The Use of Regression Statistics to Analyze Imperfect Pricing Policies

As Published: 10.1086/705553

Publisher: University of Chicago Press

Persistent URL: https://hdl.handle.net/1721.1/136641

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

The Use of Regression Statistics to Analyze Imperfect Pricing Policies

Mark R. Jacobsen

University of California San Diego and National Bureau of Economic Research

Christopher R. Knittel

Massachusetts Institute of Technology and National Bureau of Economic Research

James M. Sallee

University of California Berkeley and National Bureau of Economic Research

Arthur A. van Benthem

University of Pennsylvania and National Bureau of Economic Research

Corrective taxes can solve many market failures, but actual policies frequently deviate from the theoretical ideal because of administrative or political constraints. We present a method to quantify the efficiency costs of constraints on externality-correcting policies or, more generally, the costs of imperfect pricing, using simple regression statistics. Under certain conditions, the R^2 and the sum of squared residuals from a regression of true externalities on policy variables measure relative welfare gains from policies. We illustrate via four empirical applications: random mismeasurement of externalities, imperfect electricity pricing, heterogeneity in the longevity of energy-consuming durable goods, and imperfect spatial policy differentiation.

A prior version of this paper was circulated under the title "Sufficient Statistics for Imperfect Externality-Correcting Policies." We thank Eduardo Azevedo, Richard Blundell, Meghan Busse, Don Fullerton, Alex Gelber, Jeff Grogger, Jean-François Houde, Kelsey

Electronically published March 12, 2020
[*Journal of Political Economy*, 2020, vol. 128, no. 5]

I. Introduction

Many important public policies aim to fix market failures that create wedges between marginal social costs and benefits. Many prominent examples are externality-correcting policies, which range from taxes on cigarettes, alcohol, or sugary beverages to mandatory immunizations to the regulation of pollution. Since Pigou (1932), economists have understood that if there are no additional market failures beyond the externality, market efficiency can be fully restored when externalities are taxed directly and the marginal damage at the optimal quantity is known. Yet relatively few policies closely follow this prescription. Often it is administratively impossible, technologically too costly, or politically infeasible to price actions according to the externalities that they generate.

Consequently, externality-correcting policies are generally imperfect. Imperfection often takes the following form: the externality is dependent on a set of variables, but policy is contingent on only a subset of those variables or their imperfect proxies. For example, the external damages from sulfur dioxide depend on the amount of pollution emitted, the weather, and the location of emissions relative to population centers. But sulfur dioxide regulations are based on only emissions quantities. In transportation, congestion externalities are highly concentrated in certain times of day, but most toll prices are uniform or vary only slightly with traffic conditions. In health, the externalities associated with second-hand smoke depend on many factors, including proximity to other people, whether the smoking is indoors or outdoors, and so on. But cigarette taxes are uniform.

In this paper, we develop a model that characterizes the welfare costs of using policies that take this form. We show that when certain conditions are met, familiar statistics from simple regressions of the true externality

Jack, Damon Jones, Ravi Kanbur, Henrik Kleven, Olivia Mitchell, Juan-Pablo Montero, Joel Slemrod, Kent Smetters, Johannes Spinnewijn, and seminar participants at the 2015 American Economic Association meetings, 2015 Association of Environmental and Resource Economists conference, spring 2015 National Bureau of Economic Research Environment and Energy Economics meetings, 2015 Northeast Workshop on Energy Policy and Environmental Economics, 2015 National Tax Association conference, Catholic University of Chile, Colorado School of Mines, Cornell University, Free University Amsterdam, International Monetary Fund, Imperial College, INSEAD, Institute for Fiscal Studies, London School of Economics, Michigan State University, New York University Abu Dhabi, Stanford University, Toulouse School of Economics, University College London, University of California Berkeley, University of California Los Angeles, University of Cambridge, University of Chicago, University of Chile, University of Colorado Boulder, University of Oxford, and University of Virginia for helpful comments and suggestions. Van Benthem thanks the National Science Foundation (award SES1530494), the Kleinman Center for Energy Policy at the University of Pennsylvania, and the Wharton Dean's Research Fund for support. Sallee thanks Google for support through a gift to the Energy Institute at Haas and the US Department of Agriculture (award CA-B-AEC-0174-H). Data are provided as supplementary material online.

on the variables upon which policy is based have direct welfare interpretations. Specifically, deadweight loss scales with the sum of squared residuals, and the R^2 summarizes the fraction of the welfare gain from a Pigouvian benchmark that is achievable by the constrained (which we call second-best) policy. We demonstrate the usefulness of the method through four empirical applications.

Our theory posits a standard model of a competitive market with a representative consumer who chooses among a variety of related goods, each of which produces a different level of an externality. A vector of Pigouvian taxes on these goods can restore efficiency, but we suppose that the planner faces a constraint, so that taxes must be made contingent upon some variable that is imperfectly correlated with the externality. This induces errors in the constrained optimal tax rates. We build on Harberger (1964) in deriving a general expression that characterizes the deadweight loss of some alternative set of taxes that deviates from the Pigouvian benchmark. Evaluating the full expression requires information about all cross-price derivatives of demand, which will typically be unavailable. However, under some conditions regarding the demand matrix, second-best policies will involve a set of taxes or shadow prices under which cross-product substitution does not affect overall welfare. Intuitively, what is required is that two products that are closer substitutes for each other do not, on average, have more similar tax rate errors.

We show that when conditions on the demand matrix are met, welfare conclusions can be drawn with only limited information. Given data on the distribution of the externality and its degree of correlation with the variables upon which policy is based, one can determine the proportion of the welfare gain achievable by the Pigouvian policy that the second-best policy achieves. Unlike results in the previous literature, this policy comparison does not require an estimate of any behavioral parameters. Given an estimate of the own-price derivative for the goods and the marginal damage due to the externality, the welfare costs of employing second-best policies in lieu of the Pigouvian benchmark can be estimated directly in dollars (rather than as a proportion).

It may also be the case that the analyst will be comparing a constrained policy to a product-level benchmark that is already second best. We extend our method to provide formulas that allow our core result, based on a Pigouvian benchmark, to be adjusted using only aggregate statistics about the difference between the true externality and the imperfect benchmark being used in policy. This allows our approach to continue to provide welfare results in a number of common settings with imperfect benchmarks.

To demonstrate the power of this method, we apply it to four distinct empirical problems. The first application considers random mismeasurement—energy efficiency is measured according to laboratory test procedures that differ from in-use averages, thereby creating mismeasurement in externalities across regulated products. We take advantage of a change in the fuel economy test procedure for automobiles in the United States to quantify the efficiency cost of basing fuel economy regulation on the older, noisier test ratings. We conclude that the second-best policy is quite efficient; it obtains more than 95% of the gains achieved by the Pigouvian benchmark.

Our second application regards real-time electricity pricing. Unlike our other three applications, this does not concern an externality. Instead, there is a wedge between marginal costs and benefits due to the fact that the marginal cost of generating electricity varies hour by hour, but electricity tariffs do not vary to reflect these costs (Borenstein and Holland 2005). We apply our method to characterize the welfare gain of tariffs that vary along some time or date dimensions but fall short of the theoretical ideal of real-time pricing. We find that realistic time-varying tariffs recover only a modest fraction of the gains achieved by real-time pricing.

Our third application concerns the regulation of energy-consuming durable goods that have heterogeneous total lifetime utilization. The lifetime pollution stemming from a durable good depends on both its energy efficiency and its lifetime utilization, but policies that regulate energy efficiency ignore differences in product longevity. We use a novel data set that indicates the lifetime miles traveled for a large sample of automobiles. We quantify that average lifetime miles traveled by individual vehicles of a particular model vary substantially across different models. This implies that vehicle models with the same fuel economy rating in fact have very different levels of expected lifetime carbon dioxide emissions. We conclude that actual fuel economy policies, which treat such vehicles identically, recover only about one-quarter to one-third of the welfare gain compared with a policy that considers both fuel economy and vehicle longevity. This result is robust even when we relax key assumptions about demand.

To illustrate our results, figure 1 shows the relationship between fuel economy ratings and average lifetime carbon emissions for different types of automobiles. Each data point represents the average lifetime CO_2 emissions across a number of individual vehicles of the same model (e.g., all 2012 Toyota Camry LE observations are combined into one data point). The line is the linear best fit. Dispersion in the data comes from heterogeneity in lifetime mileage; if all vehicles had the same lifetime mileage, the data would lie on a straight line. Federal fuel economy standards impose implicit taxes on vehicles that are a linear function of each vehicle's official fuel consumption rating; they cannot be based on average lifetime mileage. Our theory shows that under some conditions, the second-best fuel economy standard creates implicit taxes equal to the ordinary least squares (OLS) prediction line and the R^2 from this regression—0.29 in

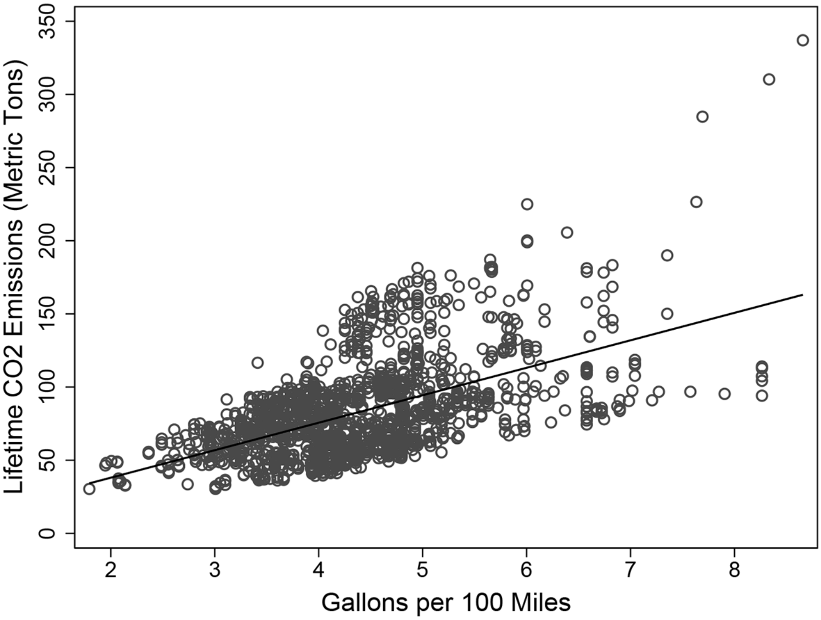

FIG. 1.—Relationship between lifetime CO_2 emissions and fuel efficiency. Each point represents a vehicle model. Solid line is an OLS regression line. The *X*-axis shows each model's fuel consumption rating: the number of gallons of gasoline per 100 miles driven. The *Y*-axis shows each model's average lifetime CO_2 emissions, calculated by dividing each model's average lifetime miles driven by its fuel economy rating to arrive at lifetime gallons of gasoline consumed and then multiplying by the tons of CO_2 per gallon of gasoline. The sample is restricted to models for which we observe at least 200 retirements from model years 1988–92. Data are described in detail in section V.

the case of figure 1—is an estimate of the fraction of the Pigouvian welfare gain that is achieved by this fuel economy policy.

A fourth application considers spatial differentiation. A given amount of pollution or energy use may have quite different health or environmental consequences depending on where it takes place, but policies often cannot differentiate their treatment by location. We use our framework to quantify the welfare costs of imperfect spatial differentiation for the case of carbon dioxide emissions resulting from the use of electric appliances. Here, differences in emissions across space are due to the fact that the emissions rate from the marginal power plant differs across regions of the country. This application also serves to demonstrate the broader applicability of regression statistics for welfare analysis because our required demand conditions will not hold for the policy we consider. Instead, we demonstrate that an alternative regression statistic, the within-R^2 from a regression with spatial fixed effects, has the desired interpretation. We conclude for this particular case that the welfare costs of failing to spatially differentiate are small.

Our method in general, and our analysis of fuel economy policy specifically, represents contributions to the evaluation of energy efficiency policies. No prior research has analyzed the implications of heterogeneity in lifetime utilization for the design of energy efficiency programs. This adds a new and apparently economically important dimension to the analysis of energy efficiency programs. In particular, it points out a new concern for the comparison between gasoline taxation and fuel economy standards as competing policies aimed to reducing greenhouse gas emissions from transportation.[1] Similar issues arise for any policy that regulates pollution control technology.

More broadly, our main contribution is to show the relationship between familiar regression statistics and second-best policies that aim to fix market failures but are constrained to be imperfect. This relates to the sufficient statistics literature in public finance, which is similar in seeking to find ways of characterizing welfare effects of policies that require information about a minimum number of parameters. Our analysis is unique in focusing on regression statistics and also adds to the small set of articles in this literature that are focused on externalities.[2]

Our analysis also connects to an important strand of literature in environmental economics that considers heterogeneity in damages from the same pollutant emitted in different locations. For example, the marginal damage from a ton of sulfur dioxide will differ depending on whether it is emitted near a densely populated city. A theoretical literature has noted

[1] For reviews of this literature for automobiles, see Parry et al. (2007), Anderson et al. (2011), and Anderson and Sallee (2016). Existing research, including Fullerton and West (2002), Fullerton and West (2010), and Feng, Fullerton, and Gan (2013), has considered how heterogeneity across consumers in driving behavior influences optimal policy design and welfare consequences, and Knittel and Sandler (2013) examine similar questions related to heterogeneity across individual automobiles in their local air pollution emissions rates. But no research considers heterogeneity in average lifetime utilization.

[2] Chetty (2009) documents a broad set of topics that have been considered by the literature on sufficient statistics in public economics, but he cites no papers focused on externalities. Recent work has included not only traditional questions in taxation (Feldstein 1999; Goulder and Williams 2003; Kleven and Kreiner 2006; Saez, Slemrod, and Giertz 2012; Hendren 2016) but also studies of social insurance (Baily 1978; Chetty 2006), health insurance (Einav, Finkelstein, and Cullen 2010), and limited rationality (Chetty, Looney, and Kroft 2009; Allcott, Mullainathan, and Taubinsky 2014). Hendren (2016) briefly notes that in order to fully assess a policy in the presence of externalities, one needs to know the effect of the policy on the externality net of many general equilibrium (cross-price) effects across a variety of related goods. But that paper does not propose a way to estimate this net effect, whereas we described conditions when they will cancel. One study that invokes the sufficient statistics tradition and does explicitly consider energy is the paper by Allcott, Mullainathan, and Taubinsky (2014), which models energy efficiency policy when heterogeneous consumers may undervalue energy efficiency due to limited rationality. They model a discrete choice between an efficient or inefficient good and derive sufficient statistics for the optimal combination of energy taxes and subsidies for energy efficient products.

that this type of spatial heterogeneity implies that uniform national policies are inefficient and suggested an efficiency gain from spatially differentiated regulation (Tietenberg 1980; Mendelsohn 1986; Baumol and Oates 1988). This type of concern has been used to study the potential benefits of spatial differentiation in policies regarding air pollution (Muller and Mendelsohn 2009; Muller, Mendelsohn, and Nordhaus 2011; Fowlie and Muller 2019), renewable energy generation (Cullen 2013; Callaway, Fowlie, and McCormick 2018), water pollution (Farrow et al. 2005), and electric vehicles (Holland et al. 2016). As we discuss in section VI, a number of these models can be understood as special cases of our general setup, and we suggest that our approach could offer a straightforward way of estimating potential gains from counterfactual policies in these contexts.

