (Goulder and Williams III 2012). It also falls within the range (3% to 7%) used by the US Department of Energy in conducting its benefit-cost analyses of appliance standards.

5.2. Identification and Preferred Specification

The product fixed effects ($\tilde{\gamma}_j$) together with the interaction between product attributes and demographic information play an important role in statistical identification. The product fixed effects capture all time-invariant preferences for product attributes for each refrigerator model. Once we control for those, the remaining variation to identify the behavioral responses to rebates, sales tax, and electricity costs is thus the variation across regions and time. One concern with exploiting the spatial variation in these variables is that consumers’ preferences for specific refrigerator attributes correlated with energy usage could also be correlated with some policy instruments. For instance, households that prefer larger appliances might live disproportionately in regions with low electricity prices and no rebates. In such a scenario, preferences for size, which is strongly correlated with overall appliance energy use, might be confounded with a response to electricity prices and rebates. Including a rich set of interaction terms between product attributes and demographics is a first way to control for region-specific preferences and ensure that the spatial variation in the various

policy instruments is exogeneous. In particular, we focus on including refrigerator attributes that are strongly correlated with energy use, such as size, door design (e.g., top freezer versus bottom-freezer), and the ice-maker option, together with a large set of demographic variables: income, education, age of the head of household, household size, and political affiliation.

We further control for region-specific preferences by interacting state dummies with the ES certification dummy. In the United States, the ES program serves as the primary means to advertise and promote energy efficient durables (other than cars). Thus, these interaction terms capture both unobserved region-specific consumer preferences for energy efficiency and the equilibrium supply-side response to the program. Controlling for the latter is particularly important given that energy efficiency subsidy programs rely primarily on the ES certification, and in regions where those programs are offered, governmental agencies, utilities, and retailers might be more likely to publicize the ES program. As a result, the awareness and understanding of ES appliance certification may vary systematically across regions due to advertising,¹⁸ which could be confounded with the responses to various subsidies relying on the ES program. The state-ES fixed effects rule out these potential confounding effects. Note that the ES dummy can also vary over time within a product, because in April 2008 the ES certification requirement became more stringent, causing a large number of products to lose the ES label.¹⁹

The state-ES fixed effects also help exploit the spatial variation in electricity operating costs in a credible manner. Suppose that all ES-certified refrigerator models had the exact same expected annual electricity use, and likewise for all of the non-ES-certified models. The distribution of electricity use in the choice sets would then have only two point masses. In this case, state variation in electricity prices could not identify the coefficient on electricity cost if state-ES fixed effects were included, because the difference in electricity costs between certified and non-certified models will be a state-specific constant that would be perfectly captured by these fixed effects. The coefficient on electricity cost is thus identified only if there is variation in electricity use within the subset of products that are ES-certified models or non-certified models. State variation in electricity prices then scales down or up the distance between products in the energy dimension of the characteristic

¹⁸Every year, the US Environmental Protection Agency (EPA) publishes a publicity “intensity” map. Figure 6 (Appendix C) shows that there is substantial variation across designed marketing areas (DMAs), but there is very little variation over time.

¹⁹Houde (2017) further discussed this natural experiment and exploited it to identify the willingness to pay for the ES label. The sample also contains a second similar natural experiment where about 21 refrigerator models lost their ES certification in 2010 because of an incorrect certification procedure.

space. In particular, in high-electricity-price regions, the distance is the greatest and electricity costs should matter more in the purchasing decision. Note that the state-ES fixed effects still capture preferences for energy efficiency correlated with high and low electricity prices.

We also account for unobserved heterogeneity in the way consumers process different pieces of energy information, by interacting the variable for electricity cost with the ES dummy. The coefficient on this interaction term represents how consumers might perceive energy operating costs for ES-certified models differently. As shown by Houde (2017), the ES label and information about energy operating costs are substitutes—consumers who have a high willingness to pay for the ES label do not value information on expected energy costs and vice versa. Houde (2017) developed a framework where (unobserved) costs in collecting and processing energy information rationalize why some consumers prefer to rely on the coarse information provided by the ES label instead of more precise energy information. In this application, the interaction term between the ES dummy and energy operating cost is a simple way to capture this unobserved heterogeneity.

To account for any systematic timing of brand marketing that could influence our estimation of the response to appliance prices, we include brand dummies interacted with time fixed effects as a robustness check. To address potential concerns about consumer sorting across time or space to take advantage of a specific fiscal instrument, we have also included a rich set of household-specific demographic controls in our empirical model.

To distinguish the various sources of variation in the sales taxes, we interact the sales tax variable with a dummy that identifies sales tax holidays that target ES models. To account for the potential for intertemporal substitution in response to the temporary CFA rebate program (Houde and Aldy 2017), we include three dummy variables, each interacted with the state rebate amount—one for two months before the program started, one for the duration of the program, and one for two months after the program ended.

Our preferred specification is thus

$$\begin{aligned}
 (11) \quad U_{ijtr} = & \gamma_{ij} - \eta_i P_{jrt} - \alpha_i T_{jrt}^S + \alpha_i^{Holiday} D_{rt}^{Holiday} \times T_{jrt}^S + \psi_i S_{rt}^{DSM} \times ES_{jt} \\
 & + \phi_i^{Before} S_{rt}^{CFA,Before} \times ES_{jt} + \phi_i^{During} S_{rt}^{CFA,During} \times ES_{jt} + \phi_i^{After} S_{rt}^{CFA,After} \times ES_{jt} \\
 & - \theta_i Elec_{jrt} - \lambda_i Elec_{jrt} \times ES_{jt} \\
 & + \nu_{ir} ES_{jt} \times D_r^{State} + \epsilon_{ijtr}
 \end{aligned}$$

5.3. Estimation Details

We estimate the conditional logit model without an outside option via maximum likelihood. We thus focus on modeling the purchasing decision of consumers that decided to buy a new appliance at a given store in a specific week. The consideration set of each consumer consists of all the refrigerator models offered in the zip code where the purchase was made.²⁰ We also restrict the consideration set of each consumer based on the overall size of the refrigerator that was purchased. Imposing such a restriction reflects the physical constraints to upsize or downsize a refrigerator due to kitchen design that most households face. We present results for two specifications. In our preferred specification, we impose a lax restriction and include in the consideration set all refrigerators within 5 cubic feet of the refrigerator volume a consumer purchased. This criterion represents a range of about 30% of an average refrigerator size. As a robustness check, we limit the consideration set to refrigerators within 0.5 cubic feet of the refrigerator volume purchased, which almost completely rules out substitution with respect to size.²¹

6. Estimation Results

Table 1 reports the results for the six income groups for our preferred specification. This specification includes interaction terms between demographics and attributes, product fixed effects, and state-ES fixed effects. The interpretation of the coefficients are graphically shown in Figure 2. Several interesting patterns emerge. Panel (a) shows that the response to electricity cost relative to the response to appliance price is increasing with income. That means that the trade-off between lifetime energy operating cost and price is rationalized by a lower implicit discount rate for higher-income groups. If we assume that the appropriate discount rate is $r=5\%$, we detect substantial microfrictions associated with energy operating costs for all income groups, but especially for lower-income groups. We also find that the interaction term between electricity cost and the ES dummy is positive, which suggests that consumers pay less attention to electricity cost for

²⁰We impute the choice set for each zip code by trimester (January-April, May-August, and September-December for each year of our sample) using sales. If we observe refrigerator model j being sold in a zip code during a given trimester, we assume that all consumers shopping at this location during that trimester could also purchase model j .

²¹When we restrict the consideration set to refrigerators within 5 cubic feet of the refrigerator volume a consumer purchased, the average size of the consideration set ranges from 60 to 85, dependently on the income group. If we limit the consideration set using 0.5 cubic feet as a cutoff, the average size of the consideration set is approximately 15 for all income groups.

ES-certified models—that is, the ES label magnifies the size of the microfrictions associated with energy operating costs.

The behavioral responses to the various types of subsidies are systematically less than the response to price (Panels (c)-(f)), with the exception of sales tax holidays for a few income groups (Panel (d)). The coefficient on sales tax is slightly more than half the coefficient on price for the second to fifth income groups. It is about 0.8 for the highest-income group, and 0.7 for the lowest-income groups (Panel (c)). Altogether, this suggests that households do not respond to a change in sales taxes the same way that they do to a change in retail prices (Chetty et al. 2009; Baker et al. 2017), but there is a heterogeneity in the degree of attention allocation to the sales tax—a result consistent with Taubinsky and Rees-Jones (2016). The response to sales tax holidays appears, however, more pronounced, with the caveat that this coefficient is imprecisely estimated. Together, these results suggest that a short-lived tax holiday for ES models might induce a larger behavioral response than a permanent tax exemption.

For the coefficients on rebates, lower-income consumers tend to respond more to rebates, and higher-income consumers less, although the difference between the estimates is not large. This result is particularly noticeable for state rebates offered during the CFA program. Interpreting the ratio of the coefficient on rebates to the coefficient on price as the probability to take the rebate, this probability ranges from 25 to 30 percent for the two lowest-income groups during the CFA program and decreases to less than 10 percent for the highest-income group. These results are consistent with the notion that hassle costs influence the decision to claim a rebate. This would partly explain why high-income consumers, who have a higher opportunity cost of time, were less likely to take time and make the effort to claim a rebate.

