The Implications of Heterogeneity for the Regulation of Energy-Consuming Durable Goods

Mark R. Jacobsen,[1,5] Christopher R. Knittel,[2,5] James M. Sallee,[3,5] Arthur A. van Benthem[4,5*]

December 8, 2014

Abstract

Many of the most important policies that aim to reduce greenhouse gas emissions and other environmental externalities do so by regulating the energy efficiency of energy-consuming durable goods. We document a hitherto unexplored connection between heterogeneity in the utilization of such durable goods and the economic efficiency of this class of policies. Inefficiency arises because products with the same energy efficiency rating have different lifetime utilizations, and hence different lifetime emissions, but are given equal policy treatment. We develop a model that characterizes sufficient statistics for the deadweight loss from using these second-best policies in lieu of efficient Pigouvian taxes in the presence of such utilization heterogeneity. Most notably, under some plausible assumptions, the R^2 from a regression of the lifetime emissions of products on their energy efficiency ratings is equal to the fraction of the first-best welfare gain that can be achieved by energy efficiency regulations that, like Corporate Average Fuel Economy standards, impose an (implicit) linear tax on energy efficiency. We explore the quantitative importance of heterogeneity for the case of automobile fuel economy regulations using data on vehicle mileage shortly before scrappage. We document significant dispersion in lifetime mileage of different types of vehicles that share a common fuel economy rating. We estimate that this heterogeneity implies that fuel-economy regulations can achieve only about one quarter of the welfare gain from an optimally designed policy.

Keywords: Corrective taxation, externalities, heterogeneity, sufficient statistic
JEL: H23, Q58, L51

*We thank seminar participants at UC Berkeley and the University of Chicago for helpful comments and suggestions.

[1]Department of Economics, University of California San Diego. [2]William Barton Rogers Professor of Energy Economics, Sloan School of Management, Director, Center for Environmental and Energy Policy Research, Massachusetts Institute of Technology. [3]Harris School of Public Policy, University of Chicago. [4]The Wharton School, University of Pennsylvania. [5]National Bureau of Economic Research. © 2014 Mark Jacobsen, Christopher Knittel, James Sallee and Arthur van Benthem. All rights reserved.

1 Introduction

The consumption of energy, which is both a tremendous engine of economic growth and the principal driver of global climate change, is nearly always achieved through the operation of some durable good. Motor vehicles combust gasoline; appliances use electricity; furnaces burn natural gas; and so on. To correct the market failures caused by the pollution that attends energy consumption, economists typically advocate the pricing of emissions. As argued by Pigou (1932), if externalities can be taxed directly, then market efficiency can be fully restored. Such policy prescriptions are indifferent to the durables that act as an intermediary between fuel inputs and emissions outputs.

Policies that directly target emissions are, however, relatively rare. Instead, a proliferation of policies focus on these durable intermediaries, often through the regulation of their energy efficiency. Examples include fuel economy regulations, appliance efficiency mandates, and building codes. Energy efficiency policies are known to suffer from inefficiencies due to the "rebound effect"; by lowering the cost of using a durable good, energy efficiency regulations induce additional use on the margin.[1] The aim of this paper is to establish another inefficiency of such policies, one that stems from heterogeneity in how durables are used.

To see the logic of our inquiry, note that two goods with the same energy efficiency rating may have very different total lifetime emissions because they may have different realized utilizations. One refrigerator may last longer than another. Or, consumers may drive certain types of cars more than others. A policy that directly taxes emissions would account for differences in realized utilization; e.g., someone who drives their car more will pay more under a carbon tax. But, policies that regulate the durables themselves, rather than emissions, are almost always limited in their ability to account for such heterogeneity in utilization. This limitation implies a substantial inefficiency, and this is the subject of our investigation.

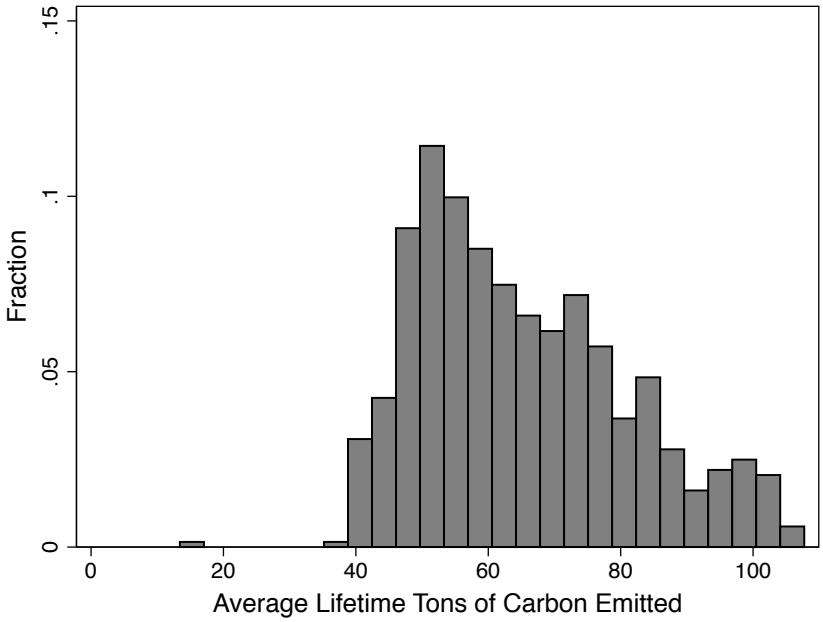
Most commonly, energy efficiency regulations must give the same policy treatment to two products that have the same official energy efficiency rating. For example, Corporate Average Fuel Economy (CAFE) standards require automakers to sell vehicles that meet or exceed a specified sales-weighted average fuel economy. This creates an implicit tax or subsidy (a shadow price) on each car sold that is a linear function of its fuel consumption rate. Thus, all vehicles that share a fuel economy rating will have a common implicit tax or subsidy.[2] But, vehicles with a common fuel economy rating in fact vary widely in their lifetime emissions, because some types of cars last longer than others.

To illustrate this heterogeneity, we plot the average lifetime carbon emissions for different types

[1] For example, fuel economy regulations that spur consumers to buy more efficient cars implicitly lower the cost of driving an additional mile, which induces people to drive more. For reviews of the literature on the efficiency of fuel economy regulation and taxation, see Anderson, Parry, Sallee, and Fischer (2011) and Sallee (2011), respectively. See Borenstein (Forthcoming) for a recent treatment of the economics of the rebound effect.

[2] Note that this is a stylized representation of CAFE, which requires each automaker to meet the standard separately. Until recently, there was no trading of compliance credits across firms, so each firm had a unique shadow price. In addition, cars and light-duty trucks are regulated separately. We also abstract from a recent policy change in which a vehicle's fuel-economy standard is conditioned on a vehicle's footprint – a measure of vehicle size determined by multiplying the vehicle's wheelbase by its average track width.

Figure 1: Distribution of Lifetime Carbon Emissions for Vehicles with EPA Average Rating of 23 Miles per Gallon



of cars that have the same fuel economy rating in Figure 1. The figure uses data, which we describe in detail below, on the odometer readings of vehicles shortly before they are scrapped, which we convert into tons of carbon using the vehicle's fuel economy rating and the average carbon content of gasoline. The graph suggests a wide dispersion in lifetime emissions among vehicles with a common 23 miles per gallon rating (the median in our data); the standard deviation is 20% of the median. At $39 per ton, which is the current federal guideline for the social cost of carbon, the standard deviation in damages across cars with the same fuel economy rating is over $600. Thus, a policy like CAFE that must give all cars with the same fuel economy rating the same shadow price is necessarily imprecise; it places the same implicit tax or subsidy on products that in fact have substantially different lifetime externalities. The goal of our paper is to understand how this imprecision affects the welfare properties of policies that regulate durable goods, as compared to efficient policies that target emissions.

To achieve this goal, we use both theory and data. We first develop a model with a representative consumer that allows us to derive sufficient statistics for the deadweight loss of using energy efficiency regulations in lieu of an efficient emissions tax. In this model, which we describe in Section 2, consumers are all alike, but there is heterogeneity in lifetime utilization of goods due to differences in product durability. In this setting, product-based policies can be fully efficient if they assign a unique (implicit) tax or subsidy to each product according to both its energy efficiency rating and its average lifetime durability. Energy efficiency policies are, however, generally limited in that they must impose regulatory treatment that is a function of only energy efficiency ratings.

Our model characterizes the deadweight loss of utilizing such restricted second-best policies

using a parsimonious set of parameters (sufficient statistics). We emphasize results relating to policies, like CAFE, that impose a linear tax on products according to their efficiency ratings. In the simplest case, the fraction of the welfare gain achieved by the first-best policy—as compared to a baseline policy that sets a single tax rate on all products—that can be achieved by a linear tax on energy efficiency is equal to the R^2 from a regression of the lifetime emissions of products on their energy efficiency ratings. The intuition for this is that the first-best policy assigns each product a tax equal to its lifetime damages, and the ability of the linear tax to mimic this differentiation is determined by the degree to which efficiency ratings predict lifetime emissions. This result links directly to graphs like Figure 1; when such figures show a greater spread in lifetime emissions for a given energy efficiency rating, the R^2 of a regression of lifetime emissions on energy efficiency ratings will be lower and the policy's deadweight loss will be greater. Under alternative assumptions, terms other than this R^2 will also influence welfare, but the R^2 remains an important factor.

