Journal of Public Economics 112 (2014) 72–88

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

Energy policy with externalities and internalities ☆

Hunt Allcott [a,b,*], Sendhil Mullainathan [b,c], Dmitry Taubinsky [c]

[a] New York University, United States
[b] NBER, United States
[c] Harvard University, United States

ARTICLE INFO

Article history:
Received 23 September 2012
Received in revised form 27 December 2013
Accepted 1 January 2014
Available online 31 January 2014

JEL classification:
D03
D04
D11
H21
H22
H23
L51
L62
L97
Q41
Q48

Keywords:
Energy efficiency
Energy-using durables
Implied discount rates
Inattention
Internality taxes
Optimal taxation
Present bias
Pigouvian taxes

ABSTRACT

We analyze optimal policy when consumers of energy-using durables undervalue energy costs relative to their private optima. First, there is an Internality Dividend from Externality Taxes: aside from reducing externalities, they also offset distortions from underinvestment in energy efficiency. Discrete choice simulations of the auto market suggest that the Internality Dividend could more than double the social welfare gains from a carbon tax at marginal damages. Second, we develop the Internality Targeting Principle: the optimal combination of multiple instruments depends on the average internality of the consumers marginal to each instrument. Because consumers who undervalue energy costs are mechanically less responsive to energy taxes, the optimal policy will tend to involve an energy tax below marginal damages coupled with a larger subsidy for energy efficient products. Third, although the exact optimal policy depends on joint distributions of unobservables which would be difficult to estimate, we develop formulas to closely approximate optimal policy and welfare effects based on reduced form "sufficient statistics" that can be estimated by using field experiments or quasi-experimental variation in product prices and energy costs.

1. Introduction

Since a seminal paper by Hausman (1979), it has frequently been asserted that consumers "undervalue" energy costs relative to purchase prices when they choose between energy-using durable goods, perhaps because they are inattentive to or imperfectly informed about these costs. Although more empirical evidence is needed, this assertion would be consistent with findings that we are inattentive to other ancillary product costs such as sales taxes (Chetty et al., 2009), shipping and handling charges (Brown, Hossein, and Morgan 2010), and the out-of-pocket costs of insurance plans (Abaluck and Gruber, 2011). In the language of Herrnstein et al. (1993), undervaluation could cause consumers to impose "internalities" on themselves as they buy goods that are less energy efficient than they would choose in their private optima. Consumer undervaluation of energy costs has become an important policy issue: along with energy use externalities such as local air pollution and climate change, many policymakers use undervaluation to justify significant regulations such as Corporate Average Fuel Economy (CAFE) standards and billions of dollars in subsidies for energy efficient durables such as air conditioners and lightbulbs.

Despite the important implications, however, there is little formal guidance on how to set and evaluate energy policy in the presence of

☆ We thank Ryan Bubb, Meghan Busse, Lucas Davis, Liran Einav, Larry Goulder, Michael Greenstone, Garth Heutel, Ryan Kellogg, Ted O'Donoghue, Jim Sallee, Josh Schwartzstein, Glen Weyl, Nathan Wozny, the two anonymous referees, and the seminar participants at the 2011 ASSA meetings, Duke, the Environmental Defense Fund, Harvard, MIT, NBER, New York University, the Stanford Institute for Theoretical Economics, UC Santa Barbara, Universidad de los Andes, and Yale for the feedback and helpful conversations. We are grateful to the Sloan Foundation for the financial support of our research on the economics of energy efficiency. MATLAB code for replicating the analysis is available from Hunt Allcott's website.

* Corresponding author at: New York University Department of Economics, 19 West 4th St., 6th Floor, New York, NY 10012, United States.
E-mail addresses: Hunt.allcott@nyu.edu (H. Allcott), mullain@fas.harvard.edu (S. Mullainathan), taubinsk@fas.harvard.edu (D. Taubinsky).

0047-2727/$ – see front matter © 2014 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jpubeco.2014.01.004

both externalities and internalities. Three questions stand out. First, how do internalities affect the traditional logic of Pigouvian externality taxation? Second, what is the optimal combination of policy instruments to address externalities and internalities? Third, what are the key empirical objects that empiricists should estimate in order to set optimal policy and measure welfare impacts?

We analyze these three questions using both a theoretical model and numerical simulations. In our model, consumers choose between two energy-using durable goods, one of which is less energy-intensive than the other. Consumers then choose utilization, incur energy costs, and consume a numeraire good. When choosing between the two energy-using durables, some consumers misoptimize: while they should be indifferent between \$1 in purchase price and \$1 in energy costs because both equally affect consumption of the numeraire good, they undervalue or overvalue energy efficiency relative to their private optima. In addition to this internality, there is a linear externality from energy use. To address these two sources of inefficiency, the policymaker has two policy instruments: "energy taxes," by which we mean carbon taxes, cap-and-trade programs, gas taxes, and other policies that change the retail energy price, and "product subsidies," by which we mean subsidies for hybrid vehicles, home weatherization, and energy efficient appliances, fuel economy standards, feebates, and other policies that affect the relative purchase prices of energy intensive durables. After deriving theoretical results, we simulate optimal policies and welfare effects for the U.S. automobile market, using a realistic representation of the choice set and driving patterns plus estimated internalities, externalities, and utility function parameters from the literature.

To answer our first question, we consider the "third best" world in which the policymaker sets only the energy tax given zero product subsidy. In the standard Pigouvian world with a homogeneous externality but no internalities, an energy tax at marginal damages achieves the first best. This increases social welfare but decreases "consumer welfare," by which we mean social welfare with zero weight on the externality. For example, a carbon tax is generally thought to be bad for the economy in the short term, even as the reduction in global warming generates positive net benefits over the long term. However, when consumers undervalue energy costs, a positive energy tax increases social welfare both by reducing externalities and by inducing more consumers to buy the energy efficient durables that they would buy in their private optima. In fact, our automobile market simulations suggest that a carbon tax set at marginal damages could in fact *increase* consumer welfare, and this effect is large enough to *more than double* the social welfare gains that an analyst would predict for the case with no internalities. This result is conceptually related to the Double Dividend hypothesis in Bovenberg and Goulder (1996), Parry (1995), and others in the basic sense that it identifies a potential additional benefit from environmental taxes other than externality reduction. We thus call this the *Internality Dividend from Externality Taxes*.

We then consider our second question, the "second best" combination of energy taxes and product subsidies. The first best policy would deliver heterogeneous corrections to decision utility that exactly offset each consumer's internality. The second best policy requires the two feasible instruments to deviate from their no-internality levels so as to best approximate the first best pattern. The problem is related to other analyses that use multiple instruments to address heterogeneous market failures, such as Fullerton and West (2002, 2010) and Innes (1996), who study the optimal combination of gas taxes and vehicle taxes to address heterogeneous pollution externalities. We derive optimal tax formulas that show that the level of each instrument is determined by the externality, the average internality of consumers marginal to that instrument, and the average internality of consumers marginal to the *other* instrument. This gives what we call the *Internality Targeting Principle*: the second best approximation to the perfectly targeted first best policy is determined by consumers' marginal internalities with respect to each of the instruments.

When our two instruments are used to address undervaluation, the Internality Targeting Principle typically gives an energy tax *below* marginal damages combined with a large product subsidy. The reason is that consumers who undervalue energy costs the most, and thus need the largest correction to decision utility, are mechanically the least sensitive to the energy tax because it works through changes in energy costs. Thus, the average marginal internality is large for the product subsidy relative to the energy tax, and a large subsidy will be used to correct the decisions of consumers that undervalue the most. This in turn causes a distortion to the less-biased consumers, which is corrected by keeping the energy price below social cost. However, the extent to which the energy tax deviates from marginal damages is tempered by the distortion that this causes on the utilization margin. In our automobile market simulations, the second best policy is a gasoline tax set 15% below marginal damages combined with a large product subsidy which would decrease the purchase price of 25 vs. 20 mile per gallon vehicles by about \$700.

In answer to our third question, we present a general but tractable approach to optimal policy design and welfare analysis in the spirit of the Chetty (2009) sufficient statistic approach. The intuition is that if consumers correctly value energy costs, their product purchases should be equally elastic to energy costs and upfront prices. The ratio of "energy cost elasticity" to price elasticity gives a valuation weight that is less than one if the average marginal consumer undervalues energy costs. Multiplying this by the energy cost savings from a more energy efficient product allows one to estimate the dollar value of the average marginal internality. This makes it possible to approximate the optimal product subsidy and its welfare effects given any fixed level of energy tax. In our auto market simulations, the "heuristic policy" of a gasoline tax at marginal damages and a product subsidy set using the sufficient statistic approach generates 94 to 99% of the welfare gains of the true second best policy, even as we simulate markets with distributions of the internality and correlations between unobservables that would be extremely difficult to identify empirically. The sufficient statistic approach offers clear benefits in setting energy policy: optimal taxes can be closely approximated and evaluated without a structural model, using reduced-form elasticities that can be estimated using variation in product prices and energy costs.

Our paper is related to a number of other analyses of public policies when agents misoptimize, including in the contexts of health care (Baicker et al. (2012), Handel (2011)), cellular phone contracts (Grubb and Osborne, 2012), drug addiction (Bernheim and Rangel (2004, 2005), Gul and Pesendorfer (2007)), taxation (Chetty et al., 2009), and many others.[1] Perhaps the most closely-related paper is O'Donoghue and Rabin (2006), who study optimal taxation on a hypothetical good ("potato chips") that some people over-consume relative to their long-run optima due to time-inconsistency. There are also several important related papers that study energy policy when consumers undervalue energy costs, including Fischer et al. (2007), Heutel (2011), Parry et al. (2010), and Krupnick et al. (2010).

Our paper makes three contributions. First, we study the optimal combination of multiple policy instruments under heterogeneous internalities, which is a natural question for energy policy and many other settings. This allows us to highlight and quantify the importance of targeting the more biased consumers based on the average internality of consumers marginal to each instrument. This discussion does not arise in most related papers because they either analyze only one policy instrument or assume that all consumers misoptimize in exactly the same way. Second, while most models assume a particular behavioral bias, we derive theoretical results that are general to many different

[1] Mullainathan et al. (2012) review this literature and discuss additional important papers too numerous to cite here. In a discussion in the Journal of Economic Literature, Kroft (2011) argues that there is much progress yet to be made: "The public finance literature is only recently beginning to consider behavioral welfare economics, and there exist few theoretical explorations of optimal policy with behavioral agents."

types of bias. Such generality seems crucial in settings such as ours, where consumers could misoptimize in multiple ways, with little empirical guidance as to which biases seem to be dominant. Third, we derive formulas that can allow policies to be designed and evaluated using reduced-form sufficient statistic that can be identified in a variety of contexts with no knowledge of the underlying structural model.

The paper proceeds as follows. In Section 2, we provide more background on undervaluation of energy costs and energy efficiency policies. Section 3 sets up our theoretical model, while Section 4 presents formal results. Section 5 details the auto market simulations, and Section 6 concludes.

2. Background

2.1. Empirical evidence on undervaluation

In this paper, we use the generic words "undervaluation" or "overvaluation" to capture multiple factors that might reduce or increase demand for energy efficient durable goods relative to consumers' private optima.[2] The first factor is naive present bias, as in Laibson (1997), Loewenstein and Prelec (1992), O'Donoghue and Rabin (1999), and Strotz (1955). Present bias results in undervaluation if the durable good's purchase price affects consumption in the present and energy costs are paid in the future. The second factor is systematically biased beliefs about the relative energy costs of different products, as studied by Allcott (2013), Attari et al. (2010), and Larrick and Soll (2008). The official cost–benefit analysis of Corporate Average Fuel Economy (CAFE) standards argues that consumers have downward-biased "perceptions" of fuel cost savings from high fuel economy vehicles (National Highway Traffic Safety Administration, NHTSA, 2010, page 2).

A third potential factor is inattention. Allcott (2011) documents suggestive evidence from the Vehicle Ownership and Alternatives Survey: 40% of Americans report that they "did not think about fuel costs at all" when buying their most recent vehicle. Inattention to energy costs would be consistent with evidence of inattention to other ancillary product costs. Consumers on eBay, for example, are less elastic to shipping and handling charges than to the listed purchase price (Brown and Hossain, 2010). Mutual fund investors appear to be less attentive to ongoing management fees than to upfront payments (Barber et al., 2005). Chetty et al. (2009) show that shoppers are less elastic to sales taxes than to prices. Seniors choosing between Medicare Part D plans place more weight on premiums than on expected out-of-pocket costs (Abaluck and Gruber, 2011).

A number of empirical papers dating to the 1970s have tested for undervaluation of energy costs. Hausman (1979) estimated that the "implied discount rate" that rationalizes consumers' tradeoffs between purchase prices and future energy costs for air conditioners was 15 to 25%, above the rates at which most consumers borrowed and invested money. His results were corroborated by Gately (1980), who showed that buyers of energy inefficient refrigerators needed to have discount rates of 45% to 300%, and by Dubin and McFadden (1984), who found that choices and utilization of home heating equipment implied a 20% discount rate. Hausman (1979) argued that consumers were making mistakes by not buying more energy efficient appliances, but that this was unsurprising because "at least since Pigou, many economists have commented on a 'defective telescopic faculty.'"

Since this early work, there have been additional tests of whether automobile consumers appear to undervalue future gasoline costs relative to purchase prices, including Allcott and Wozny (forthcoming), Austin (2008), Busse et al. (2013), Dreyfus and Viscusi (1995), Goldberg (1998), Kilian and Sims (2006), Sallee et al. (2009), Sawhill (2008), and Verboven (1999, 2002). Greene (2010) reviews 25 studies, of which 12 suggest that consumers tend to undervalue gas costs, five suggest that we overvalue gas costs, and eight indicate that the average consumer makes the tradeoff correctly.

2.2. Examples of product subsidies

The product subsidies we model are directly motivated by an important set of federal, state, and local policies. These policies include tax credits of up to $3400 for hybrid vehicles which were available for the bulk of the last decade, as well as the "gas guzzler tax," an excise tax ranging from $1000 to $7700 on low-fuel economy passenger cars. Another example is the Weatherization Assistance Program, which subsidizes weatherization for about 100,000 low-income homeowners each year. Furthermore, in many states, there are an array of rebates for weatherization and energy efficient appliances; these "Demand-Side Management programs" cost about $3.6 billion per year (U.S. EIA 2010). Importantly, our model of product subsidies also captures the effects of the Corporate Average Fuel Economy standard. This is because the standard imposes a relative shadow cost on sales of low fuel economy vehicles, which causes consumers to have to pay relatively less for high-fuel economy vehicles, just like an explicit product subsidy.[3]

3. A model of optimal policy with misoptimizing consumers

3.1. Consumer utility

We model consumers who choose between an energy inefficient durable I and an energy efficient durable E. Concretely, this could be a choice between hybrid versus non-hybrid cars, compact fluorescent lightbulbs versus incandescents, and standard versus energy efficient versions of air conditioners, washing machines, and other appliances. A durable $j \in \{I,E\}$ consumes e_j units of energy per unit of utilization m, with $e_I > e_E$. Consumers have single unit demand.

A consumer's utility from purchasing good j at price p_j and choosing utilization level m at energy price p_g is given by $u(m)+\epsilon_j + Y\text{-}p_g m\text{-}p_j$, where Y is the consumer's budget and (ϵ_I, ϵ_E) are taste shocks jointly distributed according to an atomless distribution F. We define a random variable $\epsilon = \epsilon_E\text{-}\epsilon_I$ and call its distribution G; we assume that G is also atomless and has a density function g. To ensure the existence of an interior optimum for utilization choice, we assume $u' > 0$, $u'' > 0$, $\lim_{x \to 0} u'(x) = \infty$ and $\lim_{x \to \infty} u'(x) = 0$. We also assume that $|xu''(x)/u'(x)| > 1$ to ensure that the price elasticity of utilization is less than one in absolute value, consistent with empirical estimates such as Davis (2008), Davis and Kilian (2011), Gillingham et al. (2010), Hughes et al. (2008), and Small and Van Dender (2007). This implies that consumers use less energy when they purchase the more energy efficient durable.