Finally, not shown in Table 1 but important for the policy analysis is the role of the ES label in the adoption of energy-efficient refrigerators. In the present specification, we have estimated state-specific dummies for ES-certified models. We interpret the ratio of these coefficients with the marginal utility of income as the willingness to pay (WTP) for ES products. The average WTP taken across all states²² is large for all income groups, and increasing with income: \$101, \$103, \$131, \$168, \$174, \$136, for income groups one to six, respectively. These large estimates are consistent with prior work on the ES program that suggests that consumers value certified products well

²²We obtain the average WTP across regions by estimating a simpler model where we use a single dummy for the ES label that turns on when a product is certified and does not vary across states.

beyond their associated energy savings (Houde 2017; Ward et al. 2011; Sahoo et al. 2015; Newell and Siikamäki 2015).

6.1. Extensions and Additional Robustness Tests

Adding time dummies interacted with brand dummies fixed effects should capture time shocks correlated with prices and intertemporal sorting. In appendix D (Table 8), we present the results from a specification with brand-month fixed effects and find that they have little impact on the coefficients on price, and on other coefficients as well. The idiosyncratic model-specific variation in prices displayed in Figure 1 thus provides credible identification.

We also show the results for a specification similar to that in Table 1, except that we restrict the consideration set of each consumer to refrigerator models that are within 0.5 cubic feet of the size actually purchased (Appendix D, Table 9). This restriction mostly affects the coefficient on electricity costs, which becomes larger (in absolute value). This holds for all six income groups. This way of defining the consideration sets effectively rules out substitution across size. These results suggest that the behavioral responses to electricity operating costs are not merely preferences for size—even within a subset of refrigerator models of similar size, relative energy use still matters to consumers.

This analysis shows that behavioral bias and transaction cost microfrictions differentially affect the response to any given fiscal instrument across the income distribution. We investigated how heterogeneity in responses could also occur through other observable characteristics in our data, including education, age, family structure, and political orientation. The results of these analyses are summarized in Appendix D (Tables 10 and 11). Beyond income, we could detect few robust demographic predictors of behavioral responses to fiscal instruments. We found that more-educated and older households tend to respond more strongly to energy operating costs, and this pattern holds for most income groups.

7. Policy Analysis

We use our theoretical framework to investigate the implications for optimal policy design resulting from the substantial heterogeneity in microfrictions across income groups and fiscal instrument types. First, we need to take a stand on the nature of the microfrictions affecting each fiscal

instrument—that is, a line-drawing argument in performing behavioral welfare economics (Bernheim and Rangel 2009). In particular, we need to distinguish between behavioral responses reflecting transaction costs versus behavioral biases, since only the latter type of microfrictions justifies a departure from the Pigouvian principle. Throughout the policy analysis, we maintain the following three assumptions:

Assumption 1. *The behavioral responses to energy operating costs are subject to microfrictions that are behavioral biases. Without behavioral biases, the coefficients on energy operating costs should imply a discount rate in line with other investment/borrowing decisions. We assume $r = 5\%$ for all income groups.*

Assumption 2. *The behavioral responses to sales taxes and sales tax holidays are subject to microfrictions that are behavioral biases.*

Assumption 3. *The behavioral responses to various types of rebates are subject to microfrictions that are transaction costs.*

Based on these assumptions, the optimal Pigouvian tax should be adjusted to account for behavioral biases affecting the perceptions of tax-inclusive energy operating costs. How we interpret the impacts of ES certification on consumers’ responses to operating costs plays an important role in this adjustment. In Table 1, we show consumers in all income groups responding less to energy costs for ES-certified products (positive λ coefficient estimates). Moreover, we found a large WTP for ES-certified products. Houde (2017) argues that this could result from warm glow and other *truly experienced* preferences associated with the purchase of more energy-efficient products; or it could also reflect biases in the perception of quality for ES-certified products, a phenomenon sometimes referred as the halo effect (Boatwright et al. 2008). From a welfare standpoint, if we assume that the large WTP for the ES label and the positive λ estimates reflect behavioral biases, then subsidies for ES-certified products will be hard to justify. Subsidizing ES products will increase the size of the bias consumers are subject to, a source of welfare loss. For our policy analysis, we will consider two scenarios: one for which the behavioral responses to the ES certification are treated as preferences and another scenario for which they are treated as biases.

We also make a number of simplifying assumptions in order to illuminate the key mechanisms at play. First, we assume the price of electricity equals the US national average during our sample period: \$0.11/kWh, and it is constant in nominal terms over the lifetime of the appliance. Second, we assume the marginal external cost is \$0.02/kWh for all electricity consumers. This corresponds

to a carbon tax of approximately \$30/ton of CO₂. Third, we assume consumers expect a common 18-year lifetime for newly purchased refrigerators. Fourth, we abstract from any public finance redistribution motive across income groups by setting the marginal cost of public funds (MCPF) to one. For a tax τ , the objective function used to maximize social welfare is thus

$$(12) \quad SW = \sum_k \alpha_k CS_k + (\tau - \phi) \sum_k t^L \rho^t \sum_j \alpha_k \sigma_{kj} E_j,$$

where CS_k is given by Equation 2, and α_k is the proportion of income group k in the population. The discount factor, ρ , is based on a 5% discount rate. Finally, in our base case scenario, we assume no sales tax, no rebates for energy-efficient products, and no externality (Pigouvian) tax.

Estimating social welfare should account for the transaction costs affecting rebate instruments. Although we have not estimated these costs, we can bound the welfare change using a revealed preference argument. For all consumers who responded to a policy subject to transaction costs, the gain in gross consumer surplus must be at least as large as the size of policy’s transaction costs. We will present estimates of social welfare before subtracting transaction costs, so they will represent upper bounds on welfare changes.

To simulate counterfactual policies, we first focus on the point estimates obtained in Section 6. For subsidies, we use the during-program CFA coefficient estimates and the coefficient estimate on the sales tax variable (excluding the variables interacted with the sales tax holiday dummy). We account for uncertainty in the estimates by performing exhaustive sensitivity analyses in subsection 7.2.

7.1. Main Results: Comparison of Policy Instruments

In our initial case, we assume that the willingness to pay for the ES label represents preferences and that the ES label does not impact the perception of energy operating costs (i.e., we set $\lambda = 0$). Table 2 presents the welfare ranking for eight policy scenarios that include one or more instruments calibrated to maximize social welfare: (1) the optimal policy in the absence of biases in the perception of energy operating costs, a Pigouvian tax set equal to the marginal external damage: $\tau = \phi = \$0.02/kWh$; (2) a bias-adjusted tax, which is $\tau = \$0.104/kWh$ —reflecting the average level of bias over the income groups, m_k , and the energy price ($P^e = \$0.11/kWh$); (3) the optimal, uniform CFA rebate of \$50; (4) a means-tested CFA rebate (i.e., income-group-specific CFA

rebates); (5) optimal sales tax rates of 4.62 percent for ES-rated products and 6.43 percent otherwise; (6) the optimal combination of a Pigouvian tax and a CFA rebate; and (7) the optimal combination of a Pigouvian tax and differentiated sales tax rates (based on ES-classification).

The subsidy instruments in rows 3-5 perform poorly on social welfare grounds. Moreover, the means-tested CFA rebates are regressive. Since low-income households have a higher propensity to claim a rebate, smaller rebates are necessary to change their behavior relative to high-income households. The four options with the bias-adjusted Pigouvian tax yield higher levels of social welfare relative to the other policy alternatives. The rough equivalence among these four options reflects our finding that the optimal CFA or sales tax subsidy is practically zero in the presence of the bias-adjusted tax on energy (rows 6-8). This is not a theoretical result, but specific to our estimated behavioral responses. Since consumers' responses to electricity operating costs vary over income, the non-differentiated bias-adjusted Pigouvian tax is dominated by one that would be income-group-specific with this set of taxes, $\tau_k = [0.129, 0.124, 0.106, 0.095, 0.083, 0.085]$, for income groups one to six, respectively. Thus, the uniform bias-adjusted tax in rows 6-8 is too low for low-income households, while it is too high for high-income households. For the former group, the subsidy could address the residual behavioral bias undermining welfare, but doing so would be excessive for the latter group. On average, the effects cancel out resulting in de minimis subsidies when designing the optimal combination of a Pigouvian tax and subsidies. These results illustrate the point we made earlier. For instance, although, the biases with respect to energy costs and response to CFA rebates are inversely correlated across income groups, it is still not optimal to offer a rebate. This is due to the fact the subsidy over compensate high income groups, which creates a welfare loss that offset the gains of subsidizing lower income groups. This result is driven by the level and gradient of the behavioral responses across income groups to the different instruments and the share of each income group in the population. In the next subsection, we show how varying the behavioral responses across income groups can overturn these results.

7.2. Sensitivity: Observed Heterogeneity

Our base case analysis shows that a bias-adjusted Pigouvian tax dominated scenarios that paired such a tax with subsidies. In this subsection, we explore the sensitivity of this result to our estimated parameters. We generate a new set of income-group-specific behavioral parameters where we vary both the level and the gradient of the behavioral responses for a particular fiscal instrument. In these new parameter sets, the behavioral response must fall within $[0, 1]$ and we constrain the

gradients to fall within $[-0.2, 0.2]$, where a positive gradient indicates that higher-income groups respond more to an instrument than lower-income groups.