In turn, the sum of squared residuals from this regression, along with an estimate of the marginal damages of emissions, can be used to express the deadweight loss of using second-best policies (as compared to the first-best benchmark) in terms of dollars. To demonstrate the quantitive importance of our theoretical results for actual policy, we investigate the case of automobile regulation using data on the lifetime mileage of different types of cars. We describe our data in Section 3. In Section 4 we use these data to quantify the efficiency of a linear tax on fuel consumption relative to an efficient policy. We find this CAFE-style policy recovers only about 25% of the welfare gains that are achieved by the efficient policy.[3]

We then generalize our model to account for heterogeneity across types of consumers, who may use durable goods differently, in Section 5. In this setting, the first-best allocation is achievable (assuming no rebound effect) when each consumer has a unique set of taxes on all products. Such a policy is infeasible, and we focus our study on policies that cannot discriminate across consumers and admit only a single tax rate on each product. We derive the second-best tax rates in this case. Our results demonstrate that a vast amount of information is required in order to formulate optimal policy; the second-best rates depend on a matrix of interactions between cross-product price elasticities and marginal external damages. In some sense, our formulas provide a negative result: to implement the second-best product-based policy, policymakers would need to know a great deal of information that is not readily available. An advantage of direct taxes on emissions is that they do not require this information. Our analysis in this section is related to Diamond (1973), which considers the second-best tax rates for a single consumption good that causes a different externality

[3]Throughout the paper, we focus on the question of how heterogeneity in utilization affects the welfare properties of two policies that are optimally designed, one of which faces a constraint, in the absence of a rebound effect. Actual policies, including CAFE, suffer from additional inefficiencies compared to the ideal both because of the rebound effect and because they are not necessarily second-best, even given the constraints we consider here. In particular, CAFE gets the average implicit tax on vehicles wrong because it taxes some vehicles but subsidizes others. It should instead raise the price of all cars on average (because all produce externalities), but it does not necessarily do so. This is similar to the extensive margin problem of performance standards considered in Holland, Hughes, and Knittel (2009). CAFE could be combined with a fixed tax on all cars to achieve the second-best rate. Our focus is on the role of heterogeneity, because that is our contribution in this paper, but a full assessment of real-world policies, including CAFE, must also consider these other factors.

when consumed by different individuals. Our model generalizes the results of Diamond (1973) to the case of many interrelated goods (e.g., types of cars) that each cause externalities. Doing so highlights the vast informational requirements for implementing second-best policies. We plan to develop simulations to assess how much welfare is lost when resorting to product-based standards.

Our analysis points to inefficiencies of product-based policies that have gone largely undiscussed in the energy policy literature. In some cases, product-based policies may be chosen because emissions cannot be taxed or regulated directly. In such cases, our analysis shows the cost of second-best policies compared to an unobtainable first-best. But, policies that target durables, rather than emissions, often persist even when more direct and efficient policies are feasible. For example, it would be technologically challenging to monitor the greenhouse gas emissions emitted from the tailpipe of individual automobiles. It is straightforward, however, to tax motor fuel according to its carbon content, and thus a policy very similar to an efficient tax on carbon emissions can be easily mimicked by a gasoline tax. Nevertheless, we rely on CAFE. In cases such as this, our results point to the welfare costs of preferring one form of regulation, perhaps because of political economy or equity considerations, over other feasible instruments.

2 Model with representative consumer

We begin with a representative consumer model. This allows us to establish the effects of heterogeneity in damages tied to products, rather than heterogeneity in the damages caused by two different consumers who purchase the same product. In this section, we abstract from the rebound effect for expositional clarity.[4]

In this setting, with a representative consumer and no rebound effect, product-based taxes can achieve the same efficiency as an emissions or fuel tax by taxing each product according to its ex ante expected lifetime damages. Actual policies, however, create limitations on the incentive schedule. For example, CAFE puts an implicit linear tax on a vehicle's per-mile fuel consumption.[5] These implicit taxes and subsidies will differ from the efficient policy to the extent that different vehicles that share a fuel economy rating have different average lifetime mileage (and hence different lifetime fuel consumption).

2.1 Model setup

We consider a general setup and derive a sufficient statistic to illustrate the deadweight loss of moving from a policy that charges each product according to its actual lifetime damages, to one that taxes it according to its fuel efficiency rating (or some restricted function of fuel efficiency ratings). Specifically, we consider a representative consumer who purchases a portfolio of differ-

[4] At present, the paper does not have any results that incorporate the rebound effect. Our plan is to abstract from it in the theory, but reintroduce it our empirical analysis via simulation, which will allow us to compare the magnitude of distortions from the rebound effect with the distortions due to heterogeneity.

[5] Total fuel consumption for a given number of miles traveled is linearly related to per-mile fuel consumption, but non-linearly related to fuel economy. CAFE creates a shadow price linear in per-mile fuel consumption.

ent types of goods, indexed $j = 1, ..., J$. For example, each j would be a particular automobile model. The quantities chosen of each are denoted x_j. We assume that the set of models, and their characteristics, are fixed. This marks our analysis as, in some sense, short-run.[6]

Consumer's derive utility, U, from the consumption of these products: $U(x_1, ..., x_J)$. Each product generates an externality, denoted ϕ_j, which we assume is linearly related to quantities and enters the social welfare function separably.[7] We denote the cost of production by $C(x_1, ..., x_J)$. There is an exogenous amount of income in the economy, M, and all remaining income is consumed in a quasi-linear numeraire, n. The planner maximizes social welfare, denoted W, which is just utility from the durable goods and numeraire, net of externalities. We write W substituting the budget constraint in for the numeraire:

$$W = U(x_1, ..., x_J) + M - C(x_1, ..., x_J) - \sum_{j=1}^{J} \phi_j x_j. \tag{1}$$

We assume that there are many consumers, so that they do not internalize the externality when making their choice. And, we assume that the supply side is perfectly competitive. We assume that both $U(\cdot)$ and $C(\cdot)$ are twice differentiable and are weakly convex and concave, respectively.

The consumer perceives the prices of all goods to be constants, written as p_j, and they pay taxes t_j. Thus, the consumer chooses a vector of goods to maximize:

$$\begin{aligned} \max_{x_1,...,x_J} \; & Z = U(x_1, ..., x_J) + n \\ \text{s.t.} \; & \sum_{j=1}^{J} (p_j + t_j) x_j + n \leq M. \end{aligned} \tag{2}$$

The consumer's first-order conditions imply that $\frac{\partial U}{\partial x_j} = (p_j + t_j)$. Under marginal cost pricing, this implies that $\frac{\partial U}{\partial x_j} - \frac{\partial C}{\partial x_j} = t_j$, which is just the standard tax wedge between marginal utility and marginal cost.

The planner can maximize utility by setting tax rates equal to external damages, that is, by choosing $t_j = \phi_j$ for all products. This simply imposes a Pigouvian tax on each good, and it fully internalizes the externality—the first-best is achieved. This can be seen by differentiating (1) with respect to each tax rate t_j:

$$\frac{dW}{dt_j} = \sum_{k=1}^{J} \left(\frac{\partial U}{\partial x_k} - \frac{\partial C}{\partial x_k} - \phi_k \right) \frac{\partial x_k}{\partial t_j} = \sum_{k=1}^{J} (t_k - \phi_k) \frac{\partial x_k}{\partial t_j}. \tag{3}$$

[6] Note that our model does allow for quantities to be zero.

[7] The ϕ terms can equally be thought of as the expected lifetime damages for a product that has a distribution of lifetime emissions, where uncertainty stems, for example, from random failure of a durable good. Heterogeneity in realized damages due to random failure will not influence the relative efficiency of product-based taxes as compared to fuel taxes, though it will result in different incidence on consumers who get above, or below, average lifetimes from their product.

Our aim is to characterize the welfare loss of deviating from this first-best schedule of Pigouvian taxes, where $t_j = \phi_j \ \forall j$, to some arbitrary alternative tax schedule. We can then consider particular alternative tax schedules, such as the second-best tax system where tax rates are a linear function of some other variable that is imperfectly related to damages, such as energy efficiency. To characterize the deadweight loss induced by an alternative tax schedule, we follow the sufficient statistics tradition of differentiating W with respect to the tax and integrating. But, we have many tax rates and thus employ an intermediate algebraic step to derive our expression.

2.2 Derivation of sufficient statistics

Let any generic tax schedule be denoted as $\tau_1, ..., \tau_J$. We characterize the welfare loss of moving from the tax schedule $t_j = \phi_j$ to $t_j = \tau_j$ by specifying a weighted average of the two tax schedules and then integrating the marginal welfare losses of moving the weights from ϕ_j to τ_j. That is, we specify the function $t_j = (1-\rho)\phi_j + \rho\tau_j$. We will differentiate W with respect to ρ, and then characterize the welfare loss of moving from the optimal policy (when $\rho = 0$) to the alternative policy (when $\rho = 1$).

First, we differentiate equation (1) with respect to ρ and substitute in the consumer's optimality condition. This yields:

$$\frac{dW}{d\rho} = \sum_{j=1}^{J}\sum_{k=1}^{J}\left(\frac{\partial U}{\partial x_j} - \frac{\partial C}{\partial x_j} - \phi_j\right)\frac{\partial x_j}{\partial t_k}\frac{\partial t_k}{\partial \rho} = \sum_{j=1}^{J}\sum_{k=1}^{J}(t_j - \phi_j)\frac{\partial x_j}{\partial t_k}\frac{\partial t_k}{\partial \rho}. \tag{4}$$

This term, $\frac{dW}{d\rho}$, is the incremental change in welfare as we move from the first-best rates toward the alternative tax schedule, where all rates move by an amount proportional to the difference between the first-best and the alternative taxes.[8] However, this object is not of particular interest to us; it is only an intermediate step that enables us to characterize deadweight loss in terms of demand derivatives (which are estimable) instead of the utility function (which is more difficult to recover with data).

Next note that, by our definition, $\frac{\partial t_k}{\partial \rho} = (\tau_k - \phi_k)$. This term, $(\tau_k - \phi_k)$, is the difference between the first-best and actual tax rates, which we call the "tax residual" for reasons that will become apparent. We use that substitution, as well as the definition of t_j, and simplify:

$$\begin{aligned}\frac{dW}{d\rho} &= \sum_{j=1}^{J}\sum_{k=1}^{J}\left(\{(1-\rho)\phi_j + \rho\tau_j\} - \phi_j\right)\frac{\partial x_j}{\partial t_k}(\tau_k - \phi_k) \\ &= \sum_{j=1}^{J}\sum_{k=1}^{J}\rho\left(\tau_j - \phi_j\right)\frac{\partial x_j}{\partial t_k}(\tau_k - \phi_k)\end{aligned} \tag{5}$$

[8] This procedure is related to the technique developed in Hendren (2013), who labels an object that is similar to our $\frac{dW}{d\rho}$ the "policy elasticity".

Because ρ is a constant, we can remove it from the summation, which yields the final equation. To get the change in social surplus from moving fully between the two tax schedules, we integrate from $\rho = 0$ to $\rho = 1$. If the demand derivatives are constant over the relevant range, then ρ can be pulled out in front of the summation and the integration is straightforward. This yields:

$$W(\rho = 0) - W(\rho = 1) = -\frac{1}{2} \sum_{j=1}^{J} \sum_{k=1}^{J} (\tau_j - \phi_j)(\tau_k - \phi_k) \frac{\partial x_j}{\partial t_k}. \tag{6}$$

This formula bears a resemblance to traditional tax wedge formulas, it is roughly in the form of 1/2 times a tax wedge squared times a quantity derivative.[9]

This expression can be decomposed to gain further insight. We denote $W(\rho = 1) - W(\rho = 0)$ as simply $DWL(\tau)$ and use the substitution $e_j \equiv (\tau_j - \phi_j)$ for the tax residuals. Rearranging equation (6) yields:

$$-2 \times DWL(\tau) = \underbrace{\sum_{j=1}^{J} e_j^2 \frac{\partial x_j}{\partial t_j}}_{\text{“own effects”}} + \underbrace{\sum_{j=1}^{J} \sum_{k \neq j} e_j e_k \frac{\partial x_j}{\partial t_k}}_{\text{“cross effects”}}. \tag{7}$$

Equation (7) separates the deadweight loss formula into a term related to own-price derivatives (the way that a tax affects demand for the product itself) and a term related to cross-price derivatives (the way that a tax affects demand for alternative products). Under some conditions, which we discuss in detail below, the cross-price term will be zero, or at least small compared to the own-price term. We start with this special case to build intuition and then reintroduce the cross-price term.

