

Energy Efficiency Standards Are More Regressive Than Energy Taxes: Theory and Evidence

Arik Levinson

Abstract: Economists endorse taxes as a cost-effective means of reducing pollution. But policy makers raise concerns about their regressivity, or disproportional burden on poorer families, preferring instead to regulate energy efficiency. I first show that in theory, energy efficiency standards are more regressive than energy taxes, not less. I then provide an example using data on automobiles in the United States. Taxing gas would be less regressive than regulating the fuel economy of cars if the two policies are compared on a revenue-equivalent basis.

JEL Codes: H23

Keywords: externalities, redistributive effects, environmental taxes and subsidies

FOR NEARLY 100 YEARS, economists have been explaining that a pollution tax would reduce environmental damage in the least costly way.¹ That advice is almost never followed. Instead, environmental regulations typically rely on technology or performance standards. In the United States, rather than a cost-effective carbon tax to reduce greenhouse gas emissions, we rely on energy efficiency standards for buildings and appliances. Rather than a gasoline tax, we set fuel economy standards for cars.

The failure to tax pollution in the United States has multiple roots. Some point to an American aversion to taxation of all forms. Others note that taxing gasoline in particular would be unfair to suburban or rural citizens who drive longer distances. Perhaps the most often cited objection to pollution or carbon taxes is that they would be regressive. Lower-income households spend a larger proportion of their incomes on en-

Arik Levinson is at Georgetown University and the National Bureau of Economic Research (arik.levinson@georgetown.edu). I am grateful to the Sloan Foundation for financial support, to Duncan Hobbs, Grady Killeen, and Mark Noll for research assistance, and to Sarah Aldy, Axel Anderson, Stephanie Fried, Karen Palmer, Hilary Sigman, and members of the NBER Environment and Energy Economics group for especially helpful suggestions.

1. The idea dates to Pigou (1920). There is now even a “Pigou Club” Wikipedia page listing economists and politicians who support the idea (en.wikipedia.org/wiki/Pigou_Club).

Received December 12, 2016; Accepted May 30, 2018; Published online January 23, 2019.

JAERE, volume 6, number S1. © 2019 by The Association of Environmental and Resource Economists. All rights reserved. 2333-5955/2019/06S1-0002\$10.00 <http://dx.doi.org/10.1086/701186>

S7

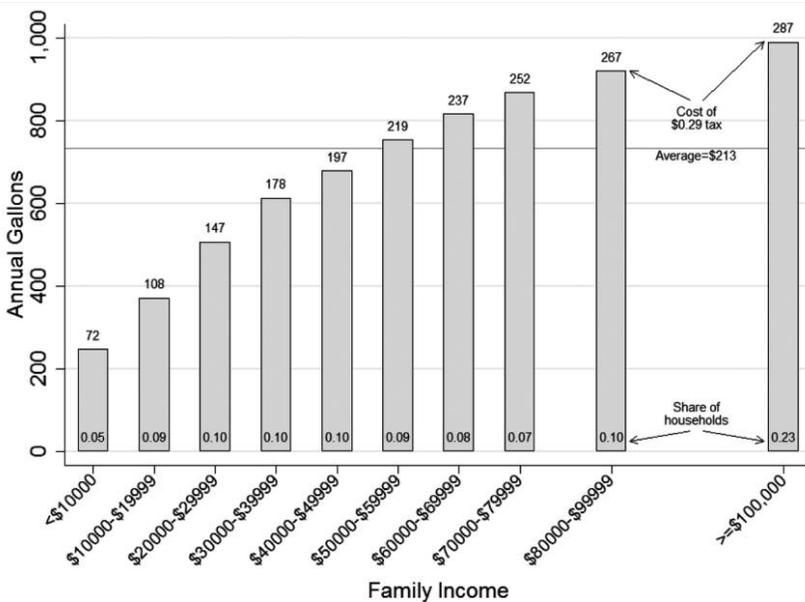


Figure 1. Annual household gasoline use by income, 2009. Source: 2009 National Household Travel Survey. Color version available as an online enhancement.

ergy and goods whose production uses energy, and as a result they would pay a disproportionate share of any pollution tax on electricity or fuel.²

That basic argument can be seen in figure 1, which plots annual household gasoline use by income, from the 2009 National Household Travel Survey (NHTS).³ The poorest 5% of households use an average of 247 gallons of gas per year, including households without vehicles. The richest 23%, with incomes over 10 times higher, each use about four times as much. A \$0.29 per gallon gas tax—one estimate of the monetized climate damage from gasoline—would cost poor households \$72 dollars per year and rich households \$287. (This assumes no change in behavior in response to the tax, which is typical of these analyses but of course defeats its purpose.) Families with more than 10 times the income would pay only four times the tax, and

2. For examples of this objection, see Bill Chameides, “Is the Conservative-Friendly Carbon Tax a Regressive Flat Tax in Disguise?” *Huffington Post*, May 25, 2011; and “Who Pays for Cap and Trade?” *Wall Street Journal*, March 9, 2009. For recent examples, see Mark Paul, Anders Fremsted, and James K. Boyce, “Can Markets Solve Climate Change?” *Nation*, April 12, 2016; and Dan Wilson and Kyle England, “Initiative 732’s Carbon Tax Hurts Families,” *Spokesman-Review*, September 24, 2016. For more academic examples, see Grainger and Kolstad (2010) or Fried et al. (2017).

3. The 2017 NHTS is available as of March 2018. For details, see nhts.ornl.gov.

thus a gas tax would be regressive. So even though an energy tax would reduce energy use in the least costly way, opponents raise distributional concerns.

Until now, no studies have asked the obvious follow-up question. Would an energy tax be more or less regressive than the energy efficiency standards used instead? If an energy tax would be more regressive, a case might be made to forgo the tax's cost-effectiveness in exchange for the efficiency standard's distributional benefits. If the energy tax would be less regressive, that case collapses.

The answer suggested by this paper's title has a straightforward intuition. If the service provided by energy—air conditioning or miles of travel—is “normal” in an economics sense, then richer people want more of it. That gives them a greater incentive to purchase more efficient air conditioners and automobiles. So richer households will buy both more energy and more energy efficiency. A tax on energy will cost those richer households more, in absolute terms even if it is a smaller share of income. By contrast a policy that targets energy-efficient appliances, either by subsidizing efficient ones or taxing inefficient ones, will favor richer households because they already spend more on energy efficiency. Rich households would pay more in energy taxes and benefit more from an energy efficiency policy, and so energy taxes are less regressive.

After making a few preliminary points in section 1, section 2 describes the trade-off between taxes and standards in theory, in a static two-good model with no uncertainty or discounting, where consumers differ only by income. The model has only two goods, an energy service (like air conditioning or miles of travel) and a numeraire. The energy service requires two inputs, energy and the energy-using durable (the air conditioner or automobile). I contrast two public policies, a tax on energy and a tax on inefficient appliances and cars. This simple framework generates a stark result: an energy tax would be both more cost-effective and less regressive.

Section 3 turns to an empirical application using automobiles and corporate average fuel economy (CAFE) standards, with data from the 2009 NHTS. The choice of automobiles as an example stems from the fact that cars are the one energy-using durable with readily available data on energy input, gallons of gasoline used. I start by ignoring the fact that cars differ in ways other than their energy efficiency. I show that in practice the theoretical result from section 2 is borne out: a gasoline tax would be less regressive than a revenue-equivalent efficiency regulation.

At the end of section 3, I consider the fact that richer households buy larger, faster, heavier cars—features correlated with inefficiency. Without controlling for car size and acceleration, richer households' cars are less fuel efficient, a contrast with the theoretical result in section 2. Once I control for car size and engine power, richer households' cars are more fuel efficient, as predicted by the model.

As an extension, I also examine the footprint-based CAFE standards adopted by the United States in 2011. Those exacerbate the regressivity of fuel economy standards, making them even more regressive when compared to a gas tax.

Pieces of these results have been noted by others. Since economists first began studying energy efficiency, researchers have observed that lower-income households purchase less efficient appliances and vehicles, sacrificing higher future energy costs for lower current appliance costs. Some attribute that choice to liquidity constraints. Others describe behavioral explanations, such as present bias (Allcott and Mulainathan 2010) and heterogeneity in discount rates (Newell and Siikamäki 2015). In the model I outline, even unconstrained rational consumers differing only in their incomes choose energy-using appliances and vehicles with different levels of fixed costs and efficiency. Poorer households use less energy and choose less efficient, less expensive appliances and vehicles. As a consequence, the burden of energy taxes falls relatively less on poor households than the burden of efficiency standards.

Researchers also have compared CAFE standards to gas taxes on efficiency grounds. Austin and Dinan (2005) estimate that a gas tax would cost 58% to 71% less than the CAFE standards, per gallon of fuel saved. More recently Jacobsen (2013) finds that the CAFE standards cost three to 10 times as much as a gasoline tax per ton of carbon dioxide avoided. So efficiency standards are, ironically, inefficient. That is easy to show in my simpler model as well.

A few researchers have analyzed energy efficiency standards on distributional grounds. For appliances, Jacobsen (2018) shows that tax incentives and rebates are regressive. For cars, Jacobsen (2013) shows that in the long run welfare losses from CAFE standards are twice as large for the poorest households as for the wealthiest, per dollar of income. Davis and Knittel (2019, in this issue) use nationally representative data on car registrations to show the same thing: CAFE standards disproportionately burden lower-income households.

Existing work has thus demonstrated three things: energy taxes like carbon and gas taxes are regressive; energy taxes are more cost-effective than energy efficiency standards; and efficiency standards are also regressive. Until now what has not been shown is whether energy efficiency standards are less or more regressive than an energy tax would be. If efficiency standards are less regressive, their extra costs might be worth the distributional benefits. If efficiency standards are more regressive, then energy taxes would be both more efficient and more equitable.