We let $m_j^* = \text{argmax}\{u(m) - p_g m e_j\}$ and set $v(e_j,p_g) \equiv u(m_j^*) - p_g m_j^* e_j$. We call $V(e_E,e_I,p_g) \equiv v(e_E,p_g) - v(e_I,p_g)$ the "gross utility gain" from energy efficiency, and we let $\xi = (e_E,e_I,p_g)$ denote all the parameters that determine this gross utility gain. The gross utility gain from energy efficiency reflects both the energy cost savings and the utility from increased utilization for the energy efficient good relative to the energy inefficient good. A fully optimizing consumer chooses durable E if and only if

$$V(\xi) + \epsilon > p_E - p_I. \tag{1}$$

[2] See DellaVigna (2009) for a review of the psychology and economics literature, which includes evidence on these biases in other contexts. Gillingham and Palmer (2012) review behavioral biases related to choices between energy-using durables.

[3] Our study is therefore related to other studies of CAFE standards and other potential policies to decrease the relative purchase prices of energy efficient vehicles, including queryAnderson, Parry, Sallee, and Fischer (2010), Austin and Dinan (2005), Fischer et al. (2007), Fullerton and West (2010), Gallagher and Muehlegger (2011), Goldberg (1998), Greene et al. (2005), Jacobsen (2010), Kleit (2004), and Sallee (2011a).

Misoptimizing consumers do not correctly value how differences in energy efficiency will impact their future utility. They choose E if and only if

$$\Gamma(V,\xi)V(\xi)+\epsilon > p_E - p_I. \quad (2)$$

where $\Gamma(V,\xi)$ is the (possibly endogenous) valuation weight. We assume that Γ is differentiable and that ΓV is strictly increasing in p_g.

A key feature of this framework is its generality. The valuation weights could be constant, or they could be endogenous to various factors. As we show formally in Appendix II, the framework is flexible enough to incorporate a number of different psychological biases, including (but not limited to):

1. *Salience bias*: As in the simple model in DellaVigna (2009), the gross utility gain from energy efficiency might be an 'opaque' component of the decision that is processed only partially, or the upfront product prices might be especially salient. Thus $\Gamma \equiv \gamma$, where $\gamma \in (0, 1)$ is the degree of attention to the opaque component, and $\gamma = 1$ gives consumers' privately-optimal choice. Alternatively, energy efficiency could also be overly salient to the consumer, which would be captured by $\gamma < 1$.
2. *Biased beliefs*: Consumers might misestimate the energy intensity difference and think that the energy intensities of I and E are $\hat{e}_I \neq e_I$ or $\hat{e}_E \neq e_E$, respectively. This could lead to over- or undervaluation.
3. *Endogenous inattention*: Attention might be endogenous to the stakes of the decision, as in Gabaix (2012), Koszegi and Szeidl (2013), or Sallee (2012). In particular, the higher the potential gross utility gain from purchasing E over I, the more consumers will pay attention to energy cost savings. The exogenous attention model can be modified to reflect this by allowing Γ to be an increasing function of $V(\xi)$.
4. *Present bias*: Assuming that purchase prices reduce consumption in the present and energy costs reduce consumption in the future, present-biased consumers will weight energy efficiency gains by a factor $\beta < 1$. In our model, this would be reflected by setting $\Gamma \equiv \beta$ for some $\beta \in (0, 1)$.

We allow for $k = 1,\dots,K$ decision utility types, with a type k consumer having bias Γ_k. We do not rule out the possibility that the discrete distribution of decision utility types could be correlated with the taste shock ϵ_j. For example, it is plausible that "green" consumers who derive warm glow from purchasing the more energy efficient product (high ϵ) might give more weight to energy efficiency (high Γ_k). We will let G_k denote the distribution of ϵ conditional on a consumer having bias Γ_k.

When $\Gamma_k(V,\xi) < \Gamma_{k'}(V,\xi)$ for all V, ξ, we say that $\Gamma_k < \Gamma_{k'}$. For simplicity, we assume that consumers can be perfectly ranked according to their bias: for any k and k', either $\Gamma_k > \Gamma_{k'}$ or $\Gamma_k < \Gamma_{k'}$. We index the bias so that $\Gamma_1 < \Gamma_2 < \cdots < \Gamma_K$, meaning that "high types" assign greater weight to energy efficiency.

3.2. Producers and the policymaker

Products $j \in \{E,I\}$ are produced in a competitive economy at a constant marginal cost c_j. Similarly, energy is produced in a competitive market at constant marginal cost c_g. The government chooses a subsidy τ_E for good E and an energy tax τ_g.[4] Prices are then given by $p_I = c_I$, $p_E = c_E - \tau_E$, $p_g = c_g + \tau_g$. Throughout the paper, $\mathbf{p}$ refers to the price vector (p_I,p_E,p_g) and $\sigma(\epsilon,k,\mathbf{p})$ denotes the consumer's choice of durable I or E (at prices $\mathbf{p}$). We use $\boldsymbol{\tau}$ to refer to the tax policy vector (τ_E,τ_g), and we use $T(\boldsymbol{\tau})$ to refer to the revenue generated by subsidizing E and taxing energy, which could be negative.

The government maintains a balanced budget through lump-sum taxes or transfers. Thus taxing or subsidizing durables purchases or energy use has no distortionary effects on other dimensions of consumption.

Define ϕ as the marginal damage per unit of energy used and $Q(\mathbf{p})$ as the amount of energy used at prices $\mathbf{p}$. The policymaker's objective function is to set $\boldsymbol{\tau}$ to maximize consumer utility net of the damage caused by energy use:

$$W = \sum_{k=1}^{K} \alpha_k \int \left[v\left(e_{\sigma(\epsilon_I,\epsilon_E,k,\mathbf{p})}, p_g\right) + Y + T(\tau) - p_{\sigma(\epsilon_I,\epsilon_E,k,\mathbf{p})} \right] dF_k(\epsilon_I,\epsilon_E) - \phi Q(\mathbf{p}). \quad (3)$$

We assume that the support of ϵ is wide enough such that for each bias type Γ_k, there is always a consumer indifferent between purchasing E and I.[5]

We use the following notation for demand for energy and demand for product E: The derivatives of energy demand Q with respect to the energy tax and the subsidy are denoted Q_{τ_g} and Q_{τ_E}, respectively. Total demand for product E is denoted D, and its derivatives with respect to the energy tax and the subsidy are denoted D_{τ_g} and D_{τ_E}, respectively. Total demand of type k consumers is denoted D^k, with respective derivatives $D^k_{\tau_g}$ and $D^k_{\tau_E}$.

3.3. Basic intuition for the policymaker's objective

Before moving on to formal results, we illustrate how the policymaker's objective function differs from consumers' objective functions. Consider a completely untaxed market: $\tau_E = 0$ and $\tau_g = 0$. A type k consumer in this untaxed market chooses durable $j \in \{I,E\}$ and utilization m to maximize the following intensive and extensive margin decision utilities:

$$m_j^* = \operatorname{argmax}\left\{u(m) - c_g e_j m\right\} \quad (4)$$

$$j^* = E \text{ iff } \Gamma_k V + \epsilon - (c_E - c_I) > 0 \quad (5)$$

In contrast, the policymaker wants the consumer to make choices according to the following criteria:

$$m_j^* = \operatorname{argmax}\left\{ \overbrace{u(m) - c_g e_j m}^{\text{Consumer's objective}} - \overbrace{\phi e_j m}^{\text{Externality}} \right\} \quad (6)$$

$$j^* = E \text{ iff } \underbrace{\Gamma_k V + \epsilon - (c_E - c_I)}_{\text{Consumer's objective}} + \underbrace{\phi\left(e_I m_I^* - e_E m_E^*\right)}_{\text{Externality}} + \underbrace{(1-\Gamma_k)V}_{\text{Internality}} > 0 \quad (7)$$

The consumer's intensive margin objective function differs from the policymaker's by the amount of externality that the consumer's choice produces. The consumer's extensive margin objective function differs from the policymaker's by both the size of the externality and the size of the internality. A comparison of Eqs. (5) and (7) reveals that the internality is a wedge between the consumer's objective function and the policymaker's objective that is almost exactly analogous to the wedge that the externality creates. The only substantive difference between the externality and the internality is that is that the externality is present at both margins of choice, while the internality is present only at the extensive margin.[6] Eq. (7) provides the intuition for why the welfare impacts of a policy will have three components: the distortion to the consumer's objective function, the externality reduction, and the internality reduction.

[4] Because there is no outside option, we do not lose any generality by not considering a tax or subsidy for I. In our model, subsidies τ_I' and τ_E' for products I and E, respectively, are choice and welfare equivalent to subsidies $\tau_I = 0$, $\tau_E = \tau_E' - \tau_I'$.

[5] Formally: for each k, the (possibly infinite) support $\left[\underline{\epsilon}^k, \bar{\epsilon}^k\right]$ of G^k is such that $\bar{\epsilon}^k > c_E - c_I$ and $\underline{\epsilon}^k < -(\Gamma_K V - (c_E - c_I))$.

[6] One potential avenue for future work would be to consider mistakes in utilization. This might arise for consumers who are on increasing block electricity pricing and choose durable goods and utilization based on average cost, as in Ito (2014) and Liebman and Zeckhauser (2004).

4. Optimal tax policy

4.1. Energy taxes for internalities and externalities

We begin by considering a situation in which the policy uses only the energy tax to address internalities and externalities. This one-dimensional optimal policy problem is similar to O'Donoghue and Rabin's (2006) analysis of sin taxes for present-biased consumers in a market with no externalities and one margin of choice.

Proposition 1 characterizes the optimal energy tax in terms of the internality, the externality, and the distortion to consumers' decision utility.

Proposition 1. *Suppose that $\tau_E = 0$. Let*

$$\mathcal{I}_{\tau_g} = \frac{\sum_k (1-\Gamma_k) V D^k_{\tau_g}}{D_{\tau_g}}$$

denote the average internality of consumers marginal to the energy tax. Then

$$\frac{d}{d\tau_g} W = \underbrace{\mathcal{I}_{\tau_g} D_{\tau_g}}_{\text{Internality change}} \quad \underbrace{-\phi Q_{\tau_g}}_{\text{Externality change}} \quad \underbrace{+\tau_g Q_{\tau_g}}_{\substack{\text{Distortion to}\\ \text{consumer decision utility}\\ \text{net of revenue recycling}}} \tag{8}$$

and if τ_g^* maximizes W, then

$$\tau_g^* = \phi + \frac{\mathcal{I}_{\tau_g} D_{\tau_g}}{-Q_{\tau_g}} \tag{9}$$

Eq. (8) consists of three parts: the marginal internality, the marginal externality, and the change in energy demand. For intuition, imagine first that all consumers optimize perfectly and that $\phi = 0$. With perfect optimization, Eq. (8) states that the product of τ_g and the change in energy demand is a sufficient statistic for the distortionary effect of taxing or subsidizing energy. This is simply a generalization of Harberger's (1964) efficiency cost of taxation formula to a setting with both extensive and intensive margin choice.[7]

In the special case with externalities and no internalities, Proposition 1 gives the Pigouvian externality tax $\tau_g^* = \phi$. When internalities are included, the optimal tax is additively separable in the internality and externality components.[8] The optimal energy tax will be above marginal damages when the average marginal internality is positive, meaning that marginal consumers tend to undervalue energy efficiency, and it will be below marginal damages when the average marginal internality is negative, meaning that the marginal consumers tend to overvalue energy efficiency.

A key insight obtained from Eq. (8) is that when the average marginal consumer undervalues energy efficiency, an increase in τ_g increases welfare through two channels: externality reduction and internality reduction. Undervaluation therefore implies higher welfare gains from increasing the energy tax than what would be expected in the externality-only case. We call this the *Internality Dividend from Externality Taxes*.

An important case is when no consumers overvalue energy efficiency, but at least some undervalue it. This is the plausible result of some behavioral models, such as present bias. In this case, Eq. (9) implies that even if an arbitrarily large portion of all consumers are perfect optimizers, the optimal energy tax will still be above marginal damages as long as *some* undervalue energy efficiency.

Corollary 1. *Suppose that $\Gamma_k \le 1$ for all k and $\Gamma_k < 1$ for some k. Then*

1. *$\frac{d}{d\tau_g} W > 0$ at $(\tau_E, \tau_g) = (0, \phi)$.*
2. *If τ_g^* is an optimal tax policy that maximizes W given that $\tau_E = 0$, then $\tau_g^* > \phi$. Thus even if $\phi = 0$, we still have that $\tau_g^* > 0$*

The first part of this corollary is a local statement about marginal changes, while the second part characterizes the global optimum. The intuition for both parts is that even when consumers pay energy costs that include the externality, some consumers that undervalue energy efficiency misoptimize by buying *I* instead of *E*. Increasing the energy tax above marginal damages induces some of these consumers to choose good *E*, as they would in their private optimum.

To see the logic behind Corollary 1 more formally, notice that when $\tau_g = \phi$ in Eq. (8), the last two terms of the equation cancel each other out, making the change in social welfare directly proportional to the marginal internality. When the energy price reflects social cost, the efficiency loss from changing the choices of optimizing consumers who are close to indifferent between *E* and *I* is only second-order. On the other hand, the gain from inducing misoptimizing consumers to purchase *E* is first-order positive. This logic is analogous to the basic result that in a previously-untaxed market with rational consumers and no externalities, the deadweight losses due to marginal tax increases are first-order zero. O'Donoghue and Rabin (2006) make a similar point when analyzing sin taxes when some consumers are rational and some misoptimize.

A second corollary to Proposition 1 considers the illustrative case of linear demand functions and constant Γ. This helps to illustrate the connection between consumers' bias and their elasticity to the energy tax.

Corollary 2. *Suppose that Γ_k is constant for all k; that is, $\Gamma_k \equiv \gamma_k$ for $\gamma_k > 0$. Suppose, moreover, that the distribution of ϵ is uniform and is distributed independently of Γ_k. Then*

$$\tau_g^* = \phi + \frac{lV(e_I m_I^* - e_E m_E^*)}{Q_{\tau_g}} \sum_k \alpha_k (1-\gamma_k)\gamma_k,$$

where $l > 0$ is the width of the support of G. Thus $\tau_g^* > \phi$ if and only if $\sum_k \alpha_k (1-\gamma_k)\gamma_k > 0$.

Under the assumptions of this corollary, a type *k* consumer has internality $(1-\gamma_k)V$ and tax elasticity proportional to γ_k. Because the elasticity depends on the internality, the average internality in the population does not equal the elasticity-weighted average marginal internality, and thus the average internality is not a sufficient statistic for the welfare impact of a corrective energy tax. Suppose, for example, that consumers are equally proportioned into two types, with $\gamma_1 = 0.2$ and $\gamma_2 = 1.2$. The average bias is $(0.2 + 1.2)/2 = 0.7$, implying substantial undervaluation on average. Increasing the energy tax above marginal damages will reduce welfare, however. The average internality of consumers marginal to the tax is now $[(0.2)(1-0.2) + (1.2)(1-1.2)]/2 = -0.08$. This means that the group of consumers whose choices will actually be changed by the increased energy tax will, on average, overvalue energy efficiency. Thus, it is actually optimal to lower the energy tax below marginal damages despite the fact that

---

[7] The simplest version of Harberger's formula applies to an economy where consumers have utility function $u(x_1,\ldots x_J) + (Z - p \cdot x - tx_1)$, where $x_1,\ldots x_J \in \mathbb{R}$ are the consumptions of the *J* different consumption goods at prices *p*, *t* is the tax on good x_1 and $Z - p \cdot x - tx_1$ is the numeraire good. Social welfare at tax *t* is given by $W(t) = \{max_x u(x) + Z - tx_1 - \sum c_i x_i\} + tx_1$, where c_i is the marginal cost of producing good *i* in a competitive market. Harberger's formula states that $W'(t) = t\frac{dx_1(t)}{dt}$; that is, the welfare impact is proportional to the product of the tax and the change in demand for x_1. This is analogous to our result that the welfare impact of increasing the energy tax is proportional to the product of the tax and the change in energy demand.

[8] The reason that the internality and the externality are additively separable is because we assume that the externality is homogeneous, and thus mechanically independent of the externality. In principle, the analysis could be generalized further to allow for heterogeneous externalities, as in Diamond (1973). Our analysis of heterogeneous *internalities* parallels Diamond's (1973) analysis of heterogeneous externalities. Just as Diamond (1973) shows that the optimal tax must equal the elasticity weighted average of the marginal externalities, our formula incorporates an elasticity-weighted average of the marginal internalities.

the average consumer undervalues energy efficiency.[9] This sets up a point that we return to in the conclusion, which is that it is not sufficient to justify energy efficiency policies based only on the proposition that the average consumer tends to undervalue energy efficiency.

