Pigou Creates Losers: On the Implausibility of Achieving Pareto Improvements from Pigouvian Taxation

James M. Sallee*
University of California, Berkeley and the NBER

December 6, 2018

Abstract

Economic theory shows that efficiency-enhancing policies, like a Pigouvian tax, generate more than enough surplus to compensate those who initially lose through transfers or tax reforms. Compensating losers, however, will frequently be impossible in the presence of substantial preference heterogeneity, as it may be impossible to adequately target transfers or reforms. This paper asks under what conditions it is possible to compensate losers from externality-correcting taxes, argues that compensating losers should be understood as a prediction problem, and demonstrates that a large fraction of US households would be net losers from corrective policies, even under generous assumptions. The main empirical application is to a gasoline tax to correct carbon emissions, though related results are presented for other sin taxes on home energy consumption, alcohol and tobacco. Results indicate that it is infeasible to create a Pareto improvement from the taxation of these goods, and moreover that plausible policies are likely to leave a large fraction of households as net losers. It is argued that the unavoidable creation of losers may explain political resistance to the adoption of efficient policies. The broader point of this exercise is to point out inherent limits in the ability of a planner, even one unconstrained by politics, to control the distributional impacts of a corrective policy.

Keywords: Corrective taxation, externalities, equity
JEL: H23, Q58, L51

*The author thanks Karl Dunkle-Warner and Catherine Wright for excellent research assistance, and Severin Borenstein, Lucas Davis, Don Fullerton, Shanjun Li and seminar participants at ASSA, Berkeley, and HERE Conference for helpful comments. The Sloan Foundation and the Hellman Fund at the University of California, Berkeley provided generous support. © 2017 James M. Sallee. All rights reserved.

1 Introduction

Why do efficient policies so often fail to gain political traction? Efficient policies, like a carbon tax, are frequently eschewed in favor of higher cost regulatory alternatives or even inaction. There are numerous possible explanations, but a common argument regards the distribution of welfare burdens imposed by the policy. There are two varieties of this argument. In one, a policy is disliked because it is regressive and disproportionately affects low-income households, as is the case for a gas or carbon tax. In the other case, a policy imposes a substantial burden on a particular firm, set of firms, or group of voters who mobilize to block the policy.

In either case, economic theory provides a potential reply, which is that such losers can be compensated. Any efficiency-enhancing policy, by definition, creates enough new surplus to compensate all losers. That is, any Kaldor-Hicks efficiency gain can be made into a Pareto improvement, if the right transfers are made in the background. A regressive tax could be combined with tax reform so as to preserve the desired income distribution, or any firms facing lost profits can be made whole.

The task of designing and implementing the right background transfers or reforms, however, may often be impossible when the planner lacks requisite information about the distribution of policy burdens. This paper asks whether or not it is possible to compensate the losers from Pigouvian taxes on externality-causing goods among US households, with a focus on motor fuels. I conclude that Pigouvian taxes cannot create Pareto improvements, and instead that a substantial fraction of households will be made net losers from externality-correcting policies. In brief, Pigouvian taxes, in fact, will inevitably create losers.

The intuition behind this is straightforward. The initial incidence of corrective taxes will be widely dispersed because consumption and production of goods that create externalities is quite diffuse. Some individual households or firms will be greatly impacted by a tax. A Pareto improvement requires that some transfer be made to those affected individuals. Given the wide diffusion in externality creation, there are no exogenous characteristics that are sufficiently strongly correlated with externality creation to allow for transfers that will keep all agents whole. That is, the failure to create a Pareto improvement is due to a *prediction* problem; lump-sum transfers can only undo the distribution of burdens if they can be targeted precisely. Barring near perfect predictability, the only way to achieve a true Pareto-improving transfer scheme is to base transfers on the externality-generating activity itself. But, making transfers contingent upon the targeted activity will inevitably create incentive problems that erode the efficiency of the corrective tax.

This paper proceeds by laying out a theoretical framework that derives a necessary con-

dition for a Pareto-improvement to be possible. I show that, for a Pareto improvement to be "technologically feasible," the variables that are used to determine the transfer program must precisely predict baseline consumption of the externality-causing good, with the degree of precision related to the size of the surplus gain created from externality reduction. This in turn depends on the marginal social damages and the own-price demand derivative for the good. To check whether this necessary condition holds for an empirical application, one therefore needs (1) an estimate of the distribution of baseline consumption of an externality-creating good, (2) knowledge of the correlation between baseline consumption and covariates that can be used in a transfer scheme, (3) an estimate of the own-price derivative, and (4) an estimate of the size of the externality.

To take the theory to data, I use the Consumer Expenditure Survey (CEX) to estimate the distribution of consumption of externality-creating goods and the correlation between consumption and covariates that could be used in transfer schemes. I combine this with estimates from the literature of the size of externalities and price derivatives. I initially focus on a gasoline tax used to correct carbon-related externalities. There is wide dispersion in consumption of gasoline across households, and only a modest fraction of this variation is correlated with variables that are likely to influence a transfer scheme, namely household structure, geographic location and income. Only about one-third of intrahousehold variation in annual gasoline expenditures is predictable by those variables, based on OLS and Lasso models. Using conventional estimates of the externality gain achieved by a carbon tax, I conclude that the transfer scheme is nowhere close to precise enough to create a Pareto improvement (i.e., the necessary condition does not hold). Instead, in the most saturated model, I find that more than one-third of households are net losers.

Additional variation can be explained with correlated endogenous variables. Specifically, vehicle ownership variables predict some of the remaining variation in gasoline expenditures, as one would expect. Conditioning transfers to recycle the gasoline tax on vehicle ownership is clearly problematic in terms of incentives, but even using these variables only pushes the explanatory power up by a modest amount and leaves a large number of losers. (Note that, even the preferred schemes involving demographics and income create some incentive problems, as those characteristics will, to varying degrees, respond to the transfer scheme. In abstracting from these distortions, I am painting the best possible picture for targeting, which still falls far short of creating a Pareto improvement.)

I then show that the degree of predictability is no better for other externality-causing goods measured in the CEX, namely natural gas, electricity, alcohol and tobacco. I interpret this as evidence that it will be infeasible to create a Pareto improvement from corrective taxes on these goods, even when a planner uses an implausible amount of information to

create an unrealistically flexible lump-sum transfer scheme.