One other current paper does pose a question similar to the one I pose here, although it is not the main focus: Davis and Knittel (2019). That paper has one major advantage over the approach I take, in that it uses comprehensive data on all vehicle registrations in the United States in 2012, including rich data about those cars' characteristics. The trade-off is vehicles can only be matched with demographic information at the level of the US census tract in which the vehicle was registered, not at the level of the individual household. In particular, that paper cannot measure two key inputs necessary to compare the distributional effects of gas taxes and CAFE standards: how many cars each household owns and how much gas those households use. For that reason, I rely on the NHTS. Though it is only a survey, and only covers one year, 2009, it

does contain all the necessary information: household income, characteristics of the vehicles owned by the household, and gasoline consumption.

Before describing the model and empirics, I need to be clear about four points. First, I focus on the “static” incidence of taxes throughout, ignoring the demand response that is the intended result of the policies. Second, critical to any calculation of tax incidence is what happens to the tax revenue. Third, energy efficiency standards can be treated analytically as equivalent to a tax on inefficient appliances or vehicles. And fourth, to study the incidence of energy efficiency, I need to take sides in the “energy paradox” debate. Readers not concerned about those clarifications can skip to the theoretical model in section 2, and readers who only want the main empirical results can go straight to section 3.

1. PRELIMINARY POINTS

Energy efficiency regulations come in two broad categories. Most, like appliance standards and building codes, set minimum thresholds for efficiency. Underperforming appliances cannot be sold and substandard buildings cannot be constructed. The distributional consequences of these bans on inefficient products seem self-evident. If rich people would otherwise buy the banned inefficient products, the ban is progressive; if poor people would buy them the ban is regressive. Those distributional outcomes are simple to show empirically, and I will return to them in the empirical analysis in section 4.

A second type of efficiency regulation raises the price of inefficient goods. One example is the CAFE standards that have regulated automobiles in the United States since 1978 and in Europe since 2009. These types of price-based standards can be compared directly to energy taxes using graphs like figure 1. To make that comparison, I need to raise four preliminary points.

1.1. Static Incidence

First, the distributional effect that I describe here is the “static” incidence of the energy tax, a straightforward calculation based on pre-tax consumption. Figure 1 provides an example: poor households spend a larger fraction of their incomes on gasoline and therefore are assumed to bear a larger fraction of the economic burden of a tax on gasoline. To know the true economic incidence—measured by equivalent or compensating variation—we would need to know the price and income elasticities of demands for energy and for energy efficiency by each income group, and the supply elasticities of energy and energy-efficient cars and appliances.

Instead, the static incidence assumes that all the elasticities are zero. It is the idea underlying the public perception of tax burden and depicted in figure 1. Grainger and Kolstad (2010) call this a “first-order” estimate of the incidence. If the welfare loss to each group from a tax on a particular good is proportional to the group’s pre-tax consumption of that good, then this static incidence will be a good measure of the true economic incidence.

1.2. Revenue Matters

Second, critical to the distributional effects of either an energy tax or an efficiency standard is whether it raises revenue and how that revenue is spent. To see why that matters, take the \$0.29 gas tax as an example.⁴ The key to the tax's cost-effectiveness is the \$0.29 per gallon opportunity cost of driving. That \$0.29 per gallon cost could be achieved by nearly identical policies with vastly different distributional consequences, merely by changing the fiscal policy associated with the tax revenues.

On the simplest level, from examining figure 1, a \$0.29 per gallon gasoline tax would be regressive, costing poor families \$72 dollars per year and rich families \$287, a \$215 difference. Rich families pay four times the gas tax even as they earn over 10 times the income.

Suppose, however, that we divided up the total tax revenue evenly among all households, as a lump-sum rebate. If figure 1 represents all US households, each would receive a check for \$213. The poorest group would pay \$72 each in gas taxes and come out ahead by \$141. The richest group would pay \$287 each in gas taxes for a net loss of \$74. Unlike the tax alone, this rebate policy would be progressive.

Alternatively, consider a policy that subsidizes gas conservation at \$0.29 per gallon. Give every household an imaginary allotment of, say, 991 gallons of gasoline per year, the amount used by the richest 23% of households in 2009. Then pay each household a \$0.29 subsidy for each of those 991 gallons not consumed. Looking again at figure 1, with no change in behavior, rich families would get nothing, while the poorest families would receive a check for \$215.

The three policies are equally cost-effective. Drivers would face the same \$0.29 opportunity cost per gallon, providing exactly the same cost-effective incentive to conserve.⁵ But the policies have very different distributional consequences for reasons unrelated to their cost-effectiveness or environmental efficacy. While the \$0.29 gas tax raises revenue, disproportionately from poor households, the \$0.29 subsidy disburses funds, disproportionately to poor households. Unless we know how the tax revenues will be spent or the subsidy funds raised, we cannot assess their distributional consequences. In the theoretical and empirical examples that follow, I take care to compare taxes and standards with equivalent financing.

1.3. Efficiency Standards Are Equivalent to Inefficiency Taxes

An important related point is that efficiency standards are economically equivalent to a tax on inefficient appliances or vehicles. An energy efficiency standard for air conditioners is equivalent to a tax levied on the unit's watt-hours of electric power used to create one British thermal unit (BTU) of cooling. An automobile fuel economy standard is

4. The exact tax rate does not matter. It only serves as an example here.

5. Setting aside income effects on driving and on the decision to own a car.

equivalent to a tax levied on the car's gallons per mile (gpm), the inverse of the common miles per gallon (mpg) measure of efficiency.

In the case of automobiles, CAFE has the same cost-effectiveness at the margin as an outright tax on inefficient cars with high gpm or a pure subsidy to low-gpm cars (setting aside the revenue consequences, per the previous section). Appendix 1 (apps. 1 and 2 are available online) demonstrates that result in a partial equilibrium model based on Kwoka (1983) and Kleit (1988). Anderson and Sallee (2016) and Gillingham (2013) similarly model the CAFE standard as a "feebate"—a set of taxes and subsidies based on the fuel economy of the vehicle rather than on actual gas consumption.

Framing efficiency standards as taxes on inefficient appliances or vehicles makes them easier to compare to an energy tax. Both raise revenue. The energy tax targets fuel consumption directly. The inefficiency tax targets appliance or vehicle characteristics and could be paid either at the time of purchase or annually as part of registration. As with an energy tax, the distributional effects of an inefficiency tax would depend on the collection and disbursement of the revenues. Again, it will be important to compare taxes and standards with equivalent financing.

1.4. The Energy Paradox

Finally, note that by describing efficiency standards as a tax on inefficient appliances or vehicles, I am implicitly taking sides in the "energy paradox" debate. Some policy advocates support efficiency standards on the grounds that market failures or consumer errors lead to the purchase of less efficient appliances and vehicles than would be privately optimal. In that view, an efficiency standard improves private welfare even without considering the environmental benefits. The US Department of Energy estimated that its 2011 energy efficiency standards for refrigerators would save owners \$200 over the life of the appliance.⁶ Similarly, the US Department of Transportation (DOT) says that its 2016 CAFE standards will save drivers over \$4,000 in gasoline costs, more than offsetting the higher up-front vehicle costs.⁷

The case of the privately beneficial standard raises a natural question: if more energy-efficient cars would make consumers better off, why don't consumers demand them and manufacturers sell them without being regulated? This "energy paradox" applies to cars, appliances, and home construction and forms the basis for an enormous amount of research. Greene (2010) summarized that work for the US Environmental Protection Agency, suggesting that for automotive fuel economy, the empirical literature is evenly

6. US Department of Energy, "Department of Energy Joins with Manufacturers, Environmentalists to Announce New Efficiency Standards for Home Refrigerators," press release, August 26, 2011.

7. US Environmental Protection Agency Office of Transportation and Air Quality, "EPA and NHTSA Finalize Historic National Program to Reduce Greenhouse Gases and Improve Fuel Economy for Cars and Trucks," regulatory announcement, April 2010.

divided between studies showing that consumers do undervalue fuel savings and those showing that they do not. Allcott and Greenstone (2012, 4–5) conclude from the literature that “it is difficult to substantiate claims” that consumers and firms “fail to make investments that would increase utility or profits.” Recently, several papers have added to the ranks of those finding they do not (Allcott 2013; Busse et al. 2013).

More importantly for this paper, if the regulation makes consumers better off even ignoring the environmental benefits, then the distributional comparison with a gasoline or energy tax becomes moot. If energy efficiency policies correct consumers’ mistakes and make them better off, then we should implement the policies, regardless of any environmental benefits. By contrast there is no reason to implement an energy tax absent concerns about externalities. Thus I focus on the case where both the tax and the standard reduce welfare. That is what recent empirical work seems to suggest, and the question of the gas tax’s regressivity is otherwise irrelevant.

In the bulk of what follows I present only the simplest case: a static two-good model and empirical calculations of static incidence, ignoring responses by consumers or car-makers to the taxes. I return to a discussion of these issues briefly in section 4.

2. COMPARING EFFICIENCY STANDARDS AND ENERGY TAXES IN THEORY

The model in this section has two aims. First, it describes a framework to compare revenue-equivalent efficiency standards and energy taxes. And second, the model shows theoretically, with the fewest possible assumptions, the result I find empirically in section 3: efficiency standards are more regressive than energy taxes.

Start with a representative consumer or household, with utility over two goods: a numeraire good, x , with a price of 1, and energy services, s , such as cooling, heating, refrigeration, lighting, or miles driven. Denote energy as e and efficiency as μ , energy services per unit of energy, s/e , which for cars would be mpg. Energy services are then the product of energy consumption and efficiency: $s = e\mu$.

$$U(s, x) = U(e\mu, x). \quad (1)$$

For air conditioners, $s = e\mu$ means cooling equals kilowatt-hours of electricity used (e) times BTU per kilowatt-hour (μ). For cars, miles equals gas consumption (e) times mpg (μ).