When the cross-price term is zero, equation (7) simplifies to:

$$-2 \times DWL(\tau) = \sum_{j=1}^{J} e_j^2 \frac{\partial x_j}{\partial t_j} \tag{8}$$

$$= \underbrace{\sum_{j=1}^{J} e_j^2}_{\text{SSR}} \times \underbrace{J^{-1} \sum_{j=1}^{J} \frac{\partial x_j}{\partial t_j}}_{\text{Mean Derivative}} + \underbrace{\sum_{j=1}^{J} \left(e_j^2 - \bar{e^2}\right) \left(\frac{\partial x_j}{\partial t_j} - \frac{\partial \bar{x}_j}{\partial t_j}\right)}_{\text{Covariance}}, \tag{9}$$

[9] This equation can alternatively be written in matrix notation. Note that

$$\frac{dW}{d\rho} = \begin{pmatrix} \tau_1 - \phi_1 & \dots & \tau_J - \phi_J \end{pmatrix} \begin{pmatrix} \frac{dx_1}{dt_1} & \dots & \frac{dx_1}{dt_J} \\ \dots & \dots & \dots \\ \frac{dx_J}{dt_1} & \dots & \frac{dx_J}{dt_J} \end{pmatrix} \begin{pmatrix} \tau_1 - \phi_1 \\ \dots \\ \tau_J - \phi_J \end{pmatrix} = (\tau - \phi)^T D (\tau - \phi),$$

where D is the substitution matrix and $(\tau - \phi)$ is the vector of tax residuals. We can now rewrite

$$W(\rho = 0) - W(\rho = 1) = -\frac{1}{2} (\tau - \phi)^T D (\tau - \phi).$$

.

where bars indicate means of variables. This first line, equation (8), shows that the deadweight loss is determined by the magnitude of the errors in the tax rate and the partial derivatives of demand. The error in the tax rate e_j is the tax wedge, relative to the socially efficient value. Thus, this formulation demonstrates that deadweight loss is the summation of a set of Harberger triangles across all the products. This is intuitive, as the assumption of the net zero cross-price terms implies that the sum of the partial equilibrium tax wedges are equal to the general equilibrium effect that considers all taxes.

Equation (9) provides additional insight by decomposing the Harberger triangle version into three parts, using the definition of the covariance. The first term is the sum of squared errors in the tax rate, multiplied by the average demand response (own-price derivative). We label the sum of squared errors as SSR because in an intuitive special case, discussed below, this will be precisely the sum of squared residuals from an OLS regression of lifetime damages ϕ_j on some observable metric upon which policy is based. The second term is the covariance between the errors in the tax rate and the own-price derivatives.

This formulation is particularly useful for understanding second-best linear tax schedules because it relates directly to ordinary least squares regression. Specifically, suppose that a planner is constrained to choose a tax schedule of the form $\tau_j = \alpha + \beta\theta_j$, where θ_j is some observable metric, such as a government energy efficiency rating. (CAFE, loosely defined, is such a policy.) The planner's second-best problem is then to choose α and β to minimize deadweight loss (from equation (9)). When true damages ϕ_j are uncorrelated with own-price derivatives (so that the covariance term is zero), this second-best problem collapses to the standard OLS problem of minimizing residuals, where the residuals are the errors in the tax rates:

$$\min_{\alpha,\beta} = -\sum_{j=1}^{J} \frac{\partial x_j}{\partial t_j} \times \sum_{j=1}^{J} (\phi_j - \alpha - \beta\theta_j)^2. \tag{10}$$

This is the standard OLS optimization problem, and the solution will be the standard OLS regression coefficients. In that case, the sum of squared errors in the tax schedule will be exactly the sum of squared residuals from an OLS regression of lifetime damages ϕ_j on θ_j, which is the metric that determines policy treatment.

For example, θ_j might be official fuel consumption (inverse of fuel economy) ratings. As discussed above, CAFE implicitly imposes a set of taxes and subsidies for vehicles that are a linear function of fuel consumption.[10] If there were no heterogeneity in average lifetime mileage of different models of cars, then this linear function could exactly predict lifetime damages; i.e., the R^2 of a regression of lifetime fuel consumption on fuel consumption ratings would be 1, and the sum of squared residuals would be zero. To the extent that there is utilization heterogeneity, however, the observable metric (fuel consumption) will be imperfectly correlated with true damages. The

[10]Note that CAFE is additionally constrained on the extensive margin because vehicles with fuel-economy ratings larger than the standard be are subsidized, while vehicles with fuel-economy ratings below the standard are taxed.

deadweight loss of the second-best linear policy is then directly tied to the sum of squared residuals from the linear regression, as captured in equation (10).

Intuitively, the deadweight loss from these imperfect tax rates are scaled by the size of own-price demand derivatives. As demand is more sensitive to prices, the inefficiency of getting prices wrong rises.

These equations express deadweight loss in dollars, but we can also express the welfare gain of using the optimal second-best linear policy (the OLS fitted line for the special case where damages are uncorrelated with own-price derivatives), in lieu of an efficient policy. To do so, we need some benchmark policy. One useful benchmark is a policy that puts a single tax rate on all products. This policy does not change the relative price of any of the products, but it allows the policymaker to impose the correct average tax on the extensive margin to correct the size of the market. In terms of equation (10), this baseline policy allows the planner to choose only a constant term, but forces the slope coefficient to be zero.

Compared to this baseline, the efficiency gain of using the linear policy over the baseline, divided by the efficiency gain of the first best policy over the same baseline is equal to the R^2 from the regression:

$$\frac{\text{DWL(Linear)} - \text{DWL(Constant)}}{\text{DWL(First best)} - \text{DWL(Constant)}} = \frac{ESS}{TSS} = R^2 \tag{11}$$

This formulation is particularly useful because it relates directly to data. In Section 4 we demonstrate this for the case of CAFE. There, we estimate the total lifetime damages of different vehicle models using data on lifetime mileage shortly before vehicles are scrapped. We then calculate the degree to which a CAFE-style policy is inefficient by estimating the best linear fit and calculating the R^2. We then use estimates of the social cost of carbon and of the own-price demand elasticity of cars, which we take from the literature, to convert the R^2, which characterizes welfare in proportional terms, into dollar amounts.

Actual policies, of course, need not be second-best. The formula in equation (7) is a valid equation for any alternative policy, provided that demand derivatives are constant over the relevant range of tax rates. (When demand derivatives are not constant, the formula is an approximation, which will be a good approximation for "small" changes in the way that traditional Harberger triangles are.) One useful decomposition is to separate the mean bias in tax rates from their variance, which can be seen by rewriting deadweight loss for the special case with zero net cross effects and no correlation between tax rate errors and demand derivatives, as follows:

$$\min_{\alpha,\beta} = -\sum_{j=1}^{J} \frac{\partial x_j}{\partial t_j} \times \left(\underbrace{J \times \bar{e_j}^2}_{\text{Bias}} + \underbrace{\sum_{j=1}^{J} (e_j - \bar{e_j})^2}_{\text{Variance}} \right). \tag{12}$$

This illustrates that there is a bias in the tax rates, and there is a variance in the tax rate errors, and their effects on welfare can be separated. The bias can be eliminated by a linear policy, but the variance cannot. A CAFE-style policy will be biased if there is a correlation between lifetime utilization and energy efficiency ratings that is not accounted for in policy design.

2.3 When are cross-price effects small?

The above discussion focused on the case when the cross-price effects in equation (7) were zero. The intuition for this can be seen by focusing on cases where the tax rates are unbiased, that is, where $\bar{e_j} = 0$. In that case, the cross-price term is the sum of the product of three things, cross-price demand derivatives (which will generally be positive) and two different tax rate errors, e_j and e_k. In the unbiased case, the errors are zero on average, so, on average, the cross-price terms are multiplying e_j and e_k, and these products will be sometimes positive, sometimes negative. As a result, it is possible for the cross-price effects to cancel out. In contrast, so long as all own-price derivatives are negative (i.e., there are no Giffen goods), the own-price terms in equation (7) must each be positive, so the summation will be positive.

Again, the OLS case helps with intuition. The sum of the product of all pairwise combinations of residuals from OLS is zero. Thus, when the alternative tax schedule τ is the OLS best fit, the cross-price term could be zero. If the tax rates are biased, then the cross-price effects will not be zero on average, but will have a bias term.

The cross-price effects also involve the cross-price derivatives, so the OLS residual intuition is not sufficient to tell us when this summation will net to zero. To see this further, we can decompose the cross-price effects further to characterize them as a function of covariance and bias.

$$-2 \times DWL(\tau) = \underbrace{\sum_{j=1}^{J} e_j^2 \frac{\partial x_j}{\partial t_j}}_{\text{“own effects”}} + \sum_{j=1}^{J} e_j \times \left\{ \underbrace{\sum_{k \neq j} \left((e_k - \bar{e_k}) \left(\frac{\partial x_j}{\partial t_k} - \frac{\partial \bar{x}_j}{\partial t_k} \right) \right)}_{\text{“covariance term”}} + \underbrace{J \times \bar{e_k} \frac{\partial \bar{x}_j}{\partial t_k}}_{\text{“bias term”}} \right\} \quad (13)$$

For the unbiased policy, the bias term in the bracketed expression will be zero, leaving only the covariance term. The covariance term is the covariance between the error in the tax rate and the degree of substitutability across alternative products k for a given model j. This will be zero, on average, when the errors in the tax rates are uncorrelated with the degree of substitutability across products.

We do not expect that the cross-price terms will be exactly zero in real settings for second-best policies because products with similar differences in utilization from the average are also likely to be closer substitutes than average. For example, if luxury vehicles tend to last longer than average, and also tend to be close substitutes for each other, we would expect the cross-price effects overall to be nonzero.