4.2. Targeting internalities with multiple instruments

We now turn to the question of how a policymaker should optimally combine two different price instruments to correct both internalities and externalities. We begin by characterizing the optimal tax policy in terms of the marginal internalities and externalities.

Proposition 2. *Let*

$$\mathcal{I}_{\tau_g} = \frac{\sum_k (1-\Gamma_k) V D^k_{\tau_g}}{D_{\tau_g}} \text{ and } \mathcal{I}_{\tau_E} = \frac{\sum_k (1-\Gamma_k) V D^k_{\tau_E}}{D_{\tau_E}}$$

denote the average internalities of consumers marginal to the energy tax and subsidy, respectively.

Then

$$\frac{d}{d\tau_g} W = \underbrace{\mathcal{I}_{\tau_g} D_{\tau_g}}_{\text{Internality change}} \quad \underbrace{-\phi Q_{\tau_g}}_{\text{Externality change}} \quad \underbrace{+\tau_g Q_{\tau_g} - \tau_E D_{\tau_g}}_{\text{Distortion to consumer decision utility net of revenue recycling}} \tag{10}$$

$$\frac{d}{d\tau_E} W = \underbrace{\mathcal{I}_{\tau_E} D_{\tau_E}}_{\text{Internality change}} \quad \underbrace{-\phi Q_{\tau_E}}_{\text{Externality change}} \quad \underbrace{+\tau_g Q_{\tau_E} - \tau_E D_{\tau_E}}_{\text{Distortion to consumer decision utility net of revenue recycling}} \tag{11}$$

The optimal policy τ_g^*, τ_E^* is given by

$$\tau_g^* \text{-} \phi = \frac{\frac{D_{\tau_g}}{\text{-}Q_{\tau_g}} \left(\mathcal{I}_{\tau_g} \text{-} \mathcal{I}_{\tau_E} \right)}{1 \text{-} \frac{Q_{\tau_E} D_{\tau_g}}{Q_{\tau_g} D_{\tau_E}}} \quad \propto \mathcal{I}_{\tau_g} \text{-} \mathcal{I}_{\tau_E} \tag{12}$$

$$\tau_E^* = \frac{\mathcal{I}_{\tau_E} - \mathcal{I}_{\tau_g} \frac{Q_{\tau_E} D_{\tau_g}}{Q_{\tau_g} D_{\tau_E}}}{1 - \frac{Q_{\tau_E} D_{\tau_g}}{Q_{\tau_g} D_{\tau_E}}} \quad \propto \mathcal{I}_{\tau_E} - \mathcal{I}_{\tau_g} \frac{Q_{\tau_E} D_{\tau_g}}{Q_{\tau_g} D_{\tau_E}} \tag{13}$$

The key concept illustrated by Eqs. (12) and (13) is that the magnitude of each of the tax instruments does not just depend on the internality of the agents marginal to the instrument; the magnitude depends also on the marginal internality with respect to the other instrument. Suppose, for example, that no consumers overvalue energy efficiency but some undervalue. Eq. (12) then shows that the energy tax is increasing in $\mathcal{I}_{\tau_g}$, the average undervaluation of the agents marginal to the energy tax. At the same time, Eq. (12) also shows that the energy tax is decreasing in $\mathcal{I}_{\tau_E}$, the average undervaluation of the agents marginal to the subsidy. In other words, the magnitude of the energy tax depends on how efficiently it addresses misoptimization *relative to the subsidy*. In fact, Eq. (12) shows that the optimal energy tax will be below marginal damages when the subsidy reduces the internality more efficiently than the energy tax.

Corollary 3. *Suppose that $\mathcal{I}_{\tau_E} > \mathcal{I}_{\tau_g}$ for all τ_g and τ_E. Then $\tau_g^* < \phi$.*

Corollary 3 contrasts with Proposition 1, which showed that when the subsidy is constrained to zero, the optimal energy tax increases with both the internality and the externality. Now, when the subsidy is set optimally and the condition of Corollary 3 is satisfied, the optimal energy tax must be below marginal damages no matter how much consumers undervalue energy efficiency.

The condition for when the optimal subsidy is negative (i.e., it is a tax on E) is more ambiguous. In general, while $\mathcal{I}_{\tau_g} - \mathcal{I}_{\tau_E} > 0$ is enough to ensure that $\tau_g^* > 0$, it is not enough to ensure that the optimal subsidy is negative. As we show in Lemma 5 in Appendix I, $\frac{Q_{\tau_E} D_{\tau_g}}{Q_{\tau_g} D_{\tau_E}} < 1$, and thus it possible that $\mathcal{I}_{\tau_g} - \mathcal{I}_{\tau_E} > 0$ and $\mathcal{I}_{\tau_E} - \mathcal{I}_{\tau_g} \frac{Q_{\tau_E} D_{\tau_g}}{Q_{\tau_g} D_{\tau_E}} > 0$ simultaneously. The reason that $\mathcal{I}_{\tau_g}$ does not receive a full weight of 1 in determining the optimal subsidy is because even if the energy tax is well-targeted at reducing the extensive margin internality, moving it away from $\tau_g = \phi$ distorts the intensive margin choice. The product subsidy, however, can target the internality without distorting that choice. This fact that the energy tax changes intensive margin choice while the subsidy does not is formally reflected in the fact that $\left| Q_{\tau_g} / D_{\tau_g} \right| > \left| Q_{\tau_E} / D_{\tau_E} \right|$, which means that relative to the impact on extensive margin choice, the energy tax has a greater impact on energy use than does the subsidy.

When is the condition of corollary 3 satisfied? Roughly, the condition holds if consumers with higher Γ are relatively more elastic to the energy tax than to the product subsidy. More precisely, let $\varepsilon^k_{\tau_g}$ be the product demand elasticity of type k consumers with respect to the energy tax, and let $\varepsilon^k_{\tau_E}$ be the product demand elasticity of type k consumers with respect to the subsidy. We say that a type k' consumer is "relatively more elastic to the energy tax" than a type k consumer if $\frac{\varepsilon^{k'}_{\tau_g}}{\varepsilon^{k'}_{\tau_E}} > \frac{\varepsilon^k_{\tau_g}}{\varepsilon^k_{\tau_E}}$. The condition of corollary 3 is satisfied when this elasticity ratio is highest for the highest consumer types, implying that the high types are relatively more responsive to the energy tax than the low types:

Lemma 1. *Suppose $K \geq 2$ and that $\frac{\varepsilon^1_{\tau_g}}{\varepsilon^1_{\tau_E}} \leq \frac{\varepsilon^2_{\tau_g}}{\varepsilon^2_{\tau_E}} \cdots \leq \frac{\varepsilon^K_{\tau_g}}{\varepsilon^K_{\tau_E}}$ for all τ_g and τ_E, with at least one of the inequalities strict for all τ_g and τ_E. Then $\mathcal{I}_{\tau_E} > \mathcal{I}_{\tau_g}$ for all τ_g and τ_E.*

The relative elasticity condition is an intuitive sufficient condition for $\mathcal{I}_{\tau_E} > \mathcal{I}_{\tau_g}$,[10] and is satisfied by many forms of bias in our model:

Lemma 2. *For $K \geq 2$, suppose that $\frac{d}{d\tau_g}(\Gamma_k V)$ is strictly increasing in k for all τ_g. Then $\frac{\varepsilon^k_{\tau_g}}{\varepsilon^k_{\tau_E}}$ is strictly increasing in k for all τ_g and τ_E, and thus $\mathcal{I}_{\tau_E} > \mathcal{I}_{\tau_g}$.*

The condition that $\frac{d}{d\tau_g}(\Gamma_k V)$ is increasing in k for all τ_g simply means that higher types' perceived utility from energy efficiency is more responsive to energy prices than lower types. This is true, for example, when $\Gamma_k \equiv \gamma_k$ for some constant γ_k. As we already pointed out in our discussion of Corollary 2, consumers with low γ_k assign little weight to reductions in energy costs—and precisely because they assign little weight to reductions in energy costs they are also not going to be responsive to taxes that change the price of energy. In Appendix II, we show that a number of other models fit this assumption as well.

The condition in Lemma 2 illuminates the intuition behind the construction of the optimal policy. Consider the optimal tax policy in a first-best world in which consumers could be targeted perfectly. In this first-best world, the optimal policy would be a K-dimensional

[9] This result is reminiscent of Bernheim and Rangel's (2004, 2005) modeling of "cue-triggered mistakes" in consumption of addictive substances. In Bernheim and Rangel's model, the policymaker would not tax the addictive substance just to correct its "mistaken" overconsumption because addicts in cue-triggered hot mode consume the subtance whenever possible and thus are inelastic to prices. Taxation of the addictive subtance is done for redistributive purposes in the model, and in many situations the optimal policy may actually be a subsidy for the addictive substance. The result also parallels Diamond's (1973) insight that the optimal tax must equal the marginal externality, not simply the average externality.

[10] It is not a necessary condition, however. It is still possible to have $\mathcal{I}_{\tau_E} > \mathcal{I}_{\tau_g}$ if the relative elasticity condition is true on average (in an appropriate weighted average) but does not hold exactly.

vector $X = (x_1,\ldots,x_K)$ specifying that a type k consumer's perceived relative utility from E should be changed by an amount $x_k = (1 - \Gamma_k)V$, where Γ_k and V are evaluated at $\tau_E = 0$ and $\tau_g = \phi$. By construction, $x_1 > \cdots > x_K$. In our second-best best world, the goal is to approximate the first-best K-dimensional policy with a two-dimensional policy consisting of a subsidy and an energy tax. A subsidy τ_E changes all consumers' decision utility by the same amount: τ_E. In contrast, when $\frac{d}{d\tau_g}(\Gamma_k V)$ is increasing in k, an energy tax affects higher types more. The optimal policy thus lowers the energy tax below marginal damages to approximate the "negative slope" pattern $x_1 > \cdots > x_K$, and then sets the optimal subsidy to the level of the average marginal internality (approximately a weighted average of X), net of the additional distortions to decision utility and externality reduction this creates.[11] This logic forms what we call the *Internality Targeting Principle*: the second best approximation to the perfectly targeted first best policy is determined by consumers' marginal internalities with respect to each of the instruments.

Eqs. (12) and (13) are a generalization of the intuition in the previous paragraph. The condition $\mathcal{I}_{\tau_g} < \mathcal{I}_{\tau_E}$ means that in some overall sense, the decision utility of high Γ_k consumers is more affected by the energy tax than the decision utility of low Γ_k consumers. In this case, $\tau_g^* < 0$ to approximates the "negative slope" pattern of X. Moreover, Eqs. (12) and (13) can be shown to imply that $\tau_E^* + \left(e_I m_I^* - e_E m_E^*\right)\left(\tau_g^* - \phi\right) = \mathcal{I}_{\tau_E}$. The interpretation is that the total deviation in energy costs from the Pigouvian benchmark, i.e. $\tau_g^*(e_I m_I^* - e_E m_E^*)$, plus the total deviation in prices from the Pigouvian benchmark, i.e. τ_E^*, must sum to the average of the internality of consumers marginal to the subsidy.

Note that $\mathcal{I}_{\tau_E} > \mathcal{I}_{\tau_g}$ will *tend* to be true because consumers that value energy efficiency more are mechanically more responsive to increases in the energy tax. But $\mathcal{I}_{\tau_E} > \mathcal{I}_{\tau_g}$ is not always guaranteed to hold. We give three examples of such cases. First, this condition will not hold if $\frac{d\Gamma_k}{d\tau_g}$ is decreasing in k sufficiently quickly, such that $\frac{d}{d\tau_g}(\Gamma_k V)$ is decreasing in k. We discuss this possibility in Appendix III for the case of endogenous partial attention. Second, notice that all intuitions for $\mathcal{I}_{\tau_E} > \mathcal{I}_{\tau_g}$ rely on consumers having heterogeneous Γ. It is straightforward to show that when consumers have homogeneous bias in our model, $\mathcal{I}_{\tau_E} > \mathcal{I}_{\tau_g}$ and the optimal policy sets $\tau_g^* = \phi$.[12] Of course, it seems unlikely that all consumers would misoptimize in exactly the same way. Finally, note that our theoretical model assumes for simplicity that utilization need is homogeneous. But if Γ_k is close to homogeneous relative to the variation in utilization needs, then we could have $\mathcal{I}_{\tau_g} > \mathcal{I}_{\tau_E}$ because high-utilization consumers have larger internalities and are more elastic to the energy tax. In the auto market simulations, the empirical distribution of utilization combined with moderate variation in Γ still give $\tau_g^* < \phi$.

The next corollary summarizes optimal policy when all consumers weakly undervalue energy efficiency. As with Corollary 1, as long as just a fraction of people undervalue energy efficiency, either an increase in the subsidy or the energy tax is welfare improving when the baseline policy is $(\tau_g,\tau_E) = (0,\phi)$. But when the subsidy is better targeted at reducing undervaluation, the *optimal* policy will set a positive subsidy and an energy tax below marginal damages.

Corollary 4. *Suppose that $\Gamma_k \le 1$ for all k and $\Gamma_k < 1$ for at least one k. Then*

1. $\frac{dW}{d\tau_g} > 0$ *and* $\frac{dW}{d\tau_E} > 0$ *at* $(\tau_g,\tau_E) = (\phi,0)$
2. *An optimal policy* (τ_E^*,τ_g^*) *must have either* $\tau_E^* > 0$ *or* $\tau_g^* > \phi$.
3. *If, additionally,* $\mathcal{I}_{\tau_E} > \mathcal{I}_{\tau_g}$ *for all* τ_g *and* τ_E, *then* $\tau_E^* > 0$ *and* $\tau_g^* < \phi$. *In particular,* $\tau_E^* > 0$ *and* $\tau_g^* < \phi$ *if* $\frac{d}{d\tau_g}(\Gamma_k V)$ *is increasing in k.*

[11] As Eq. (11) shows, an energy tax $\tau_g \neq \phi$ introduces a distortion to consumers' decision utility net of externality reduction, and this affects the impact of the subsidy on social welfare. The optimal subsidy will be set such that the total distortion to decision utility net of externalities caused by the subsidy *and* the energy tax is equal to the average marginal internality.

[12] Earlier working paper versions of our article included this as a proposition, and Proposition 3 in Heutel (2011) is also comparable.

Further elaborating on the targeting logic, when $\mathcal{I}_{\tau_E} > \mathcal{I}_{\tau_g}$, a straightforward consequence of the welfare calculations in Proposition 2 is that changing demand for E by using the subsidy is more efficient at targeting undervaluation than an equivalent change in demand resulting from the energy tax:

Corollary 5. *Suppose that $\Gamma_k \le 1$ for all k and $\Gamma_k < 1$ for at least one k. Suppose also that $\mathcal{I}_{\tau_g} > \mathcal{I}_{\tau_E}$ for all τ_E,τ_g. Then for $\tau_g \ge \phi$ and for marginal changes $d\tau_g$ and $d\tau_E$ in the energy tax and the subsidy satisfying $D_{\tau_E} d\tau_E = D_{\tau_g} d\tau_g$,*

$$\frac{dW}{d\tau_g}d\tau_g < \frac{dW}{d\tau_E}d\tau_E$$

The reasons for why the subsidy is more efficient at targeting undervaluation bias are twofold. First, when $\mathcal{I}_{\tau_E} > \mathcal{I}_{\tau_g}$, low Γ_k consumers are relatively more elastic to the subsidy than to the energy tax. Thus the optimal subsidy will correct the most biased consumers' severe underpurchasing of E, while at the same time not leading to "too much" overpurchasing of E by the the less biased consumers. Second, the energy tax distorts intensive margin choice, whereas the subsidy does not. However, the second reason is not sufficient to make the corollary hold in situations when $\Gamma_k > 1$ for at least some consumers. If all consumers overvalue energy efficiency, for example, then the energy tax *may* become more efficient. The reason is that with overvaluation, the consumers making the biggest mistakes (the high Γ_k consumers) will be relatively more elastic to the energy tax, and thus the energy tax can correct the most biased consumers' overpurchasing of E, while at the same time not leading to too much underpurchasing of E by the less biased consumers.