Households have income (Y) to spend on the numeraire and energy services. Rather than purchasing energy services directly, they purchase it indirectly through the cost of energy and the energy efficiency of the appliance or vehicle. More efficient appliances and cars are more expensive.⁸ For now assume that the cost of energy efficiency (μ) has a constant price (p_μ), but I relax that strong assumption shortly. The household budget is then

8. The US Department of Transportation estimates that fuel economy standards add nearly \$1,000 to the cost of a typical vehicle.

$$Y = x + p_e e + p_\mu \mu. \tag{2}$$

Figure 2 depicts the setup for two appliances: an efficient one with efficiency μ_e and an inefficient one with efficiency μ_i . A household can purchase no energy services and spend its entire income (Y) on the numeraire. Or it can purchase an inefficient appliance for $p_\mu \mu_i$. In that case, the cost of one unit of the energy service will be p_e / μ_i , the slope of the solid budget line in figure 2. Or the household can purchase a more expensive and efficient appliance (μ_e). Then its budget in figure 2 has a lower intercept ($Y - p_\mu \mu_e$) and the correspondingly shallower, dashed budget line with slope p_e / μ_e .

Figure 2 transforms an intertemporal decision (pay more up front, save fuel costs later) into a one-shot, static decision. Think of $p_e e$ as the present discounted cost of purchased energy over the life of the durable and $p_\mu \mu$ as the up-front fixed cost of its efficiency. Alternatively, think of $p_e e$ as annual fuel costs and $p_\mu \mu$ as the annual amortized fixed cost of efficiency. Or even more simply, consider the energy-using durable to be something that is leased for one period with fixed rental cost $p_\mu \mu$ and variable costs $p_e e$.

Now suppose the household chooses μ from many possible efficiencies. Maximizing utility in (1) with respect to the budget constraint in (2) yields two familiar looking first-order conditions:

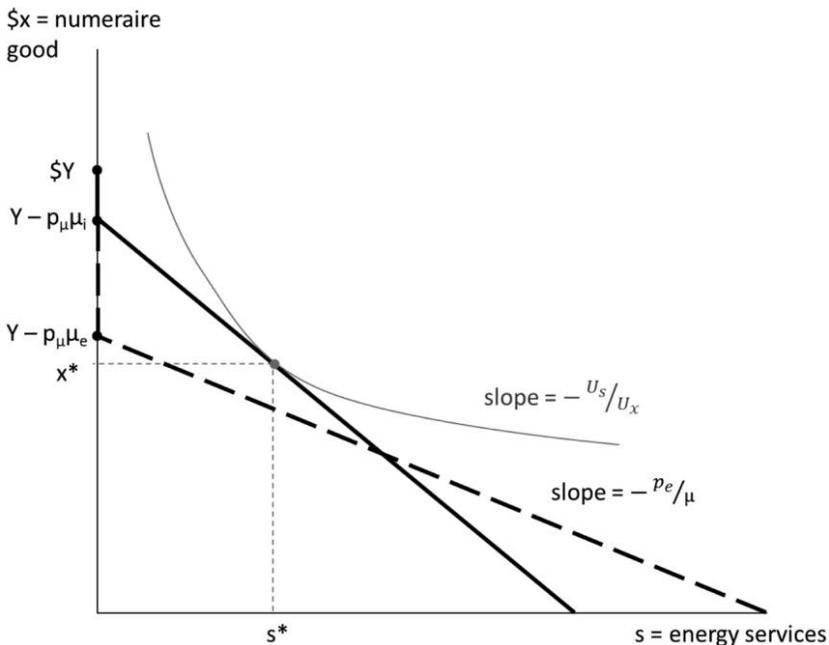


Figure 2. Energy use versus efficiency. Color version available as an online enhancement.

$$\begin{aligned} \text{(i)} \quad U_s / U_x &= p_e / \mu \\ \text{(ii)} \quad U_s / U_x &= p_\mu / e. \end{aligned} \tag{3}$$

The first comes from maximizing with respect to e . It just says that the marginal rate of substitution between energy services and the numeraire should equal the cost of energy services from purchasing more energy, for a given quantity of efficiency (μ), which is the price of energy divided by output per unit of energy (p_e/μ). The second condition in (3) comes from maximizing with respect to μ and indicates that same marginal rate of substitution equals the cost of energy services from purchasing more efficiency, for a given quantity of energy (e), which is the price of efficiency divided by energy (p_μ/e).

The conditions in (3) have a simple interpretation. There are two ways a household can purchase an extra unit of energy services (s). With given efficiency, μ , it can buy more energy at price p_e ; or with a given amount of energy, e , it can buy a more efficient appliance at price p_μ .

Now add public policies to the mix. First, a tax on energy, t_e raises the price of energy to $p_e(1 + t_e)$. That tax in turn raises the cost per unit of energy services to $p_e(1 + t_e)/\mu$, rotating the budget lines in figure 2 down. Panel *a* of figure 3 depicts such a tax for two levels of efficiency.

Next consider an energy-efficiency standard. To put the standard on the same revenue basis as the tax, imagine the efficiency standard as an “inefficiency tax” t_μ or tax on μ below some regulated amount, $\bar{\mu}$. The new budget line with both taxes t_e and t_μ added is

$$Y = x + (p_e + t_e)e + p_\mu\mu + t_\mu(\bar{\mu} - \mu). \tag{4}$$

Each unit of energy efficiency, μ , costs $p_\mu - t_\mu$. Note that the tax on inefficiency is effectively a subsidy for efficiency. It shifts the budget constraints down but shifts the efficient ones down less because they have higher μ . Panel *b* of figure 3 depicts that inefficiency tax.

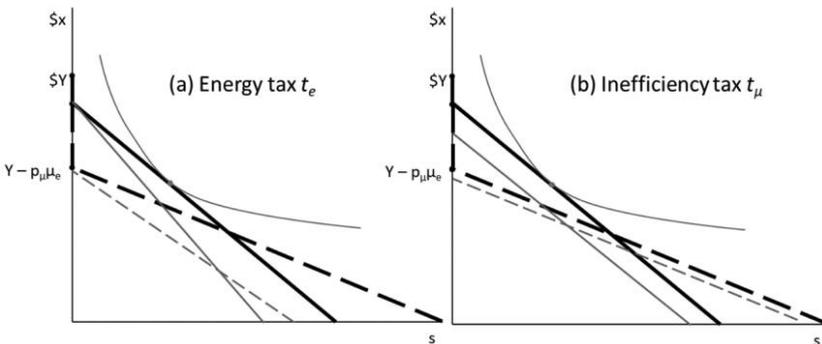


Figure 3. Energy and inefficiency taxes. Color version available as an online enhancement.

To see how taxes on energy and inefficiency would be borne by rich and poor households, turn back to the first-order conditions in equations (3). Cross multiplying, the two equations imply that $p_e e = p_\mu \mu$. Expenditures on energy equal expenditures on efficiency, and the ratio of e to μ is fixed at p_μ/p_e . Why? No matter the utility function in (1), energy services are a function of the product of e and μ . This is like having a subutility function over e and μ that is Cobb Douglas. However it decides between the two goods s and x , the household will divide expenditures on s evenly between its two components, e and μ .

What does that imply for the relative regressivity of energy and inefficiency taxes? Assume energy services s is a normal good. As income increases, demand for s increases, which means the product $e\mu$ increases. Since the ratio e/μ remains constant (from the first-order conditions in [3]), both e and μ increase proportionately. Richer households use more energy e and purchase more efficient appliances μ . Tax payments on energy $t_e e$ will increase with income, whereas tax payments on inefficiency $t_\mu(\bar{\mu} - \mu)$ will decrease with income. An energy tax will be less regressive than an efficiency standard.

As an aside, note that the energy tax is also more cost-effective than the inefficiency tax, meaning that it generates more energy savings for the same revenue and welfare change. Figure 4 depicts the scenario. An energy tax, with revenues refunded in lump

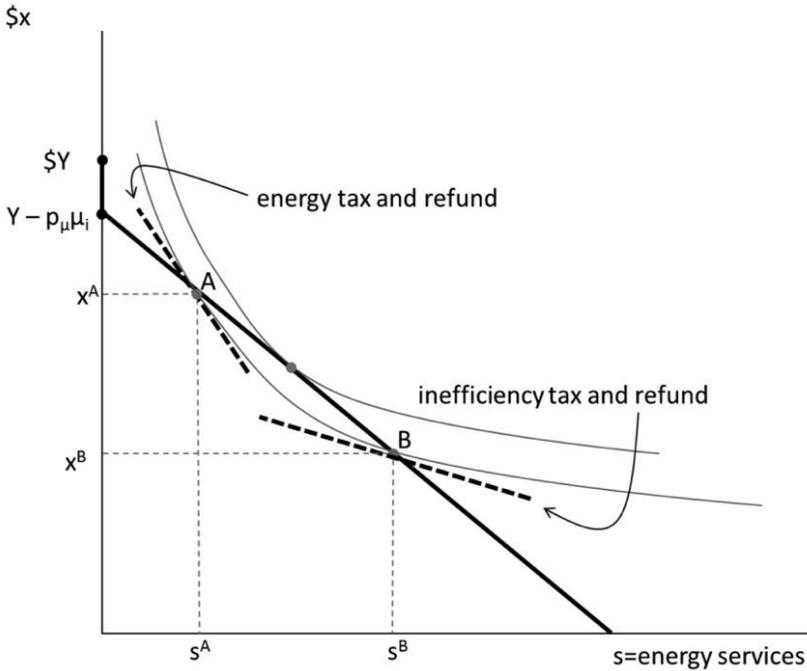


Figure 4. A revenue-equivalent efficiency subsidy saves less energy. Color version available as an online enhancement.

sums to balance the budget, leads to a steepening of the slope of the budget constraint and a reduction in energy services to s^A , at point A. An inefficiency tax with refunded revenues leads to a shallower budget constraint and increased services s^B , at point B.