The theory and prediction exercise are focused on the narrow question of whether it is possible to compensate all losers from a corrective tax. The broader point of this paper is to point out limits inherent in the economic approach that divides questions of equity and efficiency, and to point out that there is a **prediction** problem at the heart of this (potential) separation. A tradition in economics going back at least to Musgrave (1959) suggests that efficiency and equity concerns can generally be divided. Given tools that can tilt the balance between rich and poor, like a progressive income tax, a policymaker should ensure market efficiency, and then simply dial up (or down) the levers that determine the income distribution to achieve the desired resource allocation in society. This is an extremely useful modeling device, and it is favored by many who study second-best tax design (e.g., Kaplow 2008). A significant literature in public finance explores the separability of efficiency-enhancing policies, including Pigouvian taxes, in second-best constrained environments (e.g., Gauthier and Laroque 2009; Kaplow 2012). This theoretical literature has noted that preference heterogeneity is a critical assumption in their models, but there has been no accompanying empirical work to test whether or not Pareto improvements can be created when there is some degree of heterogeneity.

Should we be concerned about creating a Pareto improvement, or is it a red herring? Pareto efficiency vis-à-vis the status quo is quite distinct from social welfare maximization.[1] If one begins with the objective of maximizing social welfare, there is no reason to prioritize the status quo resource allocation in society, so fussing over Pareto improvements is largely a distraction. The motivation for seeking Pareto improvements in this paper is instead a practical one. The political process tends to favor the status quo over changes, and as such, affecting change requires satisfying a great many people. That is, a utilitarian planner would gladly accept a policy that benefits most people, but causes modest harm to the remainder. But, in practical terms, even small numbers of losers can create substantial political obstacles, consistent with the logic of collective action (Olson 1965, 1982). Empirically, this paper suggests that even implausibly well-designed schemes will leave large fractions of households as net losers.

The paper contributes to an existing literature on the distributional impacts of gasoline taxes (e.g., Poterba 1991; West 2004) and carbon taxes (e.g., Hassett, Mathur, and Metcalf 2009; Dinan 2012; Mathur and Morris 2014; Metcalf 2009; Burtraw, Sweeney, and Walls 2008; Williams, Gordon, Burtraw, Carbone, and Morgenstern 2015). That work has been

[1] See Kaplow and Shavell (2006) for an alternative view that favors de-emphasizing fairness and focusing solely on welfare maximization. My difference in view is based on a concern for practical implementation of policy, not on a difference in normative theory.

overwhelmingly focused on measuring *average progressivity/regressivity* of taxes, whereas this paper is sharply focused on heterogeneity in policy burdens and the degree to which that heterogeneity can be controlled via a transfer scheme.

Very recently, a smaller literature has emerged that seeks to quantify heterogeneity of policy burdens. Rausch, Metcalf, and Reilly (2011) use the Consumer Expenditure Survey (CEX) to characterize the overall progressivity of carbon pricing, accounting for both consumption and income channels. Pizer and Sexton (2016) analyze the CEX and similar data from the United Kingdom and Mexico to show box plots that depict the range of energy consumption within income deciles. Fischer and Pizer (2016) explore how attention to horizontal equity influences a comparison between energy-pricing schemes and a performance standard. Cronin, Fullerton, and Sexton (2017) link the CEX to income tax data to explore a variety of revenue recycling mechanisms and quantify the variation in burdens that remains, taking into account fine-grained differences in income sources. Davis and Knittel (2016) show the heterogeneity in policy impacts of fuel-economy standards across different households in the same income decile in what is otherwise a study of average progressivity.

These papers provide several initial results that are important for the development of a full analysis of heterogeneity in the incidence of energy policies. All demonstrate that there is significant heterogeneity in baseline energy consumption within households that have similar income. In this paper, I find the same descriptive facts, but I go further in several ways. Cronin, Fullerton, and Sexton (2017) and Fischer and Pizer (2016) both mention that, when there is a great deal of heterogeneity in baseline energy usage, it will be impossible to design transfer schemes to make everyone better off. They assert this, but they do not in fact test it. I develop a model that demonstrates under what conditions this claim will hold, and I then test this claim empirically. Of this literature, only Cronin, Fullerton, and Sexton (2017) link their study of heterogeneity to revenue redistribution schemes. They model several realistic schemes for revenue redistribution, but they do not attempt to optimize redistribution schemes to minimize the number of net losers, which is the centerpiece of this project.

2 Model of Pareto transfers

A simple model is developed to fix ideas. Consider an economy with a good q that causes a negative externality and a quasi-linear numeraire m. Heterogeneous agents, indexed $i = 1, ..., N$, have exogenous incomes y_i and different utilities over q, which result in demand curves written $q_i(p+t)$, where p is the market price and t is any tax levied on the consumption of the good. The supply of q is assumed to come from a perfectly competitive market with

constant returns to scale, so any tax burden is born entirely by buyers. Demand for q_i is nonnegative, and for the purposes of the initial exposition is assumed to be strictly positive. Denote the baseline consumption of the good (demand when taxes are zero) as $\tilde{q}_i \equiv q(p)$.

The externality is a global one, in that the damages experienced by each consumer can be written as a function of the total consumption of the good, $Q \equiv \sum_i q_i$. For expositional purposes, suppose that damages are linear, but potentially heterogeneous across consumers. Total damages to consumer i are written as $\phi_i Q$, where ϕ_i represents marginal damages. Denote the marginal damage summed across consumers as $\Phi \equiv \sum_i \phi_i$. In this setting, the standard Pigouvian tax equal to marginal damages ($t = \Phi$) yields a fully efficient outcome.[2]

Our objective is to understand under what conditions, in this simple economy, can a corrective tax on q achieve a Pareto improvement. To illustrate the nature of the problem, let's consider the case of a small increase Δt in the corrective tax, starting at $t = 0$. This is the case that will be most favorable to achieving a Pareto improvement. For expositional ease, suppose that demand is locally linear.

From the point of view of a social planner, the introduction of a tax will have three initial effects. It will lower private welfare from the consumption of q, raise tax revenue, and improve welfare (for all consumers) through the reduction of the externality.

1. The private welfare change is $c_i = -\tilde{q}_i \Delta t + 1/2 q_i' \Delta t^2$. For a small change in t, the squared term will approach zero, so the private welfare gain $c_i \approx -\tilde{q}_i \Delta t$.[3]

2. The revenue raised is $r_i \approx \tilde{q}_i \Delta t - q_i' \Delta t^2$. For a small change in t, the squared term will approach zero, so that $r_i = \tilde{q}_i \Delta t$. Importantly, for small changes in the tax, revenue and the private welfare loss approximately offset.

3. The externality gain is $s_i = \Phi q_i' \Delta t$. Note that even for a small tax, there is a first-order gain.

From the point of view of an individual consumer, the introduction of the tax has two effects (before redistribution of revenue).

1. The private welfare loss, as above: $c_i = -\tilde{q}_i \Delta t + 1/2 q_i' \Delta t^2$

2. The gain in the externality, due to the change in aggregate consumption: $g_i = \phi_i \sum_j q_j' \Delta t$.[4]

[2]To be clear, one unit of consumption is assumed to have the same social damage, regardless of who consumes it. The heterogeneity is that this damage experienced may be uneven across agents. This is in contrast to the heterogeneity in social damages model of Diamond (1973).