Even when the cross-price terms are not zero it is still intuitive to expect that they may be small, as compared to the own-price terms, because they have the tendency of canceling each other out (the sum of the product of all pairwise combinations of tax residuals $e_j e_k$ is zero). But, this is ultimately an empirical question. We thus emphasize the theoretical intuition of the special case with zero cross-price terms here and then, in our empirical investigation below, determine the degree to which cross-price effects are likely to be important by modeling a range of possible correlations between tax rate errors and demand derivatives.

3 Data

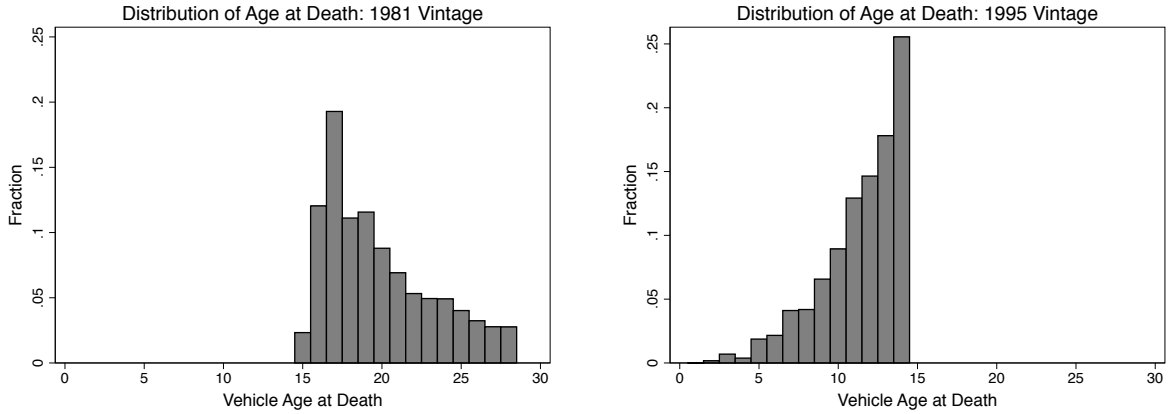
We use California smog check data matched with comprehensive registration micro data. Our analysis is primarily based upon the universe of emissions inspections from 1996 to 2010 from California's vehicle emissions testing program—the Smog Check Program, which is administered by the California Bureau of Automotive Repair (BAR). An automobile appears in the data for a number of reasons. First, vehicles more than four years old must pass a smog check within 90 days of any change in ownership. Second, in large parts of the state an emissions inspection is required every other year as a pre-requisite for renewing the registration on a vehicle that is six years or older. Third, a test is required if a vehicle moves to California from out-of-state. Vehicles that fail an inspection must be repaired and receive another inspection before they can be registered and driven in the state. There is also a group of exempt vehicles. These are: vehicles of 1975 model-year or older, hybrid and electric vehicles, motorcycles, diesel-powered vehicles, and large natural-gas powered trucks.

These data report the location of the test, the unique vehicle identification number (VIN), odometer reading, the reason for the test, and test results. We decode the VIN to obtain the vehicles' make, model, engine, and transmission. Using this information, we match the vehicles to EPA data on fuel economy. Because the VIN decoding is only feasible for vehicles made after 1981, our data are restricted to these models. This yields roughly 120 million observations.

The primary use of the smog check data is to calculate the vehicle's odometer reading at the time the vehicle was scrapped. However, vehicles may leave the smog check data because they leave California. To accurately measure when a vehicle is scrapped, we also use data obtained from CARFAX Inc. These data contain the date and location of the last record of the vehicle, regardless of state, reported to CARFAX for 32 million vehicles in the smog check data. Because the CARFAX data include import/export records, we are able to correctly classify the outcomes of vehicles which are exported to Mexico as censored, rather than scrapped, thus avoiding the issues identified in Davis and Kahn (2010). We define a vehicle as being scrapped if the vehicle is not registered anywhere in the US for two years. The CARFAX data do not report odometer readings; therefore, we restrict analysis to vehicles that were in the smog check program two years prior to being scrapped.

Together, the smog check and CARFAX data give us a measure of the total lifetime miles driven

Figure 2: The Distribution of Vehicle Age at Death for Different Vintages



by a particular unit, but we do not observe all units.[11] We observe relatively few vehicles that are under six years old because most of them are not required to do an emission test. And, we do not observe the final odometer reading for vehicles that have not yet been retired. This creates a censoring problem, which will vary across model years. To illustrate this, Figure 2 shows the age at retirement of vehicles that appear in our sample for model year 1981 and 1995 vehicles, separately. Because our data span smog check readings from 1996 to 2010, we observe 1981 vehicles that were at least 15 years old at retirement, whereas we observe retirements up to age 15 for 1995 models.

In this draft, we explore the importance of censoring for our empirical conclusions by comparing how our results vary when we use all of our data versus when we restrict attention to model years 1988 to 1992, for which censoring will be less problematic. We have plans to use additional data and extrapolation techniques to paint a more complete picture of lifetime VMT for each vehicle in our sample. Specifically, we also have access to comprehensive national registration data from R.L. Polk. These data measure the stock of registered vehicles of each VIN prefix in each year. Thus, we can use these data to measure the number of vehicles that are censored and to determine the scrappage rate of models for years (and ages) outside of our smog sheck sample. We will use this additional information, along with the pattern of mileage of related models in the data to predict lifetime VMT for all censored vehicles. Eventually, our empirical analysis will use these extrapolated data. For now, we emphasize the robustness of our results across various samples and vintage related controls to conclude that the censoring may ultimately have a limited impact on our conclusions.

[11] Note also that the R^2 statistic, which we emphasize, may suffer from bias due to sampling error. We have a large sample, but with highly long-right tailed data, sampling variation may be important. And, if there is mismeasurement in the individual product estimates of lifetime miles driven, this can cause a bias in the R^2. We are still exploring this issue.

4 Numerical Results for the Representative Consumer Model

In this section, we use our theoretical results and our automobile data to quantify the deadweight loss that results from using second-best policies. We focus first on the R^2 of regressions relating fuel-economy rating to lifetime gasoline consumption. We report the R^2 measure from a variety of specifications which take alternative approaches to dealing with the limitations of our data, namely censoring problems and sampling error. As discussed in Section 2, under certain simplifying assumptions, the R^2 alone is a sufficient statistic. Specifically, it is equal to the fraction of the welfare gain achieved by the first-best (over a baseline policy that places a common tax rate on all automobiles) that can be achieved by the optimal second-best linear policy; this fraction is an upper bound on the welfare gain from CAFE.[12] Other factors will enter the calculation that determines economic efficiency when our simplifying assumptions do not hold, or when we wish to evaluate a policy other than the second-best linear one. In those cases, the R^2 will remain a rough indicator of economic efficiency. In ongoing work, we are using simulations to determine how important these other factors may be. In Section 4.2, we use estimates of the social cost of carbon and the derivative of vehicle demand with respect to price to convert the R^2 into deadweight loss measured in dollars.

Our empirical analysis makes several important assumptions. First, we focus here on greenhouse gas emissions, in part because our data allow us to look at this externality more easily than others. Our data indicate the total miles that a vehicle has been driven, but they do not tell us *where* those miles were driven. This is not a problem for global pollutants like carbon dioxide, but it means that we can paint only an incomplete picture of the heterogeneity in damages from local air pollution. Our focus on greenhouse gases likely understates the overall importance of heterogeneity by omitting local air pollution, but it is conceivable that the two sources of heterogeneity could partially offset. We intend to incorporate local pollutants in the future.[13]

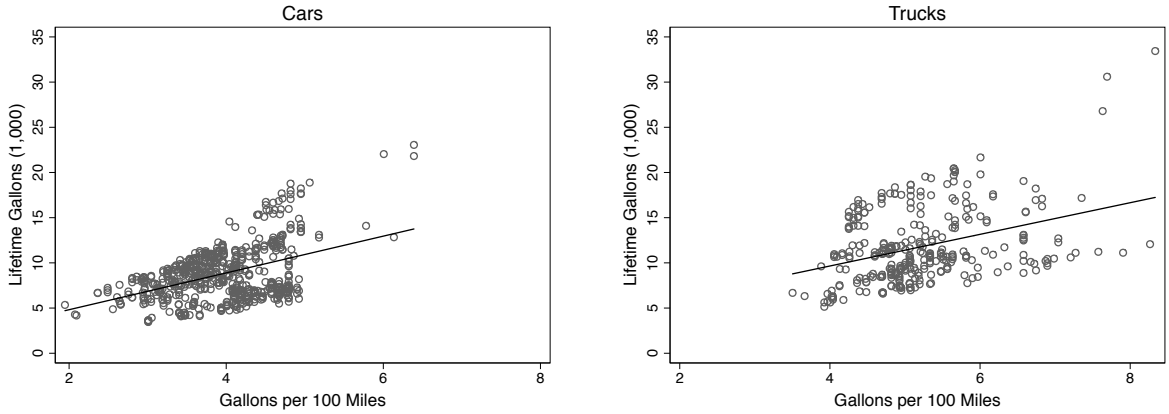
Second, we assume that in-use fuel economy matches the EPA's estimate of the combined fuel economy rating, for each type of vehicle on average across drivers and miles driven. If some vehicles are driven disproportionately in city traffic, while others are driven disproportionately on highways, or if the owners of some vehicles drive particularly aggressively, then the differences between laboratory test ratings and average in-use performance will differ across models. Either factor could be substantial, and by omitting such factors we suspect that we are understating heterogeneity. But, it is possible that heterogeneity in the deviation between on-road fuel economy and official test ratings is negatively correlated with durability, thereby offsetting some of our measured heterogeneity. We will explore this issue in the future.[14]

[12] Because a fuel-economy standard will have the wrong intercept, Holland, Hughes, and Knittel (2009) shows that there is no guarantee that welfare will increase, relative to the case of no regulation.

[13] Registration data and survey data could allow us to identify the population density and local air quality of the location where a each type of vehicle is registered. We are currently exploring the possibility of estimating heterogeneity in local air pollution by vehicle type using these data sources. See Knittel and Sandler (2013) for a related analysis of how heterogeneity in local air pollution across vehicles influences the efficiency of a fuel tax as a policy instrument.

[14] We can use a combination of registration data and survey data from the National Household Transportation

Figure 3: The Relationship Between Lifetime Gasoline Consumption and Fuel-Efficiency



Note: The unit of observation is a type of vehicle (a 10-digit VIN prefix), and gallons consumed is the average across observations for that type. Sample restricted to types for which we observe at least 500 vehicle retirements from model years 1988 to 1992. Individual observations are Winsorized to 500,000 miles. Solid lines are OLS prediction lines.