Because the budget balances at points A and B, total expenditures will be equal under each scenario. We know from (3) that expenditures on energy and energy efficiency will be equal. Therefore

$$Y = x^A + 2(p_e + t_e)e^A = x^B + 2p_e e^B, \tag{5}$$

where in each case I have replaced expenditures on the sum of energy and energy efficiency with twice the expenditures on energy. Rewriting (5),

$$x^A - x^B = 2[p_e e^B - (p_e + t_e)e^A]. \tag{6}$$

We know that $x^A > x^B$ because preferences are convex. As a result, the term in square brackets on the right side of (6) is positive, or

$$p_e(e^B - e^A) > t_e e^A, \tag{7}$$

which means in turn that $e^B > e^A$ so long as $t_e > 0$. In other words, an energy tax reduces energy demand more than a revenue-equivalent subsidy for efficient appliances. Energy taxes are more cost-effective in addition to being less regressive.

One strong assumption so far has been that the cost of energy efficiency is linear: $p_\mu \mu$. That can easily be relaxed. Suppose the fixed cost of the energy-using durable is some general function of efficiency, $c(\mu)$. Then the second first-order condition in (3) becomes

$$(ii) \quad U_s / U_x = c'(\mu) / e. \tag{8}$$

Putting it together with (i) in equation (3) gives

$$p_e e = c'(\mu) \mu. \tag{9}$$

Differentiating with respect to income, Y ,

$$p_e \frac{\partial e}{\partial Y} = \frac{\partial \mu}{\partial Y} [c'(\mu) + c''(\mu) \mu]. \tag{10}$$

If $\partial e / \partial Y$ and $\partial \mu / \partial Y$ are both positive, then richer households consume more energy (e) and more efficiency (μ). A tax on energy t_e will be less regressive than a tax on inefficient appliances or cars t_μ . Are $\partial e / \partial Y$ and $\partial \mu / \partial Y$ both positive? They will be if the term in brackets in equation (10) is positive. The first term in the brackets, $c'(\mu)$, is positive by assumption. Efficient appliances cost more. So a sufficient condition for the whole term in brackets to be positive is that $c''(\mu) > 0$, which just means that the cost of energy efficiency is increasing in efficiency, or convex.⁹ That seems intuitive. The

9. Technically all that is needed is $c''(\mu) > -c'(\mu) / \mu$. The cost of efficiency can be concave but not too concave.

more energy efficient an air conditioner or vehicle, the costlier it is to make it even more efficient.

In sum, even a simple model with few assumptions predicts that an energy tax will be both more efficient and less regressive than an efficiency standard. All it requires is that energy services be normal, the cost of energy efficiency be convex (or not too concave), and the efficiency standard be framed as a revenue-equivalent tax on inefficient appliances or cars.

The benefit of this simplification is a clear demonstration of how fuel economy choices vary with income. Richer people spend more on fuel efficiency, face a lower cost of energy services, and use more energy to generate those services.

Although it is simple—simplistic even—the framework lends itself to empirical corroboration. It predicts that richer households purchase more efficient appliances and vehicles, and use more energy operating them. In the next two sections, I look for those empirical patterns in the data.

3. COMPARING ENERGY TAXES AND EFFICIENCY STANDARDS IN PRACTICE

The logic of section 2 applies to any energy-using durables: cars, air conditioners, light bulbs, washing machines, and even entire houses. In each case, section 2 predicts that richer households will purchase more energy efficiency *and* more energy. Documenting the first is relatively easy, but obtaining household data on energy use is difficult for any particular energy-using appliances. We know the efficiency of air conditioners and light bulbs, but not how much cooling or lighting each individual one provides, nor how much electricity each consumes. An exception is vehicles. The NHTS data contain information about households and their vehicles, including how much gasoline each uses annually. So I focus mainly on cars and light trucks, returning to other appliances and building codes in section 4.

For now assume that as in section 2 cars differ only in their prices and energy efficiency. In other words, ignore the fact that richer households' cars may use more gasoline per mile because they are bigger or faster. Later, in section 3.5, I will address those differences.

3.1. Fuel Economy and Fuel Use in the National Household Travel Survey

The 2009 NHTS contains demographic information for about 150,000 households and their vehicles. I drop the households with more than five vehicles or missing incomes, and I drop cars or light trucks with missing characteristics, such as size. That leaves 101,232 households that own 74,348 cars and 71,246 light trucks. I match those NHTS vehicles to data on their characteristics in CarQuery.¹⁰ Descriptive statistics about those households and their vehicles can be found in table 1.

10. <http://www.carqueryapi.com/>.

Table 1. Household and Vehicle Characteristics

Household Income (2009 \$)	Number Households (1)	Number of Vehicles (2)	Gallons Gasoline (3)	Miles Driven (4)	Gallons per Hundred Miles (gphm)		Footprint (sq. ft.) (7)	Car Age (Years) (8)
					Average (5)	Total (6)		
<\$10,000	4,763	.60	247	5,198	3.86	2.32	47.75	9.3
\$10,000– \$19,999	8,960	.90	371	7,857	3.84	3.48	47.85	8.3
\$20,000– \$29,999	10,362	1.14	507	10,703	3.88	4.44	48.29	7.5
\$30,000– \$39,999	10,110	1.26	613	12,931	3.92	4.97	48.66	7.0
\$40,000– \$49,999	9,719	1.33	679	14,366	3.94	5.27	48.78	6.7
\$50,000– \$59,999	9,030	1.41	754	16,005	3.96	5.61	48.96	6.6
\$60,000– \$69,999	7,592	1.47	816	17,310	3.97	5.88	49.08	6.4
\$70,000– \$79,999	7,535	1.52	869	18,413	3.98	6.07	49.11	6.2
\$80,000– \$99,999	10,302	1.56	921	19,548	4.00	6.28	49.29	6.0
>=\$100,000	22,845	1.64	991	20,972	4.02	6.59	49.23	5.6

Source. 2009 National Highway Transportation Survey. Includes all households, including those without vehicles.

Note. Column 2 reports only vehicles with data on mileage and fuel economy. Self-reported total vehicle ownership is about 50% higher. Column 6 is column 2 multiplied by column 5.

Gasoline taxes do exist in the United States, but they are not designed to reduce pollution and thus do not incorporate externalities such as the social cost of carbon. The federal gas tax has been \$0.184 per gallon since 1993, and state gas taxes currently vary from \$0.12 to \$0.58 per gallon. The tax revenues mostly fund road construction and maintenance, so they serve more as user fees than a deterrent to gasoline consumption or driving.

Fuel economy regulations in the United States take the form of CAFE standards. Each carmaker has to ensure that the sales-weighted average of the vehicles it sells in the United States exceeds a minimum threshold mpg to avoid paying fines. Carmakers can sell cars with lower fuel economy but only if they are offset by sales of more efficient cars so that the carmakers' overall averages meet the minimum. For cars, that minimum started at 18 mpg in 1978 and rose to 27.5 mpg by 1989, where it remained for the next

Table 2. Revenue per Household from Three Revenue-Equivalent Taxes

Household Income (2009 \$)	Tax Revenue per Household		
	\$.29 Gas Tax (1)	gpm Tax (2)	Footprint-Based gpm Tax (3)
<\$10,000	72	92	176
\$10,000–\$19,999	108	137	220
\$20,000–\$29,999	147	175	234
\$30,000–\$39,999	178	196	224
\$40,000–\$49,999	197	208	215
\$50,000–\$59,999	219	221	201
\$60,000–\$69,999	237	232	184
\$70,000–\$79,999	252	239	203
\$80,000–\$99,999	267	248	171
≥\$100,000	287	260	238

Note. gpm = gallons per hundred miles.

two decades before rising again to 37.8 mpg in 2016. For light trucks like pickups and SUVs, the standards started at 17 mpg in 1979 and rose to 20 mpg by 1984. They stayed roughly constant until 2005, when they began rising to 28.8 mpg in 2016.¹¹

The CAFE regulations are a costly way to reduce fuel consumption, but distributional concerns are raised about increasing the gas tax. To compare the distributional consequences of the two policies, we can think of the fuel economy standard as a tax on the gpm of the vehicles owned by each household and see how its static incidence stacks up against that of the gas tax.¹²

Consider again the \$0.29 per gallon carbon tax. Its static incidence ranges from \$72 per year for the poorest households to \$287 for the richest. Those figures are listed in column 1 of table 2 and plotted as the left-most darkly shaded columns in figure 5.

To make the comparison fair, the equivalent gpm tax has to raise the same revenue as the gas tax. That works out to \$0.39 per gpm per year, or \$39 per gallon per hundred

11. The 2016 standards are averages, based on projected sales. See NHTSA, “NHTSA and EPA Establish New National Program to Improve Fuel Economy and Reduce Greenhouse Gas Emissions for Passenger Cars and Light Trucks,” http://www.nhtsa.gov/staticfiles/rulemaking/pdf/cale/CAFE-GHG_Fact_Sheet.pdf.

12. The model in section 2 subsidizes mpg, or μ . Here I tax the inverse, gpm, or $1/\mu$, because the actual US CAFE regulations target gpm rather than mpg. The US CAFE standards are expressed as miles per gallon (mpg) rather than gallons per mile, so the actual standard is a harmonic average. That complicates things needlessly. Since gpm is the inverse of mpg, the Department of Transportation takes an average of the inverse of gpm and then inverts the average (Fischer 2009). See Fullerton and West (2012) for another model of a tax on inefficient vehicles.