[3]For an infinitesimal change, this is just Roy's identity.

[4]Note that $\sum_i g_i = \sum_i s_i$.

Let us consider first the case of a small tax change, and consider also the special case where there is no heterogeneity in externality benefits across consumers. That is, suppose $g_i = \bar{g}\ \forall i$.

To achieve a Pareto improvement, revenue must be redistributed in a targeted manner, so that those who experience the largest direct welfare loss receive the largest transfer. Denote the transfer scheme as some function that maps characteristics of consumers, $\mathbf{X}_i$, into a transfer, $T(\mathbf{X}_i)$. Suppose that all of the revenue from the tax will be recycled through the transfer scheme, so that $R = \sum_i T(\mathbf{X}_i)$.

The initial (pre-transfer) welfare gain for consumer i is $c_i + g_i = -q_i + g_i$ (for a small tax). Thus, consumer i experiences a welfare gain if and only if

$$T(\mathbf{X}_i) > q_i - g_i.$$

By its nature, a Pigouvian tax will increase social welfare. Specifically, this means that $\sum_i (c_i + r_i + s_i) = \sum_i (c_i + T(\mathbf{X}_i) + g_i) \geq 0$. The challenge here is to transform this efficiency gain into a Pareto improvement, where all consumers enjoy some of the gain. To see the difficulty more clearly, consider first the case where the externality gains are uniform, so that $g_i = \bar{g}$. Then, rewrite the condition for an improvement as:

$$\bar{g} > q_i - T(\mathbf{X}_i).$$

A Pareto improvement requires that this condition be met simultaneously for all consumers. Condition 1 shows a necessary condition for this to be possible, which is that the absolute errors in the targeted transfer scheme must be smaller than the average externality gain.

Condition 1. *For a small tax increase, when ϕ_i are homogeneous and errors are symmetric, Pareto improvement is impossible whenever $N\bar{g} \leq \sum_i |q_i - T(\mathbf{X}_i)|$.*

To interpret Condition 1, consider a transfer scheme that aims to return the revenue back to individuals exactly in line with the amount that they were harmed by the price increase. In that case, one can think of $q_i - T(\mathbf{X}_i)$ as a prediction error. For a marginal tax, total revenue will equal total initial costs, so the average error is zero. To achieve a Pareto improvement, the absolute errors can be no larger than the externality gain, otherwise at least one person must be a loser. This condition is only made more stringent for a larger tax, since the revenue will be smaller than the initial welfare cost due to behavioral responses.

This first condition assumes that errors are symmetric around zero, which will be the case if the errors are the result of a least squares prediction equation. A closely related

alternative condition can be derived that relaxes the assumption of symmetry. This shows that a necessary condition for a Pareto improvement, with the same conditions other than symmetry, is that $(N-1)/N \times 2\bar{g} \leq \sum_i |q_i - T(\mathbf{X}_i)|$. Thus, relaxing symmetry makes the necessary condition somewhat less restrictive, but all of the intuitive relationships remain.

This makes clear that a Pareto improvement must be built upon a **prediction exercise**. The variables that determine the transfer scheme $\mathbf{X}_i$ must be able to predict baseline consumption with a sufficient degree of accuracy in order for a Pareto improvement to be possible. The degree of accuracy required depends upon the size of the welfare gain. The larger is $\bar{g}$, the more room there is for error. Note also that this prediction accuracy is in absolute errors, not squared errors, which will motivate an empirical approach that uses least absolute deviation regressions, rather than ordinary least squares.

Note also that as the revenue raised is large relative to the externality gain (which will be the case when the corrective tax changes the total percentage of consumption by a small amount), the accuracy of the target function needs to be improved in percentage terms. This motivates attention to the ratio of the revenue raised versus the size of the externality gain in empirical applications.

In sum, the welfare framework makes clear that the design of a Pareto-improving transfer scheme is essentially a prediction problem, where the variables that are inputs into the transfer scheme are used to predict the welfare loss from the tax. Assuming a uniform distribution of externality benefits, the prediction problem boils down to predicting baseline (pre-tax) consumption. For a non-marginal increase in the tax, heterogeneity in behavioral responses becomes relevant, because welfare losses depend on baseline consumption and changes in consumption (i.e., demand derivatives). When benefits from the externality reduction are heterogeneous, the variable to be predicted is the private welfare loss net of the idiosyncratic benefit.

3 Consumption data on externality-generating goods

This paper uses data from the interview portion of the Consumer Expenditure Survey (CEX), which is a nationally representative sample of U.S. households, from 1996 to 2016.[5] The CEX defines a unit of observation as a consumer unit, which is a set of individuals who reside together and are either related by blood or marriage, or who make financial decisions together.

[5]The CEX is actually two distinct surveys drawn from different samples of households, called the interview and diary surveys. The diary survey is less appropriate for current purposes where the focus is on variation across households because it records expenditures over a single week, so that variation in diary expenditures might simply reflect variation in timing of spending, rather than overall levels.

Table 1: Household expenditure statistics by category

	Mean	Median	St. Dev	CV	Pct 0
Motor fuels	\$1,746	\$1,353	\$1,529	0.9	9%
Electricity	\$1,143	\$984	\$913	0.8	9%
Natural gas	\$413	\$162	\$611	1.5	42%
Alcohol	\$219	\$15	\$401	1.8	47%
Tobacco	\$296	\$0	\$648	2.2	72%
All energy	\$3,306	\$2,919	\$2,226	0.7	3%
All sin goods	\$3,821	\$3,388	\$2,493	0.7	2%

Table shows annualized expenditures by category for all households in sample (N=197,668). Dollar amounts are in \$2015. All energy sums motor fuels, electricity and natural gas. All sin goods includes all five individual categories summed.

Interviews consist of retrospective questions that ask about the consumer unit's total expenditures on various items over the prior three months. Units are interviewed four times, once each quarter, but not all units complete all four rounds of interviews. For the analysis below, expenditure categories are averaged over however many interviews are completed by a consumer unit, and then scaled to represent annual consumption amounts.

Table 1 shows summary statistics on expenditures. Key for this paper is that there is wide variability in the consumption of all variables. For example, average consumer unit expenditures on motor fuels is \$1,746, but the standard deviation is nearly as large as the mean, at \$1,529.