Let us illustrate this approach with CFA subsidies. Suppose we set the coefficient on rebates for the lowest income to zero, which implies that the probability of claiming a rebate is zero, and define the coefficients for the other income groups relative to the lowest-income group. We vary the gradient by increasing the behavioral response of the highest-income group and then linearly interpolate the behavioral responses of the second to fifth income groups. For example, if we set the probability of claiming a rebate to one for the highest-income group and zero for the lowest, the gradient would be $(1 - 0)/(6 - 1) = 0.2$, and the probability of claiming a rebate for each income group would be $[0, 0.4, 0.6, 0.8, 1]$, with an average response of 0.5. We repeat this exercise by varying the behavioral responses of the lowest- and highest-income groups, which allows us to investigate a large range of gradients and average levels of behavioral responses. For each set of behavioral parameters, we solve for the optimal combination of tax and subsidy. We report the results graphically by showing the optimal instruments for each combination of gradient and level of the behavioral responses.

Panels (a) and (b) of Figure 3 presents the optimal Pigouvian tax and CFA rebates when we vary the behavioral responses to the CFA rebates, holding all other coefficients constant. Consistent with the theoretical framework, it is never optimal to offer a subsidy when the gradient for the response to rebates is positive, which means that behavioral biases with respect to energy operating costs and responses to rebates are positively correlated. But even when the behavioral responses to each instrument are negatively correlated, which corresponds to a negative gradient on Figure 3, it is not optimal to subsidize consumers for most combinations of behavioral parameters pertaining to CFA rebates. It would only be optimal to add a small rebate to a Pigouvian tax when lower-income consumers are slightly more likely to respond to rebates and where the probability of taking advantage of rebates across all income groups is very low (less than 5%)—that is, scenarios with small negative gradients and very low average responses.

We perform a similar exercise by varying the behavioral responses to the sales tax, again holding all other coefficients constant. Panel (c) of Figure 3 presents the optimal Pigouvian tax, and Panel (d) presents subsidy offered through a sales tax exemption for ES-rated products (i.e., the difference between the ES sales tax and non-ES sales tax). The results are qualitatively similar to those of the CFA rebates; for most combinations of behavioral parameters, the Pigouvian tax alone is sufficient to maximize social welfare.

In Appendix E, we also report the results (Figure 7) where we vary the gradients and level of the behavioral responses to energy operating costs. For this exercise, we solve either for the optimal combination of a Pigouvian tax and CFA rebates or Pigouvian tax and ES sales tax subsidies. Again, we find very few combinations of behavioral parameters that justify consumer subsidies in the presence of a Pigouvian tax

Although heterogeneity in behavioral biases in the perceptions of energy costs can motivate combining fiscal instruments in theory, we found that in practice, it would be optimal to do so for very few cases. A bias-adjusted Pigouvian tax alone does best in maximizing social welfare.

7.3. The Unintended Consequences of Certification

In many contexts, subsidy or tax instruments are coupled with information instruments, such as tobacco warning labels and cigarette taxes or electric vehicle tax credits and fuel economy labels. Such information disclosure requirements could influence the design and performance of fiscal instruments. In our economic environment, we have estimated the interaction effect between ES and energy operating costs—the coefficient λ in our choice model. Our positive λ estimates illustrate variation in behavioral biases across different segments of consumers and across products. Specifically, ES-certified products are subject to larger biases relative to noncertified products. As shown in Section 3, product-specific biases (i.e., $m_{kj} \neq m_{ki}$ for some j and i) can have ambiguous effects on the direction of the adjustment for the bias-adjusted policies, even if all consumers undervalue tax-inclusive energy operating costs across all products (i.e., $m_{kj} < 1 \forall k$ and j). Since ES certification further biases the perception of energy costs, there is now a trade-off between internalizing the externality and the behavioral biases. A tax that increases the price of energy will induce a substitution toward the more energy-efficient products that tend to be ES-certified. This decreases the size of the externality, but increases the behavioral biases consumers are subject to. Depending on the relative importance of the externality versus the behavioral biases, the optimal Pigouvian tax could fall below the marginal external damage and even be positive.

In Table 3 (first panel), we present the results for the various policies described in Subsection 7.1, but now account for the effect of the ES certification on the perception of energy operating costs. The optimal bias-adjusted Pigouvian tax is now close to zero, it is never optimal to offer rebates, and the optimal sales tax for ES products is higher than for non-ES products. The relative size of the behavioral biases dominates the externality; it is optimal to use fiscal policies to induce consumers to move away from ES products, which are prone to larger biases.

This effect is further magnified if we assume that the large willingness to pay for the ES label also reflects a bias. In the second panel of Table 3, we define consumer welfare such that the WTP affects decision utility, but not experienced utility. Under this assumption, the optimal policies are set so that consumers move further away from ES-certified products. The optimal Pigouvian tax is in fact a large subsidy of \$0.072/kWh, which incentivizes consumers to purchase more energy-intensive products. The very high sales tax for ES products relative to non-ES products, a difference of more than 25 percent, also reflects that it is optimal to first use the fiscal instruments to internalize the behavioral biases instead of the externality.

In these scenarios, combining the Pigouvian tax with other fiscal instruments can increase welfare by differentially targeting the externality and the behavioral biases. A welfare-maximizing policymaker could set the Pigouvian tax to \$0.066/kWh (or \$0.033/kWh in the lower-panel case) and charge a high sales tax for ES-certified products (15.87 percent or 27.08 percent, in each panel, respectively), and a low or negative sales tax on noncertified products (-3.90 percent or 4.12 percent). The negative sales tax for noncertified products and the positive Pigouvian tax show the trade-off between internalizing the behavioral biases and the externality.

These analyses abstract from firms' decisions. By increasing the willingness to pay for energy-efficient products, the ES program creates an incentive for firms to innovate along the energy efficiency frontier (Houde 2014). Generous subsidies for ES-certified products and a positive Pigouvian tax on energy could further spur innovation. A full accounting of the supply-side and dynamic effects of the program could find that ES-based subsidies are welfare enhancing.

7.4. Distributional Effects

Across all scenarios, we found that the aggregate welfare effects tended to be small—on the order of less than \$30 per consumer.²³ Given the size of the US refrigerator market (about 10 million refrigerators sold per year), these values are still economically significant. The burden of a particular policy on consumers may, however, be very large and crucially depends on how government revenues are accounted for. For example, a tax of $\tau = \$0.02/\text{kWh}$ reduces the consumer surplus in the range of \$109 to \$115. If all tax revenues could be redistributed lump-sum to these consumers, this burden would be almost completely offset. However, a perfectly targeted lump-sum transfer to consumers that just bought a refrigerator is unlikely to be feasible in practice. If we were to redistribute the tax revenue and target refrigerator consumers, we would have to design subsidies, which as our results

²³One exception is Panel I of Table 4, where the welfare losses are about \$70/consumer.

show, would be subject to various microfrictions and perverse effects. Therefore, large subsidies could compensate consumers, but have an overall negative impact on welfare. To illustrate, consider the scenarios where we offer a \$200 CFA rebate or a sales tax exemption of 6% for ES-certified products. In both cases, this will reduce the burden to consumers, but create an aggregate welfare loss.

The quantification of the burden faced by consumers also depends on our interpretation of microfrictions. In Table 4, we present the component of the consumer surplus that corresponds to the Leggett correction—that is, the component that quantifies the welfare effects due to the biases alone. Across income groups, the Leggett correction represents about half the overall change in consumer surplus. Note that the Leggett correction is larger for lower-income groups, which implies that failing to internalize behavioral biases would lead to a regressive outcome.

8. Conclusion

In this paper, we propose the concept of microfrictions to rationalize heterogeneity in the behavioral responses to fiscal instruments. Microfrictions refer to various phenomena that imply that a \$1 change induced by a fiscal instrument does not induce the same behavioral response as a \$1 change in relative prices. We show that for the design of corrective policies, it is important to distinguish microfrictions in two broad categories: transaction costs and behavioral biases. When microfrictions reflect behavioral biases, then the classic Pigouvian result of environmental taxation must be revisited. Heterogeneity in microfrictions across instruments and the population of consumers have important implications in deciding whether multiple fiscal instruments should be combined to internalize externalities and behavioral biases.

We apply our framework to the energy context and estimate the behavioral responses to different energy fiscal policies. We quantify the microfrictions associated with energy operating costs, rebates, and sales taxes and show that there is substantial heterogeneity across both the type of instruments and consumers. Focusing on the US refrigerator market, we find that lower-income households are more prone to microfrictions, relative to higher-income households, when processing energy operating costs, but appear to face less frictions in claiming rebates. For sales taxes, the highest-income group responds more, suggesting that it is less prone to behavioral issues associated with this particular instrument (e.g., salience).

We finally simulate different optimal energy fiscal policies aimed at energy-use externalities while accounting for microfrictions. We find that although the nature and heterogeneity of microfrictions

associated with energy operating costs could justify combining a tax with a subsidy, it is rarely optimal to do so in practice. Moreover, we find that the presence of a certification program such as Energy Star interacts with energy fiscal policies and can lead to unintended consequences. In essence, the coarseness of the ES program can exacerbate behavioral biases and complicate the design of a Pigouvian instrument.