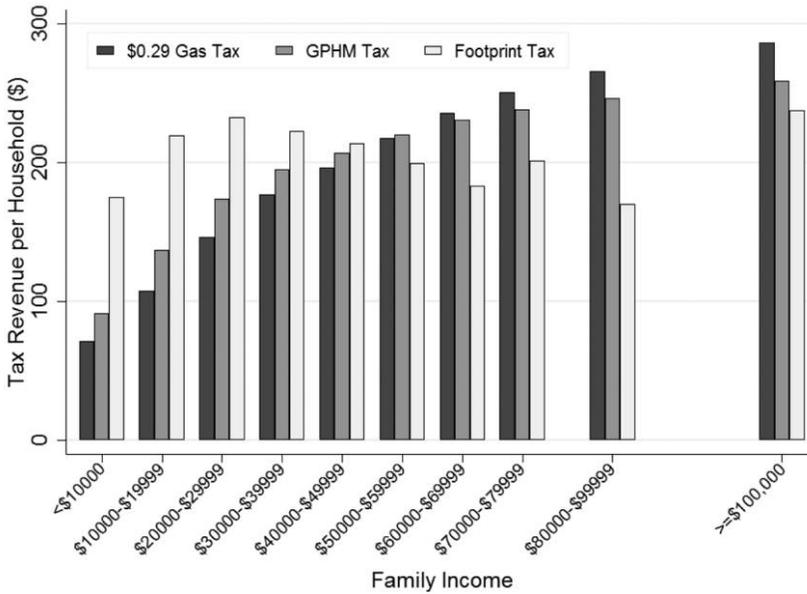


Figure 5. Revenue-equivalent gas taxes and fuel economy standards

miles (gphm) to make the units simpler. The resulting static incidence of that tax is listed in column 2 of table 2, and plotted as the middle gray columns in figure 5.¹³

Recall that it makes no sense to call either the gas tax or the gphm tax regressive or progressive, without considering what happens to the tax revenues. If the revenues are spent on government services for rich households, or wasted entirely, both taxes would be regressive. Tax payments decline in proportion to income. But if the revenues are spent on services for poor households, or refunded as an equal payment for every household, both taxes would be progressive. Poor households would pay less in taxes than they receive, and rich houses would pay more. So instead of asking whether the gas tax or the gphm tax is regressive or progressive, I ask which is more regressive than the other, or equivalently less progressive.

Like the gas tax, a gphm tax that raises the same revenue would increase with income, but more slowly than income. That can be seen in table 2. The gphm tax imposes higher costs than the gas tax on low-income households, and lower costs on high-income households. So a gphm tax would be more regressive (or less progressive) than a gas tax.

13. Using the households in table 1, I first calculate the revenue that would be raised by the \$0.29 gas tax. Then I divide that total revenue by the sum of all the gphm for all the vehicles, yielding the dollar tax per gphm that would raise the same revenue, again assuming no change in behavior. The average vehicle in the sample uses 3.98 gphm (25 mpg). A \$39 per gphm tax would cost \$155 per year for that vehicle.

The explanation for why a gphm tax would be more regressive than a gas tax can be seen back in table 1. Column 5 shows that cars owned by richer households tend to use more gas per mile driven. Column 2 shows that richer households own more cars, and fuel economy standards impose a cost on each car owned. The total gphm per household across all cars is listed in column 6 of table 1. That is the product of columns 2 and 5 and is the tax base for a gphm tax. The richest households have more than 10 times the income as poor households but would pay less than three times as much in a gphm tax.

Which is more regressive, the gas tax or the gphm tax? In this static framework, the answer depends on whether gasoline or gphm increases more with income. Column 3 of table 1 shows that rich households use four times more gasoline than poor households. Column 6 shows that rich households' cars use three times the gas per mile (in total, accounting for the fact that they have more cars). So a gas tax would increase faster with income than a gphm tax and would therefore be less regressive. Because the gphm tax is equivalent at the margin to the CAFE standards, in the United States a gas tax would be less regressive than existing fuel economy standards.

Of course, the CAFE standard is not exactly the same as a gphm tax. As currently configured, the CAFE standard effectively taxes inefficient cars and subsidizes efficient ones, while a gas tax would tax all drivers. A gas tax that refunded its revenue in some way would be fiscally equivalent to CAFE and less regressive (or more progressive).

To see that, consider a gas tax and a gphm tax that both refund their revenues in a lump sum. The resulting figure would look exactly like figure 5, but each column would be shifted down by \$215, the amount of the refund check. The poorest households would end up ahead \$141 under the gas tax, but only \$121 under the gphm tax. The richest households would end up behind \$74 under the gas tax, but only \$47 under the gphm tax. Both policies would be progressive, but a gas tax would be more progressive than a fuel economy standard. No matter how we compare the two, the gas tax is less regressive or more progressive than the gphm tax, so long as they are compared on a revenue-equivalent basis.

That is not the end of the story, however, because in 2011 the US DOT made an important change to the CAFE regulations with significant distributional implications: adjusting for the size, or "footprint," of vehicles.

3.2. Footprint-Based CAFE Standards

Before 2011, the fuel economy targets were simple averages. Each manufacturer could meet those targets by selling more small cars, which tend to consume less gasoline per mile, and fewer large cars, which consume more. Starting in 2011, the CAFE regulations shifted to footprint-based standards. A vehicle's footprint is the area under its four tires, measured in square feet.

The new footprint-based standards give larger cars a more lenient gphm target. Figure 6 illustrates the formula. A carmaker that sells only small cars in 2016, with footprints of 41 square feet or less, must meet a target of 2.43 gphm (41.09 mpg).

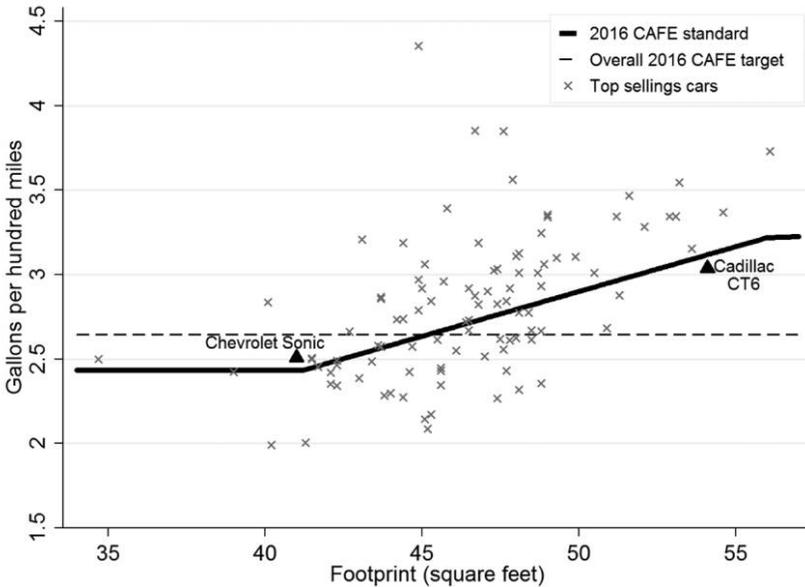


Figure 6. Footprint-based standard for cars

A carmaker selling large cars with footprints over 56 square feet needs only to achieve 3.23 gphm (30.96 mpg). Cars between those two extremes face a sliding scale.

Ito and Sallee (2017) suggest that the switch to footprint-based CAFE standards might be designed to redistribute welfare across consumers. As can be seen in figure 6, that depends on the shape of the footprint formula and who buys the larger cars. The dotted line in figure 6 plots the projected average fuel economy for cars (not light trucks) in 2016, 2.65 gphm (37.8 mpg).¹⁴ Under the simple CAFE rule that prevailed before 2011, carmakers could sell a mix of vehicles, illustrated by the top 100 bestselling cars on the graph. The weighted average of sales would have to equal or fall below 2.65 gphm. As an example, General Motors could sell more Cadillacs, priced above \$50,000, but they would have to be offset by sales of cars like Chevrolet Sonics, priced under \$20,000. Expensive Cadillacs would face an implicit tax (λR in app. 1), and cheaper Sonics an implicit subsidy (λ).

The solid line in figure 6 plots the 2016 footprint-based gphm target. Large Cadillacs meet the footprint-based standard, while smaller Sonics do not. GM can sell more non-

14. The average is 32.7 mpg for combined cars and light trucks (US EPA and NHTSA 2010, table IV.E.2-1).

compliant Sonics, but they must be offset by sales of Cadillacs. The inexpensive Sonics face the implicit tax, while the expensive Cadillacs receive the subsidy.¹⁵

The example in figure 6 illustrates the potential for the footprint-based standard to alter the regressivity of the CAFE rules. If poor people disproportionately purchase cars that are relatively inefficient for their size, like Sonics, the switch to footprint-based standards has exacerbated the regressivity of the fuel economy standards. To see whether the example in figure 6 applies in general, and whether the switch to footprint-based standards has in fact made the CAFE rules more regressive, turn back to table 1.

Table 1 lists the average vehicle footprint by household income in column 7.¹⁶ Richer households tend to own larger cars. Column 5 shows that those cars use more gas per mile.¹⁷ Under the old CAFE standard, all that mattered for fuel economy was column 5. Under the new footprint-based standard, what matters is column 5 relative to column 7.

Another way to see that the Sonic-Cadillac example in figure 6 represents a general result—and that the 2011 switch to footprint-based standards was therefore regressive—is in figure 7. The typical car driven by a household in the poorest income group is both smaller and more fuel efficient than the typical car driven by richer households. The switch to footprint-based standards reduced the advantage lower-income households had from driving smaller cars.

The next task is to measure how this switch affects the regressivity of the CAFE rules by translating the footprint-based rule into a tax, so it can be compared on a revenue-equivalent basis to the gas tax and gphm tax in figure 5. I start with the actual formula used by the CAFE rules for model years 2012–16.

$$\begin{aligned} \text{Cars:} \quad \quad \quad & \text{Target gphm} = A + \text{footprint} \times 0.05308 \\ \text{Light trucks:} & \text{Target gphm} = B + \text{footprint} \times 0.04546 \end{aligned} \tag{11}$$

15. In Levinson (2017) I provide one potential motivation for the switch to footprint-based standards: protectionism. Because imported cars are on average smaller, the switch imposed a cost on imports equivalent to a tariff of \$50 to \$200 per vehicle.