There are two important caveats to be kept in mind regarding the use of the CEX in this study. First, the analysis is concerned with variance and predictability of consumption levels across households. The survey response may mismeasure true consumption either because of sampling variability or because of inaccuracies in self-reported responses. For a discussion of CEX data quality, see Meyer, Mok, and Sullivan (2015). Throughout the paper, I winsorize all expenditure variables at 1% in order to trim the influence of outliers. Data quality remains a concern, but it turns out that the expenditure data reported here implies almost exactly the same mean gallons per year estimates as the 2009 National Household Travel Survey, which provides some reassurance.

Second, the CEX reports expenditures, not quantities. Externality-correcting policies are typically specific taxes (per unit) not ad valorem. To model an ad valorem tax on a product, only the total expenditure is required. Corrective taxes, however, will often take

Table 2: Summary statistics of demographic variables

	Mean	St. Dev.	Min	Max
Before-tax income ($2015)	59,220	61,680	-419,200	971,100
After-tax income ($2015)	55,840	57,430	-421,300	961,700
Consumer unit (CU) size	2.4	1.5	1	29
Persons < 18 in CU	0.62	1.1	0	14
Persons >64 in CU	0.28	0.6	0	8
In MSA indicator	0.83	0.38	0	1
Urban indicator	0.91	0.29	0	1
Reference person married	0.50	0.50	0	1
Year	2006	6.1	1996	2016

the form of a specific (per unit) tax. For example, a carbon tax will raise the price of gasoline by a constant amount per gallon. Thus, to model the impact of a carbon tax on gasoline consumption, we need to estimate the gallons of gasoline consumed by a household, based on their reported expenditure and prices.

For gasoline and diesel fuels, I use data from the Energy Information Administration (EIA) on the sales-weighted, tax-inclusive, retail price of all grades of each fuel type at the closest available geographic match to the consumer unit. That is, where the CEX identifies a consumer unit's metropolitan statistical area and the EIA has city-specific prices, the consumer unit is assigned prices in the past quarter that are the average EIA price for that city. In other cases, matches must be made at the state or PADD level.

For other goods, determining the price paid by consumers is more challenging. Consider alcohol. Prices will vary widely if a consumer unit is purchasing low cost beer or high-end Scotch. As a result, for goods other than motor fuels, I focus on predicting expenditures directly (rather than predicted tax burdens), which translates directly to taxes under an ad valorem tax, recognizing that this is not how a true Pigouvian tax would be designed.

The core empirical task in the paper is to determine the degree to which demographic variables that might plausibly be used in a transfer function are able to predict variation in expenditures across consumer units. Table 2 summarizes the key variables used for this purposes, which are measures of income, household size and location.

4 A gasoline tax creates losers

The primary empirical application of this paper is a gasoline tax, which is modeled here as an efficient carbon-correcting policy. This is an important policy in its own right, and

also has advantages in terms of modeling and measurement with the CEX. The conceptual goal of this analysis is to analyze an optimally designed Pigouvian tax. I thus focus on the gasoline tax as a well-targeted policy for correcting carbon externalities, but I discuss the implications of other driving-related externalities in the robustness section below.

In this section, I calculate the relative magnitude of welfare gains as compared to revenue raised from a motor fuel tax, and then demonstrate the degree to which demographic variables can predict motor fuel consumption. Specifically, I model a small tax increase of 10 cents on motor fuels (both gasoline and diesel) under the assumption that the carbon externality from motor fuel consumption is not corrected at all prior to the tax. That is, I am interpreting existing gasoline and diesel taxes as having been motivated by considerations about the optimal way to raise revenue, irrespective of a carbon externality. These assumptions are designed to be conservative against my findings, as they will maximize the implied welfare gains from carbon taxation.

In the analysis, I abstract from concerns about the tax interaction effect, through which a tax on gasoline is expected to create additional distortions by lowering real factor returns (Goulder 1995; Bovenberg 1999). I again interpret this abstraction as conservative against my results, as the tax interaction effect will imply that net welfare gains from correcting the externality are smaller.

4.1 What are the carbon externality gains from motor fuel taxation?

As described in the model, the welfare gain from a small tax on gasoline will be equal to the change in gasoline consumption induced by the tax times the externality per gallon. I assume that in the long run a gasoline tax will be born completely by consumers so that prices will rise by 10 cents per gallon.[6]

The gasoline demand literature typically estimates elasticities, so I translate the 10 cent gasoline hike into a percentage price change using the average retail gasoline price facing the consumer unit at the time of the survey in its geographic location. I then use a gasoline price elasticity of -0.4, which is interpreted as a long-run price elasticity, to translate this

[6]Existing studies find evidence of high pass through rates for state gasoline taxes, with many studies consistent with full pass through Chouinard and Perloff (2004, 2007); Doyle and Samphantharak (2008); Marion and Muehlegger (2011). Fewer studies consider the federal gas tax, perhaps because it has changed much less often, which impedes econometric investigation. (Chouinard and Perloff 2004) conclude that only half of a federal tax increase is born by consumers. If true, it would be important to consider the incidence on U.S. households through the producer side in interpreting the estimates. I return to this issue when discussing the empirical results.

price change into a change in gallons of fuel consumed.[7] By its very nature, it is challenging to estimate the long-run price elasticity of gasoline. I experiment with alternative values below.[8]

I use the EPA's conversion factor to determine the tons of carbon emitted per gallon of gasoline consumed (17.6 pounds per gallon / 2205 pounds per metric ton for E10, or 22.5 pounds per gallon / 2205 pounds per metric ton for diesel) and then multiply by $40 for the social cost of carbon.

All of the assumptions here are designed to be generous in favor of creating larger externality benefits, and using the global social cost of carbon is foremost in that generosity. Climate benefits are largely realized in the future, and the majority of benefits will be realized outside of the U.S. Indeed, the current administration advocates use of a domestic social cost of carbon ranging from $1 to $6 in rule making. Thus, while there is vigorous debate about the right estimate of the social cost of carbon, it is exceedingly likely that $40 per ton exaggerates the benefits that accrue to current U.S. drivers.

4.2 Externality gains are much smaller than the initial burden and revenue raised

Because I am modeling a small gasoline tax, the initial burden (loss of consumer surplus from the higher price) will be approximately equal to the revenue, both of which are simply the price increase times the number of gallons of gasoline consumed by the consumer unit. But, to be more precise, I use the elasticity estimate to calculate the final quantity consumed, and use that to calculate revenue. The welfare loss is calculated using a linear approximation. Specifically, revenue raised from each household is equal to 10 cents times the new consumption level, which is equal to the current observed of consumption (from data) minus the elasticity (-0.4) times the implied change in price (current price plus 10 cents divided by the current price, all minus 1). The initial private welfare loss is calculated as the new consumption level (as described above), plus the triangle, which is the change in consumption (as described above) times 1/2 times the tax (10 cents).