Third, at present, we ignore the *timing* of emissions. That is, we sum total miles driven and do not discount them into the present-value at the time when a car is new. We do so both for simplicity (our Smog Check data will allow us to estimate the average time path of VMT, so we can relax this assumption) and because initial analysis suggests that the time path of the social cost of carbon roughly offsets discounting. That is, as the social cost of carbon is set to rise at roughly the rate of interest, accounting for the time path of the social cost of carbon and the time path of vehicle emissions will yield a result that is quite similar to a baseline that ignores discounting and uses the current social cost of carbon and full lifetime mileage without discounting.

4.1 The Degree of Heterogeneity in Durability Across Models

To begin, we show scatterplots of the relationship between lifetime gasoline consumption and official fuel consumption ratings. Figure 3 shows the relationship between a model's total lifetime externality (gallons of gasoline) and its fuel economy, for both cars and trucks. A point in the figure corresponds to the average gasoline consumption at the 10-digit VIN prefix level. Therefore, the figure ignores within-VIN variation in gasoline consumption. The sample in the figure is restricted to model years 1988 to 1992, the years for which censoring is least problematic (we observe retirements between 1996 and 2009), and to vehicle types for which we have at least 500 observed

Survey to assess the degree to which vehicles are registered and driven in rural versus highway settings. This analysis will have limitations, but it will likely allow us to make an order of magnitude determination about this source of heterogeneity. Regaring heterogeneity due to driving style, see Langer and McRae (2014) for evidence that the variation across consumers in on-road fuel economy is vast. Note that what matters for our analysis is not the variance across individuals, but the degree to which the average consumer of a particular model differs from the average consumer of another model. Their analysis is based on data from a single type of vehicle, so it offers little guidance for our purposes.

retirements. This mitigates sampling error and reduces the number of observations for visual clarity. In addition, we Winsorize observations above 500,000 miles (we recode them as having exactly 500,000 miles) to limit the influence of outliers and provide a conservative view of the underlying heterogeneity. The figures also show an OLS fitted line for reference.

The figure illustrates that there is, as expected, a strong, positive correlation between fuel consumption (the inverse of fuel economy) and lifetime fuel consumption. But, there is also a great deal of dispersion. Vehicles have significantly different average lifetime mileage, and this translates into variation in lifetime fuel consumption, conditional on the official fuel consumption rating. The R^2 for cars and trucks is only 0.16 and 0.10, respectively, in this sample. (The R^2 from a combined sample regression is 0.23.) This implies that the second-best CAFE-style policy captures only 16% and 10% of the welfare gains for cars and trucks that would be achievable with an efficient set of product-based taxes that varies not only with fuel economy, but also with vehicle durability.

The figures are drawn with a particular subset of the data. To illustrate how the implied efficiency of second-best policies varies with different sample restrictions, Table 1 reports the R^2 from a set of regressions that take the form:

$$\text{Average Lifetime Gasoline Consumption}_j = \alpha + \beta \text{Gallons per Mile}_j + \varepsilon_j, \tag{14}$$

where j indexes a vehicle type (10-digit VIN prefix). We report a range of estimates in order to assess the importance of censoring and sampling error. As discussed above, we are in the process of using registration data and extrapolation to generate a more complete set of estimates regarding vehicle lifetime, and we are still analyzing the nature of bias that could arise from sampling error. For now, we consider censoring by comparing results that include all model years to those that limit the data to model years 1988 to 1992. To assess the importance of sampling error, we compare results from OLS to weighted least squares (WLS), which puts more weight on observations according to the number of underlying observations, and we check how our results change when we limit the sample to vehicles for which we observe relatively many retirements. In all cases, we Winsorize observations with reported mileage above one million (this affects 1,525 observations out of roughly 4 million, or less than 0.05%).

Table 1 shows that our estimate of the R^2 remains small in all specifications, ranging from a low of 0.17 to a high of 0.28. Importantly, our estimates change very little when we restrict the sample to include only 1988 to 1992. This leads us to suspect that our welfare conclusions will be broadly robust to additional measures that account for censoring in the data. The R^2 also changes only modestly when moving between OLS and WLS, and when restricting the sample to VIN-prefixes for which we have at least 200 observed retirements. This suggests that our qualitative findings are not overly sensitive to sampling considerations. For reference, we also report the R^2 from the OLS regression on our underlying micro data, rather than on the data collapsed to VIN-prefix averages, which is 0.08. This regression includes heterogeneity in mileage across different individual units of the same car type. As such, it includes differences in how individual drivers depreciate their vehicles, including accidents that lead newer vehicles to be scrapped.

Table 1: Regression R^2

	OLS	WLS
VIN-pre averages, all MYs		
All models	.20	.20
Models with $N \geq 200$	.21	.17
VIN-pre averages, MYs 1988-1992		
All models	.21	.26
Models with $N \geq 200$	.28	.22
Microdata, all MYs		
All models	.08	

Note: Table shows R^2 from regressions of average lifetime gallons consumed on fuel consumption rating. The unit of observation is a VIN-prefix, except for the last row, which includes individual vehicles.

Our estimates of the R^2 could also be sensitive to outliers, even when restricting to vehicles with relatively large sample sizes. To check whether or not this is a concern, we Winsorize the underlying micro data at different levels and calculate new estimates, which are reported in Table 2. For example, in the second row, any observation that has a reported odometer rating above 600,000 miles is recoded as having exactly 600,000 miles. Its gasoline consumption is recalculated assuming the new odometer reading, and the observation is then averaged along with all other observations from the same VIN-prefix. The table reports OLS and WLS results, restricting the sample to model years 1988 to 1992 and to VIN-prefixes with at least 200 observed retirements.

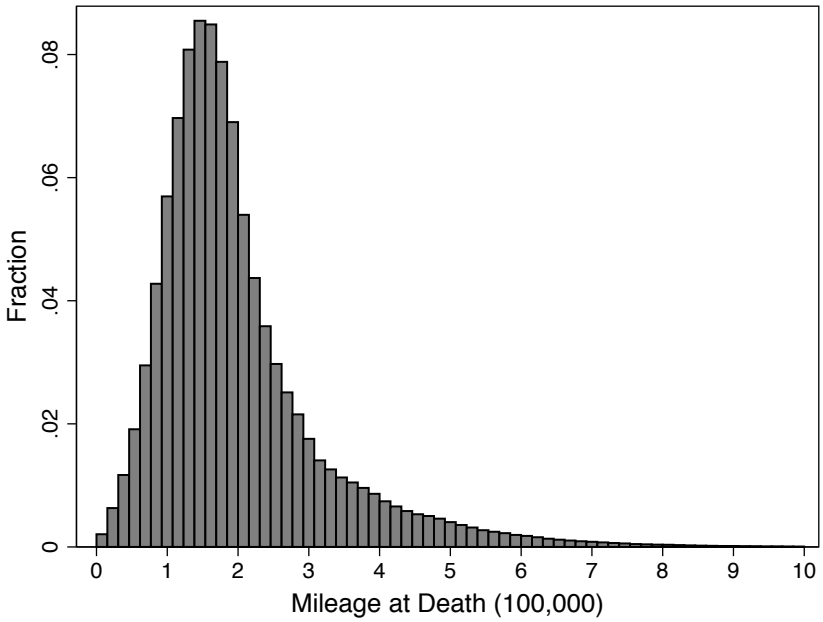
Table 2: Regression R^2

VIN-pre averages, MYs 1988-1992, Models with $N \geq 200$	**OLS**	**WLS**
All odometer readings	.28	.22
Odometer $\leq 600,000$	.30	.23
Odometer $\leq 500,000$	.32	.25
Odometer $\leq 400,000$	.36	.30

Note: Table shows R^2 from regressions of average lifetime gallons consumed on fuel consumption rating. The unit of observation is a VIN-prefix.

Table 2 shows that the R^2 increases, but only modestly, when we Winsorize the highest VMT vehicles. The OLS R^2 rises from a baseline of 0.28 to a maximum of 0.36 when we limit the influence of data over 400,000 miles. The WLS R^2 rises from a baseline of 0.22 to a maximum of 0.30 for the same restriction. We again interpret this as a sign that our qualitative conclusions are robust. Furthermore, note that Winsorizing the data at a level as low as 400,000 miles is truly restrictive. We plot the full histogram of lifetime odometer readings across all of our microdata in Figure 4. The data have a mode around 150,000 miles, but there is a long right tail. Just under 7% of vehicles in our data are scrapped with over 400,000 miles.

Figure 4: Distribution of Lifetime VMT Across All Observations

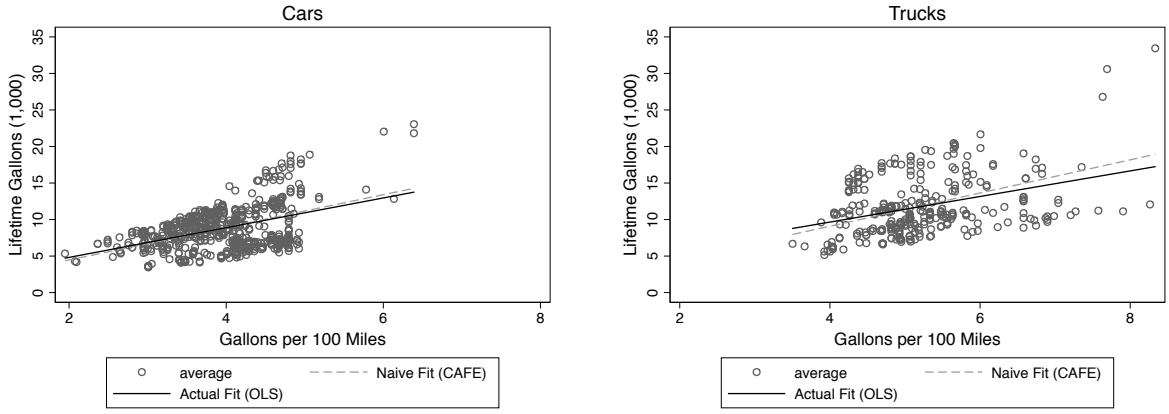


Overall, the relatively low R^2 statistics suggest that there is substantial variation in lifetime gasoline consumption, conditional on fuel economy ratings. We have also explored a number of alternative estimates that treat cars and trucks separately and that estimate the R^2 for each model year separately. In all cases, the qualitative conclusion remains that there is substantial variation in lifetime consumption that is not explained by fuel economy, which implies that policies based only of fuel economy ratings, not on average product durability, will raise welfare by significantly less than would an efficient policy (including a carbon tax or a gasoline tax). As noted throughout, these results are preliminary, as we are both working to more fully account for our data limitations using registration data and extrapolation and to use simulations to determine how much the results might change when we relax the assumptions that make the R^2 the sole determinant of policy efficiency.