16. The CarQuery data contain cars' wheelbase (length from front to rear tires) but not track width (distance between left and right tires). Footprint is track width times wheelbase. I approximate track width based on regression estimates from the top 100 best-selling cars and trucks. The estimates are track width (in millimeters) = $911.2 + 0.365 \times \text{width}$ for cars, and $799.1 + 0.502 \times \text{width}$ for light trucks.

17. In the theoretical model in section 2, rich households purchase more energy efficiency, but here rich households purchase less efficient cars. The difference is that in section 2, energy-using durables differ only in their efficiency and price. In practice, cars differ along multiple dimensions correlated with both price and efficiency. Richer families purchase larger, faster-accelerating cars. I control for those other vehicle features at the end of section 3.

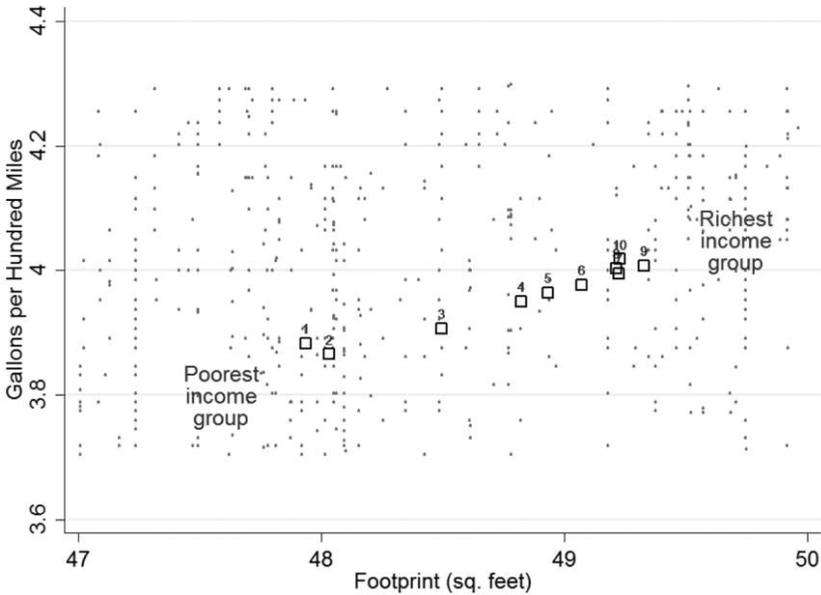


Figure 7. Fuel economy by car size, by household income. Color version available as an online enhancement.

For each square foot in a vehicle’s footprint, required fuel consumption is allowed to increase by 0.053 gphm for cars and 0.045 for light trucks. That is the slope of the solid line in figure 6. The constants, *A* and *B*, decrease each year as the regulation tightens, from 0.5842 in 2012 to 0.2406 in 2016 for cars, and from 1.46 to 1.0413 for light trucks.¹⁸

To calculate the tax base for each vehicle, I use the difference between its fuel economy and the target for a vehicle with its footprint. I then sum that number across all vehicles in the 2009 NHTS to find the economy-wide tax base. The units of this tax base are gphm in excess of the target, where larger cars and light trucks have higher gphm targets. To ensure the gphm tax raises the same revenue as the \$0.29 gas tax, I divide the gas tax revenues calculated earlier by the economy-wide footprint-based gphm tax base. That division yields a per-unit gphm tax rate that I apply to each household.

Tax revenues by income class for this footprint-based gphm tax are in column 3 of table 2. The poorest households pay more than under either the gas tax or the straight gphm tax, and the richest households pay less. To show the distinction visually, figure 5 plots those revenues as the third and lightest-shaded set of columns. The footprint-based CAFE standards are even more regressive than the flat pre-2011 CAFE standards, and both are more regressive than a gas tax would be.

18. US EPA and NHTSA 2010, 25612–15.

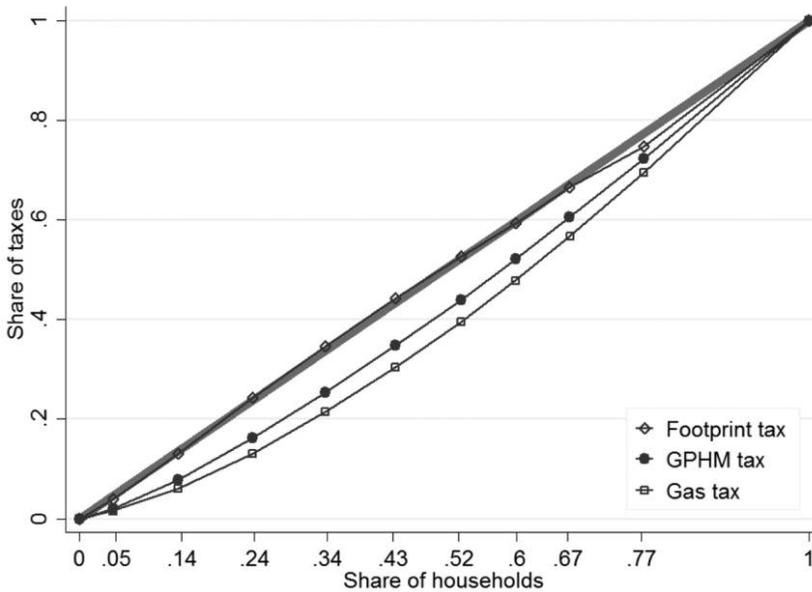


Figure 8. Tax regressivity. Color version available as an online enhancement.

3.3. Concentration Curves Comparing Energy Taxes and Efficiency Standards

For an alternative visual representation of the distributional consequences, see figure 8 where I plot concentration curves for each of the taxes under consideration.¹⁹ The bottom axis plots the share of all households, from 0 to 1, with tick marks denoting fractions in each of the NHTS income categories. The left axis plots the share of taxes collected. The 45-degree line represents a per-household tax, where each household pays the same amount regardless of income—the household equivalent of a head tax.

The bottom line (open square markers), denotes the tax shares paid by each household under the \$0.29 gas tax from table 2 and figure 5. It is bowed downward because poor households pay less in gas taxes than rich households. It still denotes a regressive tax because a proportional income tax would be even more bowed below the 45 degree line.

The second line (solid circles) denotes the gphm tax from column 2 of table 2. It is closer to the 45-degree line, thus more regressive. The third line (open diamonds) depicts the footprint-based gphm tax. It essentially mirrors the 45-degree line. The footprint-based gphm tax is as regressive as a household-level head tax.

19. Concentration curves are similar to Lorenz curves, in that the population on the horizontal axis is sorted by income, but differ in that some other variable of interest is plotted on the vertical axis—tax expenditures in this case. Concentration curves can thus lie above the 45 degree line. See Borenstein (2017) for another recent application in this journal.

3.4. Vehicle Age

As Jacobsen (2013) points out, CAFE standards apply to new cars, and poor households tend to buy used cars. As a result, poor households do bear some of the burden of the CAFE standards but only after those new cars trickle into the used car market—and only to a depreciated extent relative to rich households that purchase new cars. So it is worth considering the distributional consequences of a gphm tax that is disproportionately borne by owners of newer vehicles, even though a true gphm tax could in principle be levied on all cars, new and used.

I begin by assuming that the burden of the CAFE standard declines proportionally to the price of used cars, from the National Automobile Dealers Association. I regress the used car price on age and age squared and calculate weights equal to the ratio of predicted prices for cars of each vintage to the price of new cars. The weights vary from 1.0 for new vehicles down to 0.49 for 21-year-old vehicles.²⁰ I then scale the gphm tax so that total tax revenues equal those collected from a \$0.29 gas tax, but where owners of older cars pay less than owners of newer cars, proportional to the weights from the vehicle age regression. For comparison, I also calculate a footprint-based version of that depreciation-adjusted gphm tax, again ensuring equal total revenues.

The results are plotted in figure 9. For reference, the gas tax from figure 8 is plotted as the dashed line with open square markers. The depreciation-adjusted gphm tax lies almost on top of the gas tax, suggesting that even if we account for the fact that CAFE standards fall disproportionately on owners of new cars, the CAFE standards are just as regressive as the gas tax. Moreover, the footprint-based version (denoted by open diamonds) is more regressive than the gas tax, appearing about halfway between the gas tax and a household-level head tax represented by the 45-degree line.

On a revenue-equivalent basis, a gasoline tax would be regressive, fuel economy standards more regressive, and the new footprint-based fuel economy standards that the United States switched to in 2011 the most regressive of the three. Even if we skew the fuel economy standards toward new cars, the gas tax is no more regressive than the older standards and is less regressive than footprint-based standards. Economic theory tells us that the gas tax would be the most cost-effective policy. This analysis suggests that it would also be the least regressive.

3.5. Automobile Characteristics

The model in section 2 predicted that richer households would purchase more energy-efficient appliances and vehicles. Table 1 shows that rich households drive cars that are less fuel efficient, not more. The apparent contradiction hinges on the fact that in the theoretical model in section 2, appliances or vehicles differ only in their energy efficiency (μ).

20. I also tried assuming the gpm tax declines linearly with the age of the car, from 100% down to zero for the oldest cars in the sample. That straight-line depreciation yields nearly identical results to the price-based depreciation.