We investigate four distinct empirical applications in this paper, and we believe that the methods can be applied more broadly. The key data requirement is some measure of the distribution of the externality (or other efficiency wedge) and its correlation with the variables upon which policy is contingent. For a few more examples, consider the policies we described at the beginning of this introduction. The efficiency of sulfur dioxide trading programs could be assessed using estimates of the spatial distribution of marginal damages generated by Muller and Mendelsohn (2009) and Muller, Mendelsohn, and Nordhaus (2011).[3] The efficiency of various congestion pricing policies could be estimated using existing traffic data, such as the high-frequency records from thousands of locations in the California highway system (Caltrans 2016). Data on second-hand smoke exposure at home and in the workplace from the National Adult Tobacco Survey for the United States or the Global Adult Tobacco Survey could be used to estimate the efficiency of cigarette taxes as tools for mitigating externalities from second-hand smoke.[4]

The balance of the paper is as follows. In section II, we develop the theory for deriving sufficient statistics. In section III, we apply our method to the case of random mismeasurement in externalities, using a recent change in fuel economy testing procedures for automobiles. Section IV shows how our method applies to mispricing in electricity markets. In section V, we apply our results to heterogeneity in the longevity of automobiles. Section VI considers spatial heterogeneity in emissions from identical products used in different locations, using carbon emissions from refrigerators as an example. Section VII concludes.

[3] Spatial heterogeneity is not the only factor that determines the efficiency of SO_2 trading. Montero (1999), e.g., demonstrates that adverse selection in voluntary opt-in to the SO_2 trading program in the United States had significant efficiency impacts in the program's early years.

[4] For details on those data sources, see http://www.cdc.gov/tobacco/data_statistics/surveys/nats/index.htm and http://www.who.int/tobacco/surveillance/gats/en/.

II. Theory for Deriving Sufficient Statistics

The goal of our model is to facilitate analysis of the efficiency costs of policies that correct an externality or another wedge between marginal costs and benefits but that deviate from the theoretical ideal. Actual policies may be less efficient than an ideal policy for a variety of reasons, including political constraints, technological cost, and administrative feasibility. After presenting our model setup and notation, we first derive a general expression for the welfare loss from using some alternative, constrained policy in lieu of the ideal. We then specify sufficient conditions under which this general expression collapses so that simple regression statistics have welfare interpretations. Finally, we describe what can be learned from simple regression statistics even when our sufficient conditions are not met, and we present formulas for adjusting the welfare results in three common settings in which we compare a constrained policy against a benchmark that is already second best.

A. Model Setup

We emphasize a simple model in which there is only one market failure. We model a representative consumer in a perfectly competitive market. The economy has products indexed $j = 1, \ldots, J$. The consumer chooses quantities of each, denoted x_j. The consumer derives utility, U, from the consumption of these products according to the function $U(x_1, \ldots, x_J)$, which we assume is twice differentiable, increasing, and weakly concave in each argument. We denote the cost of production by $C(x_1, \ldots, x_J)$, which we assume is twice differentiable, increasing, and weakly convex in each argument. There is an exogenous amount of income in the economy, M, and all remaining income is consumed in a quasilinear numeraire, n. We assume no technological change and do not model the endogenous entry and exit of products into the market.[5] As such, ours is a short-run model, although one could allow for zero quantities so that the product vector represents potential products.

We posit that there is some market failure that leads the market, absent policy, to choose quantities so that there is a wedge, denoted ϕ_j, between the marginal private benefit and the marginal private cost of a unit of x_j. Our first assumption is that ϕ_j is fixed and unchanging with respect to policy intervention and that the total social inefficiency is the sum of these wedges across goods, multiplied by quantities: $\phi = \sum_{j=1}^{J} \phi_j x_j$. The simplest interpretation is that ϕ_j is an externality, as is the case in three of our four applications. In one of our applications, ϕ_j is a gap between

[5] For a treatment of how product redesigns can influence the design of a tax system that is limited in its ability to assign unique tax rates to each product, see Gillitzer, Kleven, and Slemrod (2017).

marginal cost and marginal benefit due to coarse pricing, where the cost of producing a good varies over time but the price is constrained to be constant.[6] Inspired by the externality interpretation, we refer to ϕ_j as the marginal social damage per unit of x_j.

Assumption 1. Marginal social damages from each product, ϕ_j, are fixed with respect to the tax vector t.

A natural way to think of our setup is that it models a sector of the economy—for example, j indexes types of refrigerators, and n is a separable bundle that represents all other goods. Each of the goods in the sector contributes varying amounts, ϕ_j, to a common externality—for example, the use of each refrigerator over its lifetime leads to a different amount of carbon dioxide, discounted to the present. The consumer ignores the externality when making choices, and the goal of the planner is to use taxes to internalize the externality.

The planner can impose product taxes, denoted t_j. We describe policies as taxes on products, but this is equivalent to regulatory policies that create implicit taxes (shadow prices). We assume that consumers remit taxes, so that the price to consumers is $p_j + t_j$. Revenue is recycled lump sum to consumers through a grant D. The consumer acts as a price taker. The consumer's optimization problem is

$$\max_{x_1, \dots, x_J} Z = U(x_1, \dots, x_J) + n \text{ such that } \sum_{j=1}^{J}(p_j + t_j)x_j + n \leq M + D. \quad (1)$$

The consumer's first-order conditions imply that $\partial U/\partial x_j = p_j + t_j$, which we assume holds at an interior solution.

Social welfare W is the utility from the product bundle, the numeraire (substituted out for the budget constraint), and the externality:

$$W = U(x_1, \dots, x_J) + M - C(x_1, \dots, x_J) - \sum_{j=1}^{J}\phi_j x_j. \quad (2)$$

We say the planner is unconstrained when she can set a unique tax rate on each product. In this case, the planner's problem is

$$\max_{t_1, \dots, t_J} W = U(x_1, \dots, x_J) + M - C(x_1, \dots, x_J) - \sum_{j=1}^{J}\phi_j x_j. \quad (3)$$

The first-order condition for product j is

$$\frac{dW}{dt_j} = \sum_{k=1}^{J}\left(\frac{\partial U}{\partial x_k} - \frac{\partial C}{\partial x_k} - \phi_k\right)\frac{\partial x_k}{\partial t_j} = \sum_{k=1}^{J}(t_k - \phi_k)\frac{\partial x_k}{\partial t_j} = 0, \quad (4)$$

[6] The wedge could come from other sources, such as market power, but our derivation assumes that ϕ_j is fixed with respect to the policy vector. Markups will generally shift with policy intervention, so application of our framework to market power would require modifications.

where the second equality follows from substituting the consumer's first-order condition and from our assumption of marginal cost pricing.

For wedges other than an externality, the same expression will arise as long as the wedge satisfies assumption 1. For example, one of our applications relates to coarse pricing—electricity prices are constant at all hours of the day, whereas marginal cost varies. In this case, ϕ_j is the gap between the price faced by the consumer (which equals marginal utility) and the true marginal cost. To consider that case, drop the externality term from equation (2). Then, differentiating W with respect to t_j yields the same $\Sigma_{k=1}^{J}(t_k - \phi_k)(\partial x_k/\partial t_j)$.

Equation (4) shows that all J first-order conditions for the planner will be met if and only if $t_j = \phi_j \; \forall \; j$. That is, the planner's optimum is a vector of Pigouvian taxes; each product's tax rate is set equal to its marginal external damage. Establishing this Pigouvian benchmark empirically will require a measure of the externality associated with each product. We derive the core theoretical results assuming that such data exist and then relax this in section II.E to show how our results can be adjusted when the benchmark policy itself is not Pigouvian.

We wish to characterize how welfare under the Pigouvian benchmark compares with that under a policy subject to some constraint. The difference represents the cost of the constraint on policy design. The constraint is a restriction on the vector of taxes that the planner can choose. We write this constraint as a function $t_j = g(f_j; \theta)$, where f is some vector of exogenous attributes of the products, θ are parameters to be chosen by the planner, and g is some function. The planner's problem can now be written as

$$\begin{aligned} \max_{\theta} W = U(x_1, \dots, x_J) + M - C(x_1, \dots, x_J) \\ - \sum_{j=1}^{J}\phi_j x_j \text{ such that } t_j = g(f_j; \theta) \; \forall \; j. \end{aligned} \tag{5}$$

We call the solution to this policy, denoted $t_j = g(f_j; \theta^*)$, the second-best, or constrained, tax vector. Recall that our goal is to provide welfare interpretations of regression statistics. Motivated by this, we restrict attention to situations where $g(f_j; \theta)$ can be written as linear in parameters, noting that this is no more restrictive than it is in any application of (multivariate) linear regression, where variables can be transformed and interacted. For example, in our third application, we consider fuel economy regulations that impose a shadow tax on vehicles that is an affine function of their fuel economy ratings. Thus, $g(f_j; \theta) = t_j = \alpha + \beta f_j$, where θ consists of two parameters, α and β, and f_j is the fuel economy rating. Our four applications demonstrate a variety of policy design constraints that fit into this framework.

Our objective is to describe the welfare cost of such policy constraints relative to the Pigouvian benchmark. We note that the Pigouvian benchmark itself is not necessarily first best in the presence of other market failures or margins of adjustment that product-based taxes cannot correct.[7] For example, taxes on new vehicles cannot induce optimal scrappage behavior. Therefore, our Pigouvian new vehicle tax vector falls short of a first-best tax on gasoline. We discuss this in more detail in section V. In such cases, our method considers the welfare gain along a particular dimension of interest that is targeted directly by the tax, assuming that other distortions are held constant. We return to this point in the applications.

To describe the welfare consequences of such policy constraints, we now proceed to deriving a generic expression that characterizes the loss of social welfare caused by moving from the Pigouvian benchmark policy to some arbitrary tax vector. We then use this expression to relate second-best policies that would arise given a particular constraint.

B. *Characterizing Deadweight Loss*

Let a generic tax schedule be denoted as $\tau_1, \ldots, \tau_J$. We characterize the welfare loss of moving from the Pigouvian benchmark $t_j = \phi_j$ 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. We denote the difference in welfare between the two schedules as $DWL(t = \tau)$.[8] To do so, we assume local linearity.

Assumption 2. Demand derivatives $\partial x_j / \partial t_k$ are constant between ϕ_j and τ_j for all j and k.

Under the assumption of constant demand derivatives, the efficiency loss incurred from imposing any arbitrary tax schedule τ in lieu of the Pigouvian tax schedule can be written as

$$\begin{aligned} W(t = \phi) - W(t = \tau) &\equiv DWL(t = \tau) \\ &= -\frac{1}{2}\sum_{j=1}^{J}\sum_{k=1}^{J}(\tau_j - \phi_j)(\tau_k - \phi_k)\frac{\partial x_j}{\partial t_k}. \end{aligned} \tag{6}$$

The proof, along with all others, is in appendix A (apps. A–D are available online). This formula is in the form of a set of Harberger triangles, and indeed the same result (although without externalities) is in Harberger (1964). When $\tau_j = \phi_j$, each term in the summation will be zero.

[7] For simplicity, even though the Pigouvian benchmark will not be first best in all settings, we refer to the constrained policy as second best.

[8] We consider relative benefits of policies, but the final policy choice also depends on relative costs (e.g., administrative or technology costs). Our method provides a bound on the costs that would make the less precise policy better overall. Note that in our four empirical examples, the Pigouvian benchmark is technically feasible and requires only a set of tax rates or prices, where taxes and prices are already being charged. It seems unlikely then that administrative costs would be a major factor in comparing policies in our settings, but this may not be true in other situations.

In line with the traditional use of Harberger triangles, we assume that demand derivatives are constant over the relevant range of taxes. In our discussion here, we also assume that producer prices are unchanged, which implies constant marginal cost. In this case, $\partial x_j/\partial t_k$ represents only a demand derivative, not a combined effect of supply and demand. Where marginal cost is increasing but linear, our mathematical results are all the same, but $\partial x_j/\partial t_k$ is interpreted as the combined response of supply and demand (see app. A).

Where demand or marginal costs are convex, our results represent a local approximation in the same way that Harberger triangles normally do. Thus, our derivations can also be understood as indicating incremental welfare losses from small movements away from the Pigouvian benchmark. Note that we relax both the linearity and the constant marginal cost assumptions in our electricity pricing application, but we preserve them here for the exposition.

To better understand the content of equation (6), we substitute $e_j \equiv \tau_j - \phi_j$, where e_j is the error in the tax rate, and decompose the own and cross effects:

$$-2 \times DWL(t = \tau) = \sum_{j=1}^{J}\sum_{k=1}^{J} e_j e_k \frac{\partial x_j}{\partial t_k} \tag{7}$$

$$= \underbrace{\sum_{j=1}^{J} e_j^2 \frac{\partial x_j}{\partial t_j}}_{\text{own effects}} + \underbrace{\sum_{j=1}^{J}\sum_{k\neq j} e_j e_k \frac{\partial x_j}{\partial t_k}}_{\text{cross effects}}. \tag{8}$$

Equations (7) and (8) are quite general expressions. But using these formulas to evaluate policy alternatives requires knowledge of the complete demand matrix, including all cross-price derivatives. This information will frequently be unavailable.

Under some conditions, however, the expression will simplify further and policy evaluation will require less information.[9] Specifically, the cross effects in equation (8) will be zero when there is no substitution between goods, so that cross-price derivatives are all zero. Alternatively, the cross effects will be proportional to the own effects when the errors in the tax rates are mean zero and products of errors are uncorrelated with cross-price derivatives. Note that we refer to $\partial x_j/\partial t_k$ as cross-price derivatives and the contribution of the $e_j e_k \partial x_j/\partial t_k$ to deadweight loss as cross effects. Zero cross-price derivatives are a sufficient condition for the cross effects to simplify, but so are the alternative conditions listed below.

[9] Goulder and Williams (2003) also build from the general Harberger formula and present a simplified expression for the excess burden of taxation that does not require estimates of all cross-price derivatives. They study interactions between commodity and labor taxes, a very different setting from ours.

We state these possibilities formally as assumption 3. We then proceed to derive results under the case where assumption 3 holds before returning to a detailed discussion of these conditions.

ASSUMPTION 3.

a. Tax errors e_j are uncorrelated with own-price derivatives: $\mathrm{cov}(e_j, \partial x_j/\partial t_j) = 0$.
b. Products of tax errors $e_j e_k$ are uncorrelated with cross-price derivatives: $\mathrm{cov}(e_j e_k, \partial x_j/\partial t_k) = 0 \ \forall\ j \neq k$. (A stronger version of this, b', assumes that cross-price derivatives are zero: $\partial x_j/\partial t_k = 0 \ \forall\ j \neq k$.)

Here $\mathrm{cov}(e_j e_k, \partial x_j/\partial t_k)$ is calculated for all nondiagonal elements of the demand matrix ($j \neq k$). Version b′ of this assumption holds if there is no substitution across products. Version b assumes that cross-price derivatives between each pair of products are uncorrelated with the product of their tax errors. This holds if externalities, conditional on policy, are orthogonal to substitutability. As we discuss further below, this is a plausible property of second-best policies. To provide a more intuitive economic interpretation, we note that one way that this assumption can be satisfied is if (1) for each product j the errors of its substitutes are uncorrelated with the cross-price derivatives and (2) across products j the tax errors are uncorrelated with average cross-price derivatives.

In our empirical applications, we provide examples where part b′ is likely to hold by approximation (electricity pricing) as well as cases in which part b is reasonable (fuel economy standards and noisy energy efficiency ratings). Nevertheless, assumption 3 will not hold in all cases, so we review what can be learned when the conditions do not hold after establishing our primary results that obtain under assumption 3. Moreover, we provide numerous robustness checks throughout our empirical applications.