Corollary 5 suggests that a policymaker who thinks that consumers undervalue energy efficiency but does not have sufficient information or regulatory flexibility to set the optimal combination of product subsidies and energy taxes might use a "heuristic policy" that sets the energy tax at marginal damages and then determines the optimal subsidy only to address internalities. Analogous to the results in Section 4.1, the optimal subsidy in this heuristic policy must equal the average internality of consumers marginal to the subsidy. This motivates our next section, which provides optimal subsidy and welfare formulas when the energy tax is fixed.

4.3. A sufficient statistic approach to energy policy with biased consumers

To implement our analysis thus far, a policymaker would need perfect knowledge of the distribution of consumers' bias. We now provide a formula for optimal product subsidies and welfare effects using only demand elasticities that can be observed in market data. Proposition 3 states the central result:

Proposition 3. *Suppose that $\frac{d}{d\tau_g}\Gamma_k \ge 0$. Then*

$$\mathcal{I}_{\tau_E} \ge p_g \frac{\left|Q_{\tau_E}\right| - D_{\tau_g}}{D_{\tau_E}} \tag{14}$$

and thus

$$\frac{d}{d\tau_E}W \ge \underbrace{\left(\left|Q_{\tau_E}\right| - D_{\tau_g}\right)p_g}_{\text{Lower bound for internality change}} \underbrace{-\phi Q_{\tau_E}}_{\text{Externality change}} \underbrace{+\tau_g Q_{\tau_E} - \tau_E D_{\tau_E}}_{\text{Distortion to consumer decision utility net of revenue recycling}} \tag{15}$$

and

$$\tau_E^* \ge \frac{\left(\left|Q_{\tau_E}\right| - D_{\tau_g}\right)p_g + \left(\tau_g - \phi\right)Q_{\tau_E}}{D_{\tau_E}}. \tag{16}$$

The basic intuition for this result can be seen by rewriting Eq. (14) as follows:

$$\mathcal{I}_{\tau_E} \geq p_g \frac{\left|Q_{\tau_E}\right| - D_{\tau_g}}{D_{\tau_E}} = \left(1 - \frac{D_{\tau_g}}{D_{\tau_E}(e_I m_I^* - e_E m_E^*)}\right)(e_I m_I^* - e_E m_E^*) p_g. \quad (17)$$

The term $(e_I m_I^* - e_E m_E^*) p_g$ corresponds to the energy cost savings from purchasing the more energy efficient product, which is a lower bound for the gross utility gain V. The term $\frac{D_{\tau_g}}{D_{\tau_E}(e_I m_I^* - e_E m_E^*)}$ is the ratio of the "energy cost elasticity" to the price elasticity — i.e., how much demand for E responds to a \$1 increase in relative energy costs compared to a \$1 increase in relative purchase price. Optimizing consumers should be equally responsive to the two changes, because they both affect consumption of the numeraire good by \$1. When consumers undervalue (overvalue) energy costs, this ratio is smaller (larger) than one, and we show more formally below that it is a lower bound for the average valuation weight for marginal consumers. Letting $\widetilde{\Gamma} = \sum_k \frac{D_{\tau_E}^k}{D_{\tau_E}} \Gamma_k$ denote the valuation weight of the marginal consumers, Eq. (17) thus constitutes a lower bound for $\left(1 - \widetilde{\Gamma}\right)V$, the average marginal internality.

Eq. (17) provides only a lower bound for several reasons. First, consumers who purchase the more energy efficient durable choose higher levels of utilization, and thus obtain higher levels of utilization utility. Biased consumers, however, underweight this utility gain the same way that they underweight differences in energy costs. Second, a change in τ_g may do more than just change future energy costs. When Γ_k is increasing in energy costs, increasing τ_g will also increase Γ_k. As a consequence, D_{τ_g} reflects both the level of the bias and how much the bias responds to changes in energy costs. This in turn implies that the ratio of the energy cost elasticity to the price elasticity reflects not only the level of the bias, but also how much the bias responds to changes in energy costs.

More formally, the change in energy demand Q_{τ_E} is simply the change in demand for good E multiplied by how much less energy is used by a consumer owning E rather than I. Thus:

$$Q_{\tau_E} = D_{\tau_E}(e_E m_E^* - e_I m_I^*). \quad (18)$$

Next consider D_{τ_g}. Increasing the energy tax has two first-order effects. First, it decreases the relative price of owning good E in proportion to how much less energy a consumer owning good E will use: $(e_I m_I^* - e_E m_E^*)$. Consumers with valuation weight Γ_k value this at $\Gamma_k(e_I m_I^* - e_E m_E^*)$, however. Second, increasing the energy tax may change consumer's valuation bias and thus increase type k consumers' perceived relative value of good E by $\left(\frac{d}{d\tau_g}\Gamma_k\right)V$. Thus a \$1 increase in an energy tax translates into a *perceived* increase in the relative value of good E of $\Gamma_k(e_I m_I^* - e_E m_E^*) + \left(\frac{d}{d\tau_g}\Gamma_k\right)V$. Since a \$1 increase in the subsidy translates into a \$1 increase in the perceived relative price of good E, we have that

$$D_{\tau_g} = \sum_k \left[(e_I m_I^* - e_E m_E^*)\Gamma_k + \left(\frac{d}{d\tau_g}\Gamma_k\right)V\right] D_{\tau_E}^k \quad (19)$$

Eqs. (18) and (19) thus imply that

$$\left(\left|Q_{\tau_E}\right| - D_{\tau_g}\right) p_g < p_g \sum_k (1 - \Gamma_k)(e_I m_I^* - e_E m_E^*) D_{\tau_E}^k < \sum_k (1 - \Gamma_k) V D_{\tau_E}^k \quad (20)$$

The last inequality of Eq. (20) follows from the fact that $V = (e_I m_I^* - e_E m_E^*) p_g + u(m_E^*) - u(m_I^*)$. In words, the gross utility gain from the energy efficient product has two parts, decreased energy costs and increased utilization. When utilization is fairly inelastic, $V \approx (e_I m_I^* - e_E m_E^*) p_g$, and thus our lower bound is a close approximation.

When the energy tax is set equal to marginal damages, the formulas in Proposition (3) now reduce to

$$\frac{d}{d\tau_E} W = \mathcal{I}_{\tau_E} D_{\tau_E} \geq \left(\left|Q_{\tau_E}\right| - D_{\tau_g}\right) p_g - \tau_E D_{\tau_E} \quad (21)$$

and

$$\tau_E^* = \mathcal{I}_{\tau_E} \geq p_g \frac{\left|Q_{\tau_E}\right| - D_{\tau_g}}{D_{\tau_E}} \quad (22)$$

A policymaker who observes only the three empirical parameters Q_{τ_E}, D_{τ_E}, and D_{τ_g} can set a "heuristic policy" of an energy tax at marginal damages and a product subsidy equal to the above τ_E^*. Corollary (5) suggests that this heuristic policy might approximate the true second best policy, and we evaluate this approximation in the auto market simulations.

How does a change Δ_{τ_E} in the subsidy effect welfare? A second-order Taylor expansion shows that

$$\Delta W \approx \frac{dW}{d\tau_E}\Delta_{\tau_E} + \frac{1}{2}\left(\Delta_{\tau_E}\right)^2 \frac{d^2 W}{d\tau_E^2} \quad (23)$$

$$\approx \frac{dW}{d\tau_E}\Delta_{\tau_E} - \frac{1}{2}\left(\Delta_{\tau_E}\right)^2 D_{\tau_E} \quad (24)$$

under the assumption that $\frac{d}{d\tau_E}D_{\tau_g} \approx 0$, $\frac{d\Gamma_k}{d\tau_E} \approx 0$ and $\frac{d}{d\tau_E}D_{\tau_E}^k \approx 0$ for all k.[13] When the policy is initially set at $(\tau_g, \tau_p) = (\phi, 0)$, the formula shows that increasing τ_E to the optimal level $\tau_E^* = \mathcal{I}_{\tau_E}$ generates a welfare change of

$$\Delta W \approx \left(\mathcal{I}_{\tau_E} D_{\tau_E}\right)\tau_E^* - \frac{1}{2}\left(\tau_E^*\right)^2 D_{\tau_E} = \frac{1}{2}\mathcal{I}_{\tau_E}^2 D_{\tau_E}. \quad (25)$$

Fig. 1 illustrates the basic intuition when $\tau_g = \phi$ and demand is linear. The solid line through points c and a illustrates the distribution of relative experienced utility, while undervaluation shifts demand for E downward to the dashed line through points b and f. The equilibrium with no policy is at point b, and the distance (c-b) is the average marginal internality at $\tau_E = 0$. Thus, a marginal increase in the product subsidy generates welfare gains (c-b) by inducing the marginal consumer to purchase the energy efficient product. The "third best" optimal product subsidy when $\frac{\varepsilon_{\tau_g}^k}{\varepsilon_{\tau_E}^k}$ is to set k at the average marginal internality. The welfare gains from this policy are the blue triangle abc, which corresponds to the triangle in Eq. (25).

Finally, combining Eq. (24) with the lower bound (14) gives the formula for welfare change as a function of the empirical parameters:

$$\Delta W \geq \approx \left[\left(\left|Q_{\tau_E}\right| - D_{\tau_g}\right) p_E + Q_{\tau_E}\left(\tau_g - \phi\right) - \left(\tau_E + \frac{1}{2}\Delta_{\tau_E}\right) D_{\tau_E}\right]\Delta_{\tau_E} \quad (26)$$

It would be much more difficult to provide comparable formulas for the energy tax. The reason is illustrated by the result in Corollary 2, which shows that the average internality of consumers marginal to the energy tax is proportional to $\sum \alpha_k(1 - \Gamma_k)\Gamma_k$ because consumers' response to the energy tax is proportional to Γ_k. Thus estimating the average marginal valuation weight $\sum_k \alpha_k \Gamma_k$, as we do for the subsidy, would not be enough. Instead, we would have to find a third price instrument with the property that consumers' response to that price instrument is proportional to Γ_k^2. We would then be able to estimate $\sum \alpha_k \Gamma_k(1 - \Gamma_k) = \sum \alpha_k(\Gamma_k - \Gamma_k^2)$ by subtracting consumers' response to the third price instrument from their response to the energy tax.

[13] Together, these conditions imply that $\frac{d}{d\tau_E}\mathcal{I}_{\tau_E} \approx 0$. Differentiating Eq. (11) with respect to τ_E then shows that $\frac{d^2 W}{d\tau_E^2} \approx D_{\tau_E}$.

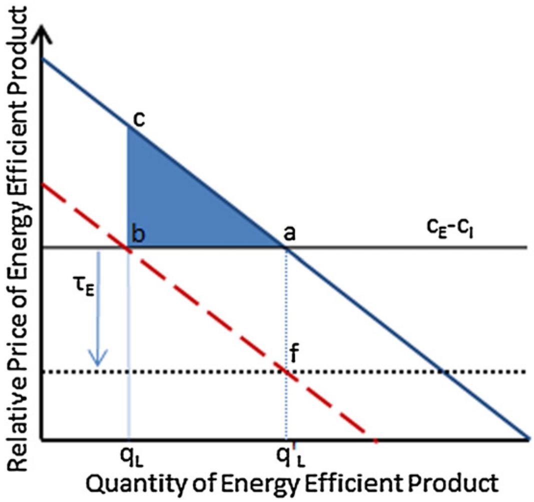

Fig. 1. Welfare effects of a product subsidy. Notes: The solid blue line is the demand curve for the energy efficient good if all consumers are rational. The dashed red line is the demand curve if consumers undervalue. (In general, these curves need not be straight lines.) Triangle abc is the consumer welfare loss from undervaluation. The dotted black line reflects the new supply curve after the product subsidy τ_E is applied.

In Appendix III, we show that these sufficient statistic-based formulas generalize to a market with multiple products and even more general variation in tastes—a setting that we also consider in our empirical simulations. In a market with multiple products, we consider a subsidy that is proportional to energy efficiency, such that two products with difference Δe in energy intensity experience a relative price change of $\tau_p \Delta e$. In this more general market, we now let D correspond to the total energy intensity of the products owned. That is, if κ_j consumers own a durable with energy efficiency j, then $D = \sum_j \kappa_j e_j$. And as before, we let Q correspond to the total energy used. All formulas in this section then generalize verbatim, with τ_p in place of τ_E and with $-D, -D_{\tau_E}, -D_{\tau_g}$ in place of D, D_{τ_E}, D_{τ_g}.

4.3.1. Implementing the sufficient statistic approach

There are several ways that empiricists can implement these formulas.[14] One is to directly estimate Q_{τ_E}, D_{τ_E}, and D_{τ_g} and plug them into the formulas. This requires data on product purchases and utilization, along with exogenous variation in product prices and energy prices. Importantly, one does not need variation in energy taxes and explicit energy efficiency subsidies: variation in energy prices and relative pre-tax prices of products with different energy intensities can also be used.[15]

D_{τ_g} is the easiest to estimate: prices of gasoline and other forms of energy vary over time and across regions, and datasets of product purchases can be matched to energy intensity ratings to calculate $D = \sum_j \kappa_j e_j$, as in Li et al. (2009) and Klier and Linn (2010). Empiricists interested in automobiles could estimate D_{τ_E} by testing how variation in hybrid vehicle incentives across states and time affects state-level hybrid purchases, as in Gallagher and Muehlegger (2011). Q_{τ_E} could be estimated by testing for effects of hybrid vehicle policies in gasoline consumption data or smog check data, which typically include odometer readings and fuel economy ratings.

In practice, Q_{τ_E} may be the most difficult to estimate. Household utility bills combine energy use across appliances where subsidies differ, while state-level gasoline consumption data combine energy use across vintages of vehicles that faced different hybrid vehicle subsidies, and smog check data are available in only a handful of states. Eq. (18) illustrates how Q_{τ_E} can be proxied by multiplying D_{τ_E} by the difference in expected energy use, which can come from cross sectional surveys. Allcott and Wozny (forthcoming), Busse et al. (2013), and Sallee et al. (2009) proxy in this way using vehicle utilization data from the National Household Travel Survey.

All of these approaches exploit the intuitive idea from Eq. (17) that the ratio of energy cost elasticity to price elasticity captures $\widetilde{\Gamma}$, the average valuation weight for marginal consumers. One could also use other approaches to estimating $\widetilde{\Gamma}$, and then multiply this by $(e_I m_I^* - e_E m_E^*) p_g$ to estimate the average marginal internality. For example, if misoptimization results from biased beliefs about energy costs, $\widetilde{\Gamma}$ can be estimated by eliciting beliefs through surveys, as in Allcott (2013) or Attari et al. (2010). If misoptimization results from biased beliefs or inattention, one could measure how an intervention to provide information or draw attention to energy costs affects demand for energy efficient products. Allcott and Taubinsky (2013) use randomized control trials to test informational interventions with buyers of energy efficient lightbulbs, and the sales tax labeling field experiment in Chetty et al. (2009) is analogously used to measure internalities due to inattention to sales taxes.

5. Automobile market simulation model

In this section we complement our theoretical analysis with discrete choice simulations of the US automobile market. Our simulations provide evidence on several questions that arise from the theory. In practice, how large is the Internality Dividend from Externality Taxes? What are the magnitudes of the second best policies, and are the welfare gains large in this context? How close do the sufficient statistic-based "heuristic policies" come to the second best optimum?

5.1. Setup

In order to simulate the theoretical model for the automobile market, we make three functional form assumptions. First, we assume that the taste shock ϵ is an identically and independently distributed type I extreme value variable that gives nested logit substitution patterns. The mean of ϵ is allowed to differ across products to reflect differences in size, power, quality, or other features, which generates differences in market shares. Second, we assume that the valuation weight for each consumer type k is an exogenous constant γ_k. This can correspond to exogenous partial inattention or to present bias. Third, we assume that utility from utilization takes the Constant Relative Risk Aversion form. We also allow consumers to have heterogenous utilization, building off of the framework in Appendix III. Appendix IV gives more information on these issues and other details of the simulations.

As in the theoretical model, the policymaker has two instruments, an energy tax and a product subsidy. In this context, the "energy tax" is simply a gasoline tax. As in the sufficient statistic analysis with more than two products, we have a product subsidy τ_p that scales linearly in each vehicle's energy intensity e_j. Because there is no substitution to an outside option and budget balance is maintained via lump sum transfers, the product subsidy can equally be interpreted as an "MPG tax," a "feebate" that combines a fee on low-MPG vehicles with a rebate for high-MPG vehicles, or an average fuel economy standard that imposes a relative shadow cost on the sale of low-MPG vehicles.