References

- Allcott, H. (2013). The welfare effects of misperceived product costs: data and calibrations from the automobile market. *American Economic Journal: Economic Policy*, 5(3):30–66.
- Allcott, H. and Knittel, C. (2017). Are consumers poorly-informed about fuel economy? evidence from two experiments. Working Paper 23076, National Bureau of Economic Research.
- Allcott, H., Mullainathan, S., and Taubinsky, D. (2014). Energy policy with externalities and internalities. *Journal of Public Economics*, 112:72–88.
- Allcott, H. and Wozny, N. (2014). Gasoline prices, fuel economy, and the energy paradox. *Review of Economics and Statistics*, 96(5):779–795.
- Anderson, S. T., Kellogg, R., Sallee, J. M., and Curtin, R. T. (2011). Forecasting gasoline prices using consumer surveys. *American Economic Review*, 101(3):110–114.
- Ariely, D., Bracha, A., and Meier, S. (2009). Doing good or doing well? image motivation and monetary incentives in behaving prosocially. *American Economic Review*, 99(1):544–55.
- Baker, S. R., Johnson, S., and Kueng, L. (2017). Shopping for lower sales tax rates. Working Paper 23665, National Bureau of Economic Research.
- Baranzini, A. and Weber, S. (2013). Elasticities of gasoline demand in switzerland. *Energy policy*, 63:674–680.
- Bernheim, B. D. and Rangel, A. (2009). Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics. *Quarterly Journal of Economics*, 124(1):51–104.
- Boatwright, P., Kalra, A., and Zhang, W. (2008). Research note—should consumers use the halo to form product evaluations? *Management Science*, 54:217–223.
- Bovenberg, A. L. and Goulder, L. H. (2002). Environmental taxation and regulation. *Handbook of Public Economics*, 3:1471–1545.
- Busse, M. R., Knittel, C. R., and Zettelmeyer, F. (2013). Are consumers myopic? evidence from new and used car purchases. *American Economic Review*, 103(1):220–256.

- Chetty, R., Looney, A., and Kroft, K. (2007). Salience and taxation: theory and evidence. Working Paper 13330, National Bureau of Economic Research.
- Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: theory and evidence. *American Economic Review*, 99(4):1145–1177.
- Congdon, W. J., Kling, J. R., and Mullainathan, S. (2009). Behavioral economics and tax policy. *National Tax Journal*, 62(3):375.
- Currie, J. (2004). The take up of social benefits. Working Paper 10488, National Bureau of Economic Research.
- Davis, L. W. and Kilian, L. (2011). Estimating the effect of a gasoline tax on carbon emissions. *Journal of Applied Econometrics*, 26(7):1187–1214.
- De Groot, O. and Verboven, F. (2016). Subsidies and myopia in technology adoption: evidence from solar photovoltaic systems. Technical report, KU Leuven, Faculty of Economics and Business, Department of Economics.
- Dubois, P., Griffith, R., and O’Connell, M. (2017). The effects of banning advertising in junk food markets. *Review of Economic Studies*, Forthcoming.
- Einav, L., Jenkins, M., and Levin, J. (2012). Contract pricing in consumer credit markets. *Econometrica*, 80(4):1387–1432.
- Einav, L., Knoepfle, D., Levin, J., and Sundaresan, N. (2014). Sales taxes and internet commerce. *American Economic Review*, 104(1):1–26.
- Farhi, E. and Gabaix, X. (2015). Optimal taxation with behavioral agents. Working Paper 21524, National Bureau of Economic Research.
- Finkelstein, A. (2009). E-ztax: Tax salience and tax rates. *Quarterly Journal of Economics*, 124(3):969–1010.
- Frey, B. S. and Oberholzer-Gee, F. (1997). The cost of price incentives: An empirical analysis of motivation crowding-out. *American Economic Review*, 87(4):746–755.
- Gallagher, K. S. and Muehlegger, E. (2011). Giving green to get green? incentives and consumer adoption of hybrid vehicle technology. *Journal of Environmental Economics and Management*, 61(1):1–15.
- Gneezy, U. and Rustichini, A. (2000). Pay enough or don’t pay at all. *Quarterly Journal of Economics*, 115(3):791–810.

- Goldin, J. and Homonoff, T. (2013). Smoke gets in your eyes: cigarette tax salience and regressivity. *American Economic Journal: Economic Policy*, 5(1):302–336.
- Goulder, L. H. and Williams III, R. C. (2003). The substantial bias from ignoring general equilibrium effects in estimating excess burden, and a practical solution. *Journal of Political Economy*, 111(4):898–927.
- Goulder, L. H. and Williams III, R. C. (2012). The choice of discount rate for climate change policy evaluation. *Climate Change Economics*, 3(4).
- Handel, B. R. and Kolstad, J. T. (2015). Health insurance for “humans”: information frictions, plan choice, and consumer welfare. *American Economic Review*, 105(8):2449–2500.
- Hausman, J. A. (1979). Individual discount rates and the purchase and utilization of energy-using durables. *The Bell Journal of Economics*, pages 33–54.
- Houde, S. (2014). Bunching with the stars: How firms respond to environmental certification. Technical report.
- Houde, S. (2017). How consumers respond to environmental certification and the value of energy information. *The RAND Journal of Economics*, Forthcoming.
- Houde, S. and Aldy, J. E. (2017). Consumers’ response to state energy efficient appliance rebate programs. *American Economic Journal: Economic Policy*, 9(4):227–255.
- Huse, C. and Koptuyug, N. (2017). Taxes vs. standards as policy instruments: Evidence from the auto market. Technical report, Working paper, Stockholm School of Economics.
- Ketcham, J. D., Kuminoff, N. V., and Powers, C. A. (2016). Estimating the heterogeneous welfare effects of choice architecture: An application to the medicare prescription drug insurance market. Working Paper 22732, National Bureau of Economic Research.
- Lanz, B., Wurlod, J.-D., Panzone, L., and Swanson, T. (2017). The behavioral effect of pigovian regulation: Evidence from a field experiment. *Journal of Environmental Economics and Management*, Forthcoming.
- Leggett, C. G. (2002). Environmental valuation with imperfect information the case of the random utility model. *Environmental and Resource Economics*, 23(3):343–355.
- Li, S., Linn, J., and Muehlegger, E. (2014). Gasoline taxes and consumer behavior. *American Economic Journal: Economic Policy*, 6(4):302–342.
- Newell, R. G. and Siikamäki, J. (2015). Individual time preferences and energy efficiency. *American Economic Review*, 105(5):196–200.

- Parry, I. W. and Bento, A. M. (2000). Tax deductions, environmental policy, and the double dividend hypothesis. *Journal of environmental economics and management*, 39(1):67–96.
- Pigou, A. C. (1920). *The economics of welfare*, 4th. London: MacMillan.
- Reynaert, M. and Sallee, J. M. (2016). Corrective policy and goodhart’s law: The case of carbon emissions from automobiles. Working Paper 22911, National Bureau of Economic Research.
- Rivers, N. and Schaufele, B. (2015). Salience of carbon taxes in the gasoline market. *Journal of Environmental Economics and Management*, 74:23–36.
- Sahoo, A., Sawe, N., et al. (2015). The heterogeneous effects of eco-labels on internalities and externalities. Technical report.
- Sexton, S. (2015). Automatic bill payment and salience effects: Evidence from electricity consumption. *Review of Economics and Statistics*, 97(2):229–241.
- Small, K. A. and Rosen, H. S. (1981). Applied welfare economics with discrete choice models. *Econometrica*, 49(1):105–130.
- Taubinsky, D. and Rees-Jones, A. (2016). Attention variation and welfare: theory and evidence from a tax salience experiment. Working Paper 22545, National Bureau of Economic Research.
- Ward, D. O., Clark, C. D., Jensen, K. L., Yen, S. T., and Russell, C. S. (2011). Factors influencing willingness-to-pay for the ENERGY STAR® label. *Energy Policy*, 39(3):1450–1458.

9. Tables and Figures




FIGURE 1. Temporal and Cross-Store Variation in Promotional Price, Brand A

Notes: The red line shows the normalized prices of the nine most popular models offered by Brand A. The gray shaded area corresponds to the 25th and 75th percentile of the normalized price distribution. The blue line is the median price after controlling for brand-week-of-sample fixed effects. Source: Houde (2017).






FIGURE 2. Interpretation of Behavioral Responses to Electricity Cost, Rebates, and Sales Tax, Table 1


Notes: Panels (a) and (b) plot the ratio $\theta_i / \left(\eta_i \cdot \rho_i \cdot \frac{1 - \rho_i^L}{1 - \rho_i} \right)$ for a discount rate equal to 5%. Panels (c)-(f) show the ratio of the coefficient for a given type of fiscal instrument divided by the marginal value of income (i.e., absolute value of the coefficient on price). A ratio that takes a value of one means that the behavioral response to the fiscal instrument is the same as the response to the price.




(a) Pigou with CFA



(b) Pigou with CFA



(c) Pigou with ES Sales Tax



(d) Pigou with ES Sales Tax

FIGURE 3. Sensitivity with Respect to Behavioral Responses to CFA Rebate or Sales Tax

Notes: Each point represents the optimal policy for a particular set of behavioral parameters where the level and the gradient of the behavioral responses have been modified. In panels (a) and (b), sensitivity with respect to the behavioral responses to CFA rebates is investigated. The two panels show the optimal Pigouvian tax combined with a CFA rebate. For most values of the behavioral parameters pertaining to CFA rebates, the optimal CFA rebate is zero. In panels (c) and (d), sensitivity with respect to the behavioral responses to sales taxes is investigated. Panel (d) shows the difference between the sales tax targeting ES products and non-ES products.