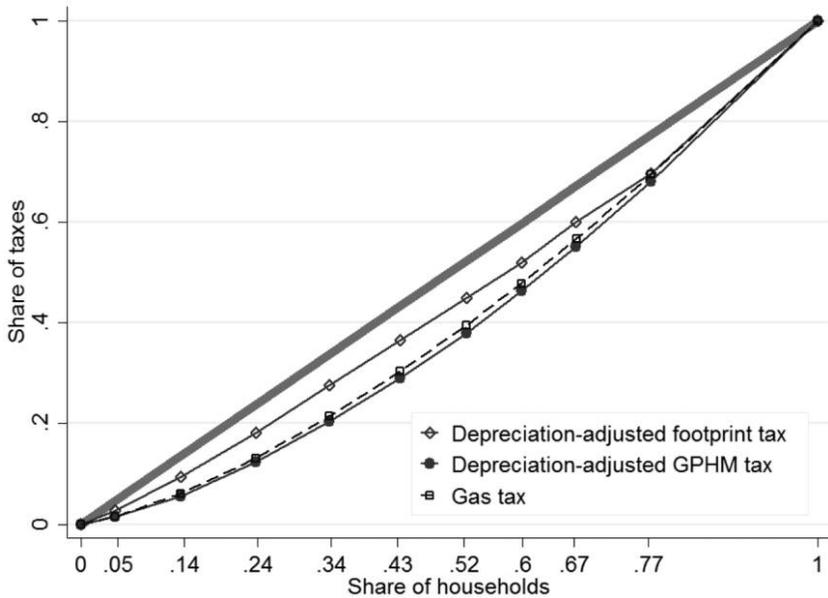


Figure 9. Depreciation-adjusted tax regressivity. Color version available as an online enhancement.

In practice, however, rich households’ cars are larger, heavier, accelerate faster, and differ from poor households’ cars in many other ways that sacrifice fuel economy.

To control for those other vehicle characteristics correlated with income and fuel economy, I take two approaches. First, I control for all other vehicle characteristics by examining models that come in two versions: an ordinary gasoline-powered version and a hybrid battery-powered version. There are 5,211 such vehicles in the 2009 NHTS, in eight different model pairs. They are listed in table 3. If the hybrid and gasoline versions of the

Table 3. Vehicles with Hybrid and Gasoline Versions

Make and Model	Proportion Hybrid in NHTS	Number
Chevrolet Blazer/Tahoe	.462	91
Ford Escape	.155	696
Honda Civic/Crx/Del Sol	.167	2,058
Mercury Mariner	.286	42
Nissan Altima	.075	226
Saturn Vue	.105	124
Toyota Camry	.183	1,408
Toyota Highlander	.276	566
Totals	.182	5,211

Table 4. Hybrid and Gasoline Characteristics

	Nonhybrids (1)	Hybrids (2)
Gallons per hundred miles	3.02 (.59)	2.26 (.52)
MPG equivalent	33.1	44.2
Annual miles driven	12,664 (7,462)	14,178 (7,915)
Car price	19,301 (4,469)	20,158 (5,279)
Gasoline price	3.07 (.14)	3.09 (.19)
Household income	76,098 (35,008)	90,746 (32,987)
Number of observations	4,262	949

Source. Car characteristics from CarQuery, car prices from National Automobile Dealers Association, and household characteristics from the 2009 National Highway Transportation Survey.

Note. All five differences are statistically significant at 5%.

vehicles are otherwise identical, then we can compare the two without concern about other car characteristics that may affect fuel economy.

Table 4 describes the differences between the hybrid and nonhybrid cars and the households that own them. As predicted by the theory in section 2, the hybrids are owned by households with higher incomes. Consistent with the model, they also get better gas mileage, are driven further each year, and cost more.

Table 5 takes a different approach, using all the vehicles in the NHTS that could be matched with CarQuery data on vehicle characteristics. I regress fuel economy (in gallons per thousand miles to make the units sensible) on vehicle characteristics and household income. Column 2 regresses fuel economy on just the log of household income, taken as the midpoint of the categorical variables in the NHTS. The coefficient on income (6.7) means that richer households' cars consume more gasoline per mile, as noted in table 1 and in contradiction to the theoretical prediction. Column 3 of table 5 regresses fuel economy on four vehicle characteristics: engine power, length, width, and height. Each adds statistically significantly to gallons of gas consumed per mile.

Column 4 of table 5 includes income and vehicle characteristics together. The coefficient on the log of household income reverses, suggesting that households with higher incomes own cars that use slightly less gasoline per mile.²¹

21. A version using income categories rather than the log of income yields the same result.

Table 5. Fuel Economy, Controlling for Car Characteristics

	Means (1)	Income Only (2)	Car Only (3)	Car and Income (4)
ln(income)	10.97* (.002)	6.66* (.36)		-3.869* (.223)
Engine power (metric horsepower)	167.6* (.16)		.363* (.004)	.371* (.004)
Wheelbase (length in mm)	2,766.7* (.62)		.063* (.001)	.062* (.001)
Width (mm)	1,830* (.32)		.054* (.002)	.051* (.002)
Height (mm)	1,590* (.53)		.225* (.001)	.227* (.001)
Constant		325.0* (3.9)	-292.6* (3.0)	-247.4* (4.0)
Observations	130,007	130,007	130,007	130,007
R-squared		.003	.621	.622

Source. Car characteristics from CarQuery and household characteristics from the 2009 National Highway Transportation Survey.

Note. Dependent variable: gallons per thousand miles.

* $p < .05$.

So although the conclusion that richer households drive more efficient cars (*ceteris paribus*) supports the results from section 2, the difference in efficiency does not by itself explain the fact that a fuel economy standard would be so much more regressive than a gasoline tax. Rather, the difference in regressivity between taxes and standards stems from car ownership increasing less steeply with income than miles driven. In the theory in section 2, there is no reason for households to purchase more than one car or appliance. Car or appliance ownership does not increase with income, and a tax on inefficiency becomes a regressive tax on poorer households that purchase inefficient cars and appliances. The next section considers a few extensions, including other energy-using durables such as appliances and houses.

4. EXTENSIONS: OTHER ENERGY-USING DURABLES;

POLLUTION-EQUIVALENT TAX RATES; LIFETIME INCIDENCE

The simple model in section 2 and the static measure of incidence in section 3 leave open a number of questions.

4.1. Other Energy-Using Durables

The model in section 2 shows that richer households will own more energy-efficient appliances and homes and use more energy to run them. Demonstrating the first part of that

Table 6. Home Energy Efficiency and Household Income

Energy Efficiency Feature	Average Household Income (2010 US\$)		
	No (1)	Yes (2)	Difference (3)
Double or triple-paned windows	\$45,063 (32,859) [5,071]	\$61,964 (37,300) [7,012]	\$16,901 (641)
Compact fluorescent bulbs installed	52,783 (35,185) [853]	59,657 (36,643) [6,567]	6,874 (1,287)
Energy Star clothes washer	52,176 (35,093) [2,197]	65,955 (36,310) [4,476]	13,779 (925)
Frontloading washer	55,471 (35,324) [8,092]	76,189 (36,218) [1,941]	20,718 (911)
Energy Star dishwasher	57,920 (35,088) [1,642]	75,121 (35,859) [3,042]	17,201 (1,083)
Energy Star fridge	46,497 (33,692) [2,981]	63,198 (37,178) [4,645]	16,701 (824)
Energy Star air conditioner	40,904 (32,088) [679]	53,332 (36,335) [1,176]	12,428 (1,625)

Source. 2009 Residential Energy Consumption Survey.

Note. Average (SD) [no. obs.].

prediction is straightforward. Table 6 reports the average incomes of households surveyed in the 2009 Residential Energy Consumption Survey (RECS).²² Each row lists a different household energy-efficiency characteristic, and the average incomes and numbers of houses with and without those characteristics. For example, the average incomes for households reporting that their homes had double or triple-paned windows was 30% higher than the incomes for homes without that energy-saving feature. A mandate requiring all homes to have double-paned windows, or a tax on homes that do not, would be borne disproportionately by low-income households. For every one of the seven energy efficiency features listed in table 6, households with the feature have significantly higher incomes. As predicted, income and energy efficiency are positively associated.

22. The RECS is collected by the US Energy Information Administration. See www.eia.gov/consumption/residential. Davis (2012) uses the RECS to show that homeowners have more energy-efficient appliances than renters.

The second part of the theory—that richer households will use more energy in those appliances and homes—is more difficult to assess. Although we can tell from table 6 that an energy efficiency mandate or tax would be regressive, directly comparing that to an energy tax is not possible without data on appliance-specific energy use. No data tell us how many loads of laundry or dishes (s in the model) were washed by the Energy Star appliances, or how much light was cast by the compact fluorescent bulbs.

4.2. Pollution-Equivalent Tax Rates

Readers of early drafts of this paper have asked whether it would be more appropriate to compare policies that abate the same amount of pollution, rather than to compare policies that raise the same amount of revenue. That is correct, in theory. In practice, however, that calculation for vehicles requires knowing the elasticity of demand for gasoline by income group, the elasticity of demand for fuel efficiency in vehicles by income group, and, to be complete, the elasticity of supply of fuel efficiency in vehicles by carmakers.²³

Incorporating those elasticities could work in either direction. Low-income households might be more responsive to either a gas tax or a gphm tax, if they are more willing to seek substitutes such as public transit or ride sharing. That would lessen the burden of either tax on low-income households. Or high-income households that own multiple vehicles may more easily reduce their gas tax liability by spending less time driving their least efficient cars. That would lessen the burden of a gas tax on high-income households.

Those calculations are beyond the scope of this paper. Jacobsen (2013) comes closest, comparing the welfare loss per dollar of gasoline abated for a gas tax and a fuel economy standard. His model could compare the welfare cost across income groups of a gas tax and an efficiency standard that reduce gas consumption by the same total amount. That approach would have one drawback. It would not be obvious how much of the different incidences across income groups would result from different elasticities—of demand for gasoline, demand for fuel efficient cars, and supply of fuel efficient cars—versus different ex ante gas consumption and vehicle choice by households with different incomes. An advantage of the approach here is transparency. All of the different incidences calculated result only from those ex ante choices.