5.1.1. Data and calibration

Table 1 presents an overview of the choice set and simulation assumptions. The choice set is the 301 new cars and trucks from model year 2007 defined at the level of a manufacturer's model name, such

[14] This discussion illustrates how the formulas can be easily applied, using approaches already in common use in the literature. However, the formulas and insights are novel — none of these papers calculate optimal policies, and only Allcott (2013) estimates welfare impacts.

[15] Differential tax salience in the sense of Chetty et al. (2009) or Finkelstein (2009) would change many parts of our results, including this statement, and this could be an interesting avenue of future research. Furthermore, in practice one must consider consumers expectations of future energy prices and taxes, as pointed out by Allcott (2011), Anderson et al. (2012), and Li et al. (2012).

as the "Honda Civic" or "Ford F-150."[16] Each model's price is the average across transactions recorded by the JD Power and Associates "Power Information Network," which collects data at 9500 dealers covering about one third of U.S. retail auto transactions. Market shares are from the National Vehicle Population Profile, a comprehensive national database of vehicle registrations obtained from R.L. Polk. Energy intensity e_j is the inverse of the U.S. Environmental Protection Agency (EPA) miles per gallon (MPG) fuel economy ratings. Different submodels within a model — for example, the manual vs. automatic transmission versions or the sedan vs. the coupe — may have different energy intensities, so we use each model's sales-weighted average energy intensity.

The most uncertain parameters in the simulations are the magnitudes of the internalities and externalities. In our base case, we assume a population average γ of 0.8, which is slightly more conservative than Allcott and Wozny's (forthcoming) $\gamma = 0.76$ and the $\gamma = 0.78$ implied by the corresponding estimates from Busse et al. (2013).[17] Lacking any empirical evidence on the distribution of valuation weights, our base case assumes a two-point discrete distribution: one half of consumers are biased, with $\gamma_b \equiv 0.6$, and the other half are unbiased, with $\gamma_u \equiv 1$. The fact that half of the consumers do not misoptimize highlights that there is significant scope for internality-targeting polices even when a large share of consumers choose according to the standard model. We experiment with the distribution of γ in alternative simulations.

We focus on externalities from carbon dioxide emissions, both because this is an externality under significant policy debate and because marginal damages per gallon from other externalities such as congestion and local air pollutant emissions are heterogeneous across consumers.[18] Consistent with estimates by Greenstone et al. (2011), we assume a marginal damage of \$20 per metric ton, which translates into $\phi = \$0.18$ per gallon of gasoline.

Appendix IV contains full details on how the model is calibrated. In brief, we set each vehicle's mean ϵ such that the baseline simulated market shares equal the observed 2007 market shares. The mean own-price elasticity of demand across all models is -5, consistent with estimates from Berry et al. (1995). The vehicle nests and nested logit substitution parameter are taken from Allcott and Wozny (forthcoming). The distribution of utilization demand parameters reflects the empirical distribution of vehicle-miles traveled from the National Household Travel Survey. The price elasticity of utilization demand at the mean VMT is -0.15, which approximates recent empirical estimates.[19] The expected future costs of gasoline and utility from driving over vehicle lifetimes are discounted to the time of purchase using empirical data on vehicle scrappage probabilities and a six percent discount rate, as calculated by Allcott and Wozny (forthcoming). We use a pre-tax gasoline price c_g of \$3 per gallon.

[16] More precisely, this is the set of model year 2007 new cars and trucks that have fuel economy ratings from the U.S. Environmental Protection Agency. We exclude vans as well as ultra-luxury and ultra-high performance exotic vehicles: the Acura NSX, Audi R8 and TT, Chrysler Prowler and TC, Cadilliac Allante and XLR Roadster, Chevrolet Corvette, Dodge Viper and Stealth, Ford GT, Plymouth Prowler, and all vehicles made by Alfa Romeo, Bentley, Ferrari, Jaguar, Lamborghini, Maserati, Maybach, Porsche, Rolls-Royce, and TVR.

[17] The average of Busse et al. (2013) implied discount rates for used vehicle markets using the corresponding assumptions for vehicle miles traveled and scrappage probabilities is 13%. Using empirical data on the average opportunity cost of capital for used vehicle buyers, this translates to $\gamma = 0.78$.

[18] Fullerton and West (2002, 2010), Innes (1996), Knittel and Sandler (2013), and queryParry and Small (2005) analyze other externalities from gasoline use and solve for optimal policies when externalities per gallon vary across consumers. Parry et al. (2007) provide an overview of the literature on automobile externalities and policies.

[19] Hughes et al. (2008) find that between 2001 and 2006, the short-run elasticity of gasoline demand was between -0.034 and -0.077. Small and Van Dender (2007) estimate that between 1997 and 2001, this elasticity was -0.022. Using data from California between 2001 and 2008, Gillingham (2013) estimates a "medium-run" elasticity of -0.22. Davis and Kilian (2011) use differences in state tax rates to estimate an elasticity of -0.46.

Table 1
Auto market simulation overview.

	Mean	Std. dev.	Min	Max
Choice set				
Number of models	301			
Price p_j (\$)	36,267	24,795	12,038	174,541
Gallons per mile e_j	0.053	0.011	0.022	0.084
2007 Quantity sold	46,459	72,078	93	616,275
Energy				
Pre-tax gasoline price p_g (\$ per gallon)	3			
Marginal damage ϕ (\$ per gallon)	0.18			
Preferences				
Valuation weight γ	0.8	0.20	0.6	1
Mean vehicle own-price elasticity	−5			
Utilization elasticity at mean utilization	0.15			
Annual discount rate	6%			

Notes: All dollars are real 2005 dollars.

5.2. Results

5.2.1. Third best energy taxes for internalities and externalities

Table 2 presents simulation results. Column 1 is the "no policy" equilibrium with zero product subsidy or energy tax. Column 2 is the first best, which could result from consumer-specific product subsidies set to address each individual consumer's level of bias. Consumer welfare is equal to social welfare minus externality damages, and all welfare calculations are relative to the no policy equilibrium.[20] All figures are present discounted totals over vehicles' lifetimes.

We also ran additional simulations to quantify the effects of externality taxation in the standard case when all consumers have $\gamma \equiv 1$. In this case, of course, the first best policy is to set $\tau_p = 0$ and $\tau_g = \phi$. The results, which are not included in Table 2, show that the carbon tax reduces consumer welfare by \$5.50 per vehicle while reducing carbon dioxide damages by \$11.10 per vehicle, thereby increasing social welfare by \$5.60 per vehicle. This illustrates the sense in which carbon taxes are "bad for the economy": reducing carbon emissions is costly, in this case because consumers purchase higher-MPG vehicles that they don't like as much.

Column 3 illustrates how these traditional results change when there are internalities. Under the same policy of $\tau_p = 0$ and $\tau_g = \phi$, consumer welfare now *increases* by \$2.00 per vehicle, as the energy tax reduces the pre-existing distortion from undervaluation. The social welfare gain is now \$12.50 per vehicle — more than twice what it would be if there were no internalities. This illustrates how the Internality Dividend from Externality Taxes could be quantitatively very important. In this case, an analyst who evaluates externality taxes without considering the Internality Dividend would understate the welfare gains by a factor of two.

Fig. 2 plots the gains in consumer welfare and social welfare at different levels of the energy tax, assuming zero product subsidy. Any energy tax between \$0 and \$0.34 per gallon increases consumer welfare through the internality dividend. As illustrated on the graph and in Column 4 of Table 2, the third best energy tax is \$0.32 per gallon. This optimal tax is almost twice the level of ϕ, which illustrates the result from Corollary 1 of Proposition 1. Not coincidentally, the socially optimal tax is slightly higher than the point at which a

[20] The changes in carbon dioxide damages are in the range of \$10 to \$23 per vehicle, which might seem small relative to the total gasoline costs of \$15,000. The reason for this is that the carbon externality is only 6% of gasoline costs, and none of the simulated policies reduce carbon dioxide emissions by more than 2.5%. Of course, if consumers were more price elastic on the purchase or utilization margin, the policies would reduce emissions more.

Table 2
Auto market simulation results.

Policy	1	2	3	4	5	6	7
	No policy	First best (Type-specific product subsidies)	$\tau_g = \phi$ $\tau_p = 0$	Third best energy tax: $\tau_p = 0$, τ_g to max social welfare	Third best product subsidy: $\tau_g = 0$, τ_p to max social welfare	Second best: τ_g and τ_p to max social welfare	Heuristic policy: $\tau_g = \phi$, τ_p from sufficient statistic formula
Optimal policies							
Gas tax τ_g ($/gal)	0	0.18	0.18	0.34	0.18	0.15	0.18
Product subsidy τ_p ($000 s/GPM)	0		0	0	66.8	68.6	55.8
Resulting allocations							
Harmonic mean MPG	19.3	19.6	19.3	19.4	19.6	19.6	19.6
Average lifetime VMT	152,840	151,950	151,590	150,550	151,950	152,140	151,890
Average PDV of gas cost	15,824	16,330	16,580	17,245	16,362	16,248	16,398
Average CO_2 tons emitted	68.9	67.1	68.1	67.5	67.2	67.3	67.4
Welfare vs. no policy							
ΔConsumer welfare/vehicle		46.9	3.1	−2.3	17.5	18.8	18.8
ΔCO_2 damages/vehicle		−25.1	−10.9	−20.1	−23.3	−22.1	−21.3
Consumer welfare loss/ton CO_2		−25.7	−3.9	1.6	−10.4	−11.7	−12.2
ΔSocial welfare/vehicle		72.0	14.0	17.8	40.8	40.9	40.1

Notes: All dollars are real 2005 dollars. Carbon emissions and damages are denominated in metric tons of carbon dioxide, and the externality is $0.18 per gallon. Welfare effects are per new vehicle sold, discounted at 6% per year over the vehicle's life.

marginal increase begins to decrease consumer welfare. To see the intuition for this, think of the first order condition from Proposition 1: the social welfare-maximizing energy tax is such that a marginal change has zero effect on the sum of externalities and consumer welfare.

5.2.2. Targeting and the second best

Column 5 of Table 2 presents the third best product subsidy, by which we mean the product subsidy that maximizes social welfare when $\tau_g = \phi$. The social welfare gains relative to no policy are $34.20 per vehicle, as compared to $15.30 with the third best energy tax. This shows that Corollary 5 is a quantitatively important result: if a policymaker has to choose between either the energy tax or the product subsidy to address internalities, the product subsidy is significantly more effective. It also confirms that in this simulated market, the corollary's basic logic is true globally, not just locally.

Column 6 presents the second best policy combination. Consistent with Corollary 3, τ_g^* is $0.15 per gallon, or about 15% less than ϕ. The socially optimal product subsidy is $66,800 per GPM. To put this in perspective, a 20 MPG vehicle, such as a Subaru Outback Wagon, uses 0.05 GPM, while a 25 MPG vehicle, such as a Toyota Corolla, uses 0.04 GPM. This τ_p^* therefore implies a relative price increase of $668 for the 20 MPG vehicle. The social welfare gains are $34.30 per vehicle, just slightly more than the third best when $\tau_g = \phi$. These gains equate to $549 million when summed over the 16 million vehicles sold in a typical year.[21]

Fig. 3 illustrates how the second best tax and subsidy rates change as γ_b varies from 0.3 to 1.6. The optimal energy tax deviates increasingly from marginal damages as γ_b deviates increasingly from one.[22] When γ_b is very low, the optimal policy is to target these highly biased consumers with a large product subsidy. This causes a large distortion to the unbiased consumers, causing them to buy more high-MPG vehicles than they would in the social optimum. The energy tax below marginal damages is well-targeted at undoing that distortion, as the unbiased consumers are more responsive than biased consumers to that change in energy prices.

When $\gamma_b > 1$, meaning that the biased consumers overvalue energy costs, the optimal policy includes both an energy tax below marginal damages and a tax on energy efficient vehicles. The product subsidy curve is slightly convex. This is because when $\gamma_b > 1$, the biased consumers are increasingly elastic to the energy tax, and deviations from marginal damages are increasingly effective at targeting the biased consumers. This is balanced against the intensive margin distortion that results when $\tau_g \neq \phi$. Fig. 3 shows that when utilization is more price elastic, optimal policy has the same features, but the energy tax stays closer to marginal damages.

Table 3 replicates the key results from Table 2, except with a three-point distribution of γ: one-quarter of the population has $\gamma_l = 0.6$, one quarter has $\gamma_h = 1.4$, and one half is still unbiased at $\gamma_u = 1$. Although the setup may or may not be realistic, the results sharply illustrate the intuition behind the Internality Targeting Principle. The mean valuation weight in the population is now $\gamma = 1$, and one might initially think that the optimal policy might be the externality-only optimum of $\tau_g = \phi$ and $\tau_p = 0$. Indeed, when τ_g is fixed at ϕ, the third best product subsidy is very close to zero. This is because all γ types are (approximately) equally elastic to the product subsidy, so a positive (negative) subsidy cannot improve the decisions of the low-γ (high-γ) types without equally distorting decisions by the opposite type.

The third best energy tax, however, deviates significantly from the externality-only case: it is only about half of marginal damages. This is because the high-γ types are much more elastic to the energy tax than the low-γ types, and an energy tax below marginal damages can reduce overconsumption of E by the high-γ types with less distortion to the low-γ types. For the same reason, the second best policy involves an

[21] Table 2 shows that the harmonic mean fuel economy rating increases by only about 0.3 in the second best compared to the no policy case. This result is driven by the combination of consumers' price elasticity and the magnitude of the internality relative to the differences in decision utility across vehicles. At the mean utilization, the true difference in lifetime discounted fuel costs between a 20 and a 25 MPG vehicle is $3272. A consumer with $\gamma = 0.8$ undervalues this by $654. This is small compared to price differences across vehicles: the standard deviation is $24,000, and the interquartile range is $21,000. Thus, correcting the internality does not induce large changes in fleet fuel economy. Allcott and Greenstone (2012) point out that this result suggests that Corporate Average Fuel Economy standards that require much larger increases in fuel economy cannot be justified based on internalities of this magnitude, and Fischer et al. (2007) make a similar argument.

[22] When $\gamma_b \approx 0.9$, so that consumers are close to homogeneous in their valuation weights, the energy tax is very slightly above marginal damages. This reflects the insight from Section 4 that the optimal energy tax can exceed marginal damages when variance in valuation weights is small compared to variance in utilization demand.

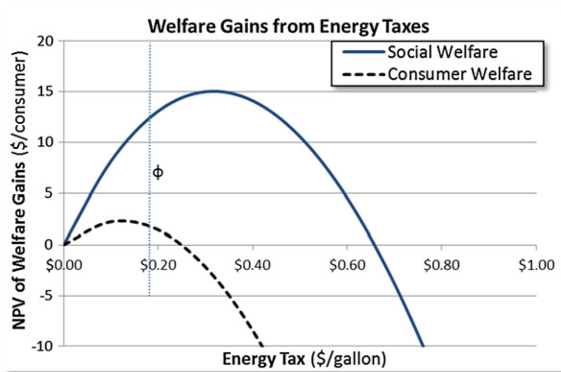

Fig. 2. The internality dividend from externality taxes. Notes: This figure shows the welfare gains from different energy taxes when the product subsidy is set to zero.

even lower energy tax that is balanced by a positive product subsidy. This follows directly from Proposition 2, because this distribution of γ implies that $\mathcal{I}_{\tau_E} \approx 0$ and $\mathcal{I}_{\tau_g} < 0$. Eq. (13) thus immediately implies that $\tau_E^* > 0$, while Eq. (12) implies that $\tau_g^* < \phi$. Therefore, even though the average internality of consumers marginal to the subsidy is near zero, the targeting logic of Eq. (13) gives a positive subsidy as part of the optimal policy combination.

A first best policy, which would combine large subsidies for E for the low-γ types with large taxes on E for the high-γ types, would produce much larger welfare gains than the second best policy. This is because although the energy tax is somewhat effective at differentially targeting the two types, the γ types still have somewhat similar elasticities to the energy tax. Furthermore, moving τ_g away from ϕ increasingly distorts utilization decisions.