TABLE 1. Estimation Results: Conditional Logit

	<\$30,000	≥\$30,000 <\$50,000	≥\$50,000 <\$75,000	≥\$75,000 <\$100,000	≥\$100,000 <\$150,000	≥\$150,000
Retail Price	-0.42*** (0.01)	-0.41*** (0.01)	-0.38*** (0.01)	-0.36*** (0.01)	-0.35*** (0.01)	-0.30*** (0.01)
Elec. Cost	-2.68*** (0.33)	-2.68*** (0.32)	-2.69*** (0.31)	-2.65*** (0.30)	-2.72*** (0.31)	-2.34*** (0.31)
Elec. Cost X ES	1.06*** (0.24)	0.93*** (0.23)	0.91*** (0.22)	0.78*** (0.22)	0.73*** (0.22)	0.47* (0.22)
Sales Tax	-0.30** (0.10)	-0.23* (0.09)	-0.25** (0.08)	-0.21** (0.08)	-0.15* (0.07)	-0.25*** (0.07)
Sales Tax X Tax Holiday	-0.01 (0.33)	-0.15 (0.29)	0.002 (0.27)	-0.19 (0.28)	-0.14 (0.24)	-0.08 (0.22)
Utility Rebate	0.03 (0.04)	0.01 (0.04)	-0.01 (0.04)	0.00 (0.04)	-0.01 (0.04)	-0.03 (0.03)
CFA Rebate: During Program	0.12*** (0.02)	0.13*** (0.02)	0.08*** (0.02)	0.06** (0.02)	0.06** (0.02)	0.02 (0.02)
CFA Rebate: 2 Months After	0.14*** (0.03)	0.04 (0.03)	0.10*** (0.03)	0.05 (0.03)	0.01 (0.03)	0.02 (0.03)
CFA Rebate: 2 Months Before	-0.05 (0.04)	-0.05 (0.04)	0.002 (0.04)	0.01 (0.04)	-0.002 (0.04)	-0.001 (0.04)

Notes: All variables are measures in hundreds of dollars. The model includes product fixed effects, state-ES fixed effects, and interaction terms between demographics and attributes. The consideration set for each consumer is restricted to refrigerator models that are within 5 cubic feet of the size purchased and offered in the same zip code as where the purchase was made. Standards errors are in parentheses. * ($p < 0.05$), ** ($p < 0.01$), *** ($p < 0.001$)

TABLE 2. Comparison of Fiscal Instruments

Policy Scenario	Optimal Policy	Social Welfare (\$/consumer)	Policy Ranking
Pigou tax: no adjustment	$\tau = 0.020\$/\text{kWh}$	1.791	2
Bias-adjusted Pigou tax	$\tau = 0.104\$/\text{kWh}$	5.049	1
CFA rebate	$S^{CFA} = \$50$	0.047	8
Mean-tested CFA rebate	$S^{CFA} = [\$32, \$42, \$61, \$74, \$57, \$36]$	0.050	4
ES Sales tax	$T_{ES}^s = 4.62\%, T_{nonES}^s = 6.43\%$	0.575	3
Pigou tax with CFA rebate	$\tau = 0.104\ \$/\text{kWh}, S^{CFA} = \1	5.049	1
Pigou tax with ES sales tax	$\tau = 0.104\ \$/\text{kWh}, T_{ES}^s = -0.01\%, T_{nonES}^s = -0.08\%$	5.049	1*

Notes: For all policy scenarios with subsidies, the ES certification is used as the eligibility criterion. All policies with a ranking of 1 differ with respect to the aggregate social welfare, but the differences are very small. The combination of the Pigouvian tax with the ES sales tax has the highest social welfare. For policies with rebates, transaction costs are not accounted for in the calculation of social welfare. The welfare estimates for these scenarios are then an upper bound.

TABLE 3. Comparison of Fiscal Instruments: Role of Welfare Interpretation of ES

Policy Scenario	Optimal Policy	Social Welfare (\$/consumer)	Policy Ranking
Welfare definition: $ES \times ElecCost \neq 0$			
Pigou tax: no adjustment	$\tau = 0.020\$/kWh$	-0.403	5
Bias-adjusted Pigou tax	$\tau = -0.004\$/kWh$	0.012	3
CFA rebate	$S^{CFA} = \$0$	0	4
ES Sales tax	$T_{ES}^s = 13.40\%$, $T_{nonES}^s = 2.10\%$	4.548	2
Pigouvian tax with CFA rebate	$\tau = -0.004\$/kWh$, $S^{CFA} = \$0$	0.012	3
Pigou tax with ES sales tax	$\tau = 0.066\$/kWh$, $T_{ES}^s = 15.87\%$, $T_{nonES}^s = -3.90\%$	6.428	1
Welfare definition: $ES \times ElecCost \neq 0$ and WTP for ES acts as a bias			
Pigou tax: no adjustment	$\tau = 0.020\$/kWh$	-2.513	5
Bias-adjusted Pigou tax	$\tau = -0.072\$/kWh$	4.483	3
CFA rebate	$S^{CFA} = \$0$	0	4
ES Sales tax	$T_{ES}^s = 25.15\%$, $T_{nonES}^s = -1.41\%$	19.044	2
Pigouvian tax with CFA rebate	$\tau = -0.072\$/kWh$, $S^{CFA} = \$0$	4.483	3
Pigou tax with ES sales tax	$\tau = 0.033\$/kWh$, $T_{ES}^s = 27.08\%$, $T_{nonES}^s = -4.12\%$	19.606	1

Notes: For all policy scenarios with subsidies, the ES certification is used as the eligibility criterion. Transaction costs are not accounted for in the calculation of social welfare. The welfare estimates for scenarios with rebates are then an upper bound. In the first panel, the welfare measure accounts for the effect of the ES label on electricity costs, which is captured by the parameter λ in the choice model and acts as a bias. The WTP for the label, which is captured by the interaction of an ES dummy with state fixed effects in the choice model, is considered to be preferences. In the second panel, both the parameter λ and the WTP for the ES label act as biases.

TABLE 4. Welfare Decomposition

	$\tau = 0.02\$/\text{kWh}$	$\tau = 0.02\$/\text{kWh}$ $S^{CFA} = \$200$	$\tau = 0.02\$/\text{kWh}$ $T_{ES}^S = -6\%, T_{nonES}^S = 0$
Gvt. Revenues	110.6	79.5	49.3
Ext. Costs	-0.5	-0.5	-0.3
With biases in energy costs and sales tax, ES \times Elec \neq 0, WTP for ES is preference			
SW	-0.6	-4.9	-6.8
CS: <\$30k	-109.8	-75.6	-66.7
CS: \$30k-\$50k	-110.7	-68.8	-60.7
CS: \$50k-\$75k	-111.6	-82.8	-56.3
CS: \$75k-\$100k	-112.5	-91.0	-52.8
CS: \$100k-\$150k	-112.7	-88.1	-50.8
CS: >\$150k	-112.9	-103.5	-52.7
Leggett: <\$30k	-64.1	-70.7	-56.6
Leggett: \$30k-\$50k	-61.7	-67.7	-41.8
Leggett: \$50k-\$75k	-58.4	-62.2	-41.4
Leggett: \$75k-\$100k	-53.1	-55.6	-29.9
Leggett: \$100k-\$150k	-47.8	-50.6	-13.3
Leggett: >\$150k	-43.0	-43.9	-35.3
With biases in energy costs and sales tax, ES \times Elec \neq 0, WTP for ES acts as a bias			
SW	-2.8	-10.2	-13.0
CS: <\$30k	-111.8	-81.7	-72.3
CS: \$30k-\$50k	-112.6	-74.9	-65.3
CS: \$50k-\$75k	-113.8	-88.2	-62.2
CS: \$75k-\$100k	-115.1	-96.5	-59.9
CS: \$100k-\$150k	-115.3	-94.3	-57.0
CS: >\$150k	-114.5	-106.3	-60.1
Leggett: <\$30k	-66.2	-76.7	-62.2
Leggett: \$30k-\$50k	-63.5	-73.8	-46.5
Leggett: \$50k-\$75k	-60.6	-67.5	-47.4
Leggett: \$75k-\$100k	-55.6	-61.1	-37.0
Leggett: \$100k-\$150k	-50.4	-56.8	-19.5
Leggett: >\$150k	-44.6	-46.7	-42.7

Notes: All values are in dollars. The government revenue is the discounted sum of the tax revenue, minus the rebate, over the lifetime of the refrigerator. We assume a lifetime of 18 years. All flows of revenues and utility are discounted at a 5% discount rate. SW refers to social welfare, CS refers to consumer surplus, and Leggett is the adjustment term that accounts for biases in the calculation of the consumer surplus (see Equation 2).

Appendix A. Theory: Additional Results and Proofs

Proof of Proposition 1:

For both cases, we define social welfare as the sum of the changes in consumer surplus, government revenue, and externality costs.