[7] Small and Van Dender (2007) estimate long-run elasticities closer to half this magnitude. Hughes, Knittel, and Sperling (2008) conclude that the elasticity has been declining over time, finding preferred estimates well below -0.4. Espey (1998) finds a range of estimates that extend well beyond -0.4 in magnitude, but this is based on a variety of studies with varying credibility of empirical strategy. There is some suggestion that demand might respond more to gasoline taxes than price variation (Davis and Kilian 2011; Li, Linn, and Muehlegger 2014), though these estimates, taken from monthly changes in consumption, may be due inflated estimates due to consumers pre-buying in anticipation of price changes (Coglianese, Davis, Kilian, and Stock 2017). This difference seems unlikely to persist in the long run.

[8] I assume a homogeneous elasticity. Simple back of the envelope calculations make clear that allowing for heterogeneity will have little impact on the qualitative results.

According to these calculations, the externality gains are $8.25 per consumer unit per year on average, while the revenue raised is $90.00 per consumer unit per year. The revenue raised is an order of magnitude larger than the externality gain. This has an important implication for the ability of the planner to create a Pareto improvement because, as shown by the theory, the externality gains represent the "error budget" available. A large amount of revenue needs to be reallocated via a transfer function, and the error budget is small relative to the revenue raised. This implies that the revenue must be allocated very precisely.

4.3 Most variation in burdens is not predictable

The key suggestion of the theoretical model is that the degree to which the initial (pre-transfer) burden of the corrective tax can be predicted by variables that are used in the transfer function will determine whether a Pareto improvement is technologically feasible. Simple regression of the household level burden on variables that constitute the transfer function thus provide the required estimates. Below, I present results where the left-hand side variable is the estimated household level initial burden of a 10 cent gas tax.[9] All values are inflation adjusted to 2015.

The theory involves non-squared errors, so I present least absolute deviation (LAD) regressions that will minimize non-squared errors. But, I also present parallel specifications from OLS because the properties of OLS and the R^2 goodness of fit statistic is most familiar. Note that LAD will, by definition, yield lower absolute errors, but OLS, by definition, will maximize the R^2.

Table 3 presents the primary estimates from this exercise. All regressions include year of sample fixed effects, which account for any time trends, though it turns out that excluding them has almost no impact on the results. Designing a transfer scheme that depends on any variables that are not strictly exogenous will create distortionary incentives. As a result, I focus attention first on the "most exogenous" variables that are likely components of a tax scheme, which are demographic indicators for household structure and geographic indicators for state and urban versus rural. Specifically, regressions include state dummies, dummies for an urban indicator and an MSA indicator, and dummy variables for the number of people in the household, as well as the number of minors, and the number over age 60. These variables predict a just under 30% of the variation in gasoline tax burdens (which is approximately the same as baseline consumption).

Column B adds a linear income control, followed by a non-parametric function of income (dummies in five-year bins) in column C. These provide a modest boost in the explanatory

[9]Because I am assuming a homogenous elasticity across households, this is equivalent to using initial baseline consumption (in gallons) as the left-hand side variable.

Table 3: Predictability of Motor Fuel Tax Burden (OLS)

	A	B	C	D
Avg. Abs. Error	\$46.6	\$45.0	\$44.2	\$39.9
R^2	.292	.331	.356	.456
N	197,668	197,668	197,668	197,668
Year FE	Y	Y	Y	Y
Demo & geo controls	Y	Y	Y	Y
Linear income		Y	Y	Y
Binned income			Y	Y
Vehicles & energy				Y

power, with the R^2 bumping up to .331 and .356, respectively. For reference, income by itself, without any demographic or geographic variables, explains only about 15% of variation (results not shown). Column C is my preferred specification. It is based on characteristics that are already part of the tax system, and could plausibly be used to design a tax reform or transfer scheme that accompanies an externality-correcting tax.

The unexplained variation in this specification is far too large to achieve a Pareto improvement. The average absolute error allows for direct comparison with the welfare gains from the externality. The residuals are around \$45 per household. This compares to the \$8.25 welfare gain. This is directly related to Condition 1: as long as the absolute average error exceeds the welfare gain, a Pareto improvement is not possible. Moreover, it is not just a matter of a few people being left as net losers. The best fitted scheme leaves more than one-third of households as net losers, even with the generous assumptions employed throughout.

Column D adds some clearly endogenous variables that would create significant distortions and are thus likely problematic variables for inclusion in a transfer scheme, including home energy consumption and dummies for the number of vehicles owned by the household, and dummies for the number leased. These variables do provide an additional boost to explanatory power, but even with vehicle ownership variables included, the variables explain less than half of the variation.

Table 4 shows LAD specifications. As expected, these lower the absolute error for identical specifications, but only by a very small amount. Figure 1 shows the distribution of net losses, accounting for both the externality gain and the targeted transfers, based on column C in Table 3. A full 37% of households remain as net losers under this scheme.

The variables chosen here are the ones that are most likely to be used for a transfer

Table 4: Predictability of Motor Fuel Tax Burden (LAD)

	A	B	C	D
Avg. Abs. Error	$45.7	$44.1	$43.2	$38.8
N	197,668	197,668	197,668	197,668
Year FE	Y	Y	Y	Y
Demo & geo controls	Y	Y	Y	Y
Linear income		Y	Y	Y
Binned income			Y	Y
Vehicles & energy				Y

scheme that operates through the tax code. The tax code is essentially a function of income and demographic structure of the household. As such, I interpret the results of Tables 3 and 4 as demonstrating that gasoline expenditures are not predicted well enough to come remotely close to enabling a Pareto improvement. A Pareto improvement is not technologically feasible.

5 Machine learning marginally improves prediction

The problem posited here is fundamentally a prediction problem. This is therefore a natural application for machine learning. Table 5 reports results from an exploration of problem of predicting gasoline tax burdens using Lasso. These results are based on a 20% sample of the data.[10] The first column repeats column C from Table 3 on the 20% sample, which delivers nearly the same average absolute error and R^2.

The second column is a Lasso analysis that uses cross validation to eliminate overfitting, run on the same set of possible covariates that are included in the first column. Lasso drops a number of explanatory variables, but in doing so barely changes the R^2 and average absolute error. This suggests that overfitting is not causing a significant bias in the main results.

The third column then allows for a large number of interactions, including year dummies interacted with the $5k bin income dummies, and geography and demographic dummies interacted with the income dummies. These additional variables provide only a modest boost to explanatory power. Other initial explorations have similarly found that significantly expanding the set of explanatory variables, based on what is available in the CEX, has limited ability to explain the remaining cross-household heterogeneity in motor fuel consumption.

[10] Additional exploration, including random forest analysis, is in progress.