5.2.2.1. Robustness to alternative assumptions. Table 4 presents the second best policies under a series of alternative assumptions. Column 1 adds a correlation between preferences and the bias, to reflect the fact that more environmentalist consumers might also be more attentive to energy costs. Specifically, we assume that unbiased consumers are willing to pay $4000 more for a vehicle with 0.01 lower fuel intensity. This has little effect on the optimal policies. Column 2 holds constant the unconditional distribution of γ but adds a positive covariance between utilization and γ. This reflects the idea that in an endogenous partial attention model, consumers who drive more might endogenously be more attentive to energy costs. This reduces the energy tax further below marginal damages, because consumers who are less responsive to the energy tax because they drive less are

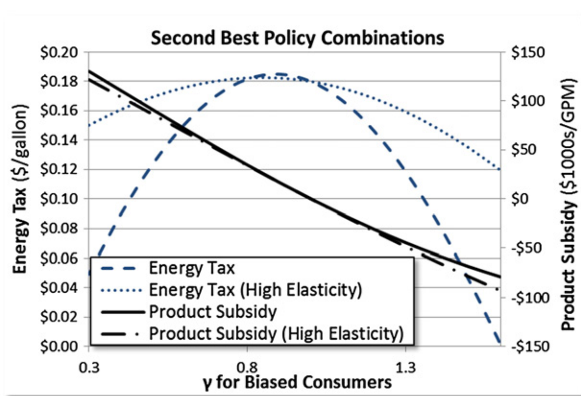

Fig. 3. Second best policy combinations. Notes: This figure shows socially-optimal policies while varying γ_b. The "High Elasticity" case assumes that the price elasticity of utilization is −0.45.

now more biased. The average marginal internality of the energy tax thus becomes even smaller relative to the product subsidy.

Columns 3–5 change key utility function parameters. Column 3 assumes that utilization is much more elastic relative to the base case, approximating estimates by Davis and Kilian (2011). As in the "High Elasticity" case in Fig. 3, this increases the intensive margin distortion as τ_g deviates from ϕ, so the optimal policy has τ_g much closer to ϕ. Columns 4 and 5 modify the marginal utility of money and the nested logit substitution parameter. As one might expect, consumption and welfare are more sensitive to taxes and subsidies when consumers are more elastic, but the optimal policies depend on marginal internalities and externalities, not substitution elasticities.

Columns 6 and 7 modify the distribution of γ. Column 6 assumes homogeneity: $\gamma = 0.8$ for all consumers. This provides a useful example that violates the assumptions of Corollary 3, giving an optimal energy tax above marginal damages. This happens because variation in the internality now derives entirely from variation in utilization demand rather than variation in γ, and the high-utilization consumers are both more elastic to the energy tax and undervalue their private gains from energy efficiency the most. The average marginal internality is thus larger for the energy tax than for the product subsidy. Column 7 does the opposite of Column 6, increasing the variation in γ_b. In this case, the energy tax is even further below ϕ, as the most biased consumers with very low γ are especially inelastic to the energy tax relative to the product subsidy.

5.2.3. Evaluating the sufficient statistic approach

The simulations show that the exact optimal policies depend at least somewhat on market features that could be difficult to estimate, such as the distribution of γ and its correlation with preferences. However, Corollary 5 suggests that if the policymaker must choose one of the two instruments to address the internality, it should be the product subsidy, and Eq. (22) provides a formula based on sufficient statistics to bound the optimal product subsidy when $\tau_g = \phi$. How close does this "heuristic policy" get to the true second best?

The sufficient statistics for the heuristic τ_p^* can be found by comparing Column 3 of Table 2 to Columns 4 and 5, which change only the energy tax and product subsidy, respectively. In this market, the heuristic policy bound is $\tau_p^* \geq$\$56,500/GPM.[23] This is 13% lower than the true third best product subsidy, which maximizes welfare with $\tau_g = 0$. Column 6 of Table 2 presents the effects of the heuristic policy. The social welfare gains are $33.80 per vehicle, only one to two percent less than the social welfare gains from the true second best policy or the third best product tax. Thus, despite the fact that the sufficient statistic-based formula is a lower bound, it appears to be a very tight lower bound, both in the sense of giving a policy that is quantitatively close to the true third best policy and in generating welfare gains that are almost exactly the same. The heuristic policy also performs quite well in the alternative simulations in Table 3: it never fails to capture less than 94% of the welfare gains from the true second best policy, even as the second best energy tax deviates more significantly from marginal damages.

Eq. (26) provides a bound on the welfare effects of a change in the product subsidy when $\tau_g = \phi$. As an example, we measure the tightness of this bound by evaluating the welfare gains from Column 5, the third best product subsidy, relative to Column 3, which has $\tau_p = 0$. The sufficient statistic-based welfare lower bound is $\Delta W \geq$\$16.90 per vehicle,

[23] The absolute value of the change in energy demand from a change in the product subsidy is $\left|Q_{\tau_p}\right| \approx \left|\frac{\Delta Q}{\Delta_{\tau_p}}\right| = \left|\frac{(\$15{,}955-\$16{,}142)/(\$3.18/gallon)}{\$65{,}000/GPM}\right| \approx 9.1 \cdot 10^{-4}$. The change in average GPM from a change in the energy price is $D_{\tau_g} \approx \frac{\Delta D}{\Delta_{\tau_g}} = \frac{(1/20.0)GPM-(1/19.9)GPM}{(\$0.32-\$0.18)/gallon} \approx 7.2 \cdot 10^{-4}$. The change in average GPM from a change in the product subsidy is $D_{\tau_p} \approx \frac{\Delta D}{\Delta_{\tau_p}} = \frac{|(1/20.2)GPM-(1/19.9)GPM|}{\$65{,}000/GPM} \approx 1.0 \cdot 10^{-8}$. Thus, the heuristic policy bound is $\tau_p^* \geq 3.18 \frac{9.1 \cdot 10^{-4} - 7.2 \cdot 10^{-4}}{1.0 \cdot 10^{-8}} \approx \$56{,}500/GPM$.

Table 3
Simulations of three-point distribution.

Policy	1	2	3	4	5	6	7
	No policy	First best (Type-specific product subsidies)	$\tau_g = \phi$ $\tau_p = \phi$	Third best energy tax: $\tau_p = \phi, \tau_g$ to max social welfare	Third best product subsidy: $\tau_g = \phi, \tau_p$ to max social welfare	Second best: τ_g and τ_p to max social welfare	Heuristic policy: $\tau_g = \phi, \tau_p$ from sufficient statistic formula
Optimal policies							
Gas tax τ_g ($/gal)	0	0.18	0.18	0.10	0.18	0.08	0.18
Product subsidy τ_p ($000 s/GPM)	0		0	0	−0.8	8.2	−0.9
Resulting allocations							
Harmonic mean MPG	19.9	19.9	19.9	19.9	19.9	19.9	19.9
Average lifetime VMT	153,660	152,390	152,400	152,960	152,400	153,140	152,400
Average PDV of gas cost	15,379	16,106	16,106	15,780	16,108	15,674	16,108
Average CO_2 tons emitted	67.0	66.2	66.2	66.5	66.2	66.5	66.2
Welfare vs. no policy							
ΔConsumer welfare/vehicle		38.9	−10.5	−4.4	−10.3	−4.2	−10.3
ΔCO_2 damages/vehicle		−11.1	−11.1	−6.2	−11.0	−6.2	−11.0
Consumer welfare loss/ton CO_2		−48.5	13.0	9.8	13.0	9.2	13.0
ΔSocial welfare/vehicle		49.9	0.6	1.8	0.6	2.1	0.6

Notes: This table replicates parts of Table 2, except with a three-point distribution of γ: one-quarter of the population has $\gamma_l = 0.6$, one-quarter has $\gamma_h = 1.4$, and one-half still has $\gamma_u = 1$. All dollars are real 2005 dollars. Carbon emissions and damages are denominated in metric tons of carbon dioxide, and the externality is $0.18 per gallon. Welfare effects are per new vehicle sold, discounted at 6% per year over the vehicle's life.

Table 4
Second best policies under alternative assumptions.

Change from base case	1	2	3	4	5	6	7
	ϵ and γ correlated	θ and γ correlated	Utilization elasticity is −0.5	High η: Average own-price elasticity is −10	Logit sub patterns	$\gamma = 0.8$ for all consumers	γ_b uniform from 0 to 1.2
Second best policies							
Gas tax τ_g($/gal)	0.15	0.12	0.18	0.13	0.16	0.21	0.08
Product subsidy τ_p ($/GPM)	66.0	60.6	64.7	68.0	66.4	63.4	70.8
Resulting allocations							
Harmonic mean MPG	20.4	20.2	20.2	20.5	20.0	20.2	20.2
Average lifetime VMT	153,110	153,130	153,440	153,440	152,560	152,490	153,460
Average PDV of gas cost	15,735	15,752	15,785	15,435	16,081	16,088	15,526
Average CO_2 tons emitted	65.2	66.0	64.9	64.4	66.4	65.4	65.9
Welfare vs. no policy							
ΔConsumer welfare/vehicle	13.0	9.8	7.6	31.0	7.9	14.7	14.7
ΔCO_2 damages/vehicle	−19.0	−16.8	−31.5	−31.0	−15.7	−22.7	−16.3
Consumer welfare loss/ton CO_2	−9.4	−8.1	−3.3	−13.8	−6.9	−8.9	−12.4
ΔSocial welfare/vehicle	32.0	26.6	39.1	62.0	23.5	23.5	30.9
Share of second best social Welfare from heuristic policy	0.99	0.97	0.98	0.99	0.99	0.99	0.94

Notes: This table presents the second best energy tax and product subsidy combination under different parameter assumptions. All dollars are real 2005 dollars. Carbon emissions and damages are denominated in metric tons of carbon dioxide. Welfare effects are per new vehicle sold, discounted at 6% per year over the vehicle's life.

which is about 78% of the actual welfare difference of $21.80 per vehicle.[24]

6. Conclusion

This paper makes three related points about energy policy with externalities and internalities. First, we identify and quantify an Internality Dividend from Externality Taxes, which could significantly improve the welfare argument for the carbon tax in the third best case when it is implemented in isolation of other policies. Second, we develop the Internality Targeting Principle, which shows how the product subsidy and energy tax work in combination to target the more biased consumers while limiting the distortions to the less biased types. The basic insight generalizes to other contexts where multiple instruments target a heterogeneous market failure: the second-best approximation to the first best policy depends on the relative internalities (or externalities) of consumers marginal to each instrument. Third, we present formulas for optimal policy and welfare analysis based on reduced form sufficient statistics that can be estimated using field experiments or quasi-experimental variation in product prices and energy prices, without knowledge of the underlying structural model.

Our analysis also motivates a theoretical and empirical research agenda with caveats to the folk wisdom that "energy efficiency policies are justified because consumers tend to undervalue energy costs." Our optimal policy formulas show that tax or subsidy levels and welfare gains are determined by the instrument's average *marginal* internality, and there are natural reasons why this could differ from the average internality. Imagine, for example, that environmentalist consumers are more attentive to energy costs than non-environmentalists. If environmentalists are also more likely to be aware of energy efficiency subsidies offered by their local utility, then the marginal consumers will be more heavily composed of these attentive types. This would similarly be the case if the

[24] This is calculated using Eq. (26), with the same $\frac{|\Delta_Q^e|}{\left|\Delta_{\tilde{\tau}_p}\right|}$, $\frac{|\Delta_D^g|}{|\Delta_{p_g}|}$, and $\frac{\Delta_D^e}{\Delta_{\tilde{\tau}_p}}$ from the footnote above, an initial τ_p of 0, and a Δ_{τ_p} of $65,000/GPM. The true ΔW is $34.21–$12.46 = $21.76 per vehicle.

policymaker subsidizes niche energy efficient products, such as LED lightbulbs, that only environmentalists tend to like. When evaluating energy efficiency policies, it matters *who* is saving energy, not just *how much* energy is saved.

Appendix I. Proofs of mathematical results

We begin with some lemmas that will be used throughout the proofs:

Lemma 3. *$em^*(p_g,e)$ is increasing in e*

Proof. We have

$$\frac{d}{de}em^*\left(p_g,e\right)=m^*\left(p_g,e\right)+e\frac{d}{de}m^*\left(p_g,e\right). \tag{27}$$

Differentiating the first order condition $u'(m^*)-p_g e=0$ with respect to e yields $u''(m^*)\frac{dm^*}{de}=p_g$. Thus

$$\frac{d}{de}em^*=m^*-\frac{ep_g}{u''(m^*)}=m^*-\frac{u'(m^*)}{u''(m^*)}.$$

But but since $u''<0$, the expression (27) is positive. ■

Lemma 4. $\frac{d}{dp_g}V=e_I m_I^*-e_E m_E^*$

Proof. By the envelope theorem,

$$\frac{d}{dp_g}v\left(\theta,e,p_g\right)=-m^*e,$$

from which the conclusion follows. ■

Lemma 5. $\frac{Q_{\tau_E}D_{\tau_g}}{Q_{\tau_g}D_{\tau_E}}<1$

Proof. Total energy demand is $Q=(1-D)e_I m_I^*+De_E m_E^*$. Thus

$$Q_{\tau_E}=D_{\tau_E}\left(e_E m_E^*-e_I m_I^*\right)$$
$$Q_{\tau_g}=D_{\tau_g}\left(e_E m_E^*-e_I m_I^*\right)+(1-D)\left(\frac{d}{d\tau_g}e_I m_I^*\right)+D\left(\frac{d}{d\tau_g}e_E m_E^*\right)<D_{\tau_g}\left(e_E m_E^*-e_I m_I^*\right)$$

where we use the fact that m_E^* and m_I^* are decreasing in τ_g. From this it easily follows that $\left|Q_{\tau_g}/D_{\tau_g}\right|>\left|Q_{\tau_E}/D_{\tau_E}\right|$. ■

Proof of Proposition 1. Let μ_I be the mean of ϵ_I. Because $T(\tau)=\tau_g Q$, social welfare can now be expressed as

$$W=\mu_I+v\left(e_I,p_g\right)-p_I+\tau_g Q$$
$$+\sum_k\alpha_k\left[\int_{\Gamma_k V(\xi)-(p_E-p_I)}^{\overline{\epsilon}_k}(V(\xi)-(p_E-p_I)+\epsilon)dG_k\right]-\phi Q$$

Differentiating with respect to τ_g, and noting that $D^k=\alpha_k[1-G_k(\Gamma_k V-(p_E-p_I))]$ by definition, yields:

$$\frac{dW}{d\tau_g}=-\phi Q_{\tau_g}-e_I m_I^*+\tau_g Q_{\tau_g}+Q$$
$$+\sum_k\alpha_k\left[-(1-\Gamma_k)V\frac{d}{d\tau_g}G(\Gamma_k V-(p_E-p_I))+\int_{\Gamma_k V(\xi)-(p_E-p_I)}^{\overline{\epsilon}_k}(e_I m_I^*-e_E m_E^*)dG_k\right]$$
$$=Q+\left(\tau_g-\phi\right)Q_{\tau_g}-e_I m_I^*+\sum_k(1-\Gamma_k)VD_{\tau_g}^k+\sum_k\left(e_I m_I^*-e_E m_E^*\right)\left(D^k\right)$$
$$=Q+\left(\tau_g-\phi\right)Q_{\tau_g}-e_I m_I^*+\sum_k(1-\Gamma_k)VD_{\tau_g}^k+\left(e_I m_I^*-e_E m_E^*\right)D$$
$$=Q-\left[(1-D)e_I m_I^*+De_E m_E^*\right]+\left(\tau_g-\phi\right)Q_{\tau_g}+\sum_k(1-\Gamma_k)VD_{\tau_g}^k$$
$$=\left(\tau_g-\phi\right)Q_{\tau_g}+\sum_k\alpha_k(1-\Gamma_k)VD_{\tau_g}^k.$$

And the expression for τ_g^* follows by setting $\frac{dW}{d\tau_g}$ equal to zero. ■

Proof of Corollary 1. By assumption, $\mathcal{I}_{\tau_g}>0$. Thus, since $Q_{\tau_g}<0$, we have that $\frac{dW}{d\tau_g}=(\tau_g-\phi)Q_{\tau_g}+\mathcal{I}_{\tau_g}D_{\tau_g}>0$ whenever $\tau_g\le\phi$. This implies that $\tau_g^*>0$. ■

Proof of Corollary 2. As we show in the proof of Lemma 2, $D_{\tau_g}^k=g(\Gamma_k V-(p_E-p_I))\left[\frac{d}{d\tau_g}(\Gamma_k V)\right]$, where g is the probability density corresponding to G. Under the assumptions of this corollary, we thus have that

$$D_{\tau_g}^k=l\gamma\left[\frac{d}{d\tau_g}V\right]=l\gamma\left(e_I m_I^*-e_E m_E^*\right),$$

which yields the statement of the Corollary when plugged into the optimal tax formula in Proposition 1. ■

Proof of Proposition 2. Because total tax revenues are now given by $\tau_g Q-\tau_E D$, we have that

$$W=\mu_I+v\left(e_I,p_g\right)-\tau_g Q-\tau_E D$$
$$+\sum_k\alpha_k\left[\int_{\Gamma_k V(\xi)-(p_E-p_I)}^{\overline{\epsilon}_k}(V(\xi)-(p_E-p_I)+\epsilon)dG_k\right]-\phi Q.$$

Differentiating with respect to τ_g, and following the computations in the proof of Proposition 1, we get that

$$\frac{dW}{d\tau_g}=-\phi Q_{\tau_g}-e_I m_I^*+\tau_g Q_{\tau_g}+Q-\tau_E D_{\tau_g}$$
$$+\sum_k\alpha_k\left[-(1-\Gamma_k)V\frac{d}{d\tau_g}G(\Gamma_k V-(p_E-p_I))\right.$$
$$\left.+\int_{\Gamma_k V(\xi)-(p_E-p_I)}^{\overline{\epsilon}_k}(e_I m_I^*-e_E m_E^*)dG_k\right]$$
$$=\left(\tau_g-\phi\right)Q_{\tau_g}-\tau_E D_{\tau_g}+\sum_k(1-\Gamma_k)VD_{\tau_g}^k$$

Notice that the only new term now, as compared to Proposition 1, is the loss in tax revenue that comes from increasing demand for E, $-\tau_E D_{\tau_g}$.