For the case where microfrictions are transaction costs, suppose that the transaction costs can take only two values: $C = 0$ or $C = \infty$. When $C = 0$, then it is optimal to set $\tau = \phi$. When the transaction costs are too high, consumers do not change their behavior in response to the tax. They pay the tax, and the tax revenue collected can fully compensate for the tax consumers pay. Consequently, for $C = \infty$, social welfare remains unaffected by the level of the tax. So long as the distribution F places some weight on $C = 0$, then $\tau = \phi$. When transaction costs fall within $0 < C < \infty$, the consumer responds only if the tax is "high enough." If a consumer responds to the tax, he will fully account for the tax, and the behavioral response to the tax is the same as in a model without transaction costs. Social welfare is thus still maximized at $\tau = \phi$, but the aggregate change in social welfare must be scaled down by the constant cost C . As long as the cost C is not a function of the level of the tax, then $\tau = \phi$.²⁴

For the case where microfrictions are behavioral biases, the expression for social welfare after a tax τ is levied is given by

$$SW(\tau) = \frac{1}{\eta_k} \cdot \left[\ln \sum_j^J \exp(U_{kj}(\tau)) + \sum_j^J \sigma_{kj}(U_{kj}^E(\tau) - U_{kj}(\tau)) \right] \\ + (\tau - \phi) \sum_j^J \sigma_{kj}(\tau) \cdot E_j,$$

where $U_{kj}(\tau) = V_{kj} - \eta_k(P_j + m(P^e + \tau)E_j)$ and $U_{kj}^E(\tau) = V_{kj} - \eta_k(P_j + (P^e + \tau)E_j)$. Taking the derivative of $SW(\tau)$ with respect to τ and rearranging, we obtain

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

²⁴Farhi and Gabaix (2015) consider a model of endogenous attention where the size of the attention cost is a function of the potential bias. They show (Proposition 6.1) that if the cost is included in the social welfare, this lowers the bias-adjusted tax. We show that if the attention cost is a constant that affects social welfare, no adjustment is required.

Proof of Lemma 1:

For a small tax τ , the aggregate change in energy is always (weakly) negative: $\Delta_k^{\tau,e} = \sum_j \frac{\partial \sigma_{kj}}{\partial \tau} E_j < 0$. Note that $m_k \leq 1$ implies $(1 - m_k) \leq 1$. Therefore, the ratio $\frac{\sum_k \alpha_k (1 - m_k) \Delta_k^{\tau,e}}{\sum_k \alpha_k \Delta_k^{\tau,e}}$ is always between 0 and 1. The denominator $1 - \frac{\sum_k \alpha_k (1 - m_k) \Delta_k^{\tau,e}}{\sum_k \alpha_k \Delta_k^{\tau,e}}$ is thus smaller than one and the ratio $\frac{\sum_k \alpha_k (1 - m_k) \Delta_k^{\tau,e}}{1 - \frac{\sum_k \alpha_k (1 - m_k) \Delta_k^{\tau,e}}{\sum_k \alpha_k \Delta_k^{\tau,e}}}$ is always positive. This implies $\tau > \phi$.

Additional lemma:

Lemma 2. *Using only data on equilibrium quantities, transaction costs can be observationally equivalent to behavioral biases when $m < 1$.*

Proof. Define $s_{kj}^{Bias}(m)$ as the market share for product j when consumer k is subject to behavioral biases m , and $s_{kj}^{TC}(F)$ as the market share when he is subject to transaction costs that follow the distribution F . For a given distribution F , the following inequalities must hold:

$$\frac{\partial s_{kj}^{Bias}(0)}{\partial I} < \frac{\partial s_{kj}^{TC}(F)}{\partial I} < \frac{\partial s_{kj}^{Bias}(1)}{\partial I}.$$

Given that $\frac{\partial^2 s_{kj}^{Bias}(m)}{\partial I \partial m} > 0 \forall m$, there must exist a m such that $\frac{\partial s_{kj}^{Bias}(m)}{\partial I} = \frac{\partial s_{kj}^{TC}(F)}{\partial I}$.

Transaction costs and behavioral biases can be observationally equivalent only if $m < 1$, because the inequality $\frac{\partial s_{kj}^{TC}(F)}{\partial I} < \frac{\partial s_{kj}^{Bias}(1)}{\partial I}$ always holds for any distribution F . \square

Appendix B. Interpretation of Coefficients on Rebates

We interpret each of the coefficients on rebates (utility or states) as the probability of claiming rebates times the marginal utility of income. In a linear model for the choice probabilities, this interpretation is fully consistent with a structural model where the decision to claim rebates is explicitly modeled. We show this below. With a nonlinear model such as the conditional logit, this interpretation of the coefficient on rebates is not fully consistent with the structural model. However, this interpretation holds locally around the estimates, as we also show below.

Consider the following general choice model that explicitly models the decision to claim rebates. Suppose again that this decision is not observed, and α_i denotes the probability that consumer i claims a rebate. The observable choice probabilities for consumer i for product j thus correspond to a latent choice model taking the following form:

$$(13) \quad F_{ij} = \alpha_i F_{ij}^R + (1 - \alpha_i) F_{ij}^{NR},$$

where F_{ij}^R is the choice probability when consumer i claims a rebate, and F_{ij}^{NR} is the choice probability when consumer i does not claim a rebate. If the functions F_{ij}^R and F_{ij}^{NR} are linear, we can readily see that the structural model corresponds to a reduced-form model where the coefficient on rebates, say ψ_i , is interpreted as the product of α_i and the behavioral response to rebates η_i . Without loss of generality, suppose that we are in a choice environment where only prices and rebates matter. We then have $F_{ij}^R = -\eta_i \cdot Price_j + \eta_i \cdot Rebate_j$, and $F_{ij}^{NR} = -\eta_i \cdot Price_j$. Therefore, we have

$$(14) \quad \begin{aligned} F_{ij} &= \alpha_i \cdot (-\eta_i \cdot Price_j + \eta_i \cdot Rebate_j) + (1 - \alpha_i) \cdot (-\eta_i \cdot Price_j) \\ &= -\eta_i \cdot Price_j + \alpha_i \cdot \eta_i \cdot Rebate_j \\ &= -\eta_i \cdot Price_j + \psi_i \cdot Rebate_j. \end{aligned}$$

In a nonlinear model, the above holds only for a local linear approximation of the structural model F_{ij} around the parameter on rebates at its estimated value.

Appendix C. Additional Information: Rebate Programs, Sales Tax Exemptions/Holidays, Energy Star Advertising






FIGURE 4. Cross-Sectional Variation in Sales Tax Rates


Notes: This map is from Einav et al. (2014) and shows the (population-weighted) average sales tax rate in the United States as of January 1, 2010.



(a) Refrigerators



(b) Clothes Washers



(c) Dishwashers

FIGURE 5. Average Price vs. Rebate Amount

Each panel shows the average price of the appliance purchased (in white) and the average rebate amount claimed (in red). States with no average price but a positive rebate amount are states where program managers did not collect price information. States where both price and rebate information are missing did not offer rebates for this particular appliance.




FIGURE 6. Energy Star and Publicity Intensity

Notes: Each panel shows the publicity intensity of the Energy Star program as classified by the EPA. All maps and analysis are taken from the “National Awareness of Energy Star” yearly reports published by EPA.

TABLE 5. Summary Statistics: Cash for Appliances

Product	# of States Offering Re- bates	# of Claims	Amount Dis- tributed (\$M)	Average Price Paid (\$)	Average Rebate Claimed (\$)	Max Rebate Claimed (\$)
Air Conditioners	30	70,781	25.6	4,511	361	3,812
Boilers	18	7,678	4.0	5,516	518	4,036
Clothes Washers	43	580,863	62.1	698	107	1,034
Dishwashers	37	316,117	26.6	543	84	47,751
Electric Water Heaters	25	3,267	1.0	1,636	307	1,816
Freezers	26	24,312	2.5	579	103	1,500
Furnaces	34	76,469	30.9	5,772	404	3,227
Gas/Propane Water Heaters	30	15,766	2.1	703	130	1,742
Gas/Propane Water Heaters (Tankless)	31	11,140	3.0	2,266	267	1,223
Heat Pumps	26	47,470	23.6	6,403	497	4,400
Refrigerators	44	613,561	78.8	1,112	128	7,085
Solar Water Heaters	15	634	0.8	7,961	1,308	2,500
Total		1,768,058	260.9			

Notes: Data collected by program administrators and provided to the Department of Energy. Excludes US territories.

TABLE 6. Rebate Amounts and Eligibility Criteria for Cash for Appliances

	Refrigerators		Clothes Washers		Dishwashers	
	Rebate	Criteria	Rebate	Criteria	Rebate	Criteria
AK	300-600	ES rural/nonrural	300-600	ES	300-600	ES rural/nonrural
AL	150	ES	100	ES	75	ES
AR	275	ES	225	ES	-	
AZ	200-300	ES	125-200	ES & Above ES	75-125	ES & Above ES
CA	200	ES	100	Above ES	100	Above ES
CO	50-100	ES	75	ES	50	Above ES
CT	50	ES	100	Above ES	-	
DE	100	ES	75	ES	75	ES
FL	20%	ES	20%	ES	20%	ES
GA	50	ES	50-99	ES & Above ES	50-99	ES & Above ES
HI	250	ES	-		-	
IA	200-500	ES	200	ES	200-250	ES & Above ES
ID	75	ES	75	ES	50	ES
IL	15%	ES	15%	ES	15%	ES
IN	-		-		-	
KS	700	ES	800	Above ES	400	Above ES
KY	50	ES	100	ES	50	ES
LA	250	ES	100	ES	150	ES
MA	200	ES & Above ES	175	ES	250	ES & Above ES
MD	50	ES & Above ES	100	ES	-	
ME	-		-		-	
MI	50-100	ES & Above ES	50	ES	25-50	ES & Above ES
MN	100	ES	200	ES	150	ES
MO	250	ES	125	ES	125	ES
MS	75	ES	100-150	ES & Above ES	75-100	ES & Above ES
MT	100	ES	100	ES	50	ES
NC	15%	ES	100	ES	75 or 15%	ES
ND	150	ES	-		-	
NE	200	ES	100-200	ES & Above ES	150	Above ES
NJ	75-100		35	ES	25-50	ES & Above ES
NM	200	ES	200	ES	-	
NV	200	ES	150	ES	100	ES
NY	75-105		75-100	ES & Above ES	165	ES
OH	100	ES	150	ES	100	ES
OK	200	ES	200	ES	-	
OR	70%	ES	70%	ES	70%	ES
RI	150	ES	-		150	ES
SC	50	ES	100	ES	50	ES
SD	150	ES	100	ES	75	ES
TX	175-315	ES	100-225	ES & Above ES	85-185	ES
UT	-		75	ES		
VA	60	ES	75-350	ES & Above ES	50-275	ES & Above ES
VT	75	ES	150	ES	-	
WA	75	ES	150	ES	75	ES
WI	75	ES	100	ES	25	ES
WV	100	ES	50-75	Above ES	50-75	ES & Above ES
WY	-		100	ES	50	ES

Notes: Above ES means that criteria more stringent than ES were used. Alaska offered different rebate amounts for rural and nonrural residents.