Figure 1: Net Loss from 10-cent Gasoline Tax with Targeted Transfer

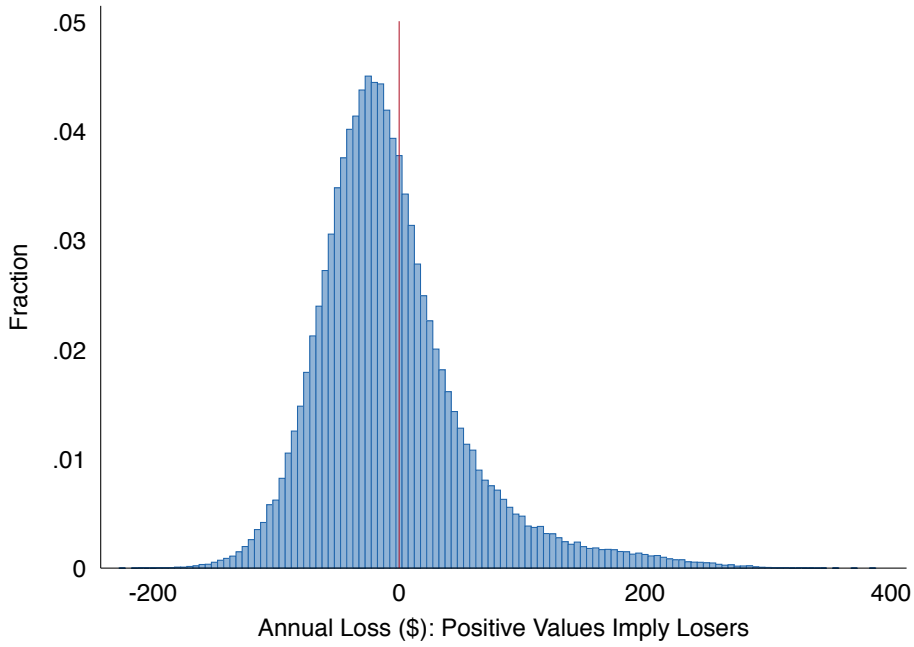


Figure shows the distribution of net impacts of a 10-cent gasoline tax, in dollars per year. A positive value implies a welfare loss. The transfer scheme used is based on the last column of Table 3, using specification C. The net impact is the private welfare loss, net of the targeted transfer scheme, net of the externality gain.

In other words, the main results are robust to (preliminary) machine-learning analysis.

6 Robustness to parameter choices

In this section, I present results that alter three assumptions about the data. First and most simply, I increase the number of data points that are winsorized. Second, I modify the elasticity of gasoline consumption from -0.4 to -0.6 and then -0.8, to reflect higher estimates from the literature. Greater elasticities are important because they will lead to greater welfare gains, which aids the elimination of losers. Third, I greatly increase the externality per gallon of gasoline consumed, from around $0.31 to $2.

The higher latter number is based on accounting for non-greenhouse gas emissions from motor fuel consumption. Harrington, Parry, and Walls (2007) survey the literature and conclude that greenhouse gas emissions externalities are quite modest compared to accident and congestion externalities. A gas tax is a very poor instrument for targeting congestion, and a mediocre at best instrument for targeting accidents or local air pollution. Nevertheless, I now show cases where the gas tax could have much larger benefits in order to compare

Table 5: Lasso regressions on gas tax burden

	OLS (C)	Lasso	Lasso
Avg. Abs. Error	\$43.5	\$43.6	\$43.2
R^2	.359	.354	.365
Vars. Supplied		178	7,844
Vars. Selected		129	433
N	39,534	39,534	39,534
Year FE	Y	Y	Y
Demog. & geog. controls	Y	Y	Y
Linear income	Y	Y	Y
Binned income	Y	Y	Y
Exogenous interactions			Y

Table 6: Fraction of losers under alternative assumptions

Elasticity	Externality per gallon	Percent Windsorized	Percent Losers
-0.4	\$0.31	1%	37.0%
-0.4	\$2	1%	15.5%
-0.6	\$2	1%	9.7%
-0.8	\$2	1%	6.2%
-0.8	\$2	10%	3.0%

results.

In arriving at a \$2 per gallon externality, I modify the values from Harrington, Parry, and Walls (2007) to account for a higher accident externality, at \$0.91 per gallon based on Anderson and Auffhammer (2014), but interpret the carbon benefits as negligible. I then subtract off the sales-weighted average gas tax in the US of \$0.48. In terms of the literature on second-best gasoline taxes, however, note that this is still a generous interpretation in that it ignores fiscal interactions that exacerbate labor market distortions. Parry and Small (2005), for instance, argue that the second-best tax is only around 60% of marginal damages due to fiscal interactions.

Table 6 uses targeted transfers from specification C from Table 3, under alternative assumptions, to calculate the number of households that are net losers. Dramatically increasing the interpreted externality gain per mile roughly halves the number of households who are net losers from a gasoline tax. Increase the elasticity of gasoline consumption to much higher rates further drives down the fraction of losers. In this scenario, the number of net losers is driven down to 6%. This is a modest number, but it should be kept in mind that there are

many generous assumptions deployed in this case, so it should be interpreted as a frontier possibility rather than a realistic point estimate. Even in this case, some households are net losers. Finally, taking all of the prior assumptions and also winsorizing a full 10% of the data drives down the number of losers to 3%.

I am in the process of conducting a number of additional robustness checks that introduce heterogeneity in the gasoline tax elasticity and heterogeneity in the price paid by different households, as well as an analysis that uses the National Household Travel Survey to corroborate the main findings.

7 Other externality-correcting taxes are similar

The focus of this paper empirically is on a gasoline tax, but the CEX enables me to make quick assessment of the degree of predictability of other consumption categories that might be the focus of sin taxes. A gas tax has the advantage that it is relatively easy to translate expenditure data into quantities using gasoline price information, and hence to estimate the impact of a specific (per gallon) gasoline tax. The impact of other sin taxes is more difficult to determine because the goods are more heterogenous (e.g., there are many types of alcohol) and are subject to non-linear prices (e.g., two-part tariffs for electricity and natural gas).

Nevertheless, a broad picture of heterogeneity and predictability can be gained by simply regressing total expenditures in these categories on the demographic variables to see how much of the baseline expenditure variation is predictable. This exercise would exactly mimic the burden of an ad valorem sin tax, and they likely come close to mimicking the scale effect of sin tax levied per unit of the sin good in question.

Note that Table 1 shows that electricity has a similar coefficient of variation with motor fuels, but that other categories have even larger variability. OLS regressions in Table 7 shows the same pattern in terms of predictability. Electricity consumption is very similar in its predictability to motor fuels, but other sin goods are substantially harder to predict.