C. Welfare Statistics When DWL Is Proportional to Squared Tax Errors

Under assumptions 1 and 2 and the strong version of assumption 3 (parts a and b′), and assuming unbiasedness on average so that $\Sigma_{j=1}^{J} e_j = 0$, the deadweight loss of an arbitrary tax vector is given by

$$DWL = -\frac{1}{2}\overline{\frac{\partial x_j}{\partial t_j}}\sum_{j=1}^{J} e_j^2, \tag{9}$$

which foreshadows a central role for minimizing a sum of squared tax errors. When we use the weaker variant of assumption 3 (parts a and b), the welfare loss expression remains very similar:

$$DWL = -\frac{1}{2}\left(\overline{\frac{\partial x_j}{\partial t_j}} - \overline{\frac{\partial x_j}{\partial t_k}}\right)\sum_{j=1}^{J} e_j^2, \tag{10}$$

where the average cross-price derivative (over the nondiagonal entries of the demand derivative matrix) $\overline{\partial x_j/\partial t_k} = \{1/[J(J-1)]\}(\Sigma_{j=1}^{J}\Sigma_{k\neq j}\partial x_j/\partial t_k)$. Note that *DWL* in equation (10) is still proportional to the sum of squared tax errors, but it is multiplied by the difference in the average own-price and the average cross-price derivative. Because the proportionality of the *DWL* is maintained, all propositions and corollaries below hold exactly using either variant of assumption 3. The algebra leading to equations (9) and (10) appears in appendix A.

When the number of goods J is large, $\overline{\partial x_j/\partial t_k}$ will become small. Thus, equation (9) will be a close approximation of the *DWL* in equation (10) even in cases where only the average own-price elasticity is known. Moreover, even for smaller values of J, $\overline{\partial x_j/\partial t_k}$ will shrink if the substitution to the outside good becomes larger.

The solution to the planner's constrained problem in equation (5) is the same as from minimizing the deadweight loss in equation (10) subject to the same constraint.[10] This makes the link between policy and regression obvious. Whenever the policy constraint $t_j = g(f_j;\theta)$ can be written as a function that is linear in parameters, minimization of deadweight loss is the same as minimizing the sum of squared residuals in a regression of the true externalities on the tax rates.

When assumptions 1–3 hold, the second-best policy will be to choose α and β to be the OLS solutions from fitting the externality to the policy variable. This is stated in proposition 1:[11]

Proposition 1. Under assumptions 1–3, the second-best policy is the OLS fit of ϕ_j to f_j, and the deadweight loss is proportional to the sum of squared residuals:

$$DWL = -\frac{1}{2}\left(\overline{\frac{\partial x_j}{\partial t_j}} - \overline{\frac{\partial x_j}{\partial t_k}}\right)SSR. \tag{11}$$

The proof is in appendix A. The intuition is as follows. When the externalities, conditional on characteristics that are in the policy function, are uncorrelated with product substitutability, then the deadweight loss is a linear function of the sum of squared tax errors and the sum of errors (bias in the tax) squared. We show that this objective function is minimized by the same line that minimizes the sum of squared tax errors: a simple OLS fit. We show below how weighted least squares provides a similar solution when own-price derivatives may be correlated with the error.

[10] Deadweight loss is just the objective function evaluated at the Pigouvian benchmark minus the objective function evaluated at an alternative tax vector, so the original objective function is just deadweight loss plus a constant term.

[11] For expositional ease, we derive results for the case where policy is contingent on one exogenous variable, denoted f_j, and the tax policy takes the form of a linear function of f_j. Then the policy choice is to choose α and β, where $t_j = \alpha + \beta f_j$. It is straightforward to modify our derivation to include many variables.

In turn, the resulting deadweight loss is the sum of squared residuals from the OLS regression scaled by the average demand derivative and an average cross-price derivative factor that is close to zero when the number of products J is large. Thus, given data on the externality and own-price derivatives and the attributes upon which policy is based, an analyst can run a simple linear regression and assign direct welfare interpretations to the regression output.

Moreover, the R^2 from this regression is a sufficient statistic that summarizes the percentage of welfare gain that could be achieved by the Pigouvian benchmark that is achievable by the second-best constrained policy. The percentage gain in welfare must be defined relative to some benchmark. The R^2 is defined relative to a benchmark policy that imposes a uniform unbiased tax rate $\bar{t}$ that is the same for all products.[12]

Corollary 1. Under assumptions 1–3, the R^2 from the OLS fit of ϕ_j to f_j represents the percentage of the welfare gain of the Pigouvian tax (relative to a baseline of a uniform unbiased tax $\bar{t}$) that is achieved by the second-best linear tax on f_j (relative to the same baseline):

$$R^2 = \frac{DWL(t = \alpha^{OLS} + \beta^{OLS} f_j) - DWL(t = \bar{t})}{DWL(t = \phi) - DWL(t = \bar{t})}. \tag{12}$$

Under our assumptions, the R^2 relaxes the information requirement of knowing own-price derivatives and also eliminates the small adjustment factor involving the average cross-price derivative. No moments of the demand system are required to calculate this sufficient statistic. This makes assessing the relative welfare gain very intuitive and easy: all that is required is running a simple OLS regression of the actual externality for each product on the variables used in the policy function.

We now relax part a of assumption 3 to allow correlation between errors and own-price derivatives. This leads to a very intuitive relationship with weighted multivariate regression:

Proposition 2. Under assumptions 1, 2, and 3 (b), the second-best policy is the weighted least squares fit of ϕ_j to a vector of attributes f_j, where the weighting matrix is diagonal with each entry equal to the own-price derivative for product j.

Given information about the own-price derivatives of each product, a researcher could calculate the weighted least squares (WLS) estimator and derive parallel welfare results for this case. The second-best policy is still a linear best fit; the deadweight loss is the weighted sum of squared residuals from that regression. The proof appears in appendix A. Further, when relaxing assumption 3 altogether, the second-best policy is

[12] The application of the deadweight loss formula in eq. (10) to this benchmark requires applying assumptions 1–3. When we relax assumption 3 in sec. II.D, we also relax its application to the benchmark policy.

the generalized least squares fit of ϕ to f where the weighting matrix is the full demand matrix. We do not emphasize this result because it requires additional information about the demand system, but in many instances this formula would be useful for robustness analysis. We demonstrate such a calculation in section V.F, using estimates on the full matrix of demand elasticities for automobiles.

Corollary 2. Suppose A and B are second-best policies with attribute vectors $f_{A,j}$ and $f_{B,j}$, as given by proposition 1 under assumptions 1–3. The welfare difference between A and B is proportional to $R_A^2 - R_B^2$.

Corollary 2 (see app. A for the derivation) compares two imperfect policies against a Pigouvian benchmark. This allows an intuitive and straightforward ranking of constrained policies by directly comparing their R^2's. Policies that better match to the true externality are more efficient, and the R^2 allows for the gains to be quantified. The welfare difference between two policies is a collection of trapezoids, the rectangle-shaped portions of which grow larger the farther the two imperfect policies are from ϕ_j.

Interpreting the assumptions about cross effects.—We now discuss the key economic implications of the assumptions needed for the results above. Part a of assumption 3 says that the strength of own-price derivatives is not correlated with a product's tax error; that is, whatever factors that determine the externality but are omitted from the policy function do not also indicate stronger or weaker own-price responses. Proposition 2 relaxes this assumption. Doing so is important in empirical applications where some products are demanded in much larger quantities than others (and so have larger own-price derivatives, all else equal).

The second part of assumption 3 has more economic content. The strong version of 3 (b′) applies to markets where products are not substitutes or complements. This assumption is unlikely to hold, although we argue in section IV that it applies, at least by approximation, to electricity pricing. Even when cross-price derivatives are not zero, corollary 1 will still apply as long as the weaker version of 3 (b) holds: this says that the difference between the errors in the tax rates between two products is no smaller or larger when the two products are closer substitutes. The errors in tax rates represent the residual variation in the externality, after conditioning on the attributes upon which policy is contingent, f. Consider the vehicle example. Two vehicles with similar externalities (ϕ) will be closer substitutes, provided that vehicle fuel economy (f) is a factor that determines vehicle choice, because ϕ is mechanically related to f. But assumption 3 can still be met if, after conditioning on fuel economy, the residual variation in the externality ϕ is not correlated with substitutability. Whether this will be true depends on the variables that are included in the policy and the source of residual variation in the externality. We discuss assumption 3 in more detail for each of our empirical applications.

The results in this section demonstrate that—under assumptions that are often plausible—the deadweight loss of deviating from the Pigouvian benchmark can be calculated with limited information about the market. The welfare gains possible in the second-best relative to those in the Pigouvian case can be calculated with even less information. In the next four sections, we demonstrate that these theoretical results have empirical relevance by illustrating four situations in which a sufficient statistic useful for evaluating policy can be derived from this framework.

D. *What Information Remains in the R^2 When the Cross Effects Do Not Simplify?*

In this subsection, we explore cases where the R^2 is biased (because our assumptions do not hold) but that bias can be signed, so the R^2 is interpretable as a bound on welfare effects. To be precise, we consider what the R^2 indicates about the welfare gains from the linear best-fit policy, showing when this overstates or understates welfare gains. When our assumptions do not hold, this linear best-fit policy may not be second best. But we still think it is the most interesting candidate policy to analyze for many situations where the policy maker lacks the detailed information about demand needed to determine how the second best deviates from the best fit. We first present a formula that highlights how different forces push the true welfare ratio away from R^2 in different directions. We then make suggestions for how empiricists might investigate the potential bias on the basis of the type of violation.[13]

When we do not impose assumption 3 so that cross effects do not simplify, we can still write out an expression for the relative gain in welfare achieved by the linear best-fit policy over a uniform tax policy, divided by the gain from the Pigouvian benchmark over the same uniform tax. We denote this welfare gain by S and compare it with the R^2:

$$S = 1 - \frac{-(1/2)\left(\overline{\partial x_j/\partial t_j}\right)SSR^{\textit{second-best}} - (1/2)\sum_j\sum_{k\neq j} e_j e_k\left(\partial x_j/\partial t_k\right)}{-(1/2)\left(\overline{\partial x_j/\partial t_j}\right)TSS^{\textit{second-best}} - (1/2)\sum_j\sum_{k\neq j}\lambda_j\lambda_k\left(\partial x_j/\partial t_k\right)}, \tag{13}$$

where λ_j are the residuals in the regression of ϕ on a constant (the uniform policy). Note that $\lambda_j = \gamma_j + e_j$, where γ_j is defined as the explained portion in the linear regression: $\gamma_j = \alpha^{OLS} + \beta^{OLS} f_j - \bar{\phi}$. Because γ is a function of f, the tax errors in the uniform policy depend on f. Thus, equation (13) allows cross-price derivatives to be correlated either with e or with f (and thus λ).

[13] Another approach is to derive bounds on the deadweight loss from eq. (7). We have constructed analytical bounds based on properties of quadratic forms and their eigenvalues, but they will be informative in only special cases. This may be a promising area of future research.

We now consider two types of correlation that determine the direction of the bias in R^2: first, correlation between products of constrained policy errors and cross-price derivatives $\mathrm{cov}(e_j e_k, \partial x_j/\partial t_k)$ (type 1), and second, correlation between products of the policy variable and cross-price derivatives $\mathrm{cov}(\gamma_j \gamma_k, \partial x_j/\partial t_k)$ (type 2).

PROPOSITION 3. Under assumptions 1 and 2, (1) $R^2 < S$ if type 1 correlation is positive (but type 2 correlation is zero) and (2) $R^2 > S$ if type 2 correlation is positive (but type 1 correlation is zero).

The proof is in appendix A. First, consider part 1 of proposition 3. If cross-price elasticities are larger for goods with similar tax errors (e.g., vehicle durability in application 3), then the true fraction of welfare recovered in the second-best policy increases relative to the R^2 measure. The intuition here is that when goods with similar tax errors are good substitutes, the Pigouvian benchmark loses some of its advantage: consumers do not substitute much along this dimension anymore, and so the two policies become more similar, acting mostly along the margin of reducing f.

Now consider part 2. If cross-derivatives are large when λ_j and λ_k (a function of the observable attribute f, such as fuel economy in application 3) are similar, then the true fraction of welfare recovered by the second-best policy S will decrease relative to R^2. The intuition for this follows from observing that correlation of substitutability with f makes the second-best policy less effective because consumers now substitute mainly among products with similar f. The Pigouvian benchmark is still based on both f and the tax error, and so its effectiveness is not damaged as much.

It is important to note that in many common settings, both types of positive correlation are likely to be present, and sometimes the bias cancels out. For example, as we discuss in section V, cars that are strong substitutes are relatively likely to have similar fuel economy and similar durability.

The results above are directional and qualitative. In many cases, as we illustrate in the applications, simulation of the true welfare gain using a range of plausible demand elasticities can be highly informative. This usually does require some knowledge on the structure of the demand matrix, for example, from existing empirical work in the literature.

E. Comparisons When the Benchmark Policy Is Imperfect

It may also be the case that the analyst will be comparing a constrained policy against a benchmark that is already second best. This happens when the data used on the left-hand side are an imperfect proxy for the true externality. In what follows, we will call the value that results from the regression using an imperfect benchmark $\hat{R}^2$ and the true fraction of welfare recovered by the constrained policy relative to a Pigouvian benchmark S. The following three scenarios span a range of data limitations an analyst could be working with, corresponding to different settings in which the benchmark policy is not Pigouvian. We show how unbiased estimates of the welfare loss from constrained policies are still recoverable using $\hat{R}^2$ and features of the scenario:

1. Measurement Error

When the analyst is facing classical measurement error in the externality, the intuition is fairly simple: the observed $\hat{\phi}_j$ is noisier than the true ϕ_j. Examples are sampling error from limited microdata, errors produced by engineering test cycles, or other errors made when the analyst produces a best guess of the product-level externality.

To implement our method in this setting, the analyst needs estimates of the externalities (or other wedges) and aggregate information on the degree of measurement error. We use the following notation: $\hat{\phi}_j \equiv \phi_j + \nu_j$, where ϕ_j is the true wedge, $\hat{\phi}_j$ is the observed or estimated wedge, and ν_j is therefore the error in measurement. Consistent with classical errors, suppose that ν_j is independent of ϕ_j (and of any regressors that are used in determining the tax scheme) and is distributed normally with mean zero and variance σ_ν^2.

Suppose that the analyst regresses $\hat{\phi}_j$ on f_j. This is a situation of errors in the dependent variable, so errors do not cause bias in the coefficients and the second-best policy is still consistently estimated. However, $\hat{R}^2$ is biased downward because of the noise from mismeasurement. A simple derivation (see Majeske, Lynch-Caris, and Brelin-Fornari 2010) shows that in expectation,

$$\hat{R}^2 = S\left(1 - \frac{\sigma_\nu^2}{\sigma_{\hat{\phi}}^2}\right), \tag{14}$$

where $\sigma_{\hat{\phi}}^2$ is the variance in $\hat{\phi}_j$. In terms of welfare interpretations, it implies that the second-best constrained policy will have larger welfare gains than indicated by the estimated statistic. In practical terms, where measurement error is a concern and errors are classical, an analyst can inflate the R^2 upward given an estimate of the signal-to-noise ratio in the data.