4.3. Lifetime Incidence

Annual incomes are distributed less equally than lifetime incomes, so policy will seem more regressive relative to annual income than to lifetime income. It is not obvious, however, whether the relative regressivity of the two taxes would be different for annual and lifetime income.

23. In fact, carmakers' responses can be strategic in complicated ways. Sallee and Slemrod (2012) show that carmakers respond to the fact that the US gas guzzler tax is a step function in fuel economy by bunching the cars they sell with fuel economy just below mpg thresholds in the tax rate. Killeen (2017) shows that carmakers responded to the footprint-based CAFE rules by expanding the footprints of their cars, holding other vehicle dimensions constant.

The one-period model in section 2 does not distinguish between lifetime and annual income. You can think of Y as lifetime income and $p_e e$ is the present discounted sum of lifetime energy expenditures. Or you can think of Y as annual income and $p_\mu \mu$ as the amortized annual cost of the appliance or car. Either way yields the same result: a tax on inefficient appliances will be more regressive than a tax on energy.

Empirically, there is plenty of evidence regarding energy taxes. Poterba (1991) compares gasoline expenditures as a share of annual income and as a share of annual expenditures, a proxy for lifetime income. Gasoline spending falls as a share of annual income but is nearly proportional as a share of expenditures. So a gas tax looks a lot less regressive relative to lifetime income than annual income. Hassett et al. (2009) take that one step further and use lifetime consumption as a proxy for lifetime income. They show that a carbon tax would be regressive under either measure of income, annual or lifetime, but that it would be less regressive under the lifetime measure. Grainger and Kolstad (2010) also examine a carbon tax. Like Poterba they use annual expenditures as a proxy for lifetime income, and they also use equivalence scales to adjust for family size. Under any of their measures, the static incidence of a carbon tax would be proportionally lower for higher-income households.

These results all demonstrate that an energy or carbon tax would be less regressive when compared to annual income than to lifetime income. I know of no evidence conducting the same analysis for efficiency standards, but I can speculate. Households with higher annual incomes purchase more energy-efficient appliances and cars. So it seems likely that households with higher lifetime incomes also purchase more energy-efficient appliances and cars. If that is the case, then a tax on inefficient appliances or cars would be borne disproportionately by lower-income households, and would be regressive in the same way that an energy or carbon tax would be regressive.

5. DISCUSSION

Pigou demonstrated in 1920 that a pollution tax would be the most cost-effective means of reducing pollution. That attribute—cost-effectiveness—should appeal to environmentalists worried about climate change. It should also appeal to those concerned about the costs associated with mitigating climate change. But in the nearly 100 years since Pigou, there have been almost no examples of pollution taxes.

Arguments against using taxes to combat climate change often invoke their distributional consequences. In the United States, we have mostly relied on energy efficiency standards rather than energy taxes, without ever making the direct distributional comparison. We know energy or pollution taxes are cost-effective and regressive, and we know energy efficiency standards are less cost-effective. Are efficiency standards also less regressive? If they are, some might view that efficiency-equity trade-off as worthwhile.

To answer this question, I demonstrate three ideas. First, it is easy to write down a simple, static, two-good, representative-household model in which an energy tax is less regressive than an energy efficiency standard. The only underlying assumptions are that

energy services are normal, the efficiency standard is representable by a tax on inefficient vehicles or appliances, and the cost of efficiency is convex, or at least not too concave. As a consequence, we should not be surprised to learn that empirically, energy taxes are both more cost-effective and more progressive than efficiency standards. There is no efficiency-equity trade-off.

Second, in the particular case of the US automobile market, a gasoline tax would be less regressive than a revenue-equivalent fuel economy standard. The key is the phrase “revenue equivalent.” Current CAFE standards do not raise government revenue and are not directly comparable to a revenue-raising gasoline tax. To evaluate the two, I compare a gasoline tax to a tax on inefficient vehicles that raises the same revenue. The gas tax is less regressive. A plausible counterargument might claim that redistribution is politically feasible through fuel economy standards, where the transfers are hidden in the prices of efficient and inefficient cars sold to rich and poor households. That argument lies in the realm of politics, outside of economics.

The relative regressivity of US CAFE standards was exacerbated by the 2011 change to footprint-based targets. The change reduced the advantage poorer households had as a consequence of their smaller cars. So not only would a gasoline tax be less regressive than a revenue-equivalent fuel economy regulation, both would be less regressive than the footprint-based fuel economy regulations promulgated since 2011.

The third idea I demonstrate is that the simple model outlined in section 2, in which an energy tax would be less regressive than an energy-efficiency standard, is supported across a wide array of home appliances and home building construction. A regulation that targets inefficient appliances and homes will cost rich households less because they are already purchasing the efficient versions of those goods.

In sum, it is true that an energy tax would be regressive, in a static model like figure 1, if the tax would be borne by income groups in proportion to their pre-tax consumption levels, and if we ignore what happens to the revenue. But this work shows that a revenue-equivalent energy efficiency standard would be even more regressive. So the fact that carbon taxes are regressive does not stand up as a coherent argument for supporting energy efficiency standards in their place.

REFERENCES

- Allcott, Hunt. 2013. The welfare effects of misperceived product costs: Data and calibrations from the automobile market. *American Economic Journal: Economic Policy* 5 (3): 30–66.
- Allcott, Hunt, and Michael Greenstone. 2012. Is there an energy efficiency gap? *Journal of Economic Perspectives* 26 (1): 3–28.
- Allcott, Hunt, and Sendhil Mullainathan. 2010. Behavior and energy policy. *Science* 327 (5970, March 5): 1204–5.
- Anderson, Soren, and James Sallee. 2016. Designing policies to make cars greener. Working paper no. 22242, National Bureau of Economic Research, Cambridge, MA.
- Austin, David, and Terry Dinan. 2005. Clearing the air: The costs and consequences of higher CAFE standards and increased gasoline taxes. *Journal of Environmental Economics and Management* 50:562–82.

- Borenstein, Severin. 2017. Private net benefits of residential solar PV: The role of electricity tariffs, tax incentives, and rebates. *Journal of the Association of Environmental and Resource Economists* 4 (S1): S85–S122.
- Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer. 2013. Are consumers myopic? Evidence from new and used car purchases. *American Economic Review* 103 (1): 220–56.
- Davis, Lucas W. 2012. Are renters less likely to have energy efficient appliances? In *Design and implementation of US climate policy*, ed. Don Fullerton and Catherine Wolfram. Chicago: University of Chicago Press.
- Davis, Lucas, and Christopher Knittel. 2019. Are fuel economy standards regressive? *Journal of the Association of Environmental and Resource Economists* 6 (S1): S37–S63.
- Fischer, Carolyn. 2009. Let's turn CAFE regulation on its head. Issue brief no. 09-06, Resources for the Future, Washington, DC.
- Fried, Stephie, Kevin Novan, and William Peterman. 2017. The distributional effects of a carbon tax on current and future generations. Finance and Economics Discussion Series Working paper no. 2016-038 Federal Reserve Board, Washington, DC.
- Fullerton, Don, and Sarah West. 2010. Tax and subsidy combinations for the control of car pollution. *B.E. Journal of Economic Analysis and Policy* 10 (1).
- Gillingham, Kenneth. 2013. The economics of fuel economy standards versus feebates. Working paper, National Energy Policy Institute, Tulsa, OK.
- Grainger, Corbett A., and Charles D. Kolstad. 2010. Who pays a price on carbon? *Environmental and Resource Economics* 46:359–76.
- Greene, David L. 2010. How consumers value fuel economy: A literature review. EPA-420-R-10-008, US Environmental Protection Agency, Washington DC.
- Hassett, Kevin, Aparna Mathur, and Gilbert Metcalf. 2009. The incidence of a U.S. carbon tax: A lifetime and regional analysis. *Energy Journal* 30 (2): 157–79.
- Ito, Koichiro, and James Sallee. 2017. The economics of attribute-based regulation: Theory and evidence from fuel-economy standards. *Review of Economics and Statistics* 100 (2): 319–36.
- Jacobsen, Grant Douglas. 2018. Examining the incidence of energy efficiency incentives. Working paper, University of Oregon.
- Jacobsen, Mark R. 2013. Evaluating US fuel economy standards in a model with producer and household heterogeneity. *American Economic Journal: Economic Policy* 5 (2): 148–87.
- Killeen, Grady. 2017. Attribute-based regulations: The case of corporate average fuel economy (CAFE) standards. Working paper Harvard Business School.
- Kleit, Andrew. 1988. The impact of automobile fuel economy standards. Working paper no. 160, Federal Trade Commission Bureau of Economics, Washington, DC.
- Kwoka, John. 1983. The limits of market-oriented regulatory techniques: The case of automotive fuel economy. *Quarterly Journal of Economics* 98 (4): 695–704.
- Levinson, Arik. 2017. Environmental protectionism: The case of CAFE. *Economics Letters* 160:20–23.
- Newell, Richard G., and Juha V. Siikamäki. 2015. Individual time preferences and energy efficiency. Working paper no. 20969, National Bureau of Economic Research, Cambridge, MA.
- Poterba, James. 1991. Is the gasoline tax regressive? *Tax Policy and the Economy* 5:145–64.
- Pigou, Arthur C. 1920. *The economics of welfare*. London: Macmillan.
- Sallee, James, and Joel Slemrod. 2012. Car notches: Strategic automaker responses to fuel economy policy. *Journal of Public Economics* 96 (11–12): 981–99.
- US EPA and NHTSA (Environmental Protection Agency and National Highway Traffic Safety Administration). 2010. Light-duty vehicle greenhouse gas emission standards and Corporate Average Fuel Economy standards; final rule. *Federal Register* 75 (88): 25324–728.