Note that our approach can easily be expanded to consider alternatives. First, the R^2 remains the relevant summary statistic for a nonlinear fuel economy policy. For example, if CAFE were more flexible so that it could put a shadow price on each model that was a quadratic function of fuel consumption ratings, we could estimate the R^2 from a quadratic regression. Second, if the linear policy used is not the second-best (that is, if the predicted line is biased), then this bias causes an additional source of welfare loss.

We illustrate this second possibility in Figure 5, which replicates Figure 3, but adds a line that represents the relationship between fuel consumption ratings and lifetime fuel consumption, if all cars (or trucks) were driven the same number of miles, which we set equal to the observed mean in our data. This line represents the best fit line that a policymaker would choose if they knew only the average mileage (separately for cars and trucks) across all vehicles, but did not know how average mileage varied with fuel consumption. This is our depiction of a "naïve" linear tax, which

Figure 5: The Relationship Between Lifetime Gasoline Consumption and Fuel-Efficiency



Note: The unit of observation is a type of vehicle (a 10-digit VIN prefix), and gallons consumed is the average across observations for that type. Sample restricted to types for which we observe at least 500 vehicle retirements from model years 1988 to 1992. Individual observations are Winsorized to 500,000 miles. Solid lines are OLS prediction lines. Dashed lines are linear fits under assumption that all vehicles are driven the mean number of miles.

gets the average shadow price right, but ignores durability completely.

In our data, there is a negative correlation between average lifetime VMT and fuel consumption ratings, within vehicle class. As a result, the naïve version of the linear tax would overtax the lifetime mileage of vehicles with high fuel-economy ratings and undertax vehicles with low fuel-economy ratings. This mispricing leads to another source of inefficiency when heterogeneity in utilization is ignored.

4.2 Estimates of Deadweight Loss

We can translate the relative gains from the first- and second-best product-based taxes, expressed above as R^2, into deadweight loss by assigning a dollar value to the externality and considering the pattern of substitution across vehicles. We narrow our focus to the 1990 model year as an example but plan a fuller simulation that explores the complete range of model years in our data. This will allow for changes over time in the amount of heterogeneity as well as analysis of mis-optimization that occurs from year to year if the policy is not continuously updated.

Evaluating the formula in equation (6) requires estimates for the externality generated by each car as well as its own- and cross-price elasticities with respect to the other vehicles in the market. We assign a value of \$39 for the social cost of carbon (Interagency Working Group on Social Cost of Carbon (2013)), leading to an external cost of 34.6 cents per gallon.[15] Using our data on lifetime fuel use for the 1990 vintage this implies an average of \$3,250 in external costs for each vehicle sold. We further impose an own-price elasticity of -5 (roughly comparable to the estimates in

[15]If the cost associated with carbon emissions has been rising approximately at the discount rate, we loosely interpret this value as being in 2011 dollars (looking retrospectively at the 1990 vintage). We plan to introduce a fuller accounting when comparing across multiple vintages.

Berry, Levinsohn, and Pakes (1995)) and cross-price elasticities distributed evenly over the full set of models. We relax both of these assumptions below, considering higher and lower own-price elasticities and cross-price elasticities that are allowed to vary with similarity in attributes.

As above we present results relative to a baseline that controls for substitution to an outside good, since this is not present under a revenue-neutral CAFE standard, and so isolate the welfare effects coming from switching among vehicles. Under our assumptions on elasticities the first-best welfare gains with a separate Pigouvian tax on each of the 356 VIN prefixes amount to $232 per car sold, or about $3.3 billion for the 1990 model year. The best linear tax on fuel use per mile, equivalent here to the optimal average fuel economy standard, generates about $0.85 billion in surplus and so leaves $2.4 billion in deadweight loss. This corresponds directly to the intuition on R^2 above: for the 1990 model year the weighted R^2 is 0.26, implying 26% of possible gains (0.85 over 3.3) are recovered with a linear policy.

Table 3 displays results allowing the demand elasticities for vehicles to vary. The case just discussed assumes an own-price elasticity of -5 with constant cross-price elasticities and appears in the first row of the table. The following two rows show the effects of altering the own-price elasticity. More elastic demand allows a larger change in the composition of the fleet and so greater welfare gains in the first best. The ratio remains fixed at the value of R^2. The next two rows allow cross-price effects to be related to the relative fuel economy of vehicles, with the two cases differing in how rapidly substitution decays as cars become more distant. A car with the same fuel economy is twice as good a substitute as a car one standard deviation away (about 1 gallon per hundred miles) in the first case, and ten times as good a substitute in the second case. When introducing this substitution pattern the gains possible in the first best become smaller. Intuitively, it is now harder to move people across vehicles with different externalities, at least to the extent the externality is correlated with fuel economy. The change in cross-price effects worsens the performance of the second-best policy even more dramatically than the first-best: this is because the only margin on which the second-best operates (fuel economy) is the margin that the cross-price effects are limiting. We plan to explore a wide range of possible cross-price effects, influencing different margins over which the first- and second-best policies operate.

The estimates above are subject to several important caveats. Perhaps the most important consideration is the role of technology: we have so far considered a static portfolio of durables offered for sale, when in the long run the products can be re-engineered according to the incentives provided by the tax schedule. In the case of cars this amounts to altering attributes, for example reductions in weight and horsepower, and introducing efficiency-enhancing technologies like hybrid drive and more sophisticated transmissions. The second-best intuition developed here will at least in part also hold for the distribution of technology: the most advanced technologies and lightest materials should be placed in the cars that have the highest expected lifetime mileage. A linear tax (or standard) based on fuel economy will encourage these technologies to be distributed much more equally across the fleet, missing welfare gains possible if the improvements could instead be targeted. We plan to explore a simple version of this point in an extended simulation. Interactions

Table 3: Welfare Effects for Model Year 1990

	Second best	**First best**	**Ratio**
Central case	0.85	3.28	0.26
Own-price elasticity			
-4	0.68	2.65	0.26
-6	1.01	3.90	0.26
Cross-price elasticities			
2	0.74	3.19	0.23
10	0.47	2.96	0.16

Note: Welfare gains are expressed in billions of 2011 dollars relative to a constant tax at the average externality. Cross-price elasticities are expressed as the factor by which substitutability decreases per standard deviation difference in fuel economy.

between the second-best effects here and a set of other distortions produced by standards (for example from attribute-basing or changes to durability induced via the used market) have the potential to further influence overall welfare.

5 Model with heterogeneous consumers

In Section 2 we present a model in which consumers are all alike, but there is heterogeneity in lifetime utilization of durable goods due to differences in product durability. We now generalize to a setting in which consumers have different intensities of use of the same good and, therefore, each consumer-good pairing can have a different lifetime externality.

5.1 The Standard Diamond Model

Diamond (1973) derives the second-best tax t^* for the case where there is a single externality-generating good x but there are $i = 1, \ldots, I$ consumers who have different own-price derivatives of demand and generate different externalities.

$$t^* = \frac{-\sum_{i=1}^{I} \sum_{k \neq i} \frac{\partial U^k}{\partial x_i} \frac{\partial x_i}{\partial t}(p+t)}{\sum_{i=1}^{I} \frac{\partial x_i}{\partial t}(p+t)} \tag{15}$$

The marginal damages (externality) is expressed as the impact of person i's consumption on others $-\sum_{k \neq i} \frac{\partial U^k}{\partial x_i} \equiv e_i$. Equation (15) has a simple interpretation: the second-best tax is the average of the externalities, weighted by the individuals' own-price derivatives. If more responsive consumers have higher damages e_i, the second-best tax rate exceeds the unweighted average externality $\bar{e}$. If demand derivatives are uniform across consumers, $t^* = \bar{e}$. Note that if consumers are

identical and have the same quasi-linear utility function ($U^i = U\ \forall i$) then the optimal demands $x_i^* = x^*$ and the externality $e_i = e$ for all consumers. Also, $\frac{\partial U^i}{\partial x_k} = \frac{\partial U^k}{\partial x_i}\ \forall i \forall k$ and $\frac{\partial x_i}{\partial t}(p+t) = x'\ \forall i$. In that case, the problem collapses to a first-best solution in which the optimal tax is equal to the total externality imposed on all other consumers: $t^* = e$.

5.2 A Diamond Model with Multiple Goods

When consumers have a different intensity of use of the same durable good, they will produce different externalities, e_{ji}, from an ex-ante identical product j. For example, different car owners may drive different amounts of miles over the lifetime of the same model.[16] Another example is that some drivers are more aggressive than others and as such realize a lower fuel economy and higher emissions (Langer and McRae (2014)). Another source of heterogeneity for local pollutants is geography; the same tailpipe emission will have a different external damages depending on population density and the starting level of pollution.

We now generalize the Diamond model to include multiple goods $j = 1, ..., J$. We continue to abstract from the rebound effect, as our aim is to compare different types of product-based policies. The setup is quite general: every consumer can have unique expected damages when paired with each product.[17] This encompasses heterogeneity due to breakdowns or accidents.[18] Each consumer i holds a continuous bundle of goods, where demand for good j by consumer i is denoted x_{ji}. In the context of vehicles, we can interpret consumer i as a "driver type" that holds a portfolio of vehicles. There are J second-best tax rates t_j^*, one for each good, whereas the first-best solution (which we assume to be infeasible) would be $J \times I$ tax rates $t_{ij}^* = e_{ji}$. Note that the latter set of taxes, which are unique to a consumer-product pairing, are equivalent to a gasoline tax as long as we abstract from the mileage rebound effect.

If a first-best (gasoline) tax were feasible, elasticities are irrelevant. The only information requirement is the marginal externality per unit of energy consumption. The first-best tax guarantees that all consumers pay exactly their total externality and choose their vehicles optimally (Pigou, 1932). In the second-best setting described here, elasticities do matter. The second-best taxes will be a function of the elasticity of substitution between vehicles, and the externality imposed by consumer i's individual use of car j: different people drive the same models differently (i.e., they have different ex-ante VMT distributions for the same vehicle at the moment of purchase).

We now derive expressions for the second-best taxes with heterogeneous consumers and multiple goods. In the Diamond framework, utility is given by:[19]

[16] To the extent that vehicle depreciation is a function of mileage, rather than time, different annual VMT is not enough to imply heterogeneity. Consumer heterogeneity comes from different accident risk or time-based depreciation.