Another relevant source of noise in our setting could arise from the use of microdata in performing the regression. Conceptually, the regression we have in mind should be run at the same level of detail over which the policy is being applied; if the benchmark policy differs with product j, then a data set containing different individual observations for each product should be collapsed to the product level before computing R^2. However, it is still possible to adjust the R^2 from a regression run on the microdata to recover the relevant welfare statistic. The average variance in the microdata across products, $\Sigma\sigma_j^2/J$, can be substituted into equation (14) in place of σ_ν^2. The resulting adjustment produces a value that is computationally equivalent to the R^2 from the product-level regression, reflecting the welfare statistic we have in mind (see app. A for details).

2. Uncertainty in the Magnitude of External Damages per Unit

Suppose that the driver of external damage is measured well, but the magnitude of the externality per unit is unknown. For example, gasoline use in cars may be measured well, while the social cost of carbon is unknown. If the researcher or policy maker uses a damage valuation above or below the true value, this introduces a slope error in the constrained policy, making it perform even worse relative to the Pigouvian benchmark. The difference in tax rate between products is not large enough if the true social cost of carbon is higher than that used in the analysis, and vice versa. As a result, $\hat{R}^2$ will overstate the welfare gain that the constrained policy produces relative to the true Pigouvian benchmark.[14]

The correction needed in this case increases nonlinearly in the distance to the true damage parameter. Suppose that the damage valuation used to compute the externality in the regression determining policy is $\hat{\delta}$, while the true damage parameter is δ. The welfare gain S of the policy being put in place (based on $\hat{\delta}$) relative to the true Pigouvian benchmark policy is

$$S = \hat{R}^2\left[1-\left(\frac{\hat{\delta}}{\delta}-1\right)^2\right]. \tag{15}$$

The derivation is in appendix A, and the intuition is as follows. When $\hat{\delta}/\delta = 1$, the benchmark policy has the right slope and $S = \hat{R}^2$. When $\hat{\delta}/\delta \in (0, 1)$, the externality estimate is too low and $S < \hat{R}^2$. The adjustment is quadratic. For example, when $\hat{\delta}/\delta = 0.5$, $\hat{R}^2$ needs to be adjusted downward by 25%. The adjustment is symmetric for the case where $\hat{\delta}$ is too high: the policy is now encouraging too much switching among products.

3. A Missing (Uncorrelated) Component of the Externality

Finally, suppose instead that there is a missing component of the externality that is uncorrelated with the observed component. For example,

[14] The welfare implications here come from relative tax errors among the J products in the setup: substitution to goods outside the model is controlled for in the benchmark.

in a setting where each j represents traffic delay on a bridge in different hours, it will be the case that random variation in weather conditions can produce more (i.e., larger ϕ_j for that hour) or less (smaller ϕ_j) delay separately from the underlying traffic pattern used to calculate $\hat{\phi}_j$. In this case, the observed $\hat{\phi}_j$ has less variance (since only one component of the externality is included) than the true ϕ_j. Given an estimate of the variance of the unobserved component, we can again recover an unbiased estimate S.

Suppose that $\phi_j \equiv \hat{\phi}_j + \xi_j$, where ϕ_j is the true externality, $\hat{\phi}_j$ is the observed externality, and ξ_j is an unobserved component independent of $\hat{\phi}_j$. Write its variance as σ_ξ^2. The analyst regresses $\hat{\phi}_j$ on f_j, producing $\hat{R}^2$. Rearranging equation (14) for this setting, we have

$$S = \hat{R}^2\left(1 - \frac{\sigma_\xi^2}{\sigma_\phi^2}\right) = \hat{R}^2\left(1 - \frac{\sigma_\xi^2}{\sigma_{\hat{\phi}}^2 + \sigma_\xi^2}\right). \tag{16}$$

Here $\hat{R}^2$ becomes an upper bound on the fraction of welfare recovered. We note that this case, combined with the one above, could also encompass more complex missing externalities that are correlated with the observed component. Consider emissions externalities from vehicles: we can measure lifetime gasoline consumption (providing a good estimate of carbon emissions) but may worry about the implications of omitting lifetime damages from local air pollution. Local emissions are correlated with weight—and therefore gasoline use—but not perfectly. If the researcher knows this correlation β (e.g., from a second data set or the engineering literature), she could add the predicted air pollution externality to the observed externality: $\hat{\phi}_{j,new} = \hat{\phi}_j + \beta \times$ gallons-per-mile. This still leaves out a—now uncorrelated by construction—piece of the local air pollution externality, which could then be corrected as above.

F. *Summary: When Is Our Theory Applicable?*

The central point of our theory is that simple regression statistics often contain intuitive information about the welfare properties of corrective policies that face some design constraint. Figure 2 provides a visual summary of the situations under which our results obtain, which is intended to serve as an initial guide for those considering our methods in other applications.

All of our theory assumes that in the absence of policy, consumption of a good deviates from the social optimum because of some wedge ϕ, such as an externality. Our base assumptions are that these wedges are fixed with respect to prices and that demand and supply are linear over the relevant range, as is generally assumed in the analysis of Harberger triangles. Under those assumptions, equation (8) expresses the deadweight loss of

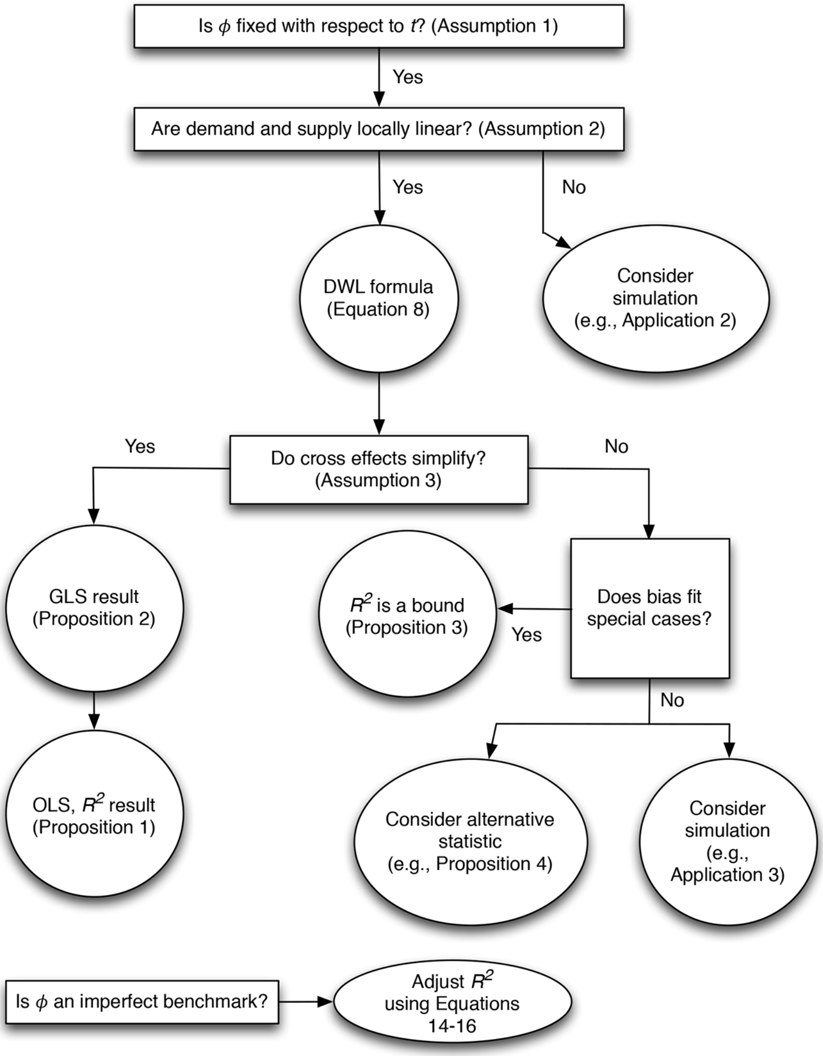

FIG. 2.—Schematic of theoretical results.

an arbitrary vector of taxes that deviates from the Pigouvian benchmark. When demand and supply are not locally linear, it is possible to amend our results through simulation, which we illustrate in our electricity pricing application.

When the conditions of assumption 3 are met, then our results about the interpretation of the sum of squared residuals and the R^2 will hold (proposition 1). We interpret our first two applications—to electricity pricing and noisy laboratory measures—as meeting these conditions most closely.

Even when violations of the assumptions are significant, the R^2 may be a useful bound. As described in proposition 3, particular types of correlations between tax errors and demand will create predictable bias in the R^2 as a measure of welfare gain. In our vehicle longevity application, we demonstrate the size of this bias after introducing correlations calibrated from the literature.

In some cases, the cross effects will not simplify, and the bias will not fit the special cases embodied in proposition 3. In that case, we suggest two approaches. One is to look for a modified relationship between regression statistics and welfare. This is what we do in application 4. There, we argue that R^2 will be substantially biased but that an alternative set of assumptions appropriate to the setting imply that the within-R^2 from a fixed effects regression has the desired interpretation (proposition 3). Other approaches that incorporate additional market failures, endogenize ϕ, or consider other relaxations of our assumptions are key topics for future research. The other approach is to use simulation to determine whether calibrated degrees of correlation between tax errors and the demand system indicate that the bias in the R^2 will be small or large. We demonstrate this approach in our vehicle longevity application.

Finally, at the bottom of figure 2, we call out a practical consideration. Here the Pigouvian benchmark is unknown, for example, because of classical mismeasurement of the wedges or uncertainty in the magnitude of the external damages, or in the case of a missing component of the externality, R^2 will be biased. This can be corrected using equations (14)–(16) when summary statistics about the difference between the true externality and the information that the analyst is using to set policy are known (or assumed in sensitivity analysis).

III. Application 1: Noisy Energy Efficiency Ratings

One reason that taxes or regulatory incentives for energy-consuming products may be imperfectly related to the true externalities that they generate is that the energy efficiency ratings themselves are imperfect. To determine the energy efficiency rating of a product, governments establish a laboratory test procedure. The government or the manufacturers themselves then test a prototype or example product. Actual performance in the field can differ from lab test results, and when it does, policies based on the official ratings will be imperfect indicators of the actual externalities associated with each product.[15] This creates inefficiencies,

[15] Such mismeasurement naturally also occurs for nonenergy goods and externalities. For example, to help prevent obesity, calorie labeling on menus will be mandatory for

and our theoretical framework can be used to quantify the consequent welfare losses.

In general, the challenge in studying this phenomenon is that it requires credible measures of average in-use energy efficiency that can be compared with the official rating. Scattered evidence of in-use performance does exist for some products, but we take a different approach here and analyze a change in the US rating system for automobiles that was meant to address mismeasurement. The Environmental Protection Agency (EPA) began measuring fuel economy of automobiles in 1978 in support of the Corporate Average Fuel Economy (CAFE) program, which mandates that each firm meet a minimum average sales-weighted fuel economy of vehicles. The ratings are based on a laboratory test during which a vehicle is driven on a dynamometer (a treadmill for cars) through a specific pattern of speeds and accelerations. The test procedure established in 1978 included two courses to represent urban and highway driving. The two ratings were averaged to determine each vehicle's rating for the CAFE program. These same ratings were presented to consumers on fuel economy labels.

In 1986, in response to consumer complaints that the ratings systematically overstated fuel economy, the EPA revised the ratings downward by simply scaling them by the same amount for all vehicles. CAFE continued to use the original values to determine automakers' compliance, but consumer labels were updated. Over time, the revised ratings were deemed to be inaccurate as well. The original test used low highway speeds, did not involve the use of air conditioning, and generally became less accurate as automobile technology and average driving patterns changed. Yet again, the EPA instituted a new test procedure in 2008 that changed the ratings substantially on average and also more for some vehicles than for others.[16]

For political reasons, however, the CAFE program continues to use the less accurate original rating system from 1978.[17] While consumers are now provided with the more accurate updated ratings, the regulation (and hence the regulatory shadow price faced by automakers) are still based on the noisy original system.

We can use our theoretical framework to quantify the welfare costs of using the old rating system in lieu of the updated one via simple linear regression. Our thought experiment is the following. We suppose that

---

many US restaurants. Using lab tests, Urban et al. (2011) found that while menus are, on average, pretty accurate, substantial variation exists. About 20% of foods purchased had at least 100 more calories than what was reported.

[16] This procedure involved five separate dynamometer tests—the original two tests and three new ones. Several tests are combined to determine the highway and city ratings that appear on fuel economy labels for consumers.

[17] Evidently, it was determined that changing the rating that entered the CAFE compliance program would require a political battle not worth waging. Changing the CAFE ratings would have created winners and losers among automakers.

(1) the new rating represents the true fuel economy rating of a vehicle; (2) after a linear adjustment, the old rating is a white noise mismeasurement of the truth; and (3) the policy maker is sophisticated and is aware of the inaccuracy in the old rating but must base policy upon it because of political or legal constraints. In other words, a sophisticated regulator can take out overall bias/tilt in measurement but does not observe car-specific mistakes. These assumptions likely hold in practice—reported on-road fuel economy is close to the 2008 EPA ratings, and the EPA explicitly presents differences between window sticker and regulatory CAFE fuel economy values.[18]

A. *The Pigouvian Benchmark versus the Constrained Policy*

In this application, we take a simplified view of the externalities associated with fuel economy and assume that the externalities associated with an automobile are proportional to its true fuel consumption per mile. This is consistent with how fuel economy standards have been designed, as such standards impose a shadow cost on each vehicle equal to a linear function of the vehicle's fuel economy rating (see, e.g., Anderson and Sallee 2016). (In application 3, we challenge this notion and discuss various complications, but here we wish to focus on only the issue of mismeasured test ratings, not other problems with fuel economy regulation.)

In terms of our model, each product j is a type of car. The externality ϕ_j is some factor ζ (e.g., the social cost of carbon times carbon emissions per gallon of gasoline times total miles driven) times true fuel economy. The shadow taxes imposed by CAFE will be a linear transformation of the old ratings, which the regulator is constrained to use in setting policy:

$$\begin{aligned}\phi_j &= \zeta \times \text{New Fuel Economy Test Rating}_j,\\ t_j &= \alpha + \beta \times \text{Old Fuel Economy Test Rating}_j.\end{aligned}$$

B. *Will Cross Effects Simplify?*

Where the noise in measurement is uncorrelated with factors that determine vehicle demand, the weaker version of assumption 3 from section II will hold, and the main theoretical results in proposition 1 and corollary 1 apply. It is logical to suppose that errors in the tax rates from an unbiased policy are uncorrelated with cross-price derivatives, as they are likely to be due to idiosyncratic aberrations from test trials or particular technologies,

[18] See http://www.epa.gov/fueleconomy/documents/420f14015.pdf and http://www.epa.gov/fueleconomy/documents/420b14015.pdf.

like stop-start systems, that are of little concern to consumers (and therefore not correlated with cross-price derivatives).

In this case, the R^2 from a regression of ϕ_j on t_j has a welfare interpretation. It indicates the fraction of the welfare gain over a flat tax (that corrects for the average externality produced by an automobile) achieved by a policy that uses the less accurate, noisy fuel economy estimates (the second best) in place of the accurate ratings (the Pigouvian benchmark). Note, however, that because ϕ_j is proportional to the old fuel economy ratings and t_j is a linear transformation of the new fuel economy estimates, the R^2 of interest is identical to the R^2 from a regression of the new fuel economy rating on the old one.