TABLE 7. State Sales Tax Holidays for Appliances, 2008-2012

State	Year	Start Date	End Date	ES Requirement	Price Cap	Sales Tax Rate
GA	2008	10/2/08	10/5/08	Energy Star Qualified	1500	4.00%
GA	2009	10/1/09	10/4/09	Energy Star Qualified	1500	4.00%
GA	2012	10/5/12	10/7/12	Energy Star & Water Sense Qualified	1500	4.00%
LA	2008	8/1/08	8/2/08	none	2500	4.00%
LA	2009	8/7/09	8/8/09	none	2500	4.00%
LA	2010	8/6/10	8/7/10	none	2500	4.00%
LA	2011	8/5/11	8/6/11	none	2500	4.00%
LA	2012	8/3/12	8/4/12	none	2500	4.00%
MA	2008	8/16/08	8/17/08	none	2500	5.00%
MA	2010	8/14/10	8/15/10	none	2500	6.25%
MA	2011	8/13/11	8/14/11	none	2500	6.25%
MA	2012	8/11/12	8/12/12	none	2500	6.25%
MD	2011	2/19/11	2/21/11	Energy Star Qualified	none	6.00%
MD	2012	2/18/12	2/20/12	Energy Star Qualified	none	6.00%
MO	2009	4/19/09	4/25/09	Energy Star Qualified	1500	4.23%
MO	2010	4/19/10	4/25/10	Energy Star Qualified	1500	4.23%
MO	2011	4/19/11	4/25/11	Energy Star Qualified	1500	4.23%
MO	2012	4/19/12	4/25/12	Energy Star Qualified	1500	4.23%
NC	2008	11/7/08	11/9/08	Energy Star Qualified	none	4.50%
NC	2009	11/6/09	11/8/09	Energy Star Qualified	none	5.75%
NC	2010	11/5/10	11/7/10	Energy Star Qualified	none	5.75%
NC	2011	11/4/11	11/6/11	Energy Star Qualified	none	5.75%
NC	2012	11/2/12	11/4/12	Energy Star Qualified	none	4.75%
SC	2008	10/1/08	10/31/08	Energy Star Qualified	2500	6.00%
TX	2009	5/23/09	5/25/09	Energy Star Qualified	6000/2000	6.25%
TX	2010	5/29/10	5/31/10	Energy Star Qualified	6000/2000	6.25%
TX	2011	5/28/11	5/30/11	Energy Star Qualified	6000/2000	6.25%
TX	2012	5/26/12	5/28/12	Energy Star Qualified	6000/2000	6.25%
VA	2008	10/10/08	10/13/08	Energy Star Qualified	2500	5.00%
VA	2009	10/9/09	10/12/09	Energy Star & Water Sense Qualified	2500	5.00%
VA	2010	10/8/10	10/11/10	Energy Star & Water Sense Qualified	2500	5.00%
VA	2011	10/7/11	10/10/11	Energy Star Qualified	2500	5.00%
VA	2012	10/5/12	10/8/12	Energy Star & Water Sense Qualified	2500	5.00%
VT	2008	7/12/08	7/18/08	Energy Star Qualified	2000	6.00%
VT	2009	8/22/09	8/22/09	none	2000	6.00%
VT	2010	3/6/10	3/6/10	none	2000	6.00%
WV	2008	9/1/08	9/7/08	Energy Star Qualified	2500	6.00%
WV	2009	9/1/09	11/30/09	Energy Star Qualified	2500	6.00%
WV	2010	9/1/10	11/30/10	Energy Star Qualified	2500	6.00%

Notes: Missouri restricted its tax holiday to clothes washers, water heaters, dishwashers, air conditioners, furnaces, refrigerators, and freezers. Maryland's restrictions apply to air conditioners, clothes washers and dryers, furnaces, heat pumps, boilers, solar water heaters, standard-size refrigerators, dehumidifiers, programmable thermostats, and compact fluorescent light bulbs. Texas' restrictions apply to air conditioners, clothes washers, ceiling fans, dehumidifiers, dishwashers, incandescent or fluorescent light bulbs, programmable thermostats, and refrigerators.

Appendix D. Additional Estimation Results

TABLE 8. Robustness Check I Conditional Logit: Brand \times Month FEs

	<\$30,000	\geq \$30,000 <\$50,000	\geq \$50,000 <\$75,000	\geq \$75,000 <\$100,000	\geq \$100,000 <\$150,000	\geq \$150,000
Retail Price	-0.42 sym*** -0.01	-0.41 sym*** -0.01	-0.38 sym*** -0.01	-0.35 sym*** -0.01	-0.34 sym*** -0.01	-0.30 sym*** -0.01
Elec. Cost	-2.64 sym*** (0.33)	-2.62 sym*** (0.32)	-2.63 sym*** (0.31)	-2.61 sym*** (0.30)	-2.69 sym*** (0.31)	-3.06 sym*** (0.32)
Elec. Cost X ES	0.96 sym*** (0.24)	0.86 sym*** (0.23)	0.84 sym*** (0.22)	0.74 sym*** (0.22)	0.66 sym** (0.22)	0.63 sym** (0.22)
Sales Tax	-0.30 sym** (0.10)	-0.24 sym** (0.09)	-0.24 sym** (0.08)	-0.21 sym** (0.08)	-0.16 sym* (0.07)	-0.25 sym*** (0.07)
Sales Tax X Tax Holiday	-0.05 (0.33)	-0.17 (0.29)	-0.01 (0.28)	-0.20 (0.28)	-0.15 (0.24)	-0.09 (0.22)
Utility Rebate	0.03 (0.04)	0.01 (0.04)	-0.01 (0.04)	0.01 (0.04)	0.00 (0.04)	-0.03 (0.03)
CFA Rebate: During Program	0.11 sym*** (0.02)	0.12 sym*** (0.02)	0.07 sym*** (0.02)	0.05 sym* (0.02)	0.06 sym** (0.02)	0.02 (0.02)
CFA Rebate: 2 Months After	0.14 sym*** (0.03)	0.04 (0.03)	0.09 sym** (0.03)	0.04 (0.03)	0.00 (0.03)	0.01 (0.03)
CFA Rebate: 2 Months Before	-0.05 (0.04)	-0.04 (0.04)	0.00 (0.04)	0.02 (0.04)	0.01 (0.04)	0.01 (0.04)

Notes: All variables are in hundreds of dollars. The model includes brand-month fixed effects in addition to product fixed effects, state-ES fixed effects, and interaction terms between demographics and attributes. The consideration set for each consumer is restricted to refrigerator models that are within 5 cu. ft. of the size purchased and offered in the same zip code as where the purchase was made. Standards errors are in parentheses. * ($p < 0.05$), ** ($p < 0.01$), *** ($p < 0.001$)

TABLE 9. Robustness Check II Conditional Logit: Restricted Consideration
Set ± 0.5 cu. ft.

	<\$30,000	\geq \$30,000 <\$50,000	\geq \$50,000 <\$75,000	\geq \$75,000 <\$100,000	\geq \$100,000 <\$150,000	\geq \$150,000
Retail Price	-0.45 sym*** (0.01)	-0.44 sym*** (0.01)	-0.40 sym*** (0.01)	-0.38 sym*** (0.01)	-0.36 sym*** (0.01)	-0.33 sym*** (0.01)
Elec. Cost	-3.86 sym*** (0.51)	-3.94 sym*** (0.48)	-2.83 sym*** (0.46)	-3.18 sym*** (0.44)	-3.27 sym*** (0.45)	-3.82 sym*** (0.46)
Elec. Cost X ES	0.47 (0.28)	0.58 sym* (0.27)	0.64 sym* (0.26)	0.74 sym** (0.26)	0.65 sym* (0.26)	0.58 sym* (0.26)
Sales Tax	-0.25 (0.14)	-0.07 (0.14)	-0.21 (0.12)	-0.26 sym* (0.11)	-0.17 (0.11)	-0.22 sym* (0.10)
Sales Tax X Tax Holiday	-0.36 (0.42)	-0.39 (0.32)	0.18 (0.31)	-0.12 (0.33)	-0.34 (0.26)	0.09 (0.29)
Utility Rebate	0.03 (0.05)	0.02 (0.05)	-0.01 (0.05)	0.02 (0.04)	0.07 (0.04)	-0.03 (0.04)
CFA Rebate: During Program	0.09 sym** (0.03)	0.13 sym*** (0.03)	0.08 sym*** (0.03)	0.05 sym* (0.03)	0.07 sym** (0.02)	0.04 (0.02)
CFA Rebate: 2 Months After	0.11 sym** (0.04)	0.04 (0.04)	0.06 (0.04)	0.02 (0.04)	0.05 (0.03)	0.03 (0.03)
CFA Rebate: 2 Months Before	-0.12 (0.05)	-0.12 (0.04)	-0.06 (0.05)	-0.01 (0.05)	-0.05 (0.05)	-0.05 (0.05)