The welfare impact of increasing the subsidy is given by

$$\frac{dW}{d\tau_E}=-\phi Q_{\tau_E}+\tau_g Q_{\tau_E}-D-\tau_E D_{\tau_E}$$
$$+\sum_k\alpha_k+\left[(1-\Gamma_k)V\frac{d}{d\tau_E}G(\Gamma_k V-(p_E-p_I))+\int_{\Gamma_k V(\xi)-(p_E-p_I)}^{\overline{\epsilon}_k}(-1)dG_k\right]$$
$$=\left(\tau_g-\phi\right)Q_{\tau_E}-\tau_E D_{\tau_E}-D+\sum_k(1-\Gamma_k)VD_{\tau_g}^k+\sum_k D_k$$
$$=\left(\tau_g-\phi\right)Q_{\tau_E}-\tau_E D_{\tau_E}+\sum_k(1-\Gamma_k)VD_{\tau_g}^k$$

Now $\frac{dW}{d\tau_g}=\frac{dW}{d\tau_E}=0$ at an optimal tax policy; thus if (τ_g^*,τ_E^*) is an optimal tax policy then it must satisfy the system of linear equations

$$0=\left(\tau_g^*-\phi\right)Q_{\tau_g}-\tau_E^* D_{\tau_g}+\mathcal{I}_{\tau_g}D_{\tau_g}$$
$$0=\left(\tau_g^*-\phi\right)Q_{\tau_E}-\tau_E^* D_{\tau_E}+\mathcal{I}_{\tau_E}D_{\tau_E}.$$

Solving this system of linear equations leads to the formulas in the proposition. Furthermore, because $D_{\tau_g}>0$ and $Q_{\tau_g}<0$, and because $\frac{Q_{\tau_E}D_{\tau_g}}{Q_{\tau_g}D_{\tau_E}}<1$ by Lemma 5, it follows that $\tau_g^*-\phi$ has the same sign as $\mathcal{I}_{\tau_g}-\mathcal{I}_{\tau_E}$ and that τ_E^* has the same sign as $\mathcal{I}_{\tau_E}-\mathcal{I}_{\tau_g}\frac{Q_{\tau_E}D_{\tau_g}}{Q_{\tau_g}D_{\tau_E}}$. ■

Proof of Corollary 3. An immediate consequence of Eq. (10). ■

Proof of Lemma 1. For each k, set $\rho_k=\frac{\epsilon_{\tau_g}^k}{\epsilon_{\tau_E}^k}$. Then by definition, for $k<k'$ we have that

$$\frac{D_{\tau_E}^k}{D_{\tau_E}^{k'}}=\frac{\rho_k}{\rho_{k'}}\frac{D_{\tau_g}^k}{D_{\tau_g}^{k'}}>\frac{D_{\tau_g}^k}{D_{\tau_g}^{k'}}. \tag{28}$$

Now define the weights $\omega^k_{\tau_g} = D^k_{\tau_g}/D_{\tau_g}$ and $\omega^k_{\tau_E} = D^k_{\tau_E}/D_{\tau_E}$. Further, set $X_k = (1 - \Gamma_k)V$. Then by definition, $I_{\tau_g} = \sum_k X_k\omega^k_{\tau_g}$ and $I_{\tau_E} = \sum X_k \omega^k_{\tau_E}$. Now X_k is decreasing in k by definition. Eq. (28) implies that $\omega^k_{\tau_E}/\omega^{k'}_{\tau_E} > \omega^k_{\tau_g}/\omega^{k'}_{\tau_g}$ for all $k < k'$, or equivalently that $\sum_{k=1}^h \omega^k_{\tau_E} > \sum_{k=1}^h \omega^k_{\tau_g}$ for all $h < K$.

Consider now probability distributions Ω_{τ_g} and Ω_{τ_E} that place probabilities $\omega^k_{\tau_g}$ and $\omega^k_{\tau_E}$ on the outcome X_k. Then the condition $\sum_{k=1}^h \omega^k_{\tau_E} > \sum_{k=1}^h \omega^k_{\tau_g}$ for all $h < K$ implies that Ω_{τ_E} first-order stochastically dominates Ω_{τ_g}, and thus that $I_{\tau_E} = \sum_k X_k\omega^k_{\tau_E} > \sum_k X_k\omega^k_{\tau_g} = I_{\tau_g}$. ■

Proof of Lemma 2. We have

$$D^k_{\tau_E} = \frac{d}{d\tau_E}[1-G(\Gamma_k V-(p_E-p_I))] = g(\Gamma_k V-(p_E-p_I))$$
$$D^k_{\tau_g} = \frac{d}{d\tau_g}[1-G(\Gamma_k V-(p_E-p_I))] = \left[\frac{d}{d\tau_g}(\Gamma_k V)\right] g(\Gamma_k V-(p_E-p_I))$$

where g is the probability density function corresponding to G. Dividing the top equation by the bottom equation shows that $D^k_{\tau_g}/D^k_{\tau_E} = \frac{d}{d\tau_g}(\Gamma_k V)$. But by assumption, $\Gamma_k V$ is increasing in k, from which the statement of the lemma follows. ■

Proof of Corollary 4. Part 1 is immediately implied by setting $(\tau_g,\tau_E) = (\phi,0)$ in Eqs. (10) and (11).

For part 2, suppose that $\tau^*_g < 0$. Then by Proposition 2, it must be that $\mathcal{I}_{\tau_g} < \mathcal{I}_{\tau_E}$ at the optimal tax policy. Moreover, because no consumers overvalue energy efficiency, $(1 - \Gamma_k) \geq 0$ for all k, and thus $\mathcal{I}_{\tau_E} > 0$. Because $\frac{Q_{\tau_E} D_{\tau_g}}{Q_{\tau_g} D_{\tau_E}} < 1$, this then implies that $\mathcal{I}_{\tau_E} - \mathcal{I}_{\tau_g} \frac{Q_{\tau_E} D_{\tau_g}}{Q_{\tau_g} D_{\tau_E}} > 0$—which then implies that $\tau^*_E > 0$ by Eq. (13).

Part 3 is just a consequence of Part 2, Corollary 4, and Lemma 2. ■

Proof of Corollary 5. As shown in Lemma 5,

$$Q_{\tau_E} = D_{\tau_E}(e_E m^*_E - e_I m^*_I)$$
$$Q_{\tau_g} = D_{\tau_g}(e_E m^*_E - e_I m^*_I) + (1-D)\left(\frac{d}{d\tau_g} e_I m^*_I\right) + D\left(\frac{d}{d\tau_g} e_E m^*_E\right) < D_{\tau_g}(e_E m^*_E - e_I m^*_I)$$

Thus for $\tau_g \geq \phi$,

$$\frac{dW}{d\tau_g} d\tau_g < \left(\mathcal{I}_{\tau_g} - \tau_E\right) D_{\tau_g} d\tau_g - \left(\tau_g - \phi\right) D_{\tau_g}(e_I m^*_I - e_E m^*_E) d\tau_g$$
$$\frac{dW}{d\tau_E} d\tau_E = \left(\mathcal{I}_{\tau_g} - \tau_E\right) D_{\tau_E} d\tau_E - \left(\tau_g - \phi\right) D_{\tau_E}(e_I m^*_I - e_E m^*_E) d\tau_E$$

But since $D_{\tau_g} d\tau_g = D_{\tau_E} d\tau_E$, the conclusion follows. ■

Proof of Proposition 3. The proof of Lemma 5 verifies that $Q_{\tau_E} = D_{\tau_E}(e_E m^*_E - e_I m^*_I)$ as in Eq. (18). The proof of Lemma 2 shows that $D^k_{\tau_g} = \frac{d}{d\tau_g}(\Gamma_k V) D^k_{\tau_E}$, which formally establishes Eq. (19). The rest is established in the text. ■

Appendix II. Specific behavioral biases and their properties

II.A. Partial (exogenous) attention and present bias

Suppose that $\Gamma_k \equiv \gamma_k$ for constants $\gamma_k > 0$. It's clear that $\gamma_k V$ is increasing τ_g. Moreover, $\frac{d}{d\tau_g}(\gamma_k V) = \gamma_k \frac{d}{d\tau_g} V$ is clearly increasing in γ_k; thus this bias satisfies the assumption in Lemma 2. Finally, since $\frac{d}{d\tau_g}\gamma_k = 0$, the bias satisfies the assumption in Proposition 3.

II.B. Incorrect beliefs about energy intensity

Suppose that type k consumers make their extensive margin choice thinking that E consumes $\hat{e}^k_E$ units of energy and I consumes $\hat{e}^k_I$ units of energy, where $\hat{e}^k_E$ is decreasing in k while $\hat{e}^k_I$ is increasing in k. For example, consumers might be right about energy efficiency on average, but not adjust sufficiently: $\hat{e}_E = (k/K)e_E + (1-k/K)e_I$ and $\hat{e}_I = (k/K)e_I + (1-k/K)e_E$.

Let $\hat{m}^k_E$ and $\hat{m}^k_I$ denote a type k consumer's *forecast* of how much he will utilize E and I, given his beliefs about the energy efficiency of these products.

Type k consumers' bias is given by

$$\Gamma_k = \frac{V\left(\hat{e}^k_E, \hat{e}^k_I, p_g\right)}{V\left(e_E, e_I, p_g\right)}.$$

Now $V\left(\hat{e}^k_E, \hat{e}^k_I, p_g\right)$ is increasing in $\hat{e}^k_E$ and decreasing in $\hat{e}^k_I$, so it's clear that $\Gamma_k < \Gamma_{k'}$ if $k < k'$.

$\Gamma_k V = V\left(\hat{e}^k_E, \hat{e}^k_I, p_g\right)$ is also increasing in p_g by Lemma 4.

Finally, Lemma 4 shows that

$$\frac{d}{dp_g}(\Gamma_k V) = \hat{e}^k_I \hat{m}^*_I - \hat{e}^k_E \hat{m}^*_E.$$

Now by Lemma 3, $\hat{e}^k_E \hat{m}^*_E$ is decreasing in k, while $\hat{e}^k_I \hat{m}^*_I$ is increasing in k. Thus $\frac{d}{dp_g}(\Gamma_k V)$ is increasing in k, and therefore satisfies the assumption in Lemma 2.

II.C. Endogenous partial attention

The most reasonable assumption to make about endogenous partial attention is that $\frac{d}{dp_g}\Gamma_k \geq 0$; that is, that energy efficiency is more salient at high energy prices. This is consistent with models such as those of Koszegi and Szeidl (2013) or Gabaix (2012), since higher energy prices imply higher relative gains from purchasing more energy efficient products. Under this assumption, it is clear that $\Gamma_k V$ is increasing in p_g.

Γ_k also satisfies the assumption of Proposition 3 by definition. It is not necessarily true, however, that $\frac{d}{dp_g}(\Gamma_k V)$ is increasing in k, because it is possible that $\frac{d}{dp_g}\Gamma_k = 0$ for $k = K$, while $\frac{d}{dp_g}\Gamma_k$ may be very high for low k.

II.D. Combinations of biases

We could also consider combinations of these biases. For example, consumers might have misperceptions of energy efficiency and also be present-biased, so that their utility is given by $\gamma V\left(\hat{e}^k_E, \hat{e}^k_I, p_g\right)$ for some $\gamma < 1$. Clearly, any combination of the above biases will satisfy the minimal assumption that $\Gamma_k V$ is increasing in p_g. Moreover, a combination of incorrect beliefs and exogenous partial attention/present bias will still have the property that $\frac{d}{dp_g}(\Gamma_k V)$ is increasing in k, and a combination of present bias and endogenous partial attention will still have the property that $\frac{d}{dp_g}(\Gamma_k V)$ is increasing in k.

Appendix III. Generalizing the sufficient statistics formulas

In the more general case, we now allow for heterogeneity in utilization utility. In particular, we consider utility functions of the form $u(\theta,m)$, where m is utilization and θ is the consumer's *utilization type*. Again, we allow for correlations between θ,k and ϵ, and we set $D^{k,\theta}$ to denote the total energy efficiency of the products owned by consumers of type (k,θ). We let H denote the joint distribution of (k,θ). We set $m^*_j(\theta)$ to be the utilization of type θ consumer owning product j, and we set $V(\theta)$ to be the gross utility gain from energy efficiency of type θ consumers.

We begin by focusing on a two-product market, and then generalize to a multi-product market. Eqs. (18) and (19) are now generalized to

$$Q_{\tau_E} = \int_{k,\theta} (e_E m^*_E(\theta) - e_I m^*_I(\theta)) D^{k,\theta}_{\tau_E} dH \tag{29}$$

$$D_{\tau_g} = \int_{k,\theta} \left[\left(e_E m_E^*(\theta) - e_I m_I^*(\theta) \right) \Gamma_k + \left(\frac{d}{d\tau_g} \Gamma_k \right) V(\theta) \right] D_{\tau_E}^{k,\theta} dH \qquad (30)$$

As before, Eqs. (30) and (29) imply

$$\left(|Q_{\tau_E}| - D_{\tau_g} \right) p_g < \int_{k,\theta} (1-\Gamma_k) V(\theta) D_{\tau_E}^{k,\theta} dH, \qquad (31)$$

with

$$\frac{\int_{k,\theta} (1-\Gamma_k) V D_{\tau_E}^{k,\theta} dH}{\int_{k,\theta} D_{\tau_E}^{k,\theta} dH}$$

being the marginal internality in this more general setting.

Generalizing these computations to multiple products is also straightforward. For any two products j and j', consider all consumers indifferent between those two products. Then an analog of Eq. (29) for j and j' tells us the change in energy use that results from changing the choices of those marginal consumers. And an analog of Eq. (30) for j and j' tells us the change in energy efficiency that results from changing the choices of those marginal consumers. Finally, for this group of marginal consumers, an analog of Eq. (31) for products j and j' provides a lower bound for the welfare gains from changing the choices of those marginal consumers.

Thus to get the total welfare gains in the market, we take the sum of Eq. (31) over all pairs of products. But this is just the absolute value of the sum of Eq. (29) over all pairs of products minus absolute value of the sum of Eq. (30) over all pairs of products. And this directly leads to the generalization because the sum of (29) over all pairs of products is the total change in energy use while the sum of Eq. (30) over all pairs of products is the change in energy efficiency.