C. Data

To estimate this R^2, we use the sample of vehicles that the EPA itself used to establish the concordance between the old and the new highway and city test ratings. In determining how to create the new system, the EPA tested a few hundred vehicles meant to represent the car market and compared the results under the new and old regimes. We obtained the data from these tests from the EPA and use them here to assess the change in ratings.[19]

D. Results

We plot these data in figure 3. The old and new ratings are highly correlated, but there is an upward bias in the old ratings (the old miles per gallon ratings were too high on average). In addition, there are noticeable differences in how the test revision affected different models—there is dispersion around the fitted line. The rating change is quantitatively important: the average difference between the old and the new estimated present-discounted fuel costs in this sample is \$1,700. The difference ranges from \$500 to \$4,250, with a standard deviation of nearly \$700.[20] So even if the bias was recognized, it still affected different vehicles to varying degrees.

The OLS regression of the new rating on the old one yields an R^2 above 0.97.[21] This indicates that along the dimension of test rating quality, the efficiency gain from removing noise is quite minor. The vast majority of

[19] These same data are used by Sallee (2014) to characterize the uncertainty faced by consumers about true lifetime fuel costs of vehicles under the old regime.

[20] This assumes a gasoline price of \$2.50 per gallon (roughly the average in 2008) for vehicles driven 12,000 miles per year for 14 years with a 5% discount rate.

[21] The R^2 changes little when modifying the sample. Adding the 13 available hybrid models to the gasoline-powered sample produces an R^2 of 0.98. The R^2 values for the subsamples of cars and trucks are 0.96 and 0.98, respectively.

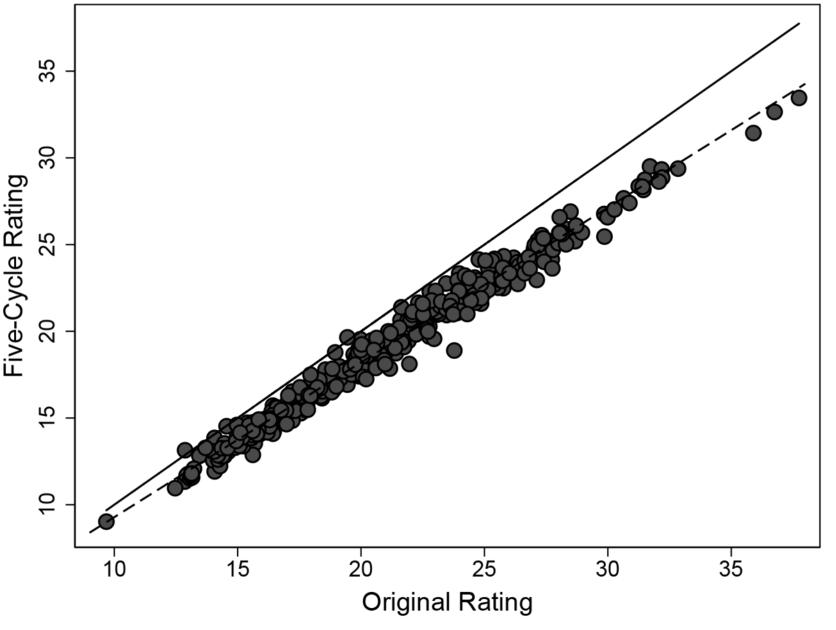

FIG. 3.—Old and new combined fuel economy ratings. Figure shows the pre-2008 (original) combined fuel economy rating and the post-2008 (five-cycle) rating in miles per gallon for a sample of vehicles. Dashed line is linear fit. Solid line is the 45° ray.

the welfare gain from an optimally designed fuel economy policy that used the new ratings can be achieved by a policy that uses the old rating system. Interestingly, this makes the lack of updating relatively innocuous, despite the fairly large differences between the two rating systems. The welfare losses from this noise, however, may be substantial if the policy maker does not take the bias in the old ratings into account and fails to make a correction (i.e., chooses a policy that is based on the assumption that the old rating system is accurate and is therefore too lax on average, causing distortions on the extensive margin). Also, note that the inefficiency from noisy energy efficiency ratings adds to a long list of existing distortions from fuel economy standards, including the welfare loss from ignoring product durability discussed in section V.

IV. Application 2: Real-Time Electricity Pricing

Our second application is to time-varying electricity prices. Electricity consumers typically pay the same price for electricity regardless of when they consume it. In contrast, the marginal cost of producing electricity varies significantly across hours of the day, days of the week, and months of the year because of variance in the marginal source of generation. At low levels of demand, marginal cost is low because only solar, wind, and so-called inexpensive base load power plants are needed. At high levels of demand, higher-cost peaker plants produce the marginal unit. As a result, the marginal cost of electricity is frequently several times higher at one hour of the day as compared with another hour in the same day.

Economists have contemplated the efficiency benefits of time-varying pricing schemes that align price and marginal cost. The theoretical ideal is called real-time pricing, which is a scheme in which the price of electricity charged to the consumer is unconstrained and is adjusted at a high frequency to reflect costs. Real-time pricing provides the right incentive to consumers at every moment and therefore achieves the efficient resource allocation, provided that no other markets failures are present (Borenstein and Holland 2005).[22]

Historically, it was infeasible to measure electricity consumption hour by hour for each end user, so this mispricing was a necessary compromise. However, with the advent and rollout of computerized electricity meters, high-frequency measurement at the customer level is already a reality in most parts of the United States. Even so, real-time electricity prices have met with considerable resistance from utilities and regulators, who fear that consumers will complain about price surges and unpredictable bills.

As a result, while the technology to implement real-time pricing is already in place, pricing reforms have been incremental. Instead of real-time pricing, utilities have experimented with peak pricing for certain times of the day, seasonal rates, or peak prices only on certain days on which demand is forecasted to be very high because of weather. A significant literature in economics has evaluated these programs, primarily with a focus on how demand responds to price variation (Jessoe and Rapson 2014; Andersen et al. 2017; Fowlie et al. 2017; Gillan 2017; Ito, Ida, and Tanaka 2018). A remaining unanswered question in this literature is whether most of the efficiency gains from real-time pricing can be achieved by these intermediate policies. If simpler rate designs can capture most of the efficiency gains of real-time pricing, then this may present a useful way forward for the industry that can accelerate reform.

We demonstrate that our model can be used to answer this policy-relevant question with readily available data and simple OLS regressions. We use wholesale pricing data from a major electricity market in the eastern United States, which provide a measure of the marginal cost of electricity at the hourly level. Using our model, we show that the R^2 from a regression of observed wholesale prices on season, day of week, or peak

[22] Although not the focus of this application, pollution externalities can be introduced to our analysis by adding environmental damages to private marginal costs and running the regressions with social marginal costs as the dependent variable.

demand periods measures the proportion of the welfare gain that an intermediate reform that allows tariffs to vary by those variables would achieve, relative to the welfare gain that would be achieved by moving all the way from a flat rate to real-time pricing. As such, our method allows us to rank a wide range of alternative policies with minimal effort, as stated in corollary 2. This can be quite valuable because the welfare gain achieved by intermediate policies will vary on the basis of the characteristics of supply and demand, for example, demand variability and capacity constraints. This increases the value of being able to quickly calculate potential welfare gains across a number of different markets in terms of both time and geographic scope.

In our application below, we find that the intermediate schemes perform relatively poorly. Fairly complex schemes are required to recover half of the welfare gains from real-time pricing, and schemes that mimic real-world policies used to date recover only a small fraction of the potential gains. These results should prove useful in the ongoing debate about electricity rate design, which is poised to undergo significant reform in the coming years.

The insights from this application will also apply to other settings that feature coarse pricing, where many related goods must be given a common price because of some exogenous constraint on the pricing policy, even though social costs differ as a result of production technologies, scarcity, or externalities. Potential examples include markets for parking, traffic congestion, taxis/ride-sharing services, or event tickets.

A. *The Real-Time Pricing Benchmark versus the Constrained Policy*

To apply our model to electricity, we interpret each product j as electricity consumed at a specific moment. Empirically, we will consider an hour to be a unique moment because this is the granularity of our wholesale pricing data. We focus on a single integrated electricity market, so we do not need to consider electricity consumed at different locations to be different goods.[23]

In our model, consumers pay $p_j + t_j$, where t_j is understood as a tax and p_j is a uniform producer price. Here, we interpret t_j as the tariff that

[23] As discussed further below, we use data from PJM (Pennsylvania, New Jersey, and Maryland), an integrated electricity market that spans multiple states. We treat this market as a single location and use PJM's reported system price. In reality, there are sometimes transmission constraints that imply that delivering electricity to one specific location has a higher cost than delivering to another location at the same time, even within the market. Most US electricity markets therefore have locational marginal prices that are specific to a particular node in the grid. We abstract from this issue, as does the bulk of the literature. In principle, however, our method could be used to gauge the welfare implications of the granularity of prices across geographical space as well as the time dimension focused on here.

applies to good j, so that the final price to consumers is just the tariff t_j (equivalent to assuming $p_j = 0$ in the original notation). Under a flat tariff, t_j is the same across all j goods. Under real-time pricing, the tariff is unique to each j. Intermediary policies will have subsets of j (such as peak demand periods) for which consumers face a common tariff.

Unlike our other applications, there is no externality. Thus, the full social cost of producing a unit of good j is just the marginal cost mc_j, which we allow to vary across time. Any mismatch between the tariff and marginal cost induces an inefficiency, where the wedge is equal to $t_j - mc_j$. This wedge plays exactly the same role in our theory as the wedge due to imperfect correction of an externality (which is denoted $t_j - \phi_j$ in the other applications). Thus, in terms of our model,

$$\phi_j = mc_j,$$
$$t_j = \alpha + \boldsymbol{\beta}' \mathbf{z}_j.$$

where $\mathbf{z}_j$ is a vector that includes tariff policy variables, such as on- versus off-peak or day of week indicators. Note that $\mathbf{z}_j$ can represent any tariff scheme that is linear in parameters, including interactions of indicator variables. The method can thus evaluate highly flexible tariffs.

If our assumptions about demand and supply hold, then the R^2 of a regression of ϕ_j on t_j will indicate the welfare fraction achieved by the constrained pricing scheme (second best) relative to the real-time pricing benchmark, where both welfare gains are calculated relative to an unbiased flat tariff.[24] Even the real-time pricing benchmark is not quite first best since it is granular on hour, ignores transmission constraints, and so on. Thus, as usual, we measure the efficiency gain from more granular pricing along the dimension that policy makers can realistically target: in this example, average hourly tariffs.

B. *Will Cross Effects Simplify?*

In this application, increasing marginal costs are essential. As discussed in section II, our results still apply in this case, but the assumptions should be interpreted in terms of combined responses of demand and supply. As detailed in appendix A, zero cross-price derivatives (the strong version [b′] of assumption 3) for demand and supply is sufficient for all our

[24] It may seem counterintuitive that we can use historical data that come from observed marginal costs, even though those realized marginal costs depend on the particular flat tariff that was in place during the sample. App. A shows that under local linearity, the R^2 of a regression of observed marginal costs (under the flat tariff) on the policy variables equals the R^2 of a regression of the benchmark marginal costs (under real-time pricing) on the policy variables. Hence, the relative efficiency gain can be computed from a regression that directly corresponds to our data.

results from section II to go through. In that case, we can characterize the deadweight loss of using an arbitrary vector of tariffs, denoted $t_j = \tau_j$, as compared with using real-time pricing, as the sum of J Harberger triangles:

$$-2 \times DWL\,(t = \tau) = \sum_{j=1}^{J}(\tau_j - mc_j)^2 \frac{\partial \tilde{x}_j}{\partial t_j},$$

where $\partial \tilde{x}_j / \partial t_j = \partial x_j / \partial t_j - \partial mc_j / \partial t_j$. Minimizing this distortion will involve fitting the tariff schedule so as to minimize the sum of squared errors between the tariff and the observed marginal cost, weighted by the derivative terms. When $\partial \tilde{x}_j / \partial t_j$ is uncorrelated with the wedges or common across all j, then the formula simplifies to its final form and our R^2 result applies.

Does it make sense to assume that cross-price derivatives are zero for supply and demand? On the demand side, the required assumption is that a change in the tariff in hour j does not affect demand in hour $k \neq j$. Substantial empirical support exists for this assumption. A consistent finding in the literature is that such cross-price derivatives are quite small and are often statistically indistinguishable from zero. In other words, the electricity tariff during hour j does not affect the demand for electricity during hour $k \neq j$.

Specifically, a substantial literature has studied experiments that raise the cost of electricity at specific hours of the day, for example, on weekdays during late afternoons in the summer, when system demand peaks because of air conditioner use in homes. A common question has been to what extent consumers will reduce electricity consumption during this high-price window and substitute this for consumption in shoulder hours around the experiment. Such studies consistently find that peak tariff schemes lower consumption during the targeted window but reveal minimal shifting of demand into off-peak hours (Jessoe and Rapson 2014; Fowlie et al. 2017; Gillan 2017; Ito, Ida, and Tanaka 2018). The exception is Andersen et al. (2017), who find that a variable pricing scheme does cause significant shifts of demand into lower-priced windows in Denmark. Therefore, the strong version (b′) of assumption 3 is likely appropriate in this application, at least in many circumstances. Bolstered by this evidence, we proceed by assuming that cross-derivatives are zero, but we also assess the performance of R^2 using estimates from the literature that quantify how large cross effects might be in section IV.E.

On the supply side, the corollary question is whether the price of electricity in time period j affects the cost of production in time period $k \neq j$. It is reasonable and indeed common in the literature to assume that production costs in different hours are separate production processes and are not directly related. Marginal costs are likely to be serially correlated, but this is because demand is serially correlated, not because production in one time causes a shift in cost in other hours.[25] (Recall that we allow the marginal cost of production to be rising at any moment *j*. The assumption is that price in one hour does not affect cost in a different hour.)

C. *Other Modeling Considerations*

Throughout our theory, we maintain the assumption that demand and supply curves are linear over the relevant range of prices. Note that we require only that each good *j* has locally linear supply and demand, not that the demand and supply curves across products *j* have the same slope. Local linearity does not seem like a problematic assumption on the demand side, but electricity supply curves can become convex when a market approaches capacity limits (although we show empirically that a linear supply assumption still fits a large part of the supply curve). When supply is convex, the R^2 statistic will provide an approximation. We investigate its accuracy via simulation, given data on the shape (convexity) of the market-level supply curve. In our case, the approximation appears to be quite good (see sec. IV.E).

D. *Data*

Our empirical application uses data from the PJM wholesale electricity market. While originally comprising the states of Pennsylvania, New Jersey, and Maryland (thus, the name PJM), the PJM market is a regional transmission organization that runs one of the largest wholesale electricity markets in the United States, stretching into 13 states in eastern and central United States plus the District of Columbia. PJM is one of five regional transmission organizations that run an active market for wholesale electricity. As in most wholesale electricity markets, PJM runs an hourly real-time auction for energy, bringing together producers and consumers (typically, utility companies) of electricity. These auctions yield hourly wholesale prices for electricity, which are a good measure of marginal costs. We use hourly pricing data for the year 2012, as this is the year for which we have data for the supply curve, which we use in the convexity simulation below.

E. *Results*

In this section, we report the R^2 for a wide variety of alternative pricing schemes. We also discuss the results from simulations that introduce

[25] Technically, this may not be true for adjacent hours because of startup and ramping costs for fossil-fueled plants. We follow much of the literature in assuming that their impact on key results is modest, although we note that Reguant (2014) and Cullen (2015) are exceptions that model startup costs explicitly.

cross-derivatives and convex supply (details are in app. B). We run variants of the following regression:

$$price_{th} = \alpha + \boldsymbol{\beta}'\boldsymbol{z}_{th} + \varepsilon_{th},$$

where t indexes date, h indexes hour, $price_{th}$ is the observed wholesale electricity price, and $\boldsymbol{z_{th}}$ is a vector that includes potential tariff policy variables, such as on- vs. off-peak indicators θ_p, hour of day fixed effects θ_h, day of week fixed effects θ_d, monthly fixed effects θ_m, season fixed effects θ_s, or their interactions.