[17] Restricted versions of this model could assume that each consumer carries a fixed amount of "utilization" (e.g., vehicle miles traveled), regardless of the product type (car) (s)he owns. In this case, the remaining individual-level heterogeneity comes from driving style and accidents. In general, the dimensions of the problem might be reducible depending on the shapes of the person-car externalities e_{ji}. If each consumer has a fixed distribution of lifetime mileage VMT_i and each vehicle has a fixed fuel economy MPG_j, then $e_{ji} = VMT_i/MPG_j$ (ignoring discounting). We can capture this with fewer than $J \times I$ taxes.

[18] The marginal damage can be interpreted as relative to the risk-adjusted mean damages.

[19] Note that in this specification, consumers derive utility from owning cars. *Services* derived from owning cars

$$U^i(x_{11}, ..., x_{ji}, ..., x_{JI}) + \mu_i, \tag{16}$$

where μ_i is income spent on other goods by consumer i. We further follow Diamond in assuming that externalities are diseconomies:

$$\frac{\partial U^i}{\partial x_{jk}} \leq 0, \tag{17}$$

for $k \neq i$. The marginal utility of consumption is independent of the demands of other consumers:

$$\frac{\partial^2 U^i}{\partial x_{ji} \partial x_{jk}} = 0, \tag{18}$$

for $k \neq i$ and for all goods j. Each consumer i maximizes utility subject to a budget constraint:

$$\underset{x_{1i},...,x_{Ji}}{max} U^i(x_{11}, ..., x_{ji}, ..., x_{JI}) + \mu_i, \tag{19}$$

subject to:

$$\sum_{j=1}^{J} (p_j + t_j) x_{ji} + \mu_i = m_i, \tag{20}$$

where m_i is total income. The first order conditions assuming an interior solution and using equation (18) are given by:

$$\frac{\partial U^i}{\partial x_{ji}}(x_{1i}, ..., x_{Ji}) = p_j + t_j = \tilde{p}_j, \tag{21}$$

for all goods j. The optimal choices for consumer i depend not only on the tax-inclusive price of the durable good $\tilde{p}_j$, but also on the consumption level – and therefore the tax-inclusive price – of the other durable goods:

$$x^*_{ji} = x_{ji}(\tilde{p}_1, ..., \tilde{p}_J). \tag{22}$$

Analogous to Diamond (1973), we derive the optimal tax rate by writing down an expression for total utility:

$$W(t_1, ..., t_J) = \sum_{i=1}^{I} U^i(x_{11}(\tilde{p}_1, ..., \tilde{p}_J), ..., x_{ji}(\tilde{p}_1, ..., \tilde{p}_J), ..., x_{JI}(\tilde{p}_1, ..., \tilde{p}_J)) + \sum_{i=1}^{I} \mu_i. \tag{23}$$

(VMT) are not explicitly modeled. However, this framework allows for different drivers of the same car to impose a different externality on other consumers. Implicitly, we can think of this as heterogeneity in VMT. In fact, in this general framework, consumer i driving car j may impose a different externality on consumer k vs. consumer l. This is relevant for local pollutants, but not for a uniformly mixed pollutant such as CO_2.

Using equation (20) and assuming that tax revenues are returned lump-sum, we can rewrite this as:

$$W(t_1,...,t_J) = \sum_{i=1}^{I} U^i(x_{11}(\tilde{p}_1,...,\tilde{p}_J),...,x_{ji}(\tilde{p}_1,...,\tilde{p}_J),...,x_{JI}(\tilde{p}_1,...,\tilde{p}_J)), \\ + \sum_{i=1}^{I}\left(m_i - \sum_{j=1}^{J} p_j x_{ji}(\tilde{p}_1,...,\tilde{p}_J)\right). \tag{24}$$

To find the second-best optimal vector of tax rates $(t_1^*,...,t_J^*)$, we need to calculate first order conditions by differentiating equation (24) w.r.t. to each of the t_j's:

$$\frac{dW}{dt_j} = \sum_{l=1}^{J}\sum_{i=1}^{I}\sum_{k=1}^{I} \frac{\partial U^i}{\partial x_{lk}} \frac{\partial x_{lk}}{\partial t_j}(\tilde{p}_1,...,\tilde{p}_J) - \sum_{l=1}^{J}\sum_{i=1}^{I} p_l \frac{\partial x_{li}}{\partial t_j}(\tilde{p}_1,...,\tilde{p}_J) = 0. \tag{25}$$

Now use the FOCs in equation (21) to rewrite this as:

$$\frac{dW}{dt_j} = \sum_{l=1}^{J}\sum_{k=1}^{I}\sum_{i\neq k} \frac{\partial U^i}{\partial x_{lk}} \frac{\partial x_{lk}}{\partial t_j}(\tilde{p}_1,...,\tilde{p}_J) + \sum_{l=1}^{J}\sum_{i=1}^{I} t_l \frac{\partial x_{li}}{\partial t_j}(\tilde{p}_1,...,\tilde{p}_J) = 0. \tag{26}$$

In sum, the formula equalizes the change in consumer surplus for both cars and the change in total externalities when t_j is raised.

This is a fairly intuitive expression. The first term in equation (26) is the total net change in the externality from an increase in t_j. If cars j and l are substitutes, this is an improvement from car j and a worsening from car l. The double summation over k and i summarizes the total externality imposed on all others for each consumer k. The marginal utility gains from reduced externalities reflect the benefits from raising the tax on car j. The second term is the reduction in consumer surplus for all cars and consumers and reflects the costs of raising t_j. In sum, we set t_j to equalize the change in net consumer surplus for all cars ("costs") and the change in net externalities ("benefits").

For the 2-good case (we are still confirming the proper notation for the J-good case), solving the system of 2 equations yields the following expression:

$$t_1^* = \left[\sum_i \frac{\partial x_{2i}}{\partial t_1}\sum_k\sum_{i\neq k}\left(\frac{\partial U^i}{\partial x_{1k}}\frac{\partial x_{1k}}{\partial t_2} + \frac{\partial U^i}{\partial x_{2k}}\frac{\partial x_{2k}}{\partial t_2}\right) - \sum_i \frac{\partial x_{2i}}{\partial t_2}\sum_k\sum_{i\neq k}\left(\frac{\partial U^i}{\partial x_{1k}}\frac{\partial x_{1k}}{\partial t_1} + \frac{\partial U^i}{\partial x_{2k}}\frac{\partial x_{2k}}{\partial t_1}\right)\right]\frac{1}{\Delta},$$
$$t_2^* = \left[\sum_i \frac{\partial x_{1i}}{\partial t_2}\sum_k\sum_{i\neq k}\left(\frac{\partial U^i}{\partial x_{1k}}\frac{\partial x_{1k}}{\partial t_1} + \frac{\partial U^i}{\partial x_{2k}}\frac{\partial x_{2k}}{\partial t_1}\right) - \sum_i \frac{\partial x_{1i}}{\partial t_1}\sum_k\sum_{i\neq k}\left(\frac{\partial U^i}{\partial x_{1k}}\frac{\partial x_{1k}}{\partial t_2} + \frac{\partial U^i}{\partial x_{2k}}\frac{\partial x_{2k}}{\partial t_2}\right)\right]\frac{1}{\Delta}, \tag{27}$$

where the arguments $(\tilde{p}_1, ..., \tilde{p}_J)$ are omitted for notational simplicity and:

$$\Delta = \sum_i \frac{\partial x_{i1}}{\partial t_1}(\tilde{p}_1, \tilde{p}_2) \sum_i \frac{\partial x_{i2}}{\partial t_2}(\tilde{p}_1, \tilde{p}_2) - \sum_i \frac{\partial x_{i2}}{\partial t_1}(\tilde{p}_1, \tilde{p}_2) \sum_i \frac{\partial x_{i1}}{\partial t_2}(\tilde{p}_1, \tilde{p}_2). \tag{28}$$

In the single good, standard Diamond model in equation (15), the second-best tax is the average of the externalities, weighted by the individuals' own-price derivatives. In the multiple goods case, this simple intuition no longer holds. t_1^* now consists of two terms. The second term still contains the weighted sum of the externalities, but now also the weighted average externalities generated by the change in consumption of the second car when the first car is taxed: the weights for the externalities from the second car are *cross*-price derivatives. Next, the entire second term is multiplied by a scale factor (the sum of the own-price elasticities of car 2 across the consumers). The first term of the expression for t_1^* is similar to the second term, but it reflects changes in externalities when the second car is taxed and its overall scale factor is a sum of cross-price derivatives. Hence, if the own-price elasticities are large relative to the cross-price elasticities, the second term dominates and the expression converges back to the simple Diamond model. Δ is a scale factor.

We conjecture that this can be simplified using matrix notation. Label two objects from equation (27) for ease of notation. First, ΔD_{jl} is the change in demand for good j when the price of good l changes. If $J = 2$, this is a 2×2 matrix. Second, ΔE_j is the total *net* change in externalities from a change in the price of good j. For example, $\sum_i \frac{\partial \alpha_{2i}}{\partial t_1} = \Delta D_{21}$ and $\sum_k \sum_{i \neq k} \left(\frac{\partial U^i}{\partial \alpha_{1k}} \frac{\partial \alpha_{1k}}{\partial t_2} + \frac{\partial U^i}{\partial \alpha_{2k}} \frac{\partial \alpha_{2k}}{\partial t_2} \right) = \Delta E_2$. For $J = 2$, we can write the vector of optimal taxes as $T^* = (\Delta D')^{-1} \Delta E$.

ΔD is the standard market-level Marshallian substitution matrix. To calculate ΔD, we need to know aggregate cross-price elasticities, but not *who* changed demand. ΔE is the *net* change in the externality resulting from a change in the price of a vehicle. Its element j is the sum of each person k's change in externality from a price increase of car j, which we denote as $\Delta e_{jk} \equiv \sum_{i \neq k} \left(\frac{\partial U^i}{\partial \alpha_{1k}} \frac{\partial \alpha_{1k}}{\partial t_j} + \frac{\partial U^i}{\partial \alpha_{2k}} \frac{\partial \alpha_{2k}}{\partial t_j} \right)$. This term can be negative or positive. For example, when $J = 2$ we expect $\frac{\partial \alpha_{11}}{\partial t_1} < 0$ and $\frac{\partial \alpha_{12}}{\partial t_1} > 0$): raising the price of car 1 will lower demand for car 1 but raise demand for car 2. The net effect on the externality depends on relative magnitude of the elasticities and externalities.