Appendix IV. Auto market simulation details

We incorporate heterogeneous utilization types using a special case of the framework in Appendix III: a type θ consumer derives utility $u(m-\theta)$ from utilization m. We assume a CRRA functional form for $u(m-\theta)$:

$$u\left(m_{ij} - \theta_i\right) = \frac{A}{1-r} \left(m_{ij} - \theta_i\right)^{1-r} \qquad (32)$$

Given this functional form, the choice of m_{ij} that maximizes utility in Eq. (36) below is:

$$m_{ij}^* = \theta_i + \left(\frac{\eta p_g e_j}{A} \right)^{-1/r} \qquad (33)$$

The parameter r is related to the price elasticity of utilization demand $\eta_{VMT} < 0$:

$$r = \frac{1}{-\eta_{VMT}} \frac{m_{ij}^* - \theta_i}{m_{ij}^*} \qquad (34)$$

We set A such that $\overline{\theta} = \frac{\overline{m}}{2}$, which ensures that elasticity does not vary too much over the support of θ. The distribution of θ is based on the empirical distribution of annual vehicle-miles traveled, using annualized odometer readings from the 2001 National Household Travel Survey (NHTS).

These annualized VMTs are re-scaled such that the average VMT over a 25-year potential lifetime matches the NHTS data. Specifically, we sum the average annual VMT $\overline{m}_a^*$ for vehicles of each age a from 1 to 25, giving $\sum_{a=1}^{25} \overline{m}_a^* \approx 236,000$. (These average annual VMTs decline from 14,500 when new to 9600 at age 12 and 4300 at age 25.)

We then multiply by a scaling factor Λ to translate this undiscounted sum over a potential lifetime to a discounted sum over an expected lifetime. Per Allcott and Wozny (forthcoming), we assume a six percent discount rate, giving a discount factor $\delta = \frac{1}{1.06}$. We use the R.L. Polk registration data for 1998 to 2007 to construct cumulative survival probabilities ϕ_a for vehicles of each age a. (A new vehicle has a 60% chance of surviving to age 12 and a 10% chance of surviving to age 25.) The scaling factor is:

$$\Lambda = \frac{\sum_{a=1}^{25} \delta^a \overline{m}_a^* \phi_a}{\sum_{a=1}^{25} \overline{m}_a^*} \approx 0.436 \qquad (35)$$

After these modifications, we now have an indirect utility function for use in the nested logit model. The indirect utility that consumer i experiences from purchasing product j, choosing optimal utilization m_{ij}^*, and receiving a transfer T is:

$$\left\{ Y_i + T - p_j - \Lambda p_g m_{ij}^* e_j \right\} + \frac{\Lambda}{\eta} u\left(m_{ij}^* - \theta_i \right) + \frac{\overline{\epsilon}_j + \epsilon_{ij}}{\eta} \qquad (36)$$

In this equation, the variable η is a scaling factor for the marginal utility of money, which is set such that the average vehicle's own-price elasticity of demand is −5. We calibrate each vehicle's mean utility $\overline{\epsilon}_j$ using the Berry et al. (1995) contraction mapping. The term in brackets is consumption of the numeraire good, while the two terms on the right represent the utility that the consumer derives from owning and using the vehicle. Welfare effects are calculated using Allcott's (2013) approach to calculating consumer surplus in logit models with biased consumers.

References

Abaluck, Jason, Gruber, Jonathan, 2011. Choice inconsistencies among the elderly: evidence from plan choice in the Medicare Part D Program. Am. Econ. Rev. 101 (4), 1180–1210 (June).

Allcott, Hunt, 2011. Consumers' perceptions and misperceptions of energy costs. Am. Econ. Rev. 101 (3), 98–104 (May).

Allcott, Hunt, 2013. The welfare effects of misperceived product costs: data and calibrations from the automobile market. Am. Econ. J. Econ. Policy 5 (3), 30–66.

Allcott, Hunt, Greenstone, Michael, 2012. Is there an energy efficiency gap? J. Econ. Perspect. 26 (1), 3–28 (Winter).

Allcott, Hunt, Taubinsky, Dmitry, 2013. The lightbulb paradox: evidence from two randomized experiments. NBER Working Paper No. 19713 (December).

Allcott, Hunt, Wozny, Nathan, 2014. Gasoline prices, fuel economy, and the energy paradox. Rev. Econ. Stat. (forthcoming).

Anderson, Soren, Kellogg, Ryan, Sallee, James, 2012. What do consumers believe about future gasoline prices? Working Paper. University of Chicago (September).

Attari, Shahzeen, DeKay, Michael, Davidson, Cliff, de Bruin, Wandi Bruine, 2010. Public perceptions of energy consumption and savings. Proc. Natl. Acad. Sci. 107 (37), 16054–16059 (September 14).

Austin, David, 2008. Effects of gasoline prices on driving behavior and vehicle markets. U.S. Congressional Budget Office Working Paper (January).

Austin, David, Dinan, Terry, 2005. Clearing the air: the costs and consequences of higher CAFE standards and increased gasoline taxes. J. Environ. Econ. Manag. 50 (3), 562–582.

Baicker, Kate, Mullainathan, Sendhil, Schwartzstein, Joshua, 2012. Choice hazard in health insurance. Working Paper. Dartmouth College (March).

Barber, Brad, Odean, Terrance, Zheng, Lu, 2005. Out of sight, out of mind: the effects of expenses on mutual fund flows. J. Bus. 78 (6), 2095–2120 (November).

Bernheim, B. Douglas, Rangel, Antonio, 2004. Addiction and cue-triggered decision processes. Am. Econ. Rev. 90 (5), 1558–1590 (December).

Bernheim, B. Douglas, Rangel, Antonio, 2005. From neuroscience to public policy: a new economic view of addiction. Swedish Econ. Pol. Rev. 12, 11–56.

Berry, Steven, Levinsohn, James, Pakes, Ariel, 1995. Automobile prices in market equilibrium. Econometrica 63 (4), 841–890 (July).

Bovenberg, Lans, Goulder, Lawrence, 1996. Optimal environmental taxation in the presence of other taxes: general equilibrium analyses. Am. Econ. Rev. 86 (4), 985–1000.

Brown, Jennifer, Hossain, Tanjim, John, Morgan, 2010. Shrouded attributes and information suppression: evidence from the field. Q. J. Econ. 125 (2), 859–876 (May).

Busse, Meghan, Knittel, Christopher, Zettelmeyer, Florian, 2013. Are consumers myopic? Evidence from new and used car purchases. Am. Econ. Rev. 103 (1), 220–256 (February).

Chetty, Raj, 2009. Sufficient statistics for welfare analysis: a bridge between structural and reduced-form methods. Annu. Rev. Econ. 1, 451–488.

Chetty, Raj, Looney, Adam, Kroft, Kory, 2009. Salience and taxation: theory and evidence. Am. Econ. Rev. 99 (4), 1145–1177 (September).
Davis, Lucas, 2008. Durable goods and residential demand for energy and water: evidence from a field trial. RAND J. Econ. 39 (2), 530–546 (Summer).
Davis, Lucas, Kilian, Lutz, 2011. Estimating the effect of a gasoline tax on carbon emissions. J. Appl. Econ. 26 (7), 1187–1214 (November-December).
DellaVigna, Stefano, 2009. Psychology and economics: evidence from the field. J. Econ. Lit. 47 (2), 315–372 (June).
Diamond, Peter, 1973. Consumption externalities and imperfect corrective pricing. Bell J. Econ. Manag. Sci. 4 (2), 526–538 (Autumn).
Dreyfus, Mark, Viscusi, Kip, 1995. Rates of time preference and consumer valuations of automobile safety and fuel efficiency. J. Law Econ. 38 (1), 79–98.
Dubin, Jeffrey, McFadden, Daniel, 1984. An econometric analysis of residential electric appliance holdings and consumption. Econometrica 52 (2), 345–362 (March).
Finkelstein, Amy, 2009. EZ-tax: tax salience and tax rates. Q. J. Econ. 124 (3), 969–1010 (August).
Fischer, Carolyn, Harrington, Winston, Parry, Ian, 2007. Do market failures justify tightening Corporate Average Fuel Economy (CAFE) Standards? Energy J. 28 (4), 1–30.
Fullerton, Don, West, Sarah, 2002. Can taxes on cars and on gasoline mimic an unavailable tax on emissions? J. Environ. Econ. Manag. 43, 135–157.
Fullerton, Don, West, Sarah, 2010. Tax and subsidy combinations for the control of car pollution. B.E. J. Econ. Anal. Policy Adv. 10 (1), 1–31 (January).
Gabaix, Xavier, 2012. A sparsity-based model of bounded rationality. Working Paper. New York University (September).
Gallagher, Kelly Sims, Muehlegger, Erich, 2011. Giving green to get green? Incentives and consumer adoption of hybrid vehicle technology. J. Environ. Econ. Manag. 61 (January), 1–15.
Gately, Dermot, 1980. Individual discount rates and the purchase and utilization of energy-using durables: comment. Bell J. Econ. 11 (1), 373–374 (Spring).
Gillingham, Kenneth, 2013. Identifying the elasticity of driving: evidence from a gasoline price shock in California. Working Paper. Yale University (May).
Gillingham, Kenneth, Palmer, Karen, 2012. Bridging the energy efficiency gap: insights for policy from economic theory and empirical analysis. Working Paper. Yale University.
Gillingham, Kenneth, Harding, Matthew, Rapson, David, 2010. Split incentives in residential energy consumption. Working Paper. Stanford University.
Goldberg, Pinelopi, 1998. The effects of the corporate average fuel efficiency standards in the US. J. Ind. Econ. 46 (1), 1–33 (March).
Greene, David, 2010. How consumers value fuel economy: a literature review. US Environmental Protection Agency Technical Report EPA-420-R-10-008 (March).
Greene, David, Patterson, Philip, Singh, Margaret, Li, Jia, 2005. Feebates, rebates, and gas-guzzler taxes: a study of incentives for increased fuel economy. Energy Policy 33 (6), 757–775 (April).
Greenstone, Michael, Kopits, Elizabeth, Wolverton, Maryann, 2011. Estimating the social cost of carbon for use in U.S. federal rulemakings: a summary and interpretation. NBER Working Paper No. 16913 (March).
Grubb, Michael, Osborne, Matthew, 2012. Cellular service demand: biased beliefs, learning, and bill shock. Working Paper. MIT (February).
Gul, Faruk, Pesendorfer, Wolfgang, 2007. Harmful addiction. Rev. Econ. Stud. 74 (1), 147–172 (January).
Handel, Benjamin, 2011. Adverse selection and switching costs in health insurance markets: when nudging hurts. NBER Working Paper No. 17459 (September).
Harberger, Arnold, 1964. The measurement of waste. Am. Econ. Rev. 54 (3), 58–76.
Hausman, Jerry, 1979. Individual discount rates and the purchase and utilization of energy-using durables. Bell J. Econ. 10 (1), 33–54.
Herrnstein, R.J., Loewenstein, George, Prelec, Drazel, Vaughan Jr., William, 1993. Utility maximization and melioration: internalities in individual choice. J. Behav. Decis. Mak. 6, 149–185.
Heutel, Garth, 2011. Optimal policy instruments for externality-producing durable goods under time inconsistency. Working Paper. UNC Greensboro (May).
Hughes, Jonathan, Knittel, Christopher, Sperling, Daniel, 2008. Evidence of a shift in the short-run price elasticity of gasoline demand. Energy J. 29 (1), 113–134.
Innes, Robert, 1996. Regulating automobile pollution under certainty, competition, and imperfect information. J. Environ. Econ. Manag. 31, 219–239.
Ito, Koichiro, 2014. Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. Am. Econ. Rev. 104 (2), 537–563.
Jacobsen, Mark, 2010. Evaluating U.S. fuel economy standards in a model with producer and household heterogeneity. Working Paper. University of California at San Diego (September).
Kilian, Lutz, Sims, Eric, 2006. The effects of real gasoline prices on automobile demand: a structural analysis using micro data. Working Paper. University of Michigan (April).
Kleit, Andrew, 2004. Impacts of long-range increases in the Corporate Average Fuel Economy (CAFE) standard. Econ. Inq. 42 (2), 279–294.
Klier, Thomas, Linn, Joshua, 2010. The price of gasoline and new vehicle fuel economy: evidence from monthly sales data. Am. Econ. J. Econ. Policy 2 (3), 134–153 (August).
Knittel, Christopher, Sandler, Ryan, 2013. The welfare impact of indirect Pigouvian Taxation: evidence from transportation. Working Paper. MIT (February).
Koszegi, Botond, Szeidl, Adam, 2013. A model of focusing in economic choice. Q. J. Econ. 128 (1), 53–107.
Kroft, Kory, 2011. Review of policy and choice: public finance through the lens of behavioral economics. J. Econ. Lit. 49 (4), 1241–1250 (December).
Krupnick, Alan, Parry, Ian, Walls, Margaret, Knowles, Tony, Hayes, Kristin, 2010. Toward a new national energy policy: assessing the options. Resources for the Future Report (November, http://www.energypolicyoptions.org/wp-content/uploads/reports/RFF-Rpt-NEPI_Tech_Manual_Final.pdf).
Laibson, David, 1997. Golden eggs and hyperbolic discounting. Q. J. Econ. 112 (2), 443–477 (May).
Larrick, Richard, Soll, Jack, 2008. The MPG illusion. Science 320 (5883), 1593–1594.
Li, Shanjun, Timmins, Christopher, von Haefen, Roger, 2009. How do gasoline prices affect fleet fuel economy? Am. Econ. J. Econ. Policy 1 (2), 113–137 (August).
Li, Shanjun, Linn, Joshua, Muehlegger, Erich, 2012. Gasoline taxes and consumer behavior. Working Paper. Harvard Kennedy School (October).
Liebman, Jeffrey, Zeckhauser, Richard, 2004. Schmeduling. Working Paper. Harvard Kennedy School.
Loewenstein, George, Prelec, Drazen, 1992. Anomalies in intertemporal choice: evidence and an interpretation. Q. J. Econ. 107, 573–597.
Mullainathan, Sendhil, Schwartzstein, Joshua, Congdon, William, 2012. A reduced form approach to behavioral public finance. Annu. Rev. Econ. 4, 511–540.
National Highway Traffic Safety Administration (NHTSA), 2010. Final Regulatory Impact Analysis: Corporate Average Fuel Economy for MY 2012-MY 2016 Passenger Cars and Light Trucks. Office of Regulatory Analysis and Evaluation, National Center for Statistics and Analysis (March).
O'Donoghue, Edward, Rabin, Matthew, 1999. Doing it now or doing it later. Am. Econ. Rev. 89, 103–121.
O'Donoghue, Edward, Rabin, Matthew, 2006. Optimal sin taxes. J. Public Econ. 90, 1825–1849.
Parry, Ian, 1995. Pollution taxes and revenue recycling. J. Environ. Econ. Manag. 29, S64–S77.
Parry, Ian, Walls, Margaret, Harrington, Winston, 2007. Automobile externalities and policies. J. Econ. Lit. 45 (2), 374–400 (June).
Parry, Ian, Evans, David, Oates, Wallace, 2010. Are energy efficiency standards justified? Resources for the Future Discussion Paper 10-59 (November).
Sallee, James, 2011. The Taxation of Fuel Economy. (Chapter 1) In: Brown, Jeffrey (Ed.), Tax Policy and the Economy. , vol. 25. National Bureau of Economic Research, Cambridge, MA.
Sallee, James, 2012. Rational inattention and energy efficiency. Working Paper. University of Chicago (September).
Sallee, James, West, Sarah, Fan, Wei, 2009. Consumer valuation of fuel economy: a microdata approach. Working Paper. University of Chicago (September, Available from http://www.ntanet.org/images/stories/pdf/proceedings/09/032.pdf).
Sawhill, James, 2008. Are capital and operating costs weighted equally in durable goods purchases? Evidence from the US automobile market. Working Paper. University of California at Berkeley (April).
Small, Kenneth, Van Dender, Kurt, 2007. Fuel efficiency and motor vehicle travel: the declining rebound effect. Energy J. 28 (1), 25–51.
Strotz, Robert, 1955. Myopia and inconsistency in dynamic utility maximization. Rev. Econ. Stud. 23, 165–180.
Verboven, Frank, 1999. The markets for gasoline and diesel cars in Europe. Centre for Economic Policy Research Discussion Paper 2069 (February).
Verboven, Frank, 2002. Quality-based price discrimination and tax incidence: evidence from gasoline and diesel cars. RAND J. Econ. 33 (2), 275–323 (Summer).