Notes: All variables are in hundreds of dollars. The model includes product fixed effects, state-ES fixed effects, and interaction terms between demographics and attributes. The consideration set for each consumer is restricted to refrigerator models that are within 0.5 cu. ft. of the size purchased and offered in the same zip code as where the purchase was made. Standards errors are in parentheses. * ($p < 0.05$), ** ($p < 0.01$), *** ($p < 0.001$)

TABLE 10. Estimation Results: Conditional Logit w. Additional Heterogeneity

	<\$30,000	≥\$30,000 <\$50,000	≥\$50,000 <\$75,000	≥\$75,000 <\$100,000	≥\$100,000 <\$150,000	≥\$150,000
Retail Price	-0.44***	-0.37***	-0.42***	-0.38***	-0.42***	0.00
Educ. College	0.04*	0.01	0.01	0.03**	0.03*	0.00
Educ. Grad. School	0.07***	0.02	0.03	0.03*	0.06***	0.00
Age 35-50	-0.03	-0.01	0.02	0.01	0.01	0.00
Age 50-65	0.03	0.00	0.04	0.01	0.04*	0.00
Age 65+	-0.04	-0.03	0.03	0.03	0.03	0.00
Family w. Children	0.00	-0.06***	0.00	-0.01	-0.01	0.00
Democrats	-0.01	-0.01	-0.02	0.00	0.04**	0.00
Independent-Others	0.02	-0.03	0.01	0.01	0.03	0.00
Elec. Cost	-2.78***	-3.24***	-3.02***	-2.39***	-1.87***	0.00
Educ. College	0.07	-0.45	-0.18	0.05	-0.59*	0.00
Educ. Grad. School	-0.32	0.18	-0.41	-0.89**	-0.75*	0.00
Age 35-50	0.27	0.10	-0.22	0.02	-0.21	0.00
Age 50-65	-0.81**	-0.56	-1.06***	-0.69*	-0.65*	0.00
Age 65+	-1.06**	-0.61	-1.05**	-0.61	-0.55	0.00
Family w. Children	0.78***	1.09***	0.96***	0.69***	0.54***	0.00
Democrats	0.02	0.47	0.80*	-0.11	-0.50	0.00
Independent-Others	0.55	0.68	0.64	-0.15	-0.32	0.00
Elec. Cost X ES	1.35***	0.88***	0.88***	0.88***	0.78**	0.00
Educ. College	-0.06	-0.08	0.06	0.05	0.05	0.00
Educ. Grad. School	0.15	0.17	0.17*	0.05	0.07	0.00
Age 35-50	-0.08	0.10	0.12	-0.14	-0.08	0.00
Age 50-65	-0.01	0.08	0.00	-0.08	-0.03	0.00
Age 65+	0.00	0.27**	0.15	0.07	-0.01	0.00
Family w. Children	-0.08	-0.08	-0.05	-0.01	-0.01	0.00
Democrats	-0.17	0.08	0.02	-0.01	-0.02	0.00
Independent-Others	-0.17	0.10	0.02	-0.02	-0.02	0.00
Sales Tax	0.01	-1.13**	0.11	-0.10	0.71*	0.00
Educ. College	-0.50*	0.04	0.13	-0.25	-0.11	0.00
Educ. Grad. School	-0.52	0.18	-0.04	-0.07	-0.35	0.00
Age 35-50	0.51	0.25	-0.12	0.16	0.00	0.00
Age 50-65	-0.19	0.18	-0.41	0.11	-0.51	0.00
Age 65+	0.46	0.48	-0.38	-0.29	-0.31	0.00
Family w. Children	-0.12	0.66**	-0.13	0.06	0.08	0.00
Democrats	0.08	0.23	0.19	-0.01	-0.66**	0.00
Independent-Others	-0.42	0.53*	-0.24	-0.09	-0.47*	0.00
Sales Tax X Tax Holiday	0.70	1.90	-1.25	0.98	-2.06	0.00
Educ. College	0.78	-0.30	0.43	0.37	0.39	0.00
Educ. Grad. School	-1.08	-1.68	-0.99	-0.66	-0.38	0.00
Age 35-50	1.92	-0.60	-0.54	-0.57	-1.23	0.00
Age 50-65	1.85	-0.83	0.01 ⁵⁸	-0.54	-0.05	0.00
Age 65+	-0.57	-0.31	0.47	-1.19	1.14	0.00
Family w. Children	-1.15	-0.44	2.06**	0.08	1.82**	0.00
Democrats	-1.13	-1.09	0.40	-0.60	0.71	0.00
Independent-Others	-2.23	-1.24	-0.37	-0.97	1.78	0.00

Notes: All variables are in hundreds of dollars. Same set of fixed effects and controls as in Table 1 in the main text. Standards errors are not displayed. * ($p < 0.05$), ** ($p < 0.01$), *** ($p < 0.001$)

TABLE 11. Estimation Results: Conditional Logit w. Additional Heterogeneity (contd.)

	<\$30,000	≥\$30,000 <\$50,000	≥\$50,000 <\$75,000	≥\$75,000 <\$100,000	≥\$100,000 <\$150,000	≥\$150,000
Utility Rebate	-0.27*	0.15	-0.05	0.00	-0.09	0.00
Educ. College	0.08	-0.04	-0.01	-0.15*	0.03	0.00
Educ. Grad. School	-0.07	-0.34***	-0.08	-0.09	0.03	0.00
Age 35-50	0.15	0.00	-0.12	0.14	0.07	0.00
Age 50-65	0.13	0.00	0.07	0.06	-0.06	0.00
Age 65+	0.19	-0.01	0.02	0.09	0.14	0.00
Family w. Children	0.10	0.00	0.04	-0.02	-0.06	0.00
Democrats	0.05	-0.08	0.09	0.01	0.09	0.00
Independent-Others	0.17	-0.12	0.05	0.02	0.10	0.00
CFA Rebate: During Program	0.17	0.15	0.03	-0.10	0.12	0.00
Educ. College	0.07	-0.06	0.04	-0.04	-0.03	0.00
Educ. Grad. School	-0.03	0.06	0.03	0.03	-0.02	0.00
Age 35-50	0.04	-0.06	0.01	0.16*	-0.02	0.00
Age 50-65	-0.07	-0.03	0.03	0.16*	0.05	0.00
Age 65+	0.03	-0.09	0.04	0.14	-0.03	0.00
Family w. Children	0.01	0.03	0.04	-0.02	0.04	0.00
Democrats	-0.09	0.02	0.00	0.02	-0.07	0.00
Independent-Others	-0.08	0.05	-0.04	0.05	-0.11*	0.00
CFA Rebate: 2 Months After	0.20	-0.04	0.24	-0.07	0.25	0.00
Educ. College	-0.07	-0.12	0.04	0.16*	0.07	0.00
Educ. Grad. School	-0.10	-0.03	-0.18*	0.13	0.16*	0.00
Age 35-50	-0.05	0.04	-0.14	0.03	-0.26	0.00
Age 50-65	-0.02	-0.12	-0.08	-0.06	-0.38*	0.00
Age 65+	-0.06	0.00	-0.10	-0.15	-0.28	0.00
Family w. Children	-0.13*	0.11	0.06	0.01	0.06	0.00
Democrats	0.14	0.19*	-0.13	0.07	0.01	0.00
Independent-Others	0.07	0.12	-0.03	0.10	-0.04	0.00
CFA Rebate: 2 Months Before	-0.13	-0.26	0.10	0.01	-0.34*	0.00
Educ. College	-0.10	-0.06	0.06	0.10	0.07	0.00
Educ. Grad. School	-0.17	-0.05	0.01	-0.01	-0.02	0.00
Age 35-50	0.20	0.21	0.12	-0.02	0.15	0.00
Age 50-65	0.05	0.22	-0.08	-0.02	0.06	0.00
Age 65+	0.04	0.27*	-0.06	-0.03	0.03	0.00
Family w. Children	0.07	0.03	-0.16	-0.03	0.10	0.00
Democrats	-0.01	-0.02	-0.06	-0.04	0.25*	0.00
Independent-Others	0.06	0.05	0.00	0.04	0.21*	0.00

Notes: All variables are in hundreds of dollars. Same set of fixed effects and controls as is Table 1 in the main text. Standards errors are not displayed. * ($p < 0.05$), ** ($p < 0.01$), *** ($p < 0.001$)

Appendix E. Additional Results: Policy Analysis




FIGURE 7. Sensitivity with Respect to Behavioral Responses to Energy Operating Costs

Notes: Each point represents the optimal policy for a particular set of behavioral parameters where the level and the gradient of the behavioral responses to energy operating costs have been modified. Panels (a) and (b) simulate the optimal combination of a tax and a CFA rebate. For most values of the behavioral parameters pertaining to energy operating costs, the optimal CFA rebate is zero. In panels (c) and (d), the optimal combination of a tax and an ES-differentiated sales tax is investigated. Panel (d) shows the difference between the sales tax targeting ES products and non-ES products. The optimal ES sales tax exemption is zero for most scenarios.