To calculate Δe_{kj} we need to know the *correlation* between marginal damages and the own- and cross-price elasticities. Absent a homogeneity assumption, we need to know this for each individual. We therefore conclude that the informational requirements to calculate the second-best taxes for the general case of heterogeneous consumers and multiple goods in equation (27) are large. We interpret this as a "negative result": the second-best tax rates require substantially more information than a first-best gasoline tax. Moreover, it may be hard to estimate these correlations without highly granular individual-level and product-level data. We are developing simulations to assess how much welfare is lost by resorting to product-based standards that ignore heterogeneity in usage intensity, and we are simultaneously investigating whether or not an analytical expression for deadweight loss can be derived for this general case.

The heterogeneous consumers, multiple goods model simplifies when additional restrictions are

imposed. We consider three special cases (see Appendix A for the derivations). First, when cross-price derivatives are all zero, the second-best taxes collapse to the single-good Diamond formula in equation (15) for each durable good j. Second, when there is no correlation between damages and elasticities, the second-best tax rate for each good j is the average externality over all consumers $\bar{e}_j$. Hence, the optimal tax formula collapses back to the durability-only case presented in Section 2. Third, consider a case in which consumers are homogeneous in that different individuals i and k drive car model j in an identical way, though cars have different fuel economy and durability (VMT). In other words, the externality from driving vehicle j is the same across each pair of consumers. The optimal tax for each vehicle reduces to its expected externality: $t_j^* = e_j$. In other words, the optimal tax in the durability only model in Section 2 is now first-best (again, maintaining the assumption of no rebound effect).

6 Conclusion

Product-based policies are a ubiquitous policy tool for regulating energy-related externalities. Heterogeneity in the lifetime utilization of durables causes product-based policies to be inefficient, as compared to emissions taxes (or, for some pollutants, equivalent fuel taxes). We show that, under intuitive conditions in a model with a representative consumer, the R^2 from a regression of lifetime emissions on energy efficiency ratings across products maps directly into welfare loss and acts as a sufficient statistic for the use of restricted policies that place a linear subsidy on products based on their energy efficiency rating, as compared to the efficient Pigouvian benchmark.

We use detailed data on vehicle scrappage, miles traveled and fuel economy to estimate welfare losses of using product-based policies versus gasoline taxes to mitigate greenhouse gas emissions in the automobile market. We estimate that CAFE-style policies recover only about one quarter of the welfare gains that could be achieved by a gasoline tax. We also derive second-best tax rates in a more general model that allows for usage heterogeneity from differences across consumers as well as differences across products. There, we demonstrate that second-best policy design requires a vast amount of information that is not likely to be available to policy makers. We plan to use simulation to assess how much welfare is lost by resorting to optimal product-based standards. An emissions tax would create superior welfare gains while simultaneously requiring much less information.

References

Anderson, Soren T., Ian W.H. Parry, James M. Sallee, and Carolyn Fischer. 2011. "Automobile Fuel Economy Standards: Impacts, Efficiency and Alternatives." *Review of Environmental Economics and Policy* 5 (1):89–108.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 64 (4):841–890.

Borenstein, Severin. Forthcoming. "A Microeconomic Framework for Evaluating Energy Efficiency Rebound and Some Implications." *Energy Journal* .

Davis, Lucas W. and Matthew E. Kahn. 2010. "International Trade in Used Vehicles: The Environmental Consequences of NAFTA." *American Economic Journal: Economic Policy* 2 (4):58–82. URL http://www.aeaweb.org/articles.php?doi=10.1257/pol.2.4.58.

Diamond, Peter A. 1973. "Consumption Externalities and Imperfect Corrective Pricing." *The Bell Journal of Economics* 4 (2):526–538.

Hendren, Nathaniel. 2013. "The Policy Elasticity." Manuscript: Harvard University.

Holland, Stephen P., Jonathan E. Hughes, and Christopher R. Knittel. 2009. "Greenhouse Gas Reductions Under Low Carbon Fuel Standards?" *American Economic Journal: Economic Policy* 1 (1):106–146.

Interagency Working Group on Social Cost of Carbon. 2013. "Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis." Http://www.whitehouse.gov/sites/default/files/omb/assets/inforeg/technical-update-social-cost-of-carbon-for-regulator-impact-analysis.pdf.

Knittel, Christopher R. and Ryan Sandler. 2013. "The Welfare Impact of Indirect Pigouvian Taxation: Evidence from Transportation." Manuscript: MIT.

Langer, Ashley and Shaun McRae. 2014. "Step on It: Evidence on the Variation in On-Road Fuel Economy." Manuscript: University of Arizona.

Pigou, Arthur C. 1932. *The Economics of Welfare*. London: Macmillan and Co., 4th edition ed.

Sallee, James M. 2011. "The Taxation of Fuel Economy." *Tax Policy and The Economy* 25:1–37.

A Appendix: Special Cases of the Multiple Goods Diamond Model

This appendix contains the derivations for the three special cases of the heterogeneous consumer, multiple good model presented in Section 5. These special cases may aid in interpretation of the general result.

A.1 No Cross-Price Effects in Demand

First, we demonstrate the our results collapse to the one-good result from Diamond (1973) when cars are neither substitutes nor complements for each other. Assume that:

$$\frac{\partial x_{ji}}{\partial t_l} = \frac{\partial x_{li}}{\partial t_j} = 0, \tag{29}$$

for all pairs of goods j and l. Demand for one good does not depend on the price of the other good. Then, equation (22) simplifies to:

$$x^*_{ji} = x_{ji}(\tilde{p}_j). \tag{30}$$

In that case, the problem is separable in the different goods. Equation (26) reduces to:

$$\frac{dW}{dt_j} = \sum_{i=1}^{I} \sum_{k \neq i} \frac{\partial U^k}{\partial x_{ji}} \frac{\partial x_{ji}}{\partial t_j}(\tilde{p}_j) + t_j \sum_{i=1}^{I} \frac{\partial x_{ji}}{\partial t_j}(\tilde{p}_j). \tag{31}$$

The second-best taxes t^*_j reduce to J separate Diamond-style formulas of externalities weighted by own-price derivatives:

$$t^*_j = \frac{\sum_{i=1}^{I} \sum_{k \neq i} \frac{\partial U^k}{\partial x_{ji}} \frac{\partial x_{ji}}{\partial t_j}(\tilde{p}_j)}{\sum_i \frac{\partial x_{ji}}{\partial t_j}(\tilde{p}_j)}. \tag{32}$$

A.2 No Correlation Between Elasticities and Externalities

Next we show what happens if if we demand elasticities are uncorrelated with the externalities. Specifically, assume that:

$$\sum_{k=1}^{I} \sum_{i \neq k} \left(\frac{\partial U^i}{\partial \alpha_{lk}} \frac{\partial \alpha_{lk}}{\partial t_j} \right) = \sum_{k=1}^{I} \frac{\partial \alpha_{lk}}{\partial t_j} \sum_{k=1}^{I} \sum_{i \neq k} \frac{\partial U^i}{\partial \alpha_{lk}} / I. \tag{33}$$

If we substitute these terms into the second-best tax rates for the 2-good case in equation (27), the result is:

$$
\begin{aligned}
t_1^* = \sum_i \frac{\partial x_{2i}}{\partial t_1} \left(\sum_k \sum_{i \neq k} \frac{\partial U^i}{\partial x_{1k}} / I \sum_{k=1}^{I} \frac{\partial x_{1k}}{\partial t_2} + \sum_k \sum_{i \neq k} \frac{\partial U^i}{\partial x_{2k}} / I \sum_{k=1}^{I} \frac{\partial x_{2k}}{\partial t_2} \right) \frac{1}{\Delta} \\
- \sum_i \frac{\partial x_{2i}}{\partial t_2} \left(\sum_k \sum_{i \neq k} \frac{\partial U^i}{\partial x_{1k}} / I \sum_{k=1}^{I} \frac{\partial x_{1k}}{\partial t_1} + \sum_k \sum_{i \neq k} \frac{\partial U^i}{\partial x_{2k}} / I \sum_{k=1}^{I} \frac{\partial x_{2k}}{\partial t_1} \right) \frac{1}{\Delta}, \\
= - \sum_k \sum_{i \neq k} \frac{\partial U^i}{\partial \alpha_{k1}} / I = \bar{e_1}.
\end{aligned} \tag{34}
$$

Analogously, $t_2^* = \bar{e_2}$. Hence, with no correlation between elasticities and externality, the second-best tax formula reduces to taxing each good at its average externality, which is the same as the second-best tax rate that we derive in the representative consumer case in section 2.

A.3 Homogeneous Consumers

Another possible step would be to assume that consumer types are homogeneous in the following way: different individuals i and k drive car model j in exactly the same way. In other words, cars may differ on characteristics such as fuel economy and durability (VMT), but there is no within-model heterogeneity of driving habits. As a consequence, the externality from driving vehicle j is the same across consumers. This restriction can be expressed as follows:

$$
\begin{aligned}
\frac{\partial U^i}{\partial x_{jk}} = \frac{\partial U^i}{\partial x_{jl}} = u_j^{'}, \\
\frac{\partial U^i}{\partial x_{jk}} = \frac{\partial U^k}{\partial x_{ji}} = u_j^{'},
\end{aligned} \tag{35}
$$

for $k \neq l \neq i$. The first condition imposes that if consumers k and l drive the same car j, they impose the same externality e_j on consumer i. The second condition says that k's externality on i from driving car j is the same as i's externality on k from driving the same car.

Substituting these constraints in the optimal tax FOCs in equation (26) yields:

$$
\begin{aligned}
\frac{dW}{dt_j} = \sum_{l=1}^{J} \sum_{k=1}^{I} \sum_{i \neq k} u_l^{'} \frac{\partial x_{lk}}{\partial t_j}(\tilde{p}_1, ..., \tilde{p}_J) + \sum_{l=1}^{J} \sum_{i=1}^{I} t_l \frac{\partial x_{li}}{\partial t_j}(\tilde{p}_1, ..., \tilde{p}_J), \\
= (I-1) \sum_{l=1}^{J} \sum_{i=1}^{I} u_l^{'} \frac{\partial x_{li}}{\partial t_j}(\tilde{p}_1, ..., \tilde{p}_J) + \sum_{l=1}^{J} \sum_{i=1}^{I} t_l \frac{\partial x_{li}}{\partial t_j}(\tilde{p}_1, ..., \tilde{p}_J) = 0,
\end{aligned} \tag{36}
$$

which implies:

$$
t_j^* = -(I-1) u_j^{'} = - \sum_{k \neq i} \frac{\partial U^k}{\partial x_{ji}} = e_j. \tag{37}
$$

Under this assumption, product-based taxes can yield the first-best policy, where the optimal taxes are given by $t_j^* = e_j$. Here, the first-best tax equals the optimal tax in the durability only model in Section 2, and the first-best is obtainable because of the assumption that all people generate a common externality when using each type of product.