This analysis is incomplete, as it does not account for the welfare gains and is based on an ad hoc assumption about how a corrective tax would impact prices. But, the results suggest that a gasoline tax is likely the easiest place to achieve broad gains, and that the other externality-creating goods are likely to create even larger numbers of losers because of the greater inability to predict variation in baseline expenditures.

Table 7: Predictability of Other Sin Expenditures (OLS)

All statistics are R^2	A	B	C
Motor Fuels	.336	.381	.403
Electricity	.281	.324	.327
Natural gas	.179	.211	.214
Alcohol	.051	.126	.129
Tobacco	.043	.046	.05
All energy	.393	.469	.486
All sin goods	.362	.438	.455
N	197,668	197,668	197,668
Year FE	Y	Y	Y
Demog. & geog. controls	Y	Y	Y
Linear income		Y	Y
Binned income			Y

8 Discussion

In conclusion, the paper lays out an analytical framework that establishes a necessary condition for transforming a Pigouvian tax into a Pareto improvement. It then shows that this condition fails to hold for a gas tax, even in the face of generous assumptions. Preliminary results on other sin goods imply that the condition is likely to fail for those cases as well. As a result, it seems clear that externality-correcting taxes, even when paired with the best-possible transfer schemes or tax reforms, will create substantial numbers of losers. This may be an important piece of the puzzle explaining the difficulty in adopting such policies. The paper now concludes with a brief discussion of a number of remaining caveats and issues of interpretation.

8.1 Pareto improvement is not the same as welfare maximization

This paper is concerned with Pareto improvements. Pareto improvements implicitly elevate the status quo allocation, whereas traditional maximization of a social welfare function does not. As such, it is critical to point out that the concern in this paper is motivated by a belief that, in practice, the status quo matters in the policy-making process. Where policy changes are difficult to enact, and can effectively be vetoed by a minority of agents through lobbying or other mechanisms, the status quo allocation takes on a greater significance. The exercise performed here is motivated, at a high level, by this idea.

In many cases, efficiency enhancing policy changes are not taken up by policymakers. There are various reasons why this might happen, but one obvious one is that the change will create losers. Economists sometimes argue that no losers need be created when efficiency is enhanced. The point of this paper is that this claim is often hollow, and that, in particular for corrective taxation, compensation of everyone is not possible, and many schemes, even when designed with the intent of minimizing the number of losers, will leave a large fraction of agents at a net loss.

8.2 Pareto improvement depends on the scope of taxation and use of revenue

The focus of this paper is on considering a single corrective tax and the transfers that can be created with the revenue raised. If such a policy struggles to create a Pareto improvement, it begs the question of whether a broader set of policies considered together can yield a different result (i.e., fewer losers). Multiple taxes taken together may create a less diffuse (or more predictable) distribution of burdens if the consumption of various goods are negatively correlated across agents. The fact that the combined raw energy expenditures are more predictable than each of the categories taken separately hints that this possibility may have empirical relevance.

In the extreme, we could ask if all policy taken together is a Pareto improvement, as compared to a Hobbesian state of nature. This may very well be the case, but the goal here is to focus more practically on a set of real economic policy challenges that are taken up as piecewise legislation.

Another issue is that the revenue could be spent on public goods and services, rather than returned lump sum. Or, it could be used to lower particular taxes. These alternatives would further alter the distribution of burdens. It could be that this creates greater concentration of burdens, but it seems more likely that it creates further dispersion, only making the problem harder.

8.3 Benefits taxes may more easily achieve Pareto improvement

If benefits from externality reduction are tightly correlated with the initial burdens of a tax, then there will less variance in the initial burden distribution, making it easier to sustain a Pareto improvement from a tax. As a result, a planner may be much better able to create a Pareto improvement from a benefits tax (i.e., a toll that funds an improvement in a highway) than from the classic externality-correcting taxes emphasized in the empirical

analysis. Benefits taxes are often perceived as more fair by voters, which may be related to the issues of compensation described in this paper.

Alcohol, cigarettes and sugary beverages are goods that may lead to externalities, but are often assumed to also be the source of internalities, such that the welfare improvements are concentrated among the heaviest users of the products. (The same is sometimes said of energy-consuming goods, though the evidence of significant behavioral biases is not clear (Allcott and Greenstone 2012).) In this case, the scope for Pareto improvements may be improved.

8.4 General equilibrium considerations matter in theory

This paper takes a partial equilibrium view. The final welfare impacts of any corrective tax depend on general equilibrium considerations, in the long run. A carbon tax, for example, is likely to have important impacts on factor prices, which will have various impacts on households (Williams et al. 2015). More generally, there are various channels through which a carbon tax will have distributional impacts (Fullerton 2011).

The partial equilibrium focus here is of course partly a practical one, but to the extent that the exercise is motivated by political economy concerns, this limited partial equilibrium view is likely the important one. If voters are unable to anticipate general equilibrium incidence effects, it is likely that they heavily discount them in forming their judgments about how a policy will affect them. The immediate impact on prices and any promised transfer scheme are likely the dominate consideration.

8.5 Difficulty in measurement does not imply political expediency

This last point segues to a broader point that the full benefits (as well as the costs) of any externality correcting tax are difficult to measure. This is a challenge for the econometrician, but the challenge is just as great for the planner and the voter, which reinforces the core point of this paper. When the planner cannot measure the benefits or costs, it is no more possible to control the final distributional effects. And, when benefits and costs are difficult for the voter to perceive, it will be difficult to convince everyone that they are in fact benefitting. Of course, if they can be fooled, then perhaps the political problem can be solved.

8.6 Distortionary instruments or historical baselines may make Pareto improvements feasible

It is possible to compensate losers in my framework by simply rebating tax revenue according to the initial expenditure levels. But, this will blunt (or completely eliminate) the incentive to correct the externality. Where only a fraction of the revenue need be rebated in this way, it may be possible to create a Pareto improvement that trades off some of the efficiency gain in order to create the Pareto improvement.

Another approach is to try to use historical baseline consumption to form the transfer scheme. This is what is often done on the firm side in carbon policies, through free permit allocation schemes. This is harder to envision for households, though not impossible, as it must accommodate entry and exit of households into the economy.

More generally, one could perhaps use other variables that are closely correlated with externality-consuming goods in ways that limited choice distortions but favor targeting. For example, vehicle ownership could be included in the transfer scheme, even if gasoline consumption is not. This creates distortions on the vehicle ownership margin, but it may improve targeting substantially. So, there may be some second-best improvements that are possible and could be explored further in subsequent work.

8.7 Loser minimization involves a different loss function

As a hint toward further investigation, I conclude by noting that if the goal was to minimize the number of losers, a planner would begin with a different loss function (neither OLS nor LAD). The LAD loss function is the correct one for checking a necessary condition for a Pareto improvement, but if the true motivation for the exercise is a political economy one, then it might be the case that the planner wishes to limit the number of losers to some politically acceptable number. What loss function could be used that usefully captures the importance of status quo bias?