Table 1 shows the efficiency gain from using increasingly flexible tariff policies. We find that simple yet commonly used tariff structures, like on- vs. off-peak prices, do not improve efficiency much. Even the highly sophisticated—and potentially hard to understand for consumers—pricing schemes that we analyze (such as tariffs that vary by weekday, hour, and month) capture less than half of the efficiency gain of real-time electricity pricing. This conclusion is similar to Borenstein (2005), who uses a more detailed simulation model of competitive electricity generation.

We test the robustness of this finding by assessing the performance of R^2 when cross-price derivatives in demand are not zero. Andersen et al. (2017) report the largest cross-price derivatives among the studies we discuss above, with substitution to shoulder hours of approximately 29% of the own-price effect. Appendix B compares the R^2 with welfare calculations that account for such substitution. Even with considerable substitution, the largest bias in the R^2 measure is approximately 2 percentage points. We also consider the cross-price effects estimated in Ata, Duran, and Islegen (2016) as well as cases with spillovers across hours, as in Jessoe and Rapson (2014). We find similarly small biases.

Finally, we evaluate how our results change when we take into account that electricity supply is convex at high levels of capacity utilization. To do this, we make use of plant-level engineering data from the same year.

TABLE 1
R^2 FROM ELECTRICITY TARIFF REGRESSIONS

Pricing Regime	R^2
On- versus off-peak fixed effects	.040
Hour of day fixed effects	.135
Hour of day and day of week fixed effects	.153
Hour of day and month of year fixed effects	.193
Hour of day, day of week, and month of year fixed effects	.211
Day of week × month of year fixed effects	.297
Hour of day × day of week × month of year fixed effects	.428
Observations	8,784

NOTE.—Dependent variable is the hourly price of electricity observed in the PJM market for 2012. Peak hours are defined as 2–6 p.m.

These data report both capacity and engineering marginal costs for each plant in the system, forming a step function supply curve.[26] We capture the convexity of supply in two ways. First, we estimate a quadratic supply curve through the aggregate engineering marginal cost curve. Second, we use the actual step function.

We then simulate the welfare gains from each of the seven pricing regimes above and compare these with their respective R^2 measures. We find that R^2 remains a reliable indicator for the efficiency gain of constrained policies even when marginal costs are convex. R^2 is always within 1% of the simulated welfare gains under our quadratic estimate of supply and generally within 10% using the step function supply curve. Furthermore, the bias in the R^2 when using the step function supply curve is relatively constant across pricing regimes, suggesting that comparisons across different pricing regimes can still be made. See appendix B for details.

V. Application 3: Automobiles and Longevity

The total externality caused by an energy-consuming durable good depends on both its energy efficiency and its lifetime utilization. For example, a car's lifetime gasoline consumption depends on fuel economy and miles driven. Were all products utilized the same amount, a set of product taxes based on only energy efficiency could accurately target lifetime externalities, thereby shifting demand across products efficiently. But heterogeneity in the longevity of products with the same energy efficiency rating implies that energy efficiency policy is necessarily imperfect.

We demonstrate the empirical importance of this issue using the case of greenhouse gas emissions from automobiles. We use a novel data set to estimate the average lifetime mileage of different car models, and we translate that into lifetime damages from greenhouse gas emissions, according to each vehicle's fuel economy and the social cost of carbon. We then use simple linear regressions motivated by our theory to evaluate second-best policies that must construct a tax vector for vehicles that depends on a vehicle's fuel economy but not its longevity. As we explain below, this constrained policy closely resembles the dominant real-world policy in this sector, fleet-average fuel economy standards, such as the US CAFE program.

Consistent with our interest in fuel economy standards, we assume that there are no taxes on used vehicles during resale. This is why the ideal new vehicle tax should capture lifetime emissions, even if an initial car buyer sells the vehicle before the end of its life. Capitalization of the new vehicle

[26] Data such as these have been used extensively to calculate hourly equilibria within electricity markets.

tax will ensure that used buyers face the same prices as they would under a direct tax on emissions.[27]

We find that the constrained policies are highly inefficient. There is a voluminous literature that explores the welfare implications of energy efficiency policies, but we know of no prior paper that has demonstrated the importance of heterogeneity in product longevity.[28] We speculate that this issue is a first-order concern not just for automobiles but also for appliances, building codes, and other efficiency programs.

A. *The Pigouvian Benchmark versus the Constrained Policy*

We study the efficiency of a fuel economy standard in correcting greenhouse gas emissions in the purchase of new vehicles. In terms of our model, each type of vehicle is a product j. We model the lifetime greenhouse gas related externality from automobiles as proportional to the total gasoline consumed by each vehicle type j. Total gasoline consumed is the total lifetime mileage of a vehicle multiplied by its miles per gallon efficiency. To translate this into a dollar externality, lifetime emissions are multiplied by a constant, ψ, which is equal to the social cost of carbon per gallon of gasoline. Note that a gasoline tax would achieve the Pigouvian benchmark, so long as consumers are aware of product durability and have a rational forward-looking valuation of fuel costs.

In contrast, a fleet-average fuel economy standard will create a shadow tax scheme where shadow taxes are equal to a linear function of a vehicle's fuel economy rating.[29] Thus, in terms of our model,

[27] Suppose, e.g., that one type of car lasts 5 years and typically has one owner, while another with the same fuel economy typically lasts 10 years and has two owners. Absent taxes on used car transactions, if each owner pays for only the emissions during their initial 5 years of ownership, market prices will fail to reflect the difference in damages between the two vehicles. To create the correct relative prices via only new vehicle taxes, we need to charge a higher tax on the long-lived car. Via tax pass through, this will raise the price of the used long-lived car so as to (imperfectly) correct the externality. Likewise, if the benchmark policy is a gasoline tax, a new car buyer who will sell the vehicle after only 5 years will account for the impact of the gasoline tax on the second owner because the asset value in the used car market will encompass the second user's gasoline taxes.

[28] Allcott and Greenstone (2012) note that differences in utilization might justify geographically differentiated appliance standards, but they do not quantify heterogeneity or calculate potential gains from differentiation.

[29] Historically, policies like CAFE were firm specific, so that the shadow price varied across firms. CAFE and most similar policies in other countries now allow trading, which fits our description here. However, there is suggestion that trading in CAFE has been thin (Leard and McConnell 2017), but this may be because trading was introduced alongside footprint-based standards, which reduces the variation in shadow costs across firms (Ito and Sallee 2018).

$$\phi_j = \psi \times \text{Fuel Economy Rating}_j \times \text{Lifetime Mileage}_j,$$
$$t_j = \alpha + \beta \times \text{Fuel Economy Rating}_j.$$

If our assumptions all hold, then the R^2 of a regression of ϕ_j on t_j will indicate the fraction of the welfare gain achieved by a fuel economy standard (second-best tax), which depends on only fuel economy ratings, as compared with the welfare gain achieved by a policy based on both fuel economy and lifetime usage (the Pigouvian benchmark), where both gains are measured compared with a policy that places a uniform tax on all cars equal to the average externality. In section V.C, we discuss additional complications related to vehicle-related externalities and ways in which fuel economy regulation differs from gasoline taxation that are not captured in this description.

Note that our unit of observation is the vehicle model (e.g., a 1995 Toyota Corolla). Because of accidents and random mechanical failure, individual units will be scrapped with variable lifetime mileage, but this is orthogonal to our welfare comparisons. That is, all of our results are robust to allowing for random product failure, with ϕ_j interpreted as the mean externality—so long as the random failure rates are not endogenous to product taxes (which we return to below). The reason is that ex ante unknowable variation in damages across identical units cannot be targeted by any new vehicle policy, so this will not affect our comparison of second-best policies to the Pigouvian benchmark.

B. *Will Cross Effects Simplify?*

Our main assumptions are about cross-price derivatives across types of automobiles. In this market, cross-price derivatives are clearly important, so the strong version of our assumption 3 (b′) will not hold. Instead, we argue that cross effects will plausibly be small and will cancel out in the R^2 ratio under a second-best tax policy so long as vehicles that have greater or lesser longevity, conditional on fuel economy ratings, are not systematically closer or further substitutes for each other. In that case, assumption 3 (b) will hold. But we relax that assumption empirically by allowing vehicles with more similar longevity and more similar fuel economy to be closer substitutes, following the derivations in section II.D. Introducing these correlations turns out to have limited impact; our results are robust.

C. *Other Modeling Considerations*

Below we find that second-best constrained policies are highly inefficient. We interpret this as evidence that fuel economy regulations are inefficient as compared with a gasoline tax. But this interpretation is generous to fuel economy regulation because it abstracts from other well-known inefficiencies in fuel economy policies. In particular, fuel economy regulations fail to incentivize abatement on the intensive margin; for example, a fuel economy standard can get people to buy the optimal vehicle, but they will not drive the optimal number of miles. Our model abstracts from that by assuming that the externality attached to each vehicle is fixed. Note, however, that we are concerned with lifetime mileage, so the intensity-of-use margin that concerns us is only the scrappage decision, not miles traveled per year.[30] In addition, revenue-neutral energy efficiency policies fail to get the average price of goods right; for example, a fuel economy standard can get the relative price of inefficient versus efficient cars right, but all cars will be too inexpensive and the car market will be too large overall.[31]

Our welfare analysis considers two alternative tax structures: a second-best scheme that imposes a tax on each vehicle that is a linear function of its fuel economy rating, and a Pigouvian benchmark that imposes taxes according to each vehicle's externality. As such, we measure welfare loss along the vehicle purchase margin, which is directly targeted by vehicle-based taxes. This abstracts from the market size effects (by assuming that both tax schedules are correct on average) and the intensive margin effect (which is omitted from both policies and therefore likely has limited impact on the proportional gains we emphasize) and bases the policy comparison on only differences related to tax rate errors driven by heterogeneity in longevity.

Thus, our R^2 results can be interpreted as an upper bound on the fraction of the welfare gain from a gasoline tax that can be achieved by a second-best fuel economy regulation. It is an upper bound because a gasoline tax would also achieve gains along the scrappage (intensity of use) margin and because a gasoline tax would correct the overall size of the car market by raising the average price of automobiles. The Pigouvian vehicle tax does neither. As discussed in section II.E, R^2 is also an upper bound if the social cost of carbon used for calculating the carbon externality is too high or too low; equation (15) quantifies this inefficiency as a function of the distance between the assumed and true social cost of carbon.

In brief, our comparison—within which CAFE performs quite poorly—understates the real welfare losses incurred from using CAFE instead of a tax on gasoline.

Another issue regards uncertainty. Knowing the future longevity of durables at the time of their purchase necessarily requires forecasting into

[30] See Jacobsen and van Benthem (2015) for evidence on how scrappage decisions influence the welfare implications of fuel economy regulations.

[31] See Holland, Hughes, and Knittel (2009) for an exploration of how performance standards create inefficiencies due to their average price effects.

the future. Where some variation in longevity across products is genuinely uncertain to all parties, no policy will be able to capture this heterogeneity in the new market. Here our results on mismeasurement can be used, as the benchmark policy will be one that leaves only random error. A closely related challenge is heterogeneous sorting, which our representative agent model abstracts from. If policy induces sorting, this could alter the expected lifetime emissions of different products not for mechanical factors but because the typical user of each product changes. Where this equilibrium sorting is predictable, it could be accounted for in the calculation of the product-level emissions rates and used to explicitly consider potential bias in R^2.[32]

Finally, recall that our model is for a representative consumer. Individual drivers may have different on-road fuel consumption rates for identical cars due to differences in driving styles and conditions (Langer and McRae 2014). Differences in maintenance or accident risk may imply that some drivers use up a vehicle faster than others, so that expected lifetime mileage for a car depends on driver behavior. We abstract from these considerations both because of data limitations and because we doubt their quantitative significance. As discussed above, we are not concerned with random failure that is unpredictable to the consumers themselves at the time of purchase. Thus, driver heterogeneity is relevant only to the extent that different types of drivers sort into different vehicles systematically in response to changing taxes (see above). Moreover, our model does permit heterogeneity in miles driven per year—all calculations are done in terms of total miles driven from new until scrappage, regardless of calendar age. Heterogeneity in annual usage matters only to the extent that faster or slower rates of utilization affect the total expected lifetime mileage of the vehicle. The fact that most cars have several owners over their life will tend to decouple any individual owner from the vehicle and will mitigate concerns related to individual heterogeneity.

D. Data

Our data come from the California Smog Check program, which records the odometer reading for all tested vehicles. We merge these data with a national registration database that identifies when a vehicle has been

[32] One avenue to explore, if parameter estimates from a sorting model are available, would be an iterative procedure to update the left-hand side externality in the spirit of Greenstone, Sunstein, and Ori (2017). In that paper, the iteration is on the actual policy as new vehicle miles traveled data come in. Here, the researcher could calculate the benchmark tax policy in the unsorted model, adjust on the basis of the sorting model, recalculate the benchmark, and so on. Construction of this sorted benchmark would allow the researcher to run a second regression, this time of the sorted benchmark on the rate of fuel consumption. The R^2 could be reported alongside the R^2 from the original regression, with similar results making it likely that sorting is an economically small factor.

retired from the US fleet and take the last observed odometer reading before a vehicle's retirement as the measure of its lifetime mileage. We aggregate individual observations to the VIN10 level (the finest distinction of a unique car type possible in our data, which delineates a vehicle by make, model, model year, engine size, and often transmission, drive type, body style, and trim) and VIN8 level (which encompasses the same vehicle characteristics as the VIN10 but aggregates across model years). We divide lifetime mileage by official fuel economy ratings to estimate lifetime gallons consumed.[33] We use this as our measure of the lifetime externality of each vehicle type (i.e., ϕ_j), multiplying by the social cost of carbon per gallon of gasoline when necessary to convert the externality into dollars.[34]

We do not observe all units, which creates the possibility of measurement error and censorship bias. Regarding the former, we are concerned with the average mileage at scrappage of cars, but we observe only a sample. To the extent that this error may be large, the R^2 value can be adjusted following the discussion in section II.E. However, we demonstrate that for our sample the bias is likely to be very small (see sec. V.E and app. C). Regarding censorship bias, we do not observe cars under 6 years old (as they are usually not required to be tested), cars that were retired before our data began, or cars that were still in the fleet when our data ended. These missing observations could bias our results in either direction, but we demonstrate through several procedures that this censorship bias is apparently small and likely causes us to slightly overstate the efficiency of fuel economy policies. We use comprehensive national registration data from R. L. Polk to address concerns related to missing data. More data details and robustness checks are included in appendix C.

E. Results

In this section, we report the R^2 from several alternative specifications. In appendix C, we explore further robustness checks. In section V.F, 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.

[33] This abstracts from 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 not only for simplicity but also because many climate models and the current federal guidelines suggest that the time path of the social cost of carbon rises at roughly the rate of interest. This means that social cost growth offsets discounting.

[34] We abstract from carbon emissions related to construction and scrappage of vehicles because standard estimates suggest that these emissions make up only 8% of life cycle emissions (National Research Council 2010). The remainder is due to gasoline consumption. Moreover, to the extent that these emissions are the same across models, incorporating them would have no effect on our welfare calculations. Only heterogeneous life cycle emissions matter.

We begin by showing the data for our preferred sample in a scatter plot. Figure 4 shows a scatter plot of the relationship between a vehicle's total lifetime externality (gallons of gasoline) and its official fuel consumption rating for both cars and trucks, along with the OLS fitted line. A point in the figure corresponds to the average lifetime gasoline consumption at the VIN10-prefix level. It ignores within-VIN10 variation in gasoline consumption. The sample in the figure is restricted to model years 1988–92, the years for which censoring is least problematic (more on this below), and to vehicle models for which we have at least 200 observed retirements.

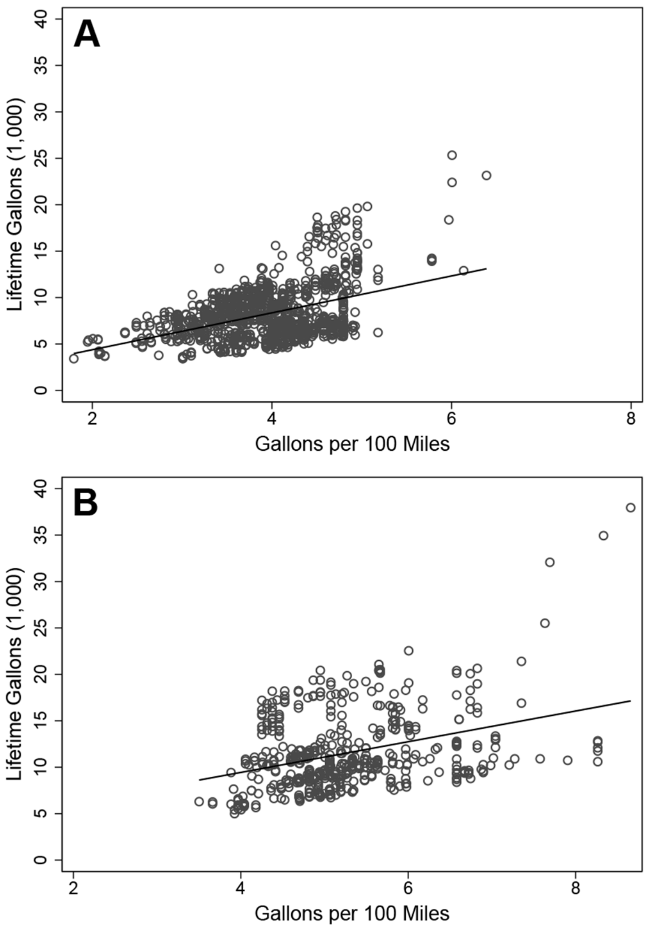

FIG. 4.—Relationship between lifetime gasoline consumption and fuel efficiency. The unit of observation is a type of vehicle (a VIN10-prefix). Gallons consumed is the average across observations for that type. The sample is restricted to models for which we observe at least 200 vehicle retirements and to model years 1988–92. Observations with vehicle miles traveled above 1 million miles are dropped. Solid lines are OLS prediction lines.

We drop observations with more than 1 million miles to limit the influence of outliers.

There is, as expected, a positive correlation between fuel consumption ratings (the inverse of fuel economy ratings) and lifetime gasoline consumption. But there is also a great deal of dispersion. Vehicles have substantially 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 in this sample is only 0.18 and 0.12, respectively. (The R^2 from a combined sample regression is 0.29.) According to our theory, this implies that the second-best linear policy captures only 18% and 12% 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.

Note that the second-best policy will undertax long-lived vehicles (observations above the regression line), and it will overtax short-lived vehicles (observations below the regression line). Some may find it counterintuitive that the efficient policy would raise taxes on long-lived vehicles. To see the intuition, consider two vehicles with the same fuel economy rating, where one lasts twice as long as the other. To drive the same number of miles (same emissions), two short-lived vehicles will be required, so that the tax will be paid twice. Thus, harmony between the tax paid and the emissions emitted requires taxing the long-lived vehicle more heavily.

Table 2 reports the R^2 from a set of regressions that take the form

$$\begin{aligned}&\text{Average Lifetime Gasoline Consumption}_j\\&\quad = \alpha + \beta\text{Gallons per Mile}_j + \varepsilon_j,\end{aligned} \tag{17}$$

where j indexes a vehicle type (VIN10-prefix or VIN8-prefix). We report a range of estimates in order to assess the importance of sample restrictions,

TABLE 2
Regression R^2

	VIN10-prefix		VIN8-prefix	
	OLS	WLS	OLS	WLS
A. All model years:				
All models	.26	.20	.23	.19
Models with $N \geq 200$	.22	.17	.27	.19
B. Model years 1988–92:				
All models	.27	.26	.28	.27
Models with $N \geq 200$	.29	.22	.34	.25

Note.—This table shows R^2 from regressions using VIN-prefix average lifetime gallons consumed on fuel consumption rating. The unit of observation is either a VIN10-prefix or a VIN8-prefix in panel A. Observations with vehicle miles traveled above 1 million miles are dropped. N is the number of observed retirements. WLS weights the regressions by N.

weighting, censoring, the level of aggregation, and sampling error. WLS results weight VIN-prefixes by the number of observed retirements N. These results are useful for assessing the effect of sampling variation, but they also approximate weighting by sales share, which leads to the preferred welfare interpretation because it aggregates to total externalities generated. In all cases, we drop observations with reported mileage above 1 million (1,525 observations out of roughly 4 million, or less than 0.05%). The unit of observation is average gasoline consumption across vehicles with the same VIN10-prefix or VIN8-prefix, consistent with figure 4.

Table 2 shows that our estimate of the R^2 remains small in all VIN10-prefix specifications, ranging from a low of 0.17 to a high of 0.29. R^2 is slightly higher when the data are collapsed at the VIN8-prefix level (0.19–0.34). Importantly, our estimates change very little when we restrict the sample to include only 1988–92 model years, which are the years in our data with the least censorship concerns and therefore our preferred specification. As these model years span the age range in which the majority of retirement happens, this provides us with a first indication that our welfare conclusions will be broadly robust to additional measures that account for censoring in the data.

As discussed above, white noise in the measurement of lifetime mileage by type (sampling error) will cause the estimated R^2 to be below the true welfare gain ratio. To assess the importance of sampling error, we compare results from OLS with WLS, which weights models by the number of vehicles scrapped. We also check how our results change when we limit the sample to vehicles for which we observe relatively many retirements ($N \geq 200$). The R^2 changes only modestly when moving between OLS and WLS and when restricting the sample to $N \geq 200$. This suggests that our qualitative findings are not overly sensitive to sampling considerations. We explore this issue further in appendix C.

Our theoretical results are focused on second-best policies—tax schedules that are set optimally against some design constraint—but actual policies may deviate from the second best. In appendix C we comment further on biased policies, which can come either because the average tax is wrong (mean bias) or because the slope is wrong (slope bias).

Summary of additional estimates.—Our approach also applies to more flexible fuel economy policies; the R^2 from the appropriate regression will have the same welfare interpretation for any policy that is linear in parameters. For example, fuel economy policy could put a shadow price on each model that was a quadratic function of fuel consumption ratings. Or tax rates could be based on fuel economy bins. Then, the R^2 from a regression of the externality on fuel consumption and fuel consumption squared, or of discrete bin dummies, would have the desired interpretation. More flexible fuel economy policy could also base shadow prices on not just fuel economy but also other attributes, like class (car versus truck), model year, or body style. The R^2 from regressions of lifetime externalities on fuel consumption and these additional attributes directly indicates the welfare gains possible from more flexible policies.

We ran a number of such regressions and summarize the results here. First, we allow for separate fleetwide average standards for cars and light-duty trucks (similar in spirit to the initial structure of US CAFE standards). In our framework, this corresponds to adding a truck indicator and its interaction with the fuel consumption rating. This adds very little explanatory power. In our preferred specifications using model years 1998–92, R^2 rises by 0.004–0.013 to a range of 0.23–0.35, depending on the specification in table 2). This strongly rejects the notion that separate regulation of cars and trucks was useful in addressing the inefficiency that we identify. Adding other policy attributes, such as body style and model year, has a similarly small impact on the R^2; none of these attributes are strongly correlated with durability (conditional on fuel economy).

We also considered attribute-based policies based on fuel consumption and vehicle size, either using the vehicle's footprint (wheelbase $\times$ track width, where wheelbase is the distance between the front and rear axles of a vehicle) or by more flexibly including wheelbase and width as separate regressors. In both specifications, we also allowed the policy to be different for cars versus light-duty trucks. This mimics current US CAFE policy, which is based on fuel consumption and footprint, with separate standards for cars versus trucks. However, including footprint in our regressions has little effect. It raises R^2 by 0.004–0.049 to 0.25–0.35. Including wheelbase and width has a somewhat larger effect: R^2 increases by 0.077–0.096 to 0.31–0.43. This suggests that there could be efficiency gains from more flexible sized-based standards because of the correlation between size and longevity, although such standards create distortionary incentives (Ito and Sallee 2018).

In all cases, the qualitative conclusion remains that there is substantial variation in lifetime consumption that is not explained by fuel economy, vehicle type, or size, which implies that policies based on only such vehicle attributes—but not on average product durability—will raise welfare by significantly less than would an efficient policy (such as a carbon tax or a gasoline tax).

F. Estimates of Deadweight Loss

We can translate the relative gains from the Pigouvian and second-best product-based taxes, expressed above as an R^2, into deadweight loss by assigning a dollar value to the externality and considering the pattern of substitution across vehicles. We begin with the 1990 model year (typical of the years in table 2), computing the possible welfare gains from a Pigouvian product-level tax and the deadweight loss from the second-best tax based on fuel economy. We then explore the influence of a range of substitution patterns across vehicles, following the theory in section II.D that allows correlation between cross-price derivatives and either the tax error or the efficiency rating. We show that when calibrating to estimates of this correlation from the literature, the R^2 remains very close to the true fraction of welfare recovered.

To evaluate the level of deadweight loss—following the formula in equation (6)—we first assign a value of \$40 for the social cost of carbon (Interagency Working Group on Social Cost of Carbon 2013), leading to an external cost of 35.5 cents per gallon.[35] Using our data on lifetime fuel use, we find that this implies an average of \$3,334 in external costs for each vehicle sold. We further impose an own-price elasticity of -5 (roughly comparable to the estimates in 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 correlated with attributes.

As above, we compute welfare results relative to a baseline that controls for substitution to an outside good (since a revenue-neutral fuel economy standard does not directly incentivize switching to an outside good) and so isolate the welfare effects coming from switching among vehicles. Under these assumptions on elasticities, the welfare gain from a Pigouvian tax on each of 356 vehicle models amounts to \$246 per car sold, or about \$3.5 billion, for model year 1990. The best linear tax on fuel use per mile, equivalent here to the optimal average fuel economy standard, generates about \$0.8 billion in surplus and so leaves \$2.7 billion in deadweight loss. This corresponds directly to the intuition on R^2 above: for the 1990 model year, the weighted R^2 is 0.24, implying that 24% of possible gains can be recovered with a single linear policy.

Table 3 presents the central case described above followed by three panels exploring sensitivity to own- and cross-price elasticities. Panel A considers changes in the own-price elasticity of demand for individual vehicle models (-5 in the central case). More elastic demand allows a larger change in the composition of the fleet, and so greater welfare gains are possible in the Pigouvian benchmark. As expected, the ratio of welfare gains in the second best versus the Pigouvian benchmark remains fixed at 0.24, the value of R^2.

Panel B turns to relaxing assumption 3, investigating how different correlations between cross-price elasticities and the residuals in the policy regressions influence the share of welfare recovered by second-best policy. We first consider the expected direction of bias building on the theory

[35] If the cost associated with carbon emissions has been rising approximately at the discount rate, we interpret this value as being in 2013 dollars (looking retrospectively at the 1988–92 vintages).

TABLE 3
Welfare Effects for Model Year 1990

	Second Best	Pigouvian Benchmark	Ratio
Central case	817	3,472	.24
A. Own-price elasticity:			
−3	501	2,128	.24
−7	1,126	4,782	.24
B. Cross-price elasticities correlated with:			
Durability	811	2,971	.27
Efficiency rating	638	3,385	.19
Price	739	3,523	.21
C. Brand and class loyalty:			
Calibrated to Bento et al. (2009)	783	3,431	.23
Doubling relative to Bento et al. (2009)	756	3,396	.22

Note.—Welfare gains are expressed in millions of 2013 dollars relative to a constant tax at the average externality. For panel B, each standard deviation reduction in attribute distance increases the cross-price elasticity by a factor of 2.

in section II.D. Specifically, we demonstrate that R^2 is biased upward or downward in this case as predicted by proposition 3. The first row in panel B makes vehicles with similar durability better substitutes, leaving all other attributes uncorrelated. As expected, this reduces the effectiveness of the Pigouvian benchmark policy, and the true fraction of welfare gained increases relative to R^2. Here, when elasticities fall in half for each standard deviation difference in durability, we see that the fraction of welfare recovered increases to 0.27. The second row examines correlation between cross effects and fuel economy rating, now making cars with similar miles per gallon better substitutes. This reduces welfare gains in both the Pigouvian benchmark and the second best; the fraction of welfare recovered falls to 0.19. Finally, the third experiment makes vehicles of similar price the best substitutes. This introduces both types of correlation together, since both miles per gallon and durability are related to price. The effects on the fraction of welfare recovered partially offset, with the fraction recovered returning toward R^2. This further supports the use of the R^2 measure.

Finally, in panel C, we calibrate cross-price derivatives using estimates of brand and class loyalty from the literature. This introduces a whole range of correlations together, with class loyalty looking most like correlation with miles per gallon and brand loyalty tending to create correlation with durability. The first line of this panel shows the net effects in our calculation when calibrating to the brand and class loyalty estimates from the demand system of Bento et al. (2009). The various effects offset almost completely, with the fraction of welfare recovered falling slightly to 0.23. The final row doubles the strength of the effects of Bento et al. (2009), doubling the fraction of buyers who substitute within brand and class, and again the effects are very close to offsetting. Across a wide range of substitution patterns, R^2 remains a robust predictor for the fraction of welfare that can be recovered in the second best.

VI. Application 4: Spatial Variation in Emissions

Externalities from pollution are typically a function of the amount of pollution emitted as well as location-specific conditions, including preexisting pollution levels, weather, and the proximity of vulnerable populations. This has long been understood as a rationale for location-specific environmental policies (Tietenberg 1980; Mendelsohn 1986; Baumol and Oates 1988). But many policies are constrained to be uniform across space because of practical or political considerations. Mendelsohn (1986) studies this issue in the context of local air pollution and provides a calibrated example showing the welfare gains from drawing two or three distinct zones that have differentiated pollution policies. Holland et al. (2016) document heterogeneity in the environmental benefit of switching from a gasoline vehicle to an electric vehicle in each US county and describe the welfare benefits of policies that vary by county or state versus a national policy.

When our assumptions hold, our model can analyze second-best environmental policies that are constrained to be uniform across space via regression statistics. As an example, we consider a tax placed on the purchase of new refrigerators that aims to correct for greenhouse gas emissions. Refrigerators vary in their energy consumption, and the externality of a given refrigerator depends on its location of use, because a unit of electricity corresponds to different amounts of greenhouse gas emissions in different locations according to what type of power is used (e.g., coal vs. renewables). We consider a constrained policy under which the tax scheme on appliances depends on only the appliance's energy consumption, not on its location.

We show that our assumptions from section II will not hold in this case and that the R^2 will be significantly biased. Instead, we offer an alternative derivation that demonstrates that the within-R^2 from a fixed effects regression has a welfare interpretation. The main purpose of this exercise is to demonstrate the promise of adapting our core framework to find alternative relationships between familiar regression statistics and the welfare properties of constrained externality-correcting policies.

Appendix D presents this application in detail, but the intuition for the within-R^2 result is as follows. Some regions have clean electricity; others are dirty. Within each region, a second-best national tax on electricity usage creates two types of mispricing. First, the tax will be systematically too low in dirty regions and too high in clean regions. Second, the single policy slope will get the relative prices wrong within each region. As long as total demand for fridges in each region is fixed by approximation, the average mispricing creates no distortion in choices, and only the second

TABLE 4
SUFFICIENT STATISTICS FOR SPATIALLY DIFFERENTIATED REFRIGERATORS

	Three Interconnections (1)	Eight NERC Regions (2)
Within-R^2, fixed effects	.96	.90
R^2, OLS	.47	.24
Sample size	4,047	10,792

NOTES.—Results are from regressing emissions on electricity consumption, with the unit of observation a refrigerator model in a particular interconnection (col. 1) or North American Electric Reliability Corporation (NERC) region (col. 2).

inefficiency—the slope errors—matter. The within-R^2 from a regression of product region–specific externalities on electricity usage and region fixed effects neutralizes the average mispricing by region. The assumption of no net substitution to the outside good will be plausible in some conditions but not others. If the latter, the within-R^2 provides an upper bound on the welfare gain.

Our empirical application uses data on the energy efficiency rates of 1,349 refrigerator models from the US Association of Home Appliance Manufacturers and carbon emissions rates per unit of electricity for each of the three major power market interconnection regions or the eight distinct regions defined by the North American Electric Reliability Corporation (Graff Zivin, Kotchen, and Mansur 2014; Holland et al. 2016). See appendix D for details.

Table 4 shows the R^2 and within-R^2 values from a regression of the carbon emissions associated with each fridge in each region on its electricity consumption rate. We find within-R^2 values of 0.96 and 0.90 (depending on geographic disaggregation), representing the fraction of the welfare gain under the Pigouvian benchmark achieved by a national policy that does no spatial differentiation. For reference, the R^2 values from OLS, which are far smaller, are also included.

These findings demonstrate a perhaps surprisingly small welfare loss from the lack of regional policy differentiation for electric appliances. This result relies on the extensive margin for refrigerator demand being zero or small. When the extensive margin response grows (and it will be larger for appliances other than refrigerators), there will be a second welfare loss due to the fact that the overall product market will be too large or small in each region.

VII. Conclusion

Externality-correcting policies rarely take on the ideal form of a direct tax on marginal damages. Actual policies are frequently constrained by administrative feasibility, technological cost, or political constraints so that they must place imperfect marginal incentives on products or actions. We demonstrate that under certain conditions, simple regression statistics have welfare interpretations that describe the efficiency costs of these constraints.

We demonstrate the usefulness of this approach through four examples. Three of our applications pertain to environmental externalities, but they span a number of distinct challenges to policy, including random mismeasurement of product attributes, spatial heterogeneity, and the implications of heterogeneity in the lifetime utilization of energy-consuming durable goods. Our other application—which is about wedges between price and marginal cost due to coarse pricing rather than externalities—suggests the potentially wider reach of our approach. These applications demonstrate the viability of our theoretical framework, but they also make contributions in their own right.

Most importantly, our study of the heterogeneity in automobile longevity points out a previously undiscussed efficiency problem with a class of energy efficiency policies that regulate new durable goods. When different products have different average lifetime utilization, energy efficiency policy—which creates explicit or implicit price incentives according to only energy efficiency ratings—is inherently imprecise. Through analysis of unique microdata on automobile mileage, we demonstrate that different types of automobiles have widely varying average lifetime mileage, which implies large inefficiencies in fuel economy policy.

We suspect that there are many additional applications that could benefit from this approach. In the introduction, we mention other possible applications in energy, environment, health, and transportation, but the possibilities extend to any setting where data are available on the distribution of an externality (or other wedge) and its correlation with the variables upon which policy is contingent. Some of our results may be relevant to settings where there is heterogeneity in the deadweight loss of taxation even in the absence of externalities. For example, labor supply elasticities differ along dimensions such as age—the young supply labor more inelastically than the old (Kleven and Schultz 2014).[36] It is generally politically infeasible to condition income or payroll taxes on age. Our findings suggest that these restrictions, while perhaps desirable on the basis of other grounds, increase the overall deadweight loss of labor taxation and provide a method to quantify the efficiency loss. In applying our model to other settings, we emphasize that it is important to consider the demand assumptions but also note that it is straightforward to conduct robustness checks that indicate the degree of error created when the assumptions do not hold.

[36] Best and Kleven (2013) show theoretically that the presence of behavioral career effects provides another reason why the contemporaneous earnings elasticity of the young is lower than of the old.

References

Allcott, Hunt, and Michael Greenstone. 2012. "Is There an Energy Efficiency Gap?" *J. Econ. Perspectives* 26 (1): 3–28.

Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky. 2014. "Energy Policy with Externalities and Internalities." *J. Public Econ.* 112:72–88.

Andersen, Laura M., Lars G. Hansen, Carsten L. Jensen, and Frank A. Wolak. 2017. "Using Real-Time Pricing and Information Provision to Shift Intra-Day Electricity Consumption: Evidence from Denmark." Manuscript, Stanford Univ. http://web.stanford.edu/group/fwolak/cgi-bin/sites/default/files/into_versus_away_paper.pdf.

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

Anderson, Soren T., and James M. Sallee. 2016. "Designing Policies to Make Cars Greener: A Review of the Literature." *Ann. Rev. Resource Econ.* 8:157–80.

Ata, Baris, A. Serasu Duran, and Ozge Islegen. 2016. "An Analysis of Time-Based Pricing in Electricity Supply Chains." Manuscript, Univ. Chicago. https://ssrn.com/abstract=2826055.

Baily, Martin N. 1978. "Some Aspects of Optimal Unemployment Insurance." *J. Public Econ.* 10 (3): 379–402.

Baumol, William J., and Wallace E. Oates. 1988. *Theory Environmental Policy*. 2nd ed. Cambridge: Cambridge Univ. Press.

Bento, Antonio M., Lawrence H. Goulder, Mark R. Jacobsen, and Roger H. von Haefen. 2009. "Distributional and Efficiency Impacts of Increased US Gasoline Taxes." *A.E.R.* 99 (3): 667–99.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 64 (4): 841–90.

Best, Michael C., and Henrik J. Kleven. 2013. "Optimal Income Taxation with Career Effects of Work Effort." Manuscript, London School Econ. http://www.henrikkleven.com/uploads/3/7/3/1/37310663/best-kleven_careers_feb2013.pdf.

Borenstein, Severin. 2005. "The Long-Run Efficiency of Real-Time Electricity Pricing." *Energy J.* 26 (3): 93–116.

Borenstein, Severin, and Stephen P. Holland. 2005. "On the Efficiency of Competitive Electricity Markets with Time-Invariant Retail Prices." *RAND J. Econ.* 36 (3): 469–93.

Callaway, Duncan S., Meredith Fowlie, and Gavin McCormick. 2018. "Location, Location, Location: The Variable Value of Renewable Energy and Demand-Side Efficiency Resources." *J. Assoc. Environmental and Resource Econ.* 5 (1): 39–75.

Caltrans. 2016. "Performance Measurement System." http://pems.dot.ca.gov/.

Chetty, Raj. 2006. "A General Formula for the Optimal Level of Social Insurance." *J. Public Econ.* 90 (10–11): 1879–901.

———. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge between Structural and Reduced-Form Methods." *Ann. Rev. Econ.* 1 (1): 451–88.

Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and Taxation: Theory and Evidence." *A.E.R.* 99 (4): 1145–77.

Cullen, Joseph A. 2013. "Measuring the Environmental Benefits of Wind Generated Electricity." *American Econ. J.: Econ. Policy* 5 (4): 107–33.

———. 2015. "Dynamic Response to Environmental Regulation in Electricity Markets." Manuscript, Washington Univ. http://www.josephcullen.com/resources/dynamicresponse.pdf.

Einav, Liran, Amy Finkelstein, and Mark R. Cullen. 2010. “Estimating Welfare in Insurance Markets Using Variation in Prices.” *Q.J.E.* 125 (3): 877–921.

Farrow, R. Scott, Martin T. Schultz, Pinar Celikkol, and George L. van Houten. 2005. “Pollution Trading in Water Quality Limited Areas: Use of Benefits Assessment and Cost-Effective Trading Ratios.” *Land Econ.* 81 (2): 191–205.

Feldstein, Martin. 1999. “Tax Avoidance and the Deadweight Loss of the Income Tax.” *Rev. Econ. and Statis.* 81 (4): 674–80.

Feng, Ye, Don Fullerton, and Li Gan. 2013. “Vehicle Choices, Miles Driven, and Pollution Policies.” *J. Regulatory Econ.* 44 (1): 4–29.

Fowlie, Meredith, and Nicholas Z. Muller. 2019. “Market-Based Emission Regulation When Damages Vary across Sources: What Are the Gains from Differentiation?” *J. Assoc. Environmental and Resource Econ.* 6 (3):593–632.

Fowlie, Meredith, Catherine Wolfram, C. Anna Spurlock, Annika Todd, Patrick Baylis, and Peter Cappers. 2017. “Default Effects and Follow-On Behavior: Evidence from an Electricity Pricing Program.” Manuscript, Univ. California Berkeley. https://www.meredithfowlie.com/s/Default-Effects-SMUD.pdf.

Fullerton, Don, and Sarah E. West. 2002. “Can Taxes on Cars and Gasoline Mimic an Unavailable Tax on Emissions?” *J. Environmental Econ. and Management* 42 (1): 135–57.

———. 2010. “Tax and Subsidy Combinations for the Control of Car Pollution.” *B. E. J. Econ. Analysis and Policy* 10 (1): 1–33.

Gillan, James. 2017. “Dynamic Pricing, Attention, and Automation: Evidence from a Field Experiment in Electricity Consumption.” Manuscript, Univ. California Berkeley. http://jamesgillan.info/s/JMP_Gillan_most_recent_draft.pdf.

Gillitzer, Christian, Henrik J. Kleven, and Joel Slemrod. 2017. “A Characteristics Approach to Optimal Taxation: Line Drawing and Tax-Driven Product Innovation.” *Scandinavian J. Econ.* 119 (2): 240–67.

Goulder, Lawrence H., and Roberton C. Williams. 2003. “The Substantial Bias from Ignoring General Equilibrium Effects in Estimating Excess Burden, and a Practical Solution.” *J.P.E.* 111 (4): 898–927.

Graff Zivin, Joshua S., Matthew Kotchen, and Erin T. Mansur. 2014. “Spatial and Temporal Heterogeneity of Marginal Emissions: Implications for Electric Cars and Other Electricity-Shifting Policies.” *J. Econ. Behavior and Org.* 107 (A): 248–68.

Greenstone, Michael, Cass R. Sunstein, and Sam Ori. 2017. “The Next Generation of Transportation Policy.” Manuscript, Univ. Chicago. https://ssrn.com/abstract=2943551.

Harberger, Arnold C. 1964. “The Measurement of Waste.” *A.E.R.* 54 (3): 58–76.

Hendren, Nathaniel. 2016. “The Policy Elasticity.” *Tax Policy and Econ.* 30 (1): 51–89.

Holland, Stephen P., Jonathan E. Hughes, and Christopher R. Knittel. 2009. “Greenhouse Gas Reductions under Low Carbon Fuel Standards?” *American Econ. J.: Econ. Policy* 1 (1): 106–46.

Holland, Stephen P., Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates. 2016. “Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors.” *A.E.R.* 106 (12): 3700–3729.

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-an.

Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2018. “Moral Suasion and Economic Incentives: Field Experimental Evidence from Energy Demand.” *American Econ. J.: Econ. Policy* 10 (1): 1–29.

Ito, Koichiro, and James M. Sallee. 2018. "The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel-Economy Standards." *Rev. Econ. and Statis.* 100 (2): 319–36.

Jacobsen, Mark R., and Arthur A. van Benthem. 2015. "Vehicle Scrappage and Gasoline Policy." *A.E.R.* 105 (3): 1312–38.

Jessoe, Katrina, and David Rapson. 2014. "Knowledge is (Less) Power: Experimental Evidence from Residential Electricity Use." *A.E.R.* 104 (4): 1417–38.

Kleven, Henrik J., and Claus Thustrup Kreiner. 2006. "The Marginal Cost of Public Funds: Hours of Work versus Labor Force Participation." *J. Public Econ.* 90 (10–11): 1955–73.

Kleven, Henrik J., and Esben A. Schultz. 2014. "Estimating Taxable Income Responses Using Danish Tax Reforms." *American Econ. J.: Econ. Policy* 6 (4): 271–301.

Knittel, Christopher R., and Ryan Sandler. 2013. "The Welfare Impact of Indirect Pigouvian Taxation: Evidence from Transportation." Working Paper no. 18849, NBER, Cambridge, MA. http://www.nber.org/papers/w18849.

Langer, Ashley, and Shaun McRae. 2014. "Step on It: Evidence on the Variation in On-Road Fuel Economy." Manuscript, Univ. Arizona. http://www.sdmcrae.com/wp-content/uploads/var_fuel_econ_langer_mcrae.pdf.

Leard, Benjamin, and Virginia McConnell. 2017. "New Markets for Pollution and Energy Efficiency: Credit Trading under Automobile Greenhouse Gas and Fuel Economy Standards." *Rev. Environmental Econ. and Policy* 11 (2): 207–26.

Majeske, Karl D., Terri Lynch-Caris, and Janet Brelin-Fornari. 2010. "Quantifying R^2 Bias in the Presence of Measurement Error." *J. Appl. Statis.* 37 (4): 667–77.

Mendelsohn, Robert. 1986. "Regulating Heterogeneous Emissions." *J. Environmental Econ. and Management* 13 (4): 301–12.

Montero, Juan-Pablo. 1999. "Voluntary Compliance with Market-Based Environmental Policy: Evidence from the U.S. Acid Rain Program." *J.P.E.* 107 (5): 1145–77.

Muller, Nicholas Z., and Robert Mendelsohn. 2009. "Efficient Pollution Regulation: Getting the Prices Right." *A.E.R.* 99 (5): 1714–39.

Muller, Nicholas Z., Robert Mendelsohn, and William Nordhaus. 2011. "Environmental Accounting for Pollution in the United States Economy." *A.E.R.* 101 (5): 1649–75.

National Research Council. 2010. *Hidden Costs of Energy: Unpriced Consequences of Energy Production and Use.* Washington, DC: Nat. Academy Press.

Parry, Ian, Margaret Walls, and Winston Harrington. 2007. "Automobile Externalities and Policies." *J. Econ. Literature* 45 (2): 373–99.

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

Reguant, Mar. 2014. "Complementary Bidding Mechanisms and Startup Costs in Electricity Markets." *Rev. Econ. Studies* 4 (81): 1708–42.

Saez, Emmanuel, Joel B. Slemrod, and Seth H. Giertz. 2012. "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review." *J. Econ. Literature* 50 (1): 3–50.

Sallee, James M. 2014. "Rational Inattention and Energy Efficiency." *J. Law and Econ.* 57 (3): 781–820.

Tietenberg, Tom H. 1980. "Transferable Discharge Permits and the Control of Stationary Source Air Pollution: A Survey and Synthesis." *Land Econ.* 56 (4): 391–416.

Urban, Lorien E., Megan A. McCrory, Gerard E. Dallal, et al. 2011. "Accuracy of Stated Energy Contents of Restaurant Foods." *J. American Medical Assoc.* 306 (3): 287–93.