A loss function that minimizes the number of losers is easy to program mathematically, but it will have impractical properties. For example, a loser minimizing program would likely take the richest person in the sample and take all of their money in the form of a negative lump sum transfer so as to enhance the budget available for others. Some other restrictions on the loss function are needed to make political sense of this problem. Thus, I believe that the next step in this research is to merge some of the issues raised here with insights from political economy models and cooperative game theory that attempt to capture different criteria that make policies acceptable to a group.

References

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

Anderson, Michael and Maximilian Auffhammer. 2014. "Pounds that Kill: The External Costs of Vehicle Weight." *Review of Economic Studies* 82 (2):535–571.

Bovenberg, A. Lans. 1999. "Green Tax Reforms and the Double Dividend: an Updated Reader's Guide." *International Tax and Public Finance* 6 (3):421–443.

Burtraw, Dallas, Richard Sweeney, and Margaret Walls. 2008. "The Incidence of US Climate Policy: Where You Stand Depends on Where You Sit." Resources for the Future Discussion Paper.

Chouinard, Hayley and Jeffrery M. Perloff. 2004. "Incidence of Federal and State Gasoline Taxes." *Economics Letters* 83 (1):55–60.

Chouinard, Hayley H. and Jeffrery M. Perloff. 2007. "Gasoline Price Differences: Taxes, Pollution Regulations, Mergers, Market Power, and Market Conditions." *The B.E. Journal of Economic Analysis & Policy* 7 (1 (Contributions)):1–26.

Coglianese, John, Lucas W. Davis, Lutz Kilian, and James Stock. 2017. "Anticipation, Tax Avoidance, and the Price Elasticity of Demand for Gasoline." *Journal of Applied Econometrics* 32 (1):1–15.

Cronin, Julie Anne, Don Fullerton, and Steven E. Sexton. 2017. "Vertical and Horizontal Redistribution from a Carbon Tax and Rebate." NBER Working Paper.

Davis, Lucas W. and Lutz Kilian. 2011. "Estimating the Effect of a Gasoline Tax on Carbon Emissions." *Journal of Applied Econometrics* 26:1187–1214.

Davis, Lucas W. and Christopher R. Knittel. 2016. "Are Fuel Economy Standards Regressive?" NBER Working Paper.

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

Dinan, Terry. 2012. "Offsetting a Carbon Tax's Costs on Low-Income Households." Https://www.cbo.gov/publication/43713.

Doyle, Jr., Joseph J. and Krislert Samphantharak. 2008. "$2.00 Gas! Studying the Effects of a Gas Tax Moratorium." *Journal of Public Economics* 92 (3-4):869–884.

Espey, Molly. 1998. "Gasoline Demand Revisited: An International Meta-Analysis of Elasticities." *Energy Economics* 20 (3):273–295.

Fischer, Carolyn and Wiliam A. Pizer. 2016. "Equity versus Efficiency in Energy Regulation." Resources for the Future Discussion Paper.

Fullerton, Don. 2011. "Six Distributional Effects of Environmental Policy." *Risk Analysis* 31 (6):923–929.

Gauthier, Stèphane and Guy Laroque. 2009. "Separability and Public Finance." *Journal of Public Economics* 93:1168–1174.

Goulder, Lawrence H. 1995. "Environmental Taxation and the Double Dividend: A Reader's Guid." *International Tax and Public Finance* 2 (2):157–183.

Harrington, Winston, Ian Parry, and Margaret Walls. 2007. "Automobile Externalities and Policies." *Journal of Economic Literature* 45 (2):373–399.

Hassett, Kevin A., Aparna Mathur, and Gilbert E. Metcalf. 2009. *The Consumer Burden of a Carbon Tax on Gasoline*. Resources for the Future Press.

Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling. 2008. "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *The Energy Journal* 29 (1).

Kaplow, Louis. 2008. *The Theory of Taxation and Public Economics*. Princeton University Press.

———. 2012. "Optimal Control of Externalities in the Presence of Income Taxation." *International Economic Review* 53 (2):487–509.

Kaplow, Louis and Steven Shavell. 2006. *Fairness versus Welfare*. Harvard University Press.

Li, Shanjun, Joshua Linn, and Erich Muehlegger. 2014. "Gasoline Taxes and Consumer Behavior." *American Economic Journal: Economic Policy* 6 (4):302–342.

Marion, Justin and Erich Muehlegger. 2011. "Fuel Tax Incidence and Supply Conditions." *Journal of Public Economics* 95 (9-10):1202–1212.

Mathur, Aparna and Adele C. Morris. 2014. "Distributional Effects of a Carbon Tax in Broader US Fiscal Reform." *Energy Policy* 66:326–334.

Metcalf, Gilbert E. 2009. "Tax Policies for Low-Carbon Technologies." *National Tax Journal* 62 (3):519–533.

Meyer, Bruce D., Wallace K.C. Mok, and James X. Sullivan. 2015. "Household Surveys in Crisis." *Journal of Economic Perspectives* 29 (4):199–226.

Musgrave, Richard. 1959. *Theory of Public Finance; A Study in Public Economy*. New York: McGraw-Hill.

Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Harvard University Press.

———. 1982. *The Rise and Decline of Nations: Economic Growth, Stagflation, and Social Rigidities*. Yale University Press.

Parry, Ian W.H. and Kenneth A. Small. 2005. "Does Britain or the United States Have the Right Gasoline Tax?" *American Economic Review* 95 (4):1276–1289.

Pizer, Wiliam A. and Steven E. Sexton. 2016. "Distributional Impacts of Energy Taxes." Manuscript: Duke University.

Poterba, James M. 1991. "Is the Gasoline Tax Regressive?" *Tax Policy and the Economy* 5:145–164.

Rausch, Sebastian, Gilbert E. Metcalf, and John M. Reilly. 2011. "Distributional Impacts of Carbon Pricing: A General Equilibrium Approach with Micro-Data for Households." *Energy Economics* 33:S20–S33.

Small, Kenneth A. and Kurt Van Dender. 2007. "Fuel Efficiency and Motor Vehicle Travel: The Declining Rebound Effect." *Energy Journal* 28 (1):25–51.

West, Sarah E. 2004. "Distributional Effects of Alterntive Vehicle Pollution Control Policies." *Journal of Public Economics* 88 (3-4):735–757.

Williams, Roberton C., III, Hal Gordon, Dallas Burtraw, Jared C. Carbone, and Richard D. Morgenstern. 2015. "The Initial Incidence of a Carbon Tax Across Income Groups." *National Tax Journal* 